Title

Cosmological Implications of Axion-Matter Couplings and Light Fields during Inflation from BOSS and Future Galaxy Surveys

Permalink

https://escholarship.org/uc/item/2qm1q4db

Author

Guo, Yi

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Cosmological Implications of Axion-Matter Couplings and Light Fields during Inflation from BOSS and Future Galaxy Surveys

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

in

Physics

by

Yi Guo

Committee in charge:

Professor Daniel Green, Chair
Professor George Fuller
Professor Julius Kuti
Professor Jeffrey Rabin

2024

The Dissertation of Yi Guo is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

I dedicate this thesis to Xinwen Hu, whose unwavering company and encouragement illuminated the darker days of my Ph.D. journey; to my parents, whose boundless love and nurturing care fortified me at every challenging turn; to my advisor, Dan, whose sagacious wisdom, enduring patience, and invaluable mentorship profoundly shaped not only this work but my entire academic trajectory.

EPIGRAPH

There is only one true heroism in the world: to see the world as it is, and to love it.

Roman Rolland

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

As I reflect on the journey of my Ph.D., I am deeply grateful for the guidance, support, and inspiration I have received from so many.

Foremost, I would like to thank my advisor, Prof. Daniel Green. His profound wisdom and steadfast patience, particularly during challenging moments of complex conceptual understanding, have been pivotal to my academic development. His unique ability to simplify intricate processes and unparalleled passion for what we study have not only clarified my understanding but has also instilled in me a motivation for continuous learning. Besides, I am fortunate to have worked alongside my colleagues Benjamin Wallisch and Jiashu Han, whose intellectual camaraderie and insightful dialogue have immeasurably enhanced my research journey.

My gratitude also extends to Prof. Julius Kuti and Prof. Javier Duarte. When I worked as the TA of the computational physics series, their mentorship was nothing short of transformative. Their mentorship ignited a deep and enduring fascination in me for N-body and quantum mechanics simulations, especially within the realm of high-performance computing. I am also thankful to Prof. George Fuller and Prof. Jeffrey Rabin for their guidance and support as members of my thesis committee. The support from staff members in the Physics Department, particularly Sharmila Poddar, Saixious Dominguez-Kilday and Catherine McConney, has been instrumental in my academic journey. Their kindness and support have been invaluable.

I cannot overlook the enduring support and friendship from my warm-hearted neighbours Weiting Kuo and Jofan Chien, my steadfast roommates Zhengdi Sun, Xiang Ji, and Dawei Li, and my dedicated schoolmates and best friends Yifan Xiang, Dachuan Lu, Yiheng Xu, Yiwen Huang, Jinlong Huang, and Xinyan Tong. Their unwavering encouragement and shared experiences, through both challenges and triumphs, have been a wellspring of comfort and inspiration.

Within the constellation of remarkable individuals who illuminated my Ph.D. journey, Xinwen Hu's radiance outshines with a unique brilliance. Her presence, radiating with gentle kindness and warm goodness, has been a beacon of love and support, constantly reminding me of the joy and balance that love brings to life. Her companionship during this academic pursuit has been a cherished gift, making every challenge easier to overcome and every accomplishment more joyful.

Lastly, but importantly, my family's unwavering support and belief in my abilities have been the foundation of my resilience. I extend heartfelt thanks to my parents for their endless sacrifices and invaluable life lessons; to my brothers and sisters, for their steadfast camaraderie and cheerful support, which have continually kept me grounded and focused; to my aunts and uncles, for their care and concern; and to my grandma, for her tenacious spirit in fighting against COVID-19 and emerging strong and healthy. Their love and encouragement have been my constant source of strength.

Chapter 2, in full, is a reprint of the material as it appears in Cosmological Implications of Axion-Matter Couplings, Green, Daniel; Guo, Yi; Wallisch, Benjamin, JCAP, 2022. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is a reprint of the material as it appears in Light Fields during Inflation from BOSS and Future Galaxy Surveys, Green, Daniel; Guo, Yi; Han, Jiashu; Wallisch, Benjamin, arXiv:2311.04882. The dissertation author was one of the primary investigators and authors of this paper.

VITA

2018	Bachelor of Science, University of California Santa Barbara
2018–2023	Teaching Assistant, Department of Physics University of California San Diego
2024	Doctor of Philosophy, University of California San Diego

PUBLICATIONS

D. Green, Y. Guo and B. Wallisch, “Cosmological Implications of Axion-Matter Couplings,” *JCAP* **02** (2022) 019, arXiv:2109.12088 [astro-ph.CO].

D. Green, Y. Guo, J. Han and B. Wallisch, “Light Fields during Inflation from BOSS and Future Galaxy Surveys,” arXiv:2311.04882 [astro-ph.CO].

FIELDS OF STUDY

Major Field: High Energy Physics

Studies in Cosmology
Professor Daniel Green

ABSTRACT OF THE DISSERTATION

Cosmological Implications of Axion-Matter Couplings and Light Fields during Inflation from BOSS and Future Galaxy Surveys

by

Yi Guo

Doctor of Philosophy in Physics

University of California San Diego, 2024

Professor Daniel Green, Chair

In this dissertation, we explore two aspects of the dark matter using cosmology. In Chapter 2, we study the interaction between axion, a famous dark matter candidate, and ordinary matter, through their trace left in the thermal history of the universe. We show that the axion-matter coupling can be constrained by the effective number of relativistic species, N_{eff}, with next-generation CMB experiments. In Chapter 3, we study the distribution of matter before inflation, by conducting massive galaxy surveys of the large-scale structure of the universe. We show that the measurement of the primordial non-gaussianity of matter distribution can be significantly optimized by various techniques and strategies for ongoing and futuristic surveys.

Chapter 1

Introduction

The universe, a vast expanse of space, time, and matter, has intrigued humans for centuries. From ancient astronomers observing the stars to modern physicists unraveling the mysteries of subatomic particles, the quest to understand our universe has been a cornerstone of human curiosity. Cosmology stands at the crossroads of various physics disciplines. It merges the knowledge of elementary particles, the forces that act upon them, and the large-scale properties of the universe. This field has provided us with insights into the early moments after the Big Bang and has led to groundbreaking theories about the nature of space and time.

One of the most perplexing discoveries in cosmology is the existence of dark matter. Unlike ordinary matter, dark matter hardly emits, absorbs, or reflects light, making it invisible and detectable only through its gravitational effects so far. Understanding dark matter is crucial for a complete picture of the universe's composition and evolution. In the journey to understand dark matter, axions emerge as a promising candidate. Originally proposed to solve a theoretical problem in quantum chromodynamics, they have gained attention in the search for dark matter. Their potential interaction with standard model fermions, the particles that make up matter as we know it, is the topic of chapter 2.

The study of primordial non-Gaussianity involves examining the very early universe's density fluctuations. These fluctuations, imprinted in the cosmic microwave background radiation and the large-scale structure of the universe, hold clues about the universe's infancy and its subsequent evolution. In chapter 3, we explore the sensitivity of current and future large-scale structure surveys to detecting primordial non-Gaussianity.

The research into axions, standard model fermions, and primordial non-Gaussianity is not just a pursuit of academic interest. It has profound implications for our understanding of the universe. It could answer fundamental questions about the nature of dark matter, the early moments of the universe, and the forces that have shaped its evolution.

Chapter 2

Cosmological Implications of Axion-Matter Couplings

2.1 Synopsis

Axions and other light particles appear ubiquitously in physics beyond the Standard Model, with a variety of possible couplings to ordinary matter. Cosmology offers a unique probe of these particles as they can thermalize in the hot environment of the early universe for any such coupling. For sub-MeV particles, their entropy must leave a measurable cosmological signal, usually via the effective number of relativistic particles, N_{eff}. In this paper, we will revisit the cosmological constraints on the couplings of axions and other pseudo-Nambu-Goldstone bosons to Standard Model fermions from thermalization below the electroweak scale, where these couplings are marginal and give contributions to the radiation density of $\Delta N_{\text{eff}} > 0.027$. We update the calculation of the production rates to eliminate unnecessary approximations and find that the cosmological bounds on these interactions are complementary to astrophysical constraints, e.g. from supernova SN 1987A. We additionally provide quantitative explanations for these bounds and their relationship.

2.1.1 Contribution of the author

The author has contributed in developing the main formalism and in writing the paper. He has performed the numerical calculations and the analysis under the direction of the supervisor.

2.2 Introduction

Light particles with very weak couplings to the Standard Model are highly-motivated experimental targets from a number of perspectives. Axions and axion-like particles have been proposed as solutions to fine-tuning problems like the strong CP problem [1, 2, 3] and the hierarchy problem [4]. Furthermore, a cold component of the axion may form a viable dark matter candidate [5, 6, 7]. Alternatively, light particles may take the form of Goldstone or pseudo-Goldstone modes that arise as a consequence of symmetry breaking, including in models of flavor [8, 9, 10, 11] (familons), neutrino masses [12, 13] (majorons) and supersymmetry (gravitino). More broadly, top-down models suggest that there could be a number of additional sectors with very weak or gravitational couplings to the Standard Model [14, 15, 16, 17, 18, 19]. The plethora of possibilities for light particles [20, 21] is mirrored in the variety of dark matter candidates and dark sectors that are being actively explored [22].

Cosmology plays a vital role in our investigations of this vast landscape. While the list of possible couplings to the Standard Model (SM) are enormous, if any of them is sufficient to thermalize one of these particles, its relic energy density is detectable through its gravitational influence [23, 24, 25, 26, 27, 28, 29, 30, 31, 32]. Thermalized particles with masses $m \ll 1\,\mathrm{eV}$ are relativistic during the radiation era and the cosmological constraints can be inferred from the effective number of relativistic species N_{eff}. Heavier particles with $m < 100\,\mathrm{eV}$ will contribute to the sum of neutrino masses and are also constrained by the evolution of the universe at low redshifts. Most importantly, when its mass is sub-MeV, the entropy carried by a new particle cannot be eliminated by its decay or annihilation and, therefore, always leaves a cosmological signal in some combination of relic abundances (big bang nucleosynthesis/BBN) [33, 34, 35], cosmic microwave background (CMB) [36, 37, 38, 39, 40, 41, 42] and/or large-scale structure [43, 44, 45, 46] observables.

The majority of couplings of light particles to the Standard Model are irrelevant. The combination of symmetries needed to protect the mass of the particle combined with the limited set of gauge-invariant operators in the Standard Model usually ensures that they are dimension five or larger [20, 21, 29]. As a result, thermalization usually occurs at high temperatures where all the particles involved are relativistic. As a result, the strongest constraints on these couplings will arise from thermalization well above the electroweak scale. This implies that their contributions to N_{eff} are diluted to their minimum value, $\Delta N_{\mathrm{eff}} = 0.027\, g_s$, where g_s is the number of degrees of freedom of the light relic (see e.g. [47, 48, 49] for reviews).

Couplings of axions (familons) to matter are an exception to this general pattern.[1] Starting from a manifestly shift-symmetric form, a scalar field ϕ can be coupled to the SM fermions ψ_i via

$$\begin{aligned}\mathscr{L}_{\phi\psi} &= -\frac{\partial_\mu \phi}{\Lambda_\psi}\bar{\psi}_i\gamma^\mu\left(g_V^{ij} + g_A^{ij}\gamma^5\right)\psi_j \\ &\to \frac{\phi}{\Lambda_\psi}\left(\mathrm{i}H\bar{\psi}_{L,i}\left[\left(\lambda_i - \lambda_j\right)g_V^{ij} + \left(\lambda_i + \lambda_j\right)g_A^{ij}\right]\psi_{R,j} + \mathrm{h.c.}\right) + \mathscr{O}\left(\phi^2\right),\end{aligned} \tag{2.1}$$

where we integrated by parts and used the equations of motion with the Higgs doublet H, the left-/right-handed spinors $\psi_{L,R} \equiv \frac{1}{2}\left(1 \mp \gamma^5\right)\psi$, the Yukawa couplings $\lambda_i \equiv \sqrt{2}m_i/v$ and the Higgs vacuum expectation value $v = 246\,\mathrm{GeV}$.[2] Importantly, we see on the second line that this interaction is effectively dimension four in the presence of a non-zero Higgs vacuum expectation value. By dimensional analysis, the production rate of ϕ is proportional to temperature (since the effective coupling is dimensionless) and, therefore, exceeds the Hubble rate at low temperatures, which scales as temperature squared during radiation domination. However, this argument is only true above the mass of the SM fermion since the production rate will again become negligibly small once the number density of the fermion is sufficiently Boltzmann suppressed. This possibility is particularly intriguing because the decoupling temperature is effectively below the mass of the associated fermion and, therefore, results in a larger contribution to N_{eff}

---

[1] For simplicity, we will refer to all scalar particles interacting with matter as axions, whether we assume specific models or independent couplings, despite the fact that the latter might be more naturally referred to as familons in most of our cases or, more generally, pseudo-Nambu-Goldstone bosons (pNGBs).

[2] For simplicity, we suppressed the $SU(2)_L$ and $SU(3)_c$ structures which take the same form as for the SM Yukawa couplings [50].

which is more easily constrained or detected with near-term CMB experiments, such as Simons Observatory (SO) and CMB-S4 [51, 52].

For the above reasons, cosmological constraints on axion-matter couplings have received significant attention in the literature [29, 32, 53, 54, 55, 56, 57, 58]. Yet, while the origin of the constraints is straightforward to estimate qualitatively, precise numerical bounds depend on a number of details that have only been partially explored in the literature. Most essentially, the thermal production rate and/or decoupling calculations are often approximated in various ways. In some cases, the estimated bounds are substantially stronger than the true bounds [32].

The result of our analysis is a more precise calculation of ΔN_{eff} for interactions with charged leptons and heavy quarks, namely the electron, muon and tau lepton, and the charm, bottom and top quarks, as a function of each of their coupling strengths. In addition, we provide an intuitive (semi-)analytic explanation for the form of each of these curves. From these results, it is possible to straightforwardly derive the constraints on axion-matter couplings from a given measurement of N_{eff}. Moreover, we compare current and future cosmological constraints to existing bounds in the literature. Of particular interest is the relation to constraints from SN 1987A which have recently been inferred for couplings to muons [59, 60, 61]. As both the cosmological and astrophysical bounds are derived from thermal production of the axion, we explore the precise relationship between these bounds in that context.

This paper is organized as follows: In Section 2.3, we summarize the cosmology and particle physics of axions, or more generally pNGBs, and their coupling to SM fermions studied in this work. This includes a qualitative explanation of the freeze-in (low-temperature re-thermalization and decoupling) and freeze-out (high-temperature decoupling) scenarios that are possible below and above the electroweak symmetry breaking scale, and how they allow measurements of N_{eff} to be translated into bounds on these interactions. In Section 2.4, we present the calculation of the axion production rates and the coupling constraints. Moreover, we infer the bounds from current and future measurements of N_{eff}, and describe the physics underlying these cosmological constraints in detail. In Section 2.5, we explore the relation between these bounds and those from astrophysical measurements, in particular the cooling rate of SN 1987A, and quantitatively compare our N_{eff}-based constraints to those and other existing bounds. In Section 2.6, we present our conclusions. A set of appendices contains technical details on the calculation of the axion production rate (Appendix A.1), and the implications of quantum statistics and the presence of the QCD phase transition on our results (Appendix A.2).

2.3 Review of Axions and Axion Cosmology

Axion-like particles and other pseudo-Nambu-Goldstone bosons arise in a variety of forms, depending on the ultraviolet completion. String theory famously contains a plenitude of axions. Alternatively, the strong CP problem suggests the need for an axion ϕ with a coupling $\phi \tilde{G}_{\mu\nu} G^{\mu\nu}$, where $G_{\mu\nu}$ is the gluon field strength tensor and $\tilde{G}_{\mu\nu}$ its dual. The key feature of the coupling of such particles is that they preserve a shift symmetry, $\phi \to \phi + c$, with constant c, such that they can be naturally light.

While the coupling of axions to gauge bosons is often what distinguishes axions from other naturally light scalars, the coupling to matter can and will arise for all such particles. In some cases, the particles are given alternate names such as *familons*. Nevertheless, given the shift symmetry, the leading interactions with the SM fermions are given by[3]

$$\mathscr{L}_{\phi\psi} = -\frac{\partial_\mu \phi}{\Lambda_\psi} \left(g_V^{ij} J_V^{ij} + g_A^{ij} J_A^{ij} \right) = -\frac{\partial_\mu \phi}{\Lambda_\psi} \bar{\psi}_i \gamma^\mu \left(g_V^{ij} + g_A^{ij} \gamma^5 \right) \psi_j \,, \tag{2.2}$$

where the couplings to the vector and axial-vector currents $J_{V,A}$ are denoted by the subscripts V and A, respectively. The diagonal vector couplings, $i = j$, vanish due to vector current conservation, i.e. diagonal couplings are only present for the axial part. This can also be seen explicitly

---

[3] In this paper, we follow the notation of [32] and parameterize the dimensionful axion couplings in terms of the effective mass scale $\Lambda_{ij} \equiv \Lambda_\psi / [(g_V^{ij})^2 + (g_A^{ij})^2]^{1/2}$. This parametrization can be straightforwardly converted to other commonly employed notations, such as the inverse scale $\tilde{g}_{ij} = 1/\Lambda_{ij}$, the dimensionless coupling constant $\tilde{\varepsilon}_i = 2m_i/\Lambda_{ii}$ or the decay constant $f_a = \Lambda_\psi$.

after integrating by parts and using the equations of motion, as performed in the second line of (2.1). While a linear combination of the axial couplings is equivalent to the coupling of axions to gauge bosons due to the chiral anomaly, we only consider the effects of the couplings to matter with no contribution from anomalies in this paper.

In the cosmological context, the impact of an axion-matter interaction is qualitatively different before and after the electroweak phase transition. Prior to the electroweak phase transition, these couplings are described by dimension-5 operators and the axion interaction rate with SM particles therefore scales as $\Gamma_\phi \sim T^3/\Lambda_\psi^2$. Meanwhile, the expansion rate of the universe scales as $H \sim T^2/M_{\rm pl}$ at those early times. This implies that the axion may be in thermal equilibrium with the rest of the SM for $T > T_F$, where freeze out at temperature T_F is defined by $H(T_F) \simeq \Gamma_\phi(T_F)$ (assuming T_F is above the electroweak scale). For $T < T_F$, axion production becomes inefficient and the axions decouple from the Standard Model. Note that a population of hot axions may exist whether or not axions form the dark matter and are therefore complementary to many of the direct detection strategies [22].

This description of decoupling is applicable to any particle coupled to the Standard Model through an irrelevant operator. This is the common origin of most (light) thermal relics in the early universe and leads to the standard contribution to the energy density in free-streaming radiation as parameterized by

$$N_{\rm eff} = \frac{8}{7}\left(\frac{11}{4}\right)^{4/3}\frac{\rho_\nu + \rho_\phi}{\rho_\gamma} \qquad \to \qquad \Delta N_{\rm eff} = \frac{8}{7}\left(\frac{11}{4}\right)^{4/3}\frac{\rho_\phi}{\rho_\gamma}\,, \tag{2.3}$$

where ρ_ν and ρ_ϕ are the energy density in neutrinos and axions (or any other light thermal relic beyond the Standard Model), respectively. Given their sub-eV masses, both are relativistic prior to recombination, which means that $N_{\rm eff}$ does not distinguish between axions, neutrinos or any other decoupling relativistic species. While $N_{\rm eff} = 3.044$ in the Standard Model due to the three neutrinos [62, 63, 64], a thermalized axion or other pNGB will contribute $\Delta N_{\rm eff} \geq 0.027$, with this bound being saturated for decoupling above all SM mass thresholds. For general T_F and

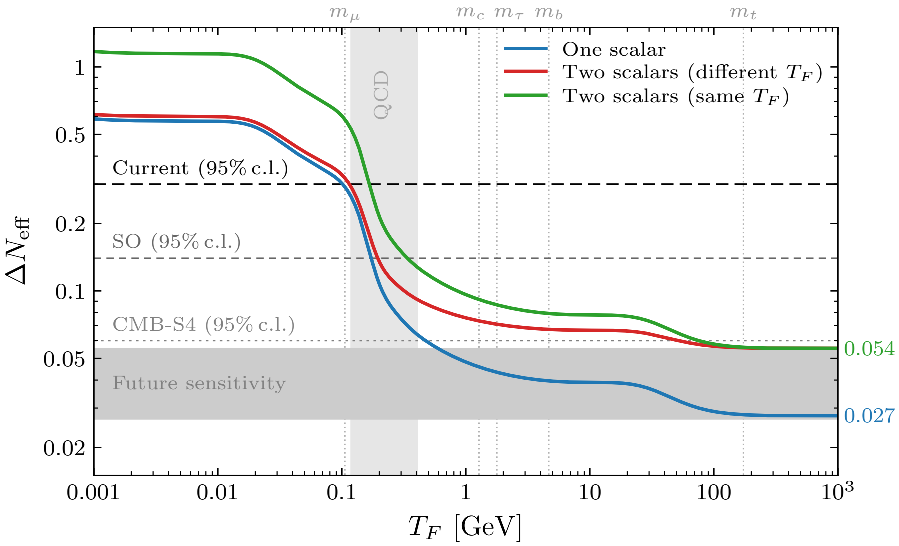

Figure 2.1. Contribution to ΔN_{eff} from a light particle that decoupled from the Standard Model at a freeze-out temperature T_F. The blue line indicates the contribution for a single real degree of freedom, such as an axion or Goldstone boson. The green and red lines show $\Delta N_{\mathrm{eff}}(T_F)$ for the case of two scalar degrees of freedom, either with the same decoupling temperature, $T_F \equiv T_{F,1} = T_{F,2}$, or two different decoupling temperatures, $T_F \equiv T_{F,1} \lesssim 10^3\,\mathrm{GeV} < T_{F,2}$. The dashed lines indicate the current bound on ΔN_{eff} at 95% c.l. from Planck 2018 and BAO data [41], and the forecasted constraints from the Simons Observatory (SO) [51] and CMB-S4 [52]. The gray band illustrates the future sensitivity that might potentially be achieved with a combination of cosmological surveys of the CMB and large-scale structure, such as CMB-HD [65], MegaMapper [66] and PUMA [67], cf. [44, 68]. We refer to [48, 49] for additional details.

g_s internal degrees of freedom, the contribution is given by[4]

$$\Delta N_{\mathrm{eff}} = g_s \left(\frac{43/4}{g_*(T_F)} \right)^{4/3} , \tag{2.4}$$

where $g_*(T)$ is the effective number of SM degrees of freedom at temperature T. The effect of lower T_F is to increase ΔN_{eff} by reducing the amount of entropy converted to photons after decoupling of the axion. This increase is shown in Fig. 2.1 since the SM particles become massive, annihilate and deposit their energy (and entropy) in the remaining thermal SM bath. This therefore provides a natural observational target (see also [48, 49], for instance).

---

[4] We make the assumption that there are no large sources of entropy beyond the Standard Model particles at or below the freeze-out temperature T_F. See e.g. [48] for a more detailed discussion.

The absence of a detection with an exclusion of $\Delta N_{\rm eff} = 0.027$ using future cosmological data would put strong constraints on the coupling strength Λ_{ij} [32] since this would exclude the presence of any thermalized relics above the electroweak scale. In this case, no additional light particles could have been in thermal equilibrium with the SM at any point in the history of the universe (including axions), back to the era of reheating at temperature T_R. Since this requires the would-be freeze-out temperature to be above the reheating temperature, $T_F(\Lambda_\psi) > T_R$, a significant exclusion of $\Delta N_{\rm eff} = 0.027$ would imply very strong constraints on the axion couplings. If we define $\Lambda_F^{(\psi)}(T)$ as the coupling Λ_ψ such that $\Gamma(T) = H(T)$ at temperature T, then our approximate bound would be $\Lambda_\psi \gtrsim \Lambda_F^{(\psi)}(T_R)$.[5] For the scales in (2.1), this constraint would imply [32][6]

$$\Lambda^I_{ij} > \begin{cases} 1.0\times 10^{11}\,\text{GeV}\,\dfrac{m_i \mp m_j}{m_\tau}\left(\dfrac{T_R}{10^{10}\,\text{GeV}}\right)^{1/2} & i,j = \text{leptons}, \\ 1.8\times 10^{13}\,\text{GeV}\,\dfrac{m_i \mp m_j}{m_t}\left(\dfrac{T_R}{10^{10}\,\text{GeV}}\right)^{1/2} & i,j = \text{quarks}, \end{cases} \tag{2.5}$$

where m_i are the SM fermion masses, with $m_\tau \approx 1.8\,\text{GeV}$ and $m_t \approx 173\,\text{GeV}$. We refer to [32] for additional details, including a comparison of the current experimental and the prospective cosmological constraints which will likely be stronger by orders of magnitude except for most interactions involving electrons.

However, after the electroweak phase transition, there exists a second scenario which is unique to the couplings of SM fermions to axions (and other pNGBs). The out-of-equilibrium axions may re-equilibrate and thermalize with the Standard Model after the Higgs acquired its

---

[5] This also assumes no dramatic increase in the number of degrees of freedom in the Standard Model or non-equilibrium evolution that could further dilute $\Delta N_{\rm eff} < 0.027$ (see [48] for more details and discussion).

[6] We directly use the results of [32] and do not include any improved calculations of the axion production rate here because the sensitivity to the (unknown) reheating temperature limits the need for a more precise calculation at this point.

non-zero vacuum expectation value. In this case, the Lagrangian (2.1) becomes

$$\mathscr{L}_{\phi\psi} = \mathrm{i}\frac{\phi}{\Lambda_\psi}\bar{\psi}_i\left[(m_i - m_j)g_V^{ij} + (m_i + m_j)g_A^{ij}\gamma^5\right]\psi_j\,, \tag{2.6}$$

which is effectively a dimension-four interaction. This implies that the interaction rate now scales as $\Gamma_\phi \sim m_\psi^2 T/\Lambda_\psi^2$, which is a weaker temperature dependence than that of the expansion rate, $H \sim T^2$. Depending on the interaction strength, the axions will eventually thermalize and decouple again at later times leaving a much larger contribution to the radiation density as displayed in Fig. 2.1. Such a contribution to ΔN_{eff} may already be ruled out with current cosmological datasets or ruled out in the (near) future. Preventing this re-thermalization of the axion to lead to a large axion density and violation of existing (or near-future) constraints on ΔN_{eff} requires the re-equilibration temperature to be smaller than the mass of the respective fermion(s) since the interaction rate Γ becomes Boltzmann suppressed in this regime. This in turn suggests that we can put limits on the axion couplings by effectively trading the reheating temperature T_R with the fermion mass m_ψ, $\Lambda_\psi \gtrsim \Lambda_F^{(\psi)}(m_\psi)$ [32].

The resulting bounds are typically weaker than those derived from freeze-out above the electroweak scale, as the higher temperatures ultimately lead to more efficient production. Nevertheless, high-temperature freeze-out is more[7] sensitive to assumptions on the reheating temperature and particle content of the universe. Moreover, these bounds will be easier (and therefore earlier) to achieve since the larger contributions to ΔN_{eff} that are generated by the freeze-in process can be more easily measured (or excluded) by a realistic cosmological survey. While order-of-magnitude estimates of the resulting bounds were provided in [32], the contributions to ΔN_{eff} are at the threshold of current and future CMB experiments and, therefore, demand a more careful treatment. Given the highly nonlinear relationship between ΔN_{eff} and the fundamental parameters of the model, even seemingly small effects can translate into large

---

[7] The freeze-in scenario (re-thermalization and decoupling) only assumes that the reheating temperature is larger than the fermion mass(es), $T_R > m_\psi$.

differences in the inference of the axion-matter couplings (see e.g. §2.4.3 and Appendix A.2).

2.4 Production Rates and Cosmological Constraints

Relating N_{eff} measurements to specific models is dependent on a reliable calculation of the production rate. The existence of a constraint usually follows from dimensional analysis, but we also have to put the cosmological constraints into the broader context of experimental probes of axions and other pNGBs. Calculating the axion production rate accurately is difficult as some approximations are unreliable when $T \approx m_\psi$ [32].[8] In this section, we will therefore recalculate these rates for the couplings of axions to SM matter particles without these approximations (§2.4.1). We will then derive predictions of ΔN_{eff} and observational bounds on these interactions, focusing on the diagonal couplings, $\Lambda_i \equiv \Lambda_\psi / g_A^{ii}$, for simplicity (§2.4.2). Moreover, we will provide a detailed discussion of the physics underlying these constraints (§2.4.3).

2.4.1 Computation of Production Rates

In the following, we summarize the calculation of the interaction rate of axions and other pseudo-Nambu-Goldstone bosons with SM fermions. While we will put an emphasis here on the conceptual steps and relevant physical processes, we refer to Appendix A.1 for the technical details.

The leading processes producing axions in the early universe with diagonal interactions described by (2.6) after electroweak symmetry breaking are shown in Fig. 2.2: (a) Compton-like scattering, $\psi_i + \{\gamma, g\} \to \psi_i + \phi$, and (b) fermion-antifermion annihilation, $\psi_i + \bar{\psi}_i \to \phi + \{\gamma, g\}$, where we denoted the photon and gluon by γ and g, respectively. The scattering amplitudes of these production channels are given by [32]

$$\sum |\mathscr{M}|^2_{(a)} = 16\pi A_\psi \, |\tilde{\varepsilon}_i|^2 \frac{t^2}{(s - m_i^2)(m_i^2 - u)} \,, \tag{2.7}$$

---

[8] Prior calculations of the axion production rate have assumed relativistic particles and a high-temperature limit [32] or used Boltzmann statistics ignoring Bose enhancement and Pauli blocking [54], for instance.

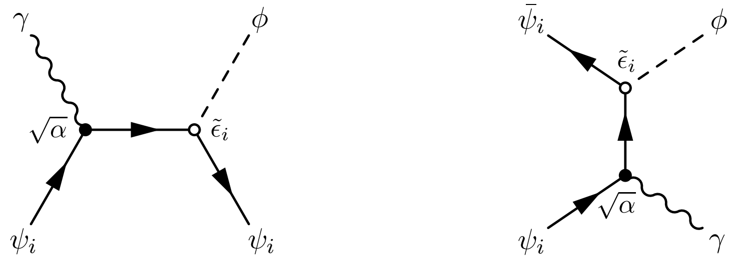

(a) Compton-like scattering. **(b)** Fermion annihilation.

Figure 2.2. Feynman diagrams for the dominant production channels of axions and other pseudo-Nambu-Goldstone bosons via the coupling to charged fermions below the electroweak scale: (a) Compton-like scattering and (b) fermion annihilation. For quarks, the coupling to photons is replaced by that to gluons. In addition to the displayed *s*- and *t*-channel diagrams, there are *u*-channel diagrams which are not shown.

$$\sum|\mathscr{M}|^2_{(b)} = 16\pi A_\psi |\tilde{\varepsilon}_i|^2 \frac{s^2}{(m_i^2-t)(m_i^2-u)}\,, \tag{2.8}$$

with $\tilde{\varepsilon}_i \equiv 2m_i/\Lambda_i$, the Mandelstam variables s, t and u, and

$$A_\psi \equiv \begin{cases} \alpha & \psi = \text{lepton}, \\ 4\alpha_s & \psi = \text{quark}. \end{cases} \tag{2.9}$$

We will neglect the weak temperature dependence of the running fine-structure constant α and approximate it by its low-energy value of $\alpha \approx 1/137$. On the other hand, the running strong coupling constant $\alpha_s(T)$ significantly depends on temperature. We include this temperature dependence by employing the five-loop corrections of the QCD beta function implemented in `RunDec` [69] for all temperatures with $\alpha_s < 1$. In the following, we will (conservatively) stop our calculation at a temperature of $T = 1\,\mathrm{GeV}$ when $\alpha_s \approx 0.5$ to avoid the strongly-coupled regime. We refer to Appendix A.2 for additional details, a discussion of the implications of this choice and a less conservative calculation.

In general, the production rate for the relevant two-to-two processes is

$$\Gamma_\phi = \frac{1}{n_\phi^{\mathrm{eq}}} \prod_{i=1}^{4} \int \frac{\mathrm{d}^3 p_i}{(2\pi)^3 2E_i} f_1(p_1)\, f_2(p_2)\, [1 \pm f_3(p_3)]\, [1 \pm f_4(p_4)] \tag{2.10}$$
$$\times (2\pi)^3 \delta^{(3)}(p_1 + p_2 - p_3 - p_4)\, (2\pi) \delta(E_1 + E_2 - E_3 - E_4) \sum |\mathscr{M}|^2 ,$$

where $n_\phi^{\mathrm{eq}}(T) = \zeta(3)\, T^3/\pi^2$ is the equilibrium number density of a relativistic scalar at temperature T, the momenta and energies of the incoming (outgoing) particles are denoted by p_i and E_i with $i = 1, 2$ $(3, 4)$, the Bose-Einstein and Fermi-Dirac distribution functions for bosons and fermions are

$$f^b(p) = \frac{1}{\mathrm{e}^{E(p)/T} - 1} , \qquad f^f(p) = \frac{1}{\mathrm{e}^{E(p)/T} + 1} , \tag{2.11}$$

and '±' indicates either Bose enhancement ('+') or Pauli blocking ('−') of the outgoing bosons and fermions, respectively. The total scattering amplitude is given by $\sum |\mathscr{M}|^2 = 2 \sum |\mathscr{M}|^2_{(a)} + \sum |\mathscr{M}|^2_{(b)}$ to account for fermions and antifermions in the Compton-like process. We show in Appendix A.1 that the rate (2.10) can be rewritten as the following five-dimensional integral:

$$\Gamma_\phi = \frac{1}{n_\phi^{\mathrm{eq}}} \int_{E_{\min}}^{\infty} \mathrm{d}E \int_0^{p_{\max}} \mathrm{d}p \int_{p_1^{\min}}^{p_1^{\max}} \mathrm{d}p_1 \int_{p_3^{\min}}^{p_3^{\max}} \mathrm{d}p_3\, \frac{p_1 p_3}{512 \pi^6 E_1 E_3} f_1(p_1) f_2(p_2) \tag{2.12}$$
$$\times [1 \pm f_3(p_3)]\, [1 \pm f_4(p_4)] \int_0^{2\pi} \mathrm{d}\phi \sum |\mathscr{M}|^2 ,$$

where E and p are the total energy and momentum, p_1 and p_3 are one of the incoming and outgoing momenta each, and ϕ is the polar angle difference between these two momenta in the plane orthogonal to p. We implicitly impose energy-momentum conservation to fix $p_{i+1} = p - p_i$ for $i = 1, 3$ and provide the integration limits in Appendix A.1.

For the specific amplitudes of Compton-like scattering (2.7) and fermion annihilation (2.8) that are of interest in this work, the integral over the angle ϕ can be conducted analytically. This means that we are left with a four-dimensional integral which we evaluate numerically using multi-dimensional adaptive quadrature. To facilitate its numerical calculation,

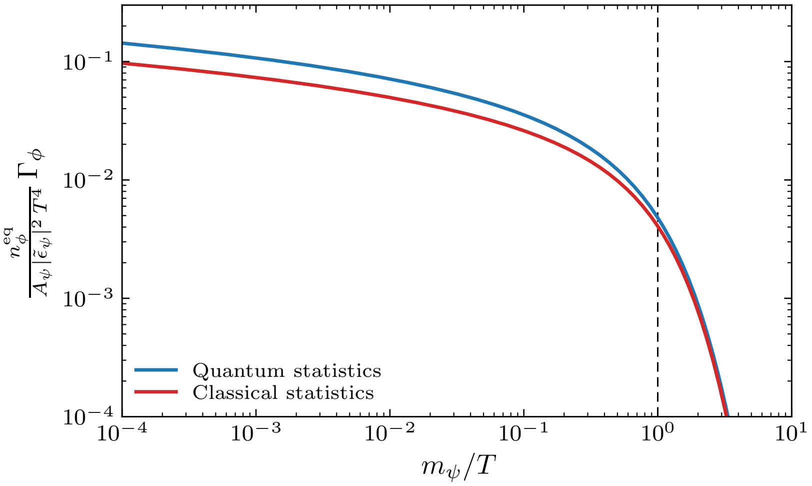

Figure 2.3. Dimensionless rescaling of the interaction rate Γ_ϕ as a function of m_ψ/T. We compare the results of our full calculation using the Bose-Einstein and Fermi-Dirac distribution functions ('quantum statistics') with the approximate result of employing the Boltzmann distribution without Bose enhancement and Pauli blocking ('classical statistics'). The vertical dashed line indicates $T = m$ which is approximately the temperature where decoupling occurs for moderate coupling strengths.

it is useful to consider the rescaling $n_\phi^{\rm eq} A_\psi^{-1} |\tilde\varepsilon_\psi|^{-2} T^{-4}\,\Gamma_\phi$ as a function of m_ψ/T since it is dimensionless and independent of both the axion-fermion coupling and the SM fermion masses. While we take these masses m_ψ to be non-zero, we assume massless axions, $m_\phi = 0$, which is a good approximation for large parts of parameter space relevant for measurements of the relativistic energy density as parameterized by $N_{\rm eff}$. The result is shown in Fig. 2.3 for the full quantum distribution functions of (2.11) with Bose enhancement and Pauli blocking, and the commonly employed classical approximation of Boltzmann statistics, $f^b(p) = f^f(p) = \exp\{-E(p)/T\}$, without Bose enhancement or Pauli blocking. As expected, the curves agree in the Boltzmann-suppressed regime of low temperatures, $T \lesssim m_\psi$ within 20%, but differ at large temperatures $T \gg m_\psi$ (e.g. about 50% for $T = 10^3\, m_\psi$). Since the difference in the range $m_\psi/T \in [1, 10]$, which is most relevant for equilibrium physics, roughly varies between 15% and 20%, we expect shifts of less than about 10% in our final predictions for the contribution to $N_{\rm eff}$, except in their tails at small couplings where the differences may be considerably larger.

2.4.2 Computation of Constraints

Having calculated the production rate as a function of temperature T, we now compute the number density n_ϕ of axions and other pNGBs, and the associated contribution to $\Delta N_{\rm eff}$ as a function of the coupling Λ_i. This subsequently allows us to provide bounds on axion-fermion interactions from current and future cosmological measurements of $\Delta N_{\rm eff}$.

Instead of assuming instantaneous annihilation of the SM fermions at $T = m_\psi$ to estimate the relevant axion abundance, we solve the Boltzmann equation for the axion number density n_ϕ,

$$\frac{\mathrm{d}n_\phi}{\mathrm{d}t} + 3H(T)\,n_\phi = \Gamma_\phi(T)\left(n_\phi^{\rm eq}(T) - n_\phi\right), \tag{2.13}$$

with the Hubble parameter $H(T)$ during radiation domination. While our calculation of the production rate $\Gamma_\phi(T)$ is general, including Bose enhancement and Pauli blocking, this equation assumes that this quantity is determined by the rate in equilibrium and is therefore independent of n_ϕ. We only expect a minor impact of these assumptions for coupling strengths Λ_i for which the axion reaches equilibrium at temperatures $T \gtrsim m_\psi$, i.e. when the expected number density of axions is near its equilibrium value. On the other hand, if the axion-fermion coupling is so small that the axion is never close to reaching equilibrium, our calculation of $\Gamma_\phi(T)$ leads to a slight overestimate of this rate because it includes the equilibrium Bose enhancement from using $n_{\rm eq}(T)$ in the final state.[9] In consequence, our calculations may overestimate the contribution to the radiation density as a function of the coupling strength, $\Delta N_{\rm eff}(\Lambda_i)$, by about 30% around $\Delta N_{\rm eff} = 0.02$ or, alternatively, the bound on the interaction strength given a $\Delta N_{\rm eff}$ measurement, $\Lambda_i(\Delta N_{\rm eff})$, by roughly 10% (with the latter being the quantity that we are more interested in).

We follow the common procedure (see e.g. [54, 53, 55]) of numerically solving this differential equation after changing variables to the dimensionless inverse temperature $x = m_\psi/T$

---

[9] We could have alternatively solved the Boltzmann equation for the distribution function $f_\phi(p,T)$, which is an integro-differential equation, instead of the respective equation (2.13) for the number density $n_\phi(T)$ to capture this effect.

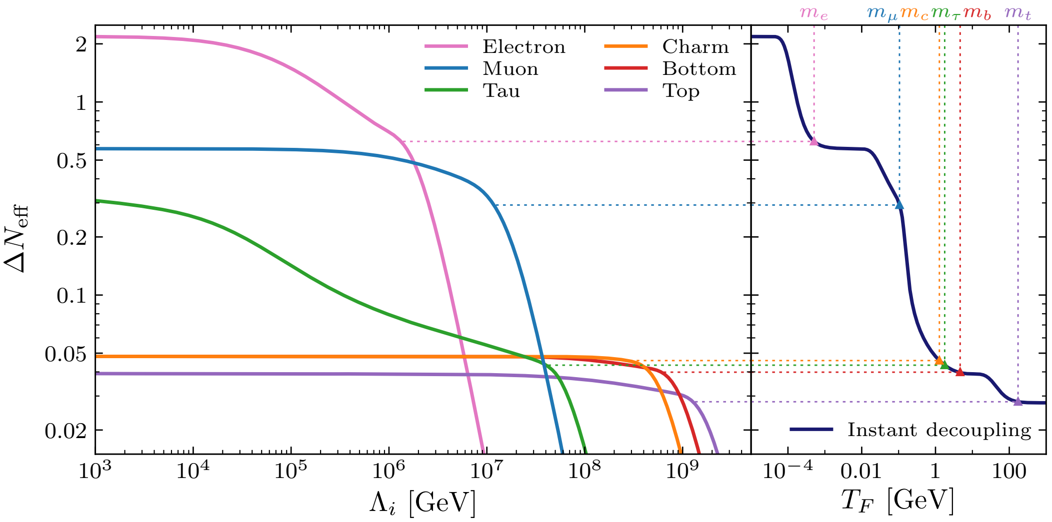

Figure 2.4. *Left:* Contribution to the radiation density as parameterized by ΔN_eff as a function of the axion-fermion coupling strength Λ_i for different SM fermions ψ_i. The displayed values for the bottom and charm couplings are conservative and may be (significantly) larger, with details of these uncertainties being discussed in Appendix A.2. *Right:* Contribution to ΔN_eff for a single thermalized (equilibrium) degree of freedom which decoupled from the Standard Model at a temperature T_F (i.e. the same as Fig. 2.1). The horizontal lines between the panels indicate the contribution to ΔN_eff expected for $T_F = m_\psi$ on the right with the appropriate value of Λ_i on the left. For larger values of Λ_i, the particle fails to reach equilibrium and, therefore, the abundance decreases rapidly.

and the dimensionless comoving number density $Y_\phi = n_\phi/s$, where $s = 2\pi^2 g_{*s} T^3/45$ is the entropy density. We adopt the effective number of degrees of freedom in entropy $g_{*s}(T)$ as numerically computed by [70], which is based on the lattice QCD calculation of [71] in the non-perturbative regime,[10] and an initial condition with no axions, $Y_{\phi,0} \equiv Y_{\phi,t=0} = 0$. Having obtained the final value for the comoving number density, $Y_{\phi,\infty}$, as a function of the SM fermion ψ_i and the axion-fermion interaction strength Λ_i, we can convert it to a prediction for the contribution to the radiation density according to $\Delta N_\mathrm{eff} \approx 74.84\, Y_{\phi,\infty}^{4/3}$. (We refer to Appendix A.1 for additional details.)

The results of this calculation are presented in Fig. 2.4, which shows the contribution

---

[10] If we used the results of [71] over the entire temperature range, our predictions for ΔN_eff would be within $\lesssim 10\%$ of the presented results.

to $\Delta N_{\rm eff}$ for each of the axion couplings to SM fermions.[11] Given the current constraint from Planck, $\Delta N_{\rm eff} < 0.30$ (95%) [41], we can constrain the axion coupling to electrons, muons and tau leptons:[12]

$$\Lambda_e > 2.5 \times 10^6\,\text{GeV} = 10^{6.4}\,\text{GeV}\,, \tag{2.16}$$

$$\Lambda_\mu > 1.1 \times 10^7\,\text{GeV} = 10^{7.1}\,\text{GeV}\,, \tag{2.17}$$

$$\Lambda_\tau > 1.7 \times 10^3\,\text{GeV} = 10^{3.2}\,\text{GeV}\,. \tag{2.18}$$

Upcoming (near-term) CMB experiments will continue to improve the measurement of $N_{\rm eff}$ which will also increase the sensitivity to these interactions. Given that the Simons Observatory [51] and CMB-S4 [52] are forecasted to reach $\Delta N_{\rm eff} < 0.14$ and 0.060 at 95% c.l., respectively, we project that SO can exclude

$$\Lambda_e > 3.7 \times 10^6\,\text{GeV} = 10^{6.6}\,\text{GeV}\,, \tag{2.19}$$

$$\Lambda_\mu > 2.2 \times 10^7\,\text{GeV} = 10^{7.3}\,\text{GeV}\,, \tag{2.20}$$

$$\Lambda_\tau > 1.0 \times 10^5\,\text{GeV} = 10^{5.0}\,\text{GeV}\,, \tag{2.21}$$

---

[11] The differences between our results obtained using full quantum statistics and calculations based on approximate classical statistics is about 5% in the plateaus, but increases when the predicted values of $\Delta N_{\rm eff}$ drop for larger values of Λ_i, reaching or even exceeding 20% at $\Delta N_{\rm eff} = 0.02$. This is because the axions decouple at higher temperatures where the difference between the classical and quantum production rate becomes more pronounced, cf. Fig. 2.3.

[12] Apart from bounds on $\Delta N_{\rm eff}$ from the CMB, we can also employ its BBN constraints (see e.g. [33, 34, 35]) which lead to

$$\Lambda_e > 1.8 \times 10^6\,\text{GeV} = 10^{6.2}\,\text{GeV}\,, \tag{2.14}$$

$$\Lambda_\mu > 1.4 \times 10^6\,\text{GeV} = 10^{6.1}\,\text{GeV}\,, \tag{2.15}$$

where we conservatively assumed $\Delta N_{\rm eff} < 0.5$ following [56]. While the electron bound is similar due to the functional dependence of $\Delta N_{\rm eff}$ on Λ_e in the Boltzmann-suppressed regime, the muon constraint is weaker by one order of magnitude. On the other hand, these BBN bounds are not limited to sub-eV axions, but extend to masses $m_\phi \lesssim 1\,\text{MeV}$ (cf. [56]).

while CMB-S4 will just fall short of the heavy quark targets,[13] but will constrain

$$\Lambda_e > 5.4 \times 10^6\,\mathrm{GeV} = 10^{6.7}\,\mathrm{GeV}\,, \tag{2.22}$$

$$\Lambda_\mu > 3.3 \times 10^7\,\mathrm{GeV} = 10^{7.5}\,\mathrm{GeV}\,, \tag{2.23}$$

$$\Lambda_\tau > 5.5 \times 10^6\,\mathrm{GeV} = 10^{6.7}\,\mathrm{GeV}\,. \tag{2.24}$$

We in particular note the improvement in Λ_τ by nearly two orders of magnitude when going from Planck to SO and from SO to CMB-S4, respectively. While the physical origins of these bounds will be discussed in §2.4.3, we clearly see that there is a highly nonlinear relationship between improvements in the measurement of N_eff and the parameters of models that produce $\Delta N_\mathrm{eff} > 0$.

These bounds hold for axions with masses $m \lesssim 1\,\mathrm{eV}$. At higher masses, the axions behave like matter (and not free-streaming radiation) during the recombination era and may also be constrained from their contribution to the effective mass of neutrinos. As we increase the axion mass, we expect that the constraints on these couplings will become more stringent due to their impact on structure formation and, eventually, over-closure of the universe. However, as we increase the mass, it is increasingly possible that the axions decay prior to recombination through a coupling to photons or neutrinos. We will therefore leave the discussion of larger masses to future work (see also [74, 75, 76]).

2.4.3 Physics of Constraints

In the previous subsection, we calculated ΔN_eff for axions and other pNGBs coupled to individual SM fermions. The shapes of the curves displayed in Fig. 2.4 vary significantly depending on the fermion. In the following, we provide qualitative explanations and describe the physics underlying this functional dependence of $\Delta N_\mathrm{eff}(\Lambda_i)$. In addition, we might hope to

---

[13] We note the relatively large uncertainty in our predictions for the coupling to the charm and bottom quarks due to the strong-coupling regime of the QCD phase transition (see Appendix A.2 for a more detailed discussion). Dedicated QCD lattice calculations may reveal that CMB-S4 and potentially even SO are sensitive to these interactions. An alternative approach is to match across the QCD phase transition using calculations for ΔN_eff before and after, similar to the calculation of [72, 73] for the QCD axion.

understand the approximate size of the constraints on the interaction strengths from dimensional analysis.

As described in detail in the previous subsection, cosmological production of axions is described by the Boltzmann equation (2.13). While the total production of axions can be determined exactly using this equation, the origin of the constraint follows from the qualitative requirement that the production becomes efficient at some temperature. Specifically, we will produce a large number density of axions if the production rate exceeds the rate of dilution due to the expansion of the universe, namely

$$\Gamma(T) > H(T) = \sqrt{\frac{\pi^2}{90} g_*(T)}\,\frac{T^2}{M_{\mathrm{pl}}}\,, \tag{2.25}$$

for some temperature T achieved in the early universe, with the reduced Planck mass M_{pl}. When this condition is met, the axion will thermalize, i.e. the number density of axions will approach the number density of photons at that temperature which can therefore yield a potentially detectable contribution to N_{eff}.

At any temperature above the mass of the fermion, we can always make the coupling sufficiently large to meet our condition in (2.25) which implies that the axions will thermalize. This is sufficient to ensure that the axions remain in thermal equilibrium with the rest of the Standard Model as the universe cools to lower temperatures (but that are still larger than the mass, $T > m_i$). However, due to the Boltzmann suppression of the fermion when $T \ll m_i$, the axions always decouple from the SM bath at sufficiently low temperatures, regardless of the coupling strength. Figure 2.5 compares the Hubble rate $H(T)$ to the production rate $\Gamma(T)$ for a range of interaction strengths to muons and illustrates both of these features. We notice that all production rates become much smaller than the Hubble rate, $\Gamma_\phi(T) \ll H(T)$, by $T \approx m_\mu/10$. As a consequence, we should expect that the left-most part of Fig. 2.4 should asymptote to ΔN_{eff} as determined by a decoupling at $T_F \approx m_i/10$ based on Fig. 2.1. This implies that the asymptotic contributions to ΔN_{eff} for large couplings to different fermions should be ordered inversely

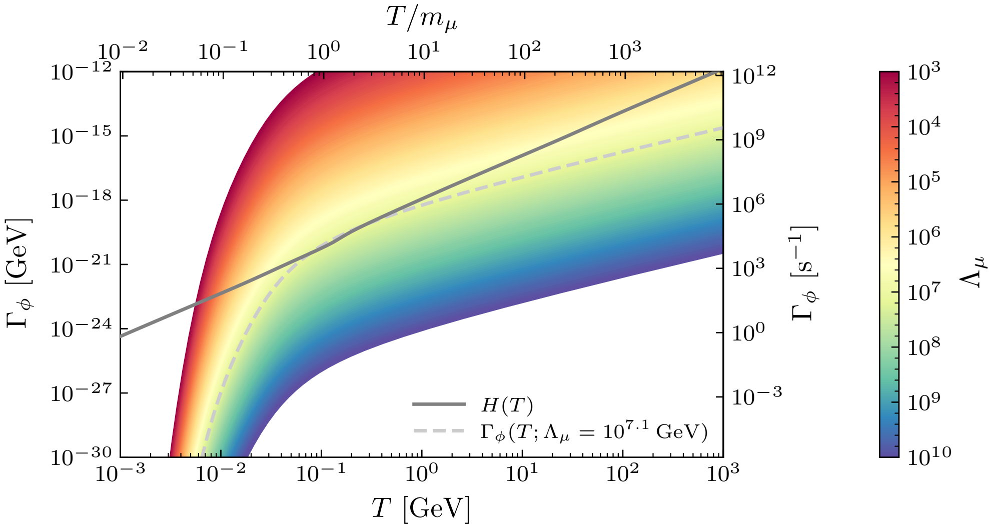

Figure 2.5. Production rate of axions and other pNGBs, Γ_ϕ, as a function of temperature T for different values of the muon coupling Λ_μ. The solid gray line shows the Hubble rate $H(T)$, i.e. $\Gamma_\phi(T) > H(T)$ indicates efficient cosmological production of axions at that temperature. The dashed line shows the production rate for the current bound on the interaction strength from N_{eff} measurements of Planck. For small values of Λ_μ, we see efficient production over many decades of $T > m_\mu$, but the axion usually decouples by $T \lesssim m_\mu/10$ because the production becomes exponentially suppressed. As Λ_μ increases, efficient production is increasingly possible only around $T \approx m_\mu$. When $\Lambda_\mu \gg 10^7$ GeV, there is no temperature where axions are efficiently produced from this interaction with muons. While the most natural units of the production rate and temperature for cosmological axion production are GeV (left) and the dimensionless ratio T/m_μ (top), production rates in s^{-1} (right) and temperatures in GeV (bottom) are useful for comparison with astrophysical axion production.

proportional to their mass, i.e. smaller masses m_i correspond to larger $\Delta N_{\rm eff}$ because $T_F \approx m_i/10$ is smaller. We see that this ordering is true for all the heavy fermions except for the ordering of the charm and tau lepton. This is a result of cutting off the axion production rate during the QCD phase transition when coupling to the charm because the strong coupling constant is no longer perturbative. This uncertainty due to QCD strong coupling is discussed in Appendix A.2 and implies that the charm and bottom curves are likely underestimated at larger couplings.

As we decrease the coupling, or increase the scale Λ_i, we see that the shape of the curves in Fig. 2.4 depends significantly on the specific fermion. Increasing the coupling changes the precise temperature at which the axion decouples and, therefore, the part of the $g_*(T)$ curve responsible for diluting the number of axions. Specifically, as the coupling decreases, the temperature where the axion production is significant becomes increasingly restricted to $T \approx m_i$. As a consequence, the decoupling temperature effectively increases from $T_F \approx m_i/10$ to $T_F \approx m_i$ as we move from stronger to weaker coupling (smaller to larger Λ_i). This implies that the contribution to $\Delta N_{\rm eff}$ from coupling to a given fermion is sensitive to the $g_*(T)$ curve in the vicinity of its mass m_i. We note that the contribution of the fermion itself is included in this change to $g_*(T)$ which means that there are more fermion-antifermion pairs present in the thermal bath when the axions decouple at higher temperatures. In all cases, we see a knee in the shape of the curve that is in good agreement with the equilibrium result for $T_F = m_i$, which is the approximate decoupling temperature when the axions just barely reach equilibrium.

Finally, as we decrease the coupling further (again, equivalent to increasing Λ_i), the production rate will eventually not reach $H(T)$ for any temperature T. Without coming into equilibrium, the number of axions is no longer tied to the number of photons and we see that the contribution to $\Delta N_{\rm eff}$ falls rapidly. Since the effective coupling is $\tilde{\varepsilon}_i = 2m_i/\Lambda_i$, this happens at smaller values of Λ_i for lighter fermions. As a consequence, the exponential falloff of the $\Delta N_{\rm eff}$ curves in Fig. 2.4 occurs in the order of increasing mass, i.e. lighter fermions lead to negligible $\Delta N_{\rm eff}$ at smaller Λ_i. The fact that the $\Delta N_{\rm eff}$ curves cross is another manifestation of the same physics.

We can gain further insights by comparing the shape of the curves relating ΔN_{eff} and the axion-fermion coupling strengths (left panel of Fig. 2.4) to the standard T_F-ΔN_{eff} curve of Fig. 2.1. This comparison is provided by the right panel of Fig. 2.4. Taking $T_F = m_i$, which is denoted by the triangles in the right panel, we see that the contributions to ΔN_{eff} are near, but slightly below the asymptotic values at strong couplings for most fermions. While this is in line with our expectations, it does not fully address how to translate the approximate inequality in (2.25) into a map between Λ_i and ΔN_{eff}. This is particularly noticeable for the coupling to muons, where the ΔN_{eff} curve begins to drop around $\Lambda_\mu = 10^6\,\text{GeV}$, but crosses $\Delta N_{\text{eff}}(T_F = m_\mu)$ only around $\Lambda_\mu = 10^7\,\text{GeV}$. More generally, the limitation of our qualitative estimates is that they do not entirely explain the shapes of the ΔN_{eff} curves.

The complication in relating the scales Λ_i to values of ΔN_{eff} is that these curves are really a combination of the axion production rate from the fermion involved in the coupling and the effective number of relativistic degrees of freedom from all the particles in the Standard Model, $g_*(T)$. We can separate these effects by removing the dependence on $g_*(T)$, as is shown in Fig. 2.6. The green curve shows the contribution to ΔN_{eff} if $g_*(T)$ was a constant over the relevant range of temperatures so that ΔN_{eff} would be a constant for any coupling reaching equilibrium. We see in Fig. 2.5 that $\Gamma(T) < H(T)$ for all temperatures T if $\Lambda_\mu > 1.5 \times 10^7\,\text{GeV} \approx 10^{7.2}\,\text{GeV}$, which is in good agreement with the value of Λ_μ at which the curve bends, indicating that the axions are never coming into equilibrium.

In contrast, even when we add the contribution to $g_*(T)$ from the muon (displayed by the red curve in Fig. 2.6), which changes between $T \gg m_\mu$ and $T \ll m_\mu$, the contribution to ΔN_{eff} begins to decrease around $\Lambda_\mu \approx 10^6\,\text{GeV}$ when the axion is still reaching equilibrium. We can understand this in terms of our simple estimate as follows: at large coupling strengths, $T_F \approx m_\mu/10$ means the muons are mostly annhilated at decoupling and do not contribute to $g_*(T_F)$. As we decrease the coupling, $T_F \to m_\mu$ so that the axions are diluted somewhat by the muon annihilation even though the axions come into thermal equilibrium.

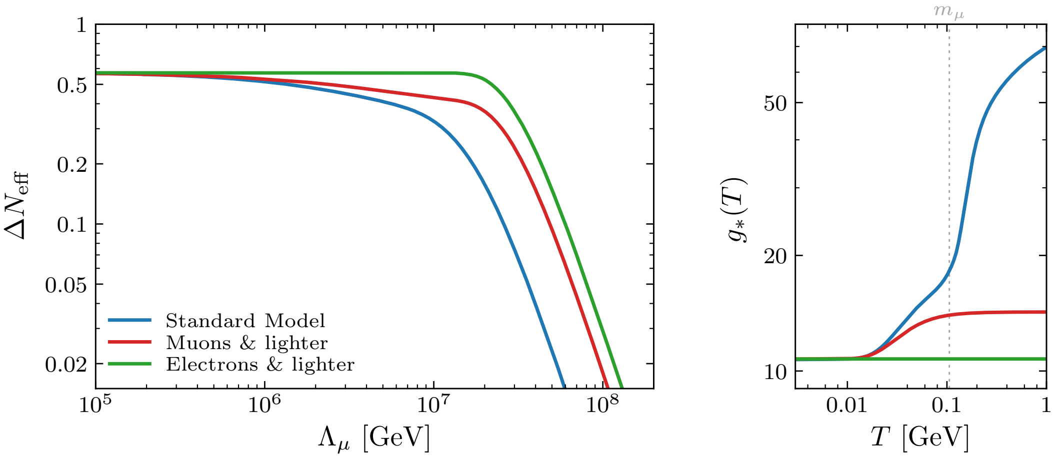

Figure 2.6. *Left*: Breakdown of the contribution to $\Delta N_{\rm eff}$ from the coupling to muons. The full Standard Model result is shown in blue, the red curve displays the predictions for a universe without the SM fermions that are heavier than the muon, i.e. a universe with photons, neutrinos, electrons and muons, and the green curve additionally removes the contribution to $g_*(T)$ from the muon. We see that the large change in $g_*(T)$ in the vicinity of $T = m_\mu$ due to the QCD phase transition has a significant impact on the resulting contribution to $\Delta N_{\rm eff}$. *Right*: The $g_*(T)$ curves underlying the calculation of $\Delta N_{\rm eff}(\Lambda_\mu)$ displayed in the left panel, matched by color.

Finally, when all the degrees of freedom of the Standard Model are included in $g_*(T)$ (blue curve in Fig. 2.6), we see a considerably larger suppression even when $\Lambda_\mu > 10^7\,$GeV because the muon mass lies on the boundary of the QCD phase transition during which the number of degrees of freedom is changing rapidly with temperature. This dependence on $g_*(T)$ depends significantly on the fermion mass and leads to the variety of shapes seen in Fig. 2.4. This is best illustrated by the coupling to the tau lepton, which has a mass that is close to the QCD phase transition which means that the $\Delta N_{\rm eff}$ curve effectively transitions from decoupling after the QCD phase transition to decoupling before the QCD phase transition as we increase Λ_τ.

2.5 Comparison with Other Probes

Cosmological constraints on axions and other pNGBs are particularly compelling as they are both easy to calculate and robust to much of the details of the model. This is largely due to thermal equilibrium which tells us the number of axions produced at a given temperature for any sufficiently large coupling. Yet, there are a wide variety of other probes of axions, both terrestrial and astrophysical, that have different strengths and weaknesses compared to cosmological probes. In this section, we will compare our results, especially our constraint on Λ_μ from Planck, to other probes of the same couplings. Of particular interest will be astrophysical constraints, such as from cooling of supernova SN 1987A, which are also the result of thermal axion production.

2.5.1 Stellar Cooling and SN 1987A

Astrophysical constraints on axion couplings [77, 78, 79] offer a useful foil for cosmological bounds. Stars also provide a controlled high-temperature environment in which the thermal production of a new light particle would be detectable. In this sense, astrophysical constraints are probing essentially the same physics as constraints from bounds on N_{eff}. It is therefore instructive to understand where differences arise and what the relative strengths of each probe are.

At a qualitative level, both probes are sensitive to large changes in the number of axions. In the case of astrophysical environments, the production of these particles is governed by

$$\frac{\mathrm{d}n_\phi}{\mathrm{d}t} = \Gamma_\phi^{(\star)}(T_\star)\left(n_\phi^{\text{eq}}(T_\star) - n_\phi\right), \tag{2.26}$$

where $T_\star$ is a temperature that is (mostly) fixed by the specific probe. This equation should be compared to the cosmological Boltzmann equation (2.13). In general, the cosmological interaction rate $\Gamma_\phi(T)$ and the astrophysical rate $\Gamma_\phi^{(\star)}(T)$ are related, but they can differ even at the same temperature due to the large chemical potentials present in astrophysical environments. This is particularly important for protons, neutrons and electrons. In contrast, muons (and the other heavier particles) are unstable and their abundance is primarily due to thermal production. As a result, we can treat $\Gamma_\phi(T) \approx \Gamma_\phi^{(\star)}(T)$ for our purposes.

At very weak coupling, the axions will escape the star after production, thus providing a new mechanism for energy to leave the system. If the number of axions produced in the timescale of observation $t_{\rm obs}$ is comparable to the number of photons,

$$\Gamma_\phi^{(\star)}(T_\star)\, n_\phi^{\rm eq}(T_\star)\, t_{\rm obs} \gtrsim n_\gamma(T_\star)\,, \tag{2.27}$$

then the energy loss due to axions is significant and would lead to detectable changes in the dynamics of the astrophysical system.

We are particularly interested in the constraints from SN 1987A because the high temperatures of the supernova can produce a large number of muons. The timescale relevant to axion cooling of the proto-neutron star is $t_{\rm obs} \approx 1\,{\rm s}$.[14] Of course, the temperature of the star depends both on the radius and time which therefore means that our estimate is necessarily approximate. Nevertheless, both the axion production and the total energy of the star are dominated by the hottest regions which therefore makes this local approximation a useful starting point. Furthermore, since $n_\phi^{\rm eq}(T_\star) \approx n_\gamma(T_\star)$ for light axions, forbidding significant cooling implies a constraint

$$\Gamma_\phi^{(\star)}(T_\star) < t_{\rm obs}^{-1} = 10^{-24.2}\,{\rm GeV} = 7\times 10^{-25}\,{\rm GeV}\,. \tag{2.28}$$

For a system in equilibrium at temperature T, we can easily determine this constraint from Fig. 2.7. We however notice that the production rate is very sensitive to the precise temperature. In particular, if we require that $\Gamma_\phi^{(\star)} = 1\,{\rm s}^{-1}$, then our bound is in the range of $\Lambda_\mu > 10^{7.7-8.5}\,{\rm GeV}$ for a temperature range of $25-45{\rm MeV}$. This is, of course, a reflection of the Boltzmann suppression of muons inside the supernova and, therefore, explains our exponential sensitivity to the temperature. These results are in agreement with the bounds found by a more detailed analysis presented in [59, 60]. Using results from a simulation with a mass-weighted, radially-averaged core temperature of $25\,{\rm MeV}$, they find $\Lambda_\mu > 10^{7.5}\,{\rm GeV}$. However, these bounds still depend sensitively on the specific simulation. Concretely, they find $\Lambda_\mu > 10^{8.0}\,{\rm GeV}$ in another simulation with

---

[14] A common description of supernovae suggests the cooling of the proto-neutron star lasts approximately ten seconds. We will conservatively take the relevant timescale to be $t_{\rm obs} \approx 1\,{\rm s}$ because approximating supernovae as constant temperature systems will break down as $t_{\rm obs} \to 10\,{\rm s}$.

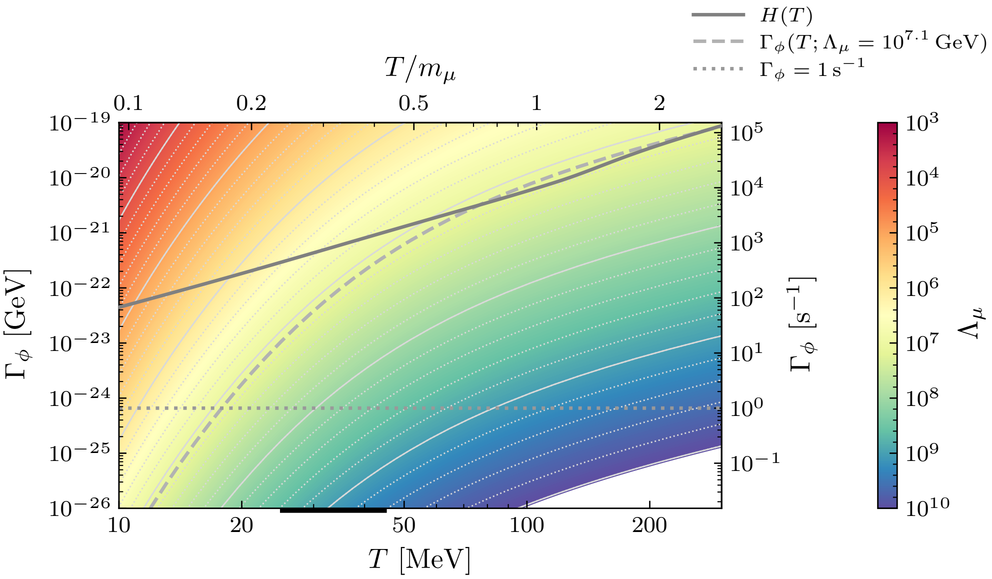

Figure 2.7. Production rate of axions and other pNGBs, Γ_ϕ, as a function of temperature T for different choices of the coupling to muons, Λ_μ. As in Fig. 2.5, which presents a larger range of temperatures, the solid gray line shows the Hubble rate $H(T)$ to indicate the region $\Gamma_\phi(T) > H(T)$ of efficient cosmological production and the dashed gray line indicates the interaction rate for the current Planck N_{eff} bound on the axion-muon coupling Λ_μ. This figure allows us to easily compare our cosmological bounds to astrophysical constraints which arise for an order one change in the number of photons during the observable timescale, namely $\Gamma_\phi^{(\star)}(T)\,t_{\text{obs}} > 1$. As SN 1987A is associated with $t_{\text{obs}} \approx 1\,\text{s}$, the inferred constraint arises approximately from the intersection of $\Gamma_\phi \approx \Gamma_\phi^{(\star)} = 1\,\text{s}^{-1}$ with the temperature in the core of the supernova. As an estimate of this temperature, the black horizontal line indicates the range of mass-weighted, radially-averaged temperatures in spherically-symmetric, one-dimensional simulations of approximately 25 – 45MeV [59].

mass-weighted, radially-averaged core temperature of 45 MeV. Importantly, the temperatures reached inside the supernova vary significantly both with position inside the core and between the simulations. As a result, we expect that the derived bounds will be exponentially sensitive to the details of the specific simulation used (see also [61] for more discussion).

From Figure 2.7, we can also compare the sensitivity of the cosmological constraint to the bound from SN 1987A. Comparing the production rate as a function of temperature T to the Hubble rate (shown by the solid gray line), we see that $\Lambda_\mu = 10^{7.2}\,\text{GeV}$ is the smallest value for which the cosmological production never becomes efficient, $\Gamma_\phi(T) > H(T)$. This is our approximate cosmological bound that essentially reproduces the exact bound from current Planck measurements of ΔN_eff of $\Lambda_\mu > 10^{7.1}\,\text{GeV}$. To compare this constraint to SN 1987A, we follow the dashed gray curve to $\Gamma_\phi = 1\,\text{s}^{-1}$ where it corresponds to the production in the supernova at $T = 18\,\text{MeV}$. In addition, we see that the temperature where $\Gamma_\phi = 1\,\text{s}^{-1}$ is nearly unchanged for the slightly smaller value of $\Lambda_\mu = 10^7\,\text{GeV}$, but with the major difference that axions are efficiently produced in the early universe over many decades in temperature. In consequence, it is more useful to treat cosmological and supernova bounds of axions and other pNGBs as complementary rather than redundant. While the constraint $\Lambda_\mu > 10^{7.5}\,\text{GeV}$ from SN 1987A is somewhat stronger, the current cosmological limit of $\Lambda_\mu > 10^{7.1}\,\text{GeV}$ (or the future CMB-S4 limit of $\Lambda_\mu > 10^{7.5}\,\text{GeV}$) can be mostly understood from equilibrium physics and is therefore quite robust. In contrast, the temperatures and dynamics inside a supernova are complex and the origin of the constraints are more uncertain (see also [80] for other potential limitations and uncertainties of the supernova-based constraints). For related reasons, the bounds in [59, 60] are theoretical constraints and not 95% c.l. limits, in contrast to the cosmological constraints derived from N_eff measurements.

2.5.2 Experimental Limits

We discussed the axion-muon coupling constraints from stellar cooling and SN 1987A in detail in the previous section. In the following, we provide a more general overview of existing bounds on couplings of axions and other pNGBs to matter and compare them to the limits that we derived in §2.4.2. We will keep our focus on the diagonal couplings and will not discuss the extensive list of existing bounds on off-diagonal couplings, in particular from SM particle decays (cf. e.g. [11, 81, 82, 83]).

Axion-Lepton Couplings

First, we will consider the couplings of axions to leptons. We focus on the axion-muon and axion-tau couplings since the ΔN_{eff}-based bounds are competitive in those cases, but will also briefly discuss the axion-electron interaction. Moreover, we are not considering model-dependent constraints that convert a bound on the coupling to electrons or photons to a limit on the other lepton couplings, e.g. by assuming these interactions to be universal or determined by a model such as DFSZ. Such bounds are generally stronger due to the tight bounds on Λ_e, but we prefer to consider the various axion couplings to be independent. This results in conservative and model-independent estimates applicable to any pNGB-fermion interaction, as mentioned above.

Coupling to Electrons.

The limit on the axion-electron interaction strength from white dwarf cooling is $\Lambda_e > 1.2 \times 10^{10}\,\mathrm{GeV}$ [84] (see also [85, 86, 87, 88, 89] for similar limits from stellar cooling and related discussions), which is considerably stronger than any limit that may be derived from upcoming cosmological ΔN_{eff} measurements. The reheating temperature-dependent freeze-out constraint from excluding $\Delta N_{\mathrm{eff}} = 0.027$ is $\Lambda_e \gtrsim 6 \times 10^7\,\mathrm{GeV}\sqrt{T_R/10^{10}\,\mathrm{GeV}}$ [32], while the current freeze-in constraint is given by $\Lambda_e > 2.5 \times 10^6\,\mathrm{GeV}$ which will improve to $5.4 \times 10^6\,\mathrm{GeV}$ with CMB-S4, cf. (2.16) and (2.22).[15] Having said that, the cosmological ΔN_{eff}-based bounds on Λ_e are less sensitive than astrophysical constraints to the physical environment where axions are produced. In addition, the environment in the early universe is quite different to the interiors of stars which means that cosmological probes can be an important complementary test [90].

---

[15] We note that BBN-based ΔN_{eff} measurements have been able to essentially close a small window in parameter space for $m_\phi \in [0.1, 1]\,\mathrm{MeV}$ [56].

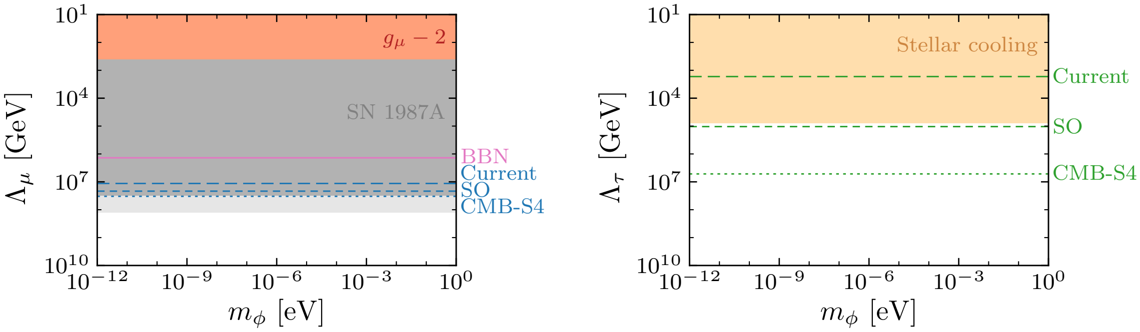

Figure 2.8. Comparison of existing and future constraints on the coupling between axions and muons (*left*) and tau leptons (*right*), respectively, as a function of the axion mass $m_\phi \leq 1\,\mathrm{eV}$, which is the mass range relevant for CMB (and large-scale structure) measurements of ΔN_{eff}. *Left*: The existing model-independent constraints for the coupling to muons come from measurements of the anomalous magnetic moment of the muon ($g_\mu - 2$), which would receive axion contributions at the loop level, and from the observed cooling rate of SN 1987A. For the latter, the dark (light) regions indicate the conservative (optimistic) bounds inferred by [59, 60] and, therefore, indicate the level of uncertainty in these constraints. The limits derived in this work from ΔN_{eff} measurements of BBN (cf. [56]), and current and future CMB experiments are complementary in nature. *Right:* The strongest, model-independent bound on the axion-tau interaction strength comes from the loop-induced coupling to electrons which is strongly constrained from stellar cooling of white dwarfs. The CMB-based limits will improve by two orders of magnitude each from Planck to the Simons Observatory and CMB-S4.

Coupling to Muons.

The main constraints on potential interactions of sub-MeV axions and muons have been derived from the cooling rate of supernova SN 1987A and from measurements of the anomalous magnetic moment of the muon, $g_\mu - 2$. While we discussed the former in detail in §2.5.1 (cf. [59, 60]), the latter allows to put laboratory bounds on the coupling strength since pNGBs contribute to $g_\mu - 2$ at the loop level. Following [91] and using the current difference between the measured and theoretically-predicted value of $a_\mu \equiv (g_\mu - 2)/2$, $\Delta a_\mu = [251 \pm 59] \times 10^{-11}$ [92, 93], at the lower 5σ limit (since the pNGB-induced contribution to the anomalous moment is negative), $\Delta a_\mu \geq -44 \times 10^{-11}$, we conservatively derive $\Lambda_\mu > 10^{2.6}$.[16]

In the left panel of Fig. 2.8, we compare these existing bounds to the current and future

---

[16] We note that the difference between the experimental value and the implied value from recent lattice calculations of the hadronic contribution to $g_\mu - 2$ is significantly smaller [94] and leads to a bound which is weaker by a factor of approximately two.

bounds studied in this work, see (2.15), (2.17), (2.20) and (2.23). For the supernova bound, we display a conservative and an optimistic estimate as derived in [59, 60] which differ in their model assumptions of the mass of the remnant neutron star in SN 1987A, with additional uncertainties possible from uncertainties in the equation of state at supernuclear densities. We clearly see the complementary nature of the displayed constraints, with current and future ΔN_{eff}-based bounds exploring the same parameter space and an exclusion of $\Delta N_{\mathrm{eff}} > 0.067$ corresponding to the conservative bound from SN 1987A.

Coupling to Tau Leptons.

In contrast to the interactions with electrons and muons, there do not appear to be model-independent, tree-level bounds on the diagonal coupling to tau leptons. However, such a coupling would induce an interaction with electrons at the loop level, i.e. the bounds on Λ_e can generally be translated into constraints on the couplings to the other leptons (and quarks) [11]. The strong bounds on the electron-axion coupling and the large masses of the third generation of SM fermions partly compensate the loop suppression which results in interesting constraints. Following [11], we infer $\Lambda_\tau \gtrsim 8 \times 10^4\,\mathrm{GeV}$ from the previously mentioned white dwarf cooling bound on Λ_e of [84] if the axion contribution to the induced coupling is dominantly proportional to m_τ^2. We compare this limit to our freeze-in constraints in the right panel of Fig. 2.8. While the current Planck bound is weaker than the loop-induced bound based on the Λ_e limit, near-term CMB experiments will strengthen this constraint, with the sensitivity of SO corresponding to this stellar bound and CMB-S4 being projected to improve upon it by about two orders of magnitude.

Overall, we notice that the cosmological constraints have the opposite strengths and weaknesses of the astrophysical bounds on axion couplings to leptons. Due to the maximum temperatures found in astrophysical settings, the implied sensitivities are much weaker for heavier leptons. In contrast, under plausible assumptions, our cosmological history reaches temperatures well above the masses of these leptons and, therefore, is also sensitive to the heavy leptons. In fact, cosmological observations can reach larger values of Λ_i for the heavier fermions because the effective coupling is proportional to their mass.

Axion-Quark Couplings

We now turn to the interaction between axions and the heavy quarks. Since current and near-term cosmological experiments are not sensitive enough to constrain these couplings,[17] we will instead estimate the required sensitivity to ΔN_{eff} to match the existing constraints. As for the lepton couplings, we will again focus on model-independent and diagonal couplings, but note that it is less clear in this case because the distinction between diagonal and off-diagonal constraints is only valid at leading order since quark flavors necessarily mix.

Coupling to Top Quarks.

The best model-independent constraints on the diagonal axion-top coupling arise from the loop-induced constraint based on the Λ_e limit, cf. [11]. Assuming that this loop contribution is dominated by m_t^2, e.g. when only coupling the axion to the right-handed top, we deduce $\Lambda_t \gtrsim 4 \times 10^9\,\mathrm{GeV}$. This corresponds to a contribution of $\Delta N_{\mathrm{eff}} \approx 0.005$ which is much smaller than the high-temperature thermal freeze-out contribution of 0.027 or near-term cosmological bounds on ΔN_{eff}.

Coupling to Bottom Quarks.

In the case of an independent axion-bottom interaction, we again follow [11] and compute the same loop-induced constraint as described for the interactions with tau leptons and top quarks: $\Lambda_b \gtrsim 2.0 \times 10^6\,\mathrm{GeV}$. Here, we assumed that the axion contribution to the electron coupling is dominated by m_b^2, which is the case if the axion only couples to the right-handed bottom, for instance. Given our conservative estimate of $\Delta N_{\mathrm{eff}}(\Lambda_b)$, we require a cosmological measurement that excludes $\Delta N_{\mathrm{eff}} \gtrsim 0.048$ to improve upon this bound, but we refer to Appendix A.2 for a discussion on the uncertainties of this estimate and potential implications for SO and CMB-S4.

---

[17] Note however our discussion in Appendix A.2 of the uncertainties in our calculation for the interactions with the charm and bottom quarks which still leave the possibility for such constraints.

Coupling to Charm Quarks.

For diagonal couplings to charm quarks, we follow the same argument while assuming that the loop contribution is dominated by m_c^2, e.g. by only coupling to the right-handed charm. In this way, we infer $\Lambda_c \gtrsim 1.4 \times 10^5\,\mathrm{GeV}$. When comparing this bound to the predictions based on our conservative estimate of $\Delta N_{\mathrm{eff}}(\Lambda_c)$, an improvement of this bound requires the exclusion of $\Delta N_{\mathrm{eff}} \gtrsim 0.048$. However, lattice QCD calculations may reveal that weaker bounds on ΔN_{eff} lead to the same bounds on the interaction strength with charm quarks, with some constraining power not only accessible for CMB-S4, but potentially even for the Simons Observatory (cf. Appendix A.2).

2.6 Conclusions

The high temperatures and densities of the early universe provide an ideal environment to test fundamental physics. Even for extremely weak couplings, new particles could be efficiently produced, potentially leaving a lasting imprint on cosmological observables. Axions and other pseudo-Nambu-Goldstone bosons provide a particularly compelling target as they are naturally light and would therefore leave a measurable imprint on cosmological observables via the effective number of relativistic species, N_{eff}.

In this paper, we calculated the predicted contributions to N_{eff} from axions and other pNGBs that are coupled to Standard Model fermions. We focused on the effectively marginal interactions that arise after electroweak symmetry breaking which can thermalize these particles beyond the Standard Model at low temperature. The axions eventually decouple when the temperature drops below the mass of the associated fermion. Since this happens below the electroweak scale, they contribute $\Delta N_{\mathrm{eff}} > 0.027$ to the radiation density in the early universe which makes them compelling targets for near-term surveys. Our main result is shown in Fig. 2.9, which provides a direct link between the measurement of N_{eff} and limits on the coupling to SM fermions.

This work includes an improved calculation of the thermal axion production rates. The

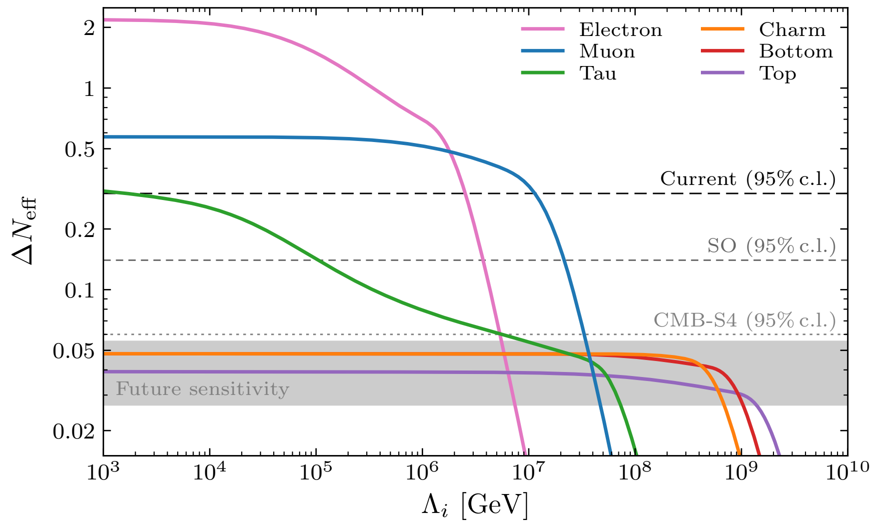

Figure 2.9. Contribution to the radiation density as parameterized by $\Delta N_{\rm eff}$ as a function of axion coupling strength Λ_i for different Standard Model fermions ψ_i (cf. Fig. 2.4). From this figure, we can translate current and future constraints on $N_{\rm eff}$ (cf. Fig. 2.1) into the equivalent bounds on Λ_i for any couplings to matter.

described method for calculating the thermal averages including the full quantum statistics is also relevant to other production rate calculations, including axion production at high temperatures. In addition, the same production rates calculated in this paper also appear in astrophysical constraints on axions. We are therefore able to use the common origin of axion production to compare the strengths and weaknesses of the cosmological and astrophysical bounds on axion couplings to matter, and saw that the cosmological production of axions is essentially determined by dimensional analysis as are the associated bounds.

Our results are particularly important in the context of ongoing cosmic surveys improving the measurement of $N_{\rm eff}$, which are on the precipice of measuring the energy density of a single scalar field decoupling prior to the QCD phase transition. Concretely, CMB-S4 is expecting to exclude $\Delta N_{\rm eff} > 0.060$ at 95% c.l. [52] and could be improved in combination with a number of large-scale structure surveys [44]. The axion couplings to matter discussed in this paper illustrate that this is a particularly compelling level of sensitivity. As several Standard Model fermions have masses around 1 GeV, the coupling of axions to these fermions naturally contributes to $\Delta N_{\rm eff}$

at this level. Furthermore, the decoupling temperature depends both on the interaction strength and the mass of the fermion, resulting in a range of ΔN_{eff}, even when the axion reaches thermal equilibrium. As a result, there are both numerous thresholds of ΔN_{eff} at this sensitivity and opportunities to continuously improve our understanding of Standard Model couplings of axions and other pseudo-Nambu-Goldstone bosons with the depth of these surveys.

Acknowledgments

The authors thank Nathaniel Craig, Peizhi Du, Peter Graham, Marilena LoVerde, Gustavo Marques-Tavares, Joel Meyers and Surjeet Rajendran for helpful discussions. The authors were supported by the US Department of Energy under Grants DE-SC0009919 and DE-SC0019035. B. W. also acknowledges support from the Simons Foundation Modern Inflationary Cosmology Initiative under Grant SFARI 560536. The completion of this work by B. W. was partially supported by a grant from the Simons Foundation and the hospitality of the Aspen Center for Physics, which is supported by National Science Foundation Grant PHY-1607611. We acknowledge the use of `FeynMP` [95], `IPython` [96] and `RunDec` [69], and the Python packages `Matplotlib` [97], `Numba` [98], `NumPy` [99] and `SciPy` [100].

Chapter 2, in full, is a reprint of the material as it appears in Cosmological Implications of Axion-Matter Couplings, Green, Daniel; Guo, Yi; Wallisch, Benjamin, JCAP, 2022. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Light Fields during Inflation from BOSS and Future Galaxy Surveys

3.1 Synopsis

Primordial non-Gaussianity generated by additional fields present during inflation offers a compelling observational target for galaxy surveys. These fields are of significant theoretical interest since they offer a window into particle physics in the inflaton sector. They also violate the single-field consistency conditions and induce a scale-dependent bias in the galaxy power spectrum. In this paper, we explore this particular signal for light scalar fields and study the prospects for measuring it with galaxy surveys. We find that the sensitivities of current and future surveys are remarkably stable for different configurations, including between spectroscopic and photometric redshift measurements. This is even the case at non-zero masses where the signal is not obviously localized on large scales. For realistic galaxy number densities, we demonstrate that the redshift range and galaxy bias of the sample have the largest impact on the sensitivity in the power spectrum. These results additionally motivated us to explore the potentially enhanced sensitivity of Vera Rubin Observatory's LSST through multi-tracer analyses. Finally, we apply this understanding to current data from the last data release of the Baryon Oscillation Spectroscopic Survey (BOSS DR12) and place new constraints on light fields coupled to the inflaton.

3.1.1 Contribution of the author

The author has contributed in developing the main formalism and in writing the paper. He has performed the numerical calculations and the analysis together with Jiashu Han.

3.2 Introduction

The statistics of the primordial density perturbations offer a window into the dynamics of the very early universe [101, 102, 103], in particular the inflationary epoch [104, 105]. While current maps of the universe are consistent with purely Gaussian fluctuations [106], non-Gaussian correlation functions encode the particles and interactions relevant to the origin of structure [107, 108, 105]. Inflation predicts a lower bound on the amount of primordial non-Gaussianity (PNG) due to gravity alone that is approximately two orders of magnitude below current constraints [109]. On the other hand, signals that are large enough to be detected in the next generation of surveys arise in many models, including examples with only Planck-suppressed interactions [110].

Due to the limited number of modes remaining to be measured to the cosmic variance limit in the cosmic microwave background (CMB), surveys of the large-scale structure (LSS) of the universe present the best opportunity to improve our understanding of the primordial statistics [105, 111, 112, 113, 102]. With the benefit of three-dimensional information, even current surveys have the raw statistical power to compete with the CMB [114]. Unfortunately, for generic non-Gaussian correlation functions, late-time nonlinearities make many modes inaccessible and present a serious obstacle to our progress [115, 116, 117, 118]. In some circumstances, however, the manifestation of the primordial signal is robust to nonlinear physics and can be measured reliably [101].

The fluctuations of extra fields beyond the inflaton present such an opportunity. In the absence of these fields, correlation functions of matter and galaxies are subject to the constraints of the single-field consistency conditions [119, 120, 121]. In practice, these conditions are similar

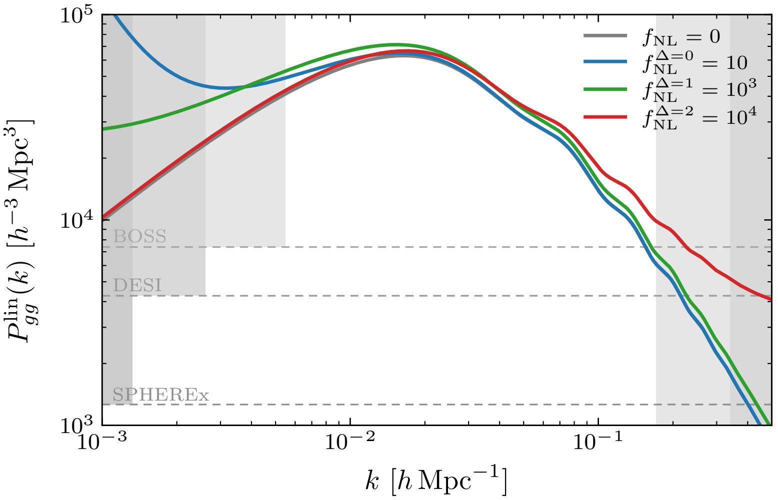

Figure 3.1. Illustration of the effect of the scale-dependent bias $b_{\rm NG}^{\Delta}$ from a non-zero non-Gaussian amplitude $f_{\rm NL}^{\Delta}$ on the linear galaxy power spectrum. We see that the enhancement of power shifts from large to small scales as we increase the exponent Δ (see §3.3.2 for its definition and details). The horizontal dashed lines indicate the effective shot noise level for BOSS, DESI and SPHEREx after scaling them to the displayed redshift $z = 0$ and linear bias $b_1 = 1.6$. The gray-shaded regions on large scales indicate the wavenumbers below the minimum wavenumber of these surveys, $k < k_{\rm min}$, computed based on their entire spherical volume. The gray-shaded regions for large wavenumbers indicate the regimes where the scales exceed the nonlinear scale, $k > k_{\rm NL}$, for the maximum redshift of BOSS and DESI, respectively.

to the equivalence principle [122] and ensure that correlations are determined by the curvature of spacetime (i.e. derivatives of the metric fluctuations). Light fields break these conditions and can lead to enhanced long-distance correlations [123, 124, 125, 126, 127, 128]. These apparent violations of the equivalence principle cannot arise from nonlinear gravity, and can therefore be distinguished from astrophysical and gravitational sources of non-Gaussianity.

From the perspective of both data analysis and survey design, the unique advantage of these non-Gaussian signals is that they can be observed in the two-point statistics of galaxies. Violations of the single-field consistency conditions introduce couplings between long- and short-wavelength modes. The short-wavelength modes control the number density of galaxies, which are then non-trivially correlated on large scales in the presence of this type of non-Gaussianity. The resulting (non-local) change to the shape of the galaxy power spectrum, shown in Figure 3.1, is known as scale-dependent bias [129] and has been searched for in many existing datasets (see e.g. [130, 131, 132, 133, 134, 135, 136, 137, 138, 139]).

The case of local primordial non-Gaussianity [140], for which the Gaussian and non-Gaussian Newtonian potentials φ and Φ are related by

$$\Phi(\vec{x}) = \varphi(\vec{x}) - f_{\mathrm{NL}}^{\mathrm{loc}} \varphi^2(\vec{x}) + \dots , \tag{3.1}$$

has been very well studied in the literature, both theoretically and observationally. Local PNG arises in the presence of multiple massless fields, such that φ is an isocurvature mode during inflation, but is converted to an adiabatic mode at later times [123, 124, 125]. Future measurements are expected to reach particularly interesting thresholds for the physics of multi-field inflation [115, 141, 142, 105].

The contribution to primordial non-Gaussianity from massive particles is much less studied observationally. It is however a classic signature of quasi-single-field inflation [126, 127, 128], which is also known as cosmological collider physics [143]. Importantly, these models are compelling targets for current and future surveys. In spite of this, little has been known about the optimal survey strategy to search for these relics from inflation up to now.

In this paper, we will examine how future surveys can best constrain these non-Gaussian signals using galaxy power spectra. For local PNG (massless fields), the signal is dominated by the largest scales and most observational strategies are designed accordingly [144, 145]. On the other hand, the signals of massive fields may arise at large or small scales (see Fig. 3.1), which means that the characteristics which best constrain these models are less clear a priori [146, 147, 148]. We therefore study this signal in detail, and investigate how redshift coverage, biasing, multi-tracer techniques and target selection impact forecasts over the full range of scaling laws induced by these particles.

We find that target selection, specifically finding highly biased objects, is the largest factor in driving current and future measurements of this effect when assuming a fixed universality relation. While large volumes help to increase the number of modes, especially in spectroscopic surveys, we observe that the benefits of going to higher redshifts are more driven by the large biases of high-redshift objects and only secondarily by the increased volume. In photometric surveys, the larger number density of galaxies also enables a significant increase in sensitivity through sample-variance cancellation. This requires splitting the sample according to the bias, which we explore in the context of Vera Rubin Observatory's Legacy Survey of Space and Time (LSST).

Using these insights, we explore current constraints on the scale-dependent bias of galaxies using the BOSS DR12 dataset [149] for the full range of scaling behaviors of light inflationary fields. We measure the contributions to the power spectrum for one scaling exponent (PNG shape) at a time and compute a correlation matrix to extrapolate between the discrete choices of scaling dimensions (masses). While our inferred constraints are less sensitive to the PNG amplitude than those from Planck, the CMB constraints for non-zero masses are driven by bispectra in equilateral configurations rather than the scaling behavior in the squeezed limit probed by the galaxy power spectrum. The only previous constraint from LSS data in this regime was made in [150]. At the same time, we forecast that galaxy power spectrum measurements in future surveys will exceed the sensitivity of Planck to these light fields for a sizable range of their masses.

This paper is organized as follows: In Section 3.3, we review the relevant theoretical background for our non-Gaussian signal. In Section 3.4, we present forecasts for a wide range of model parameters and experimental configurations. Our goal is to identify the choices that most directly impact the sensitivity to PNG beyond the local type. In Section 3.5, we apply our understanding of the signal to BOSS DR12 galaxy clustering data and present new constraints on non-Gaussianity from light fields. In Section 3.6, we discuss the role of multi-tracer analyses and astrophysical effects on the forecasts. We conclude in Section 3.7. A set of appendices contains technical details on the modeling of the galaxy power spectrum and the survey specifications (Appendix B.1), and a discussion on our ability to measure the scaling behavior of the non-Gaussian signature from the galaxy power spectrum (Appendix B.2).

3.3 Light Fields and Galaxies

Our goal in this paper is to explore the signal of additional fields that manifest themselves in the galaxy power spectrum via the scale-dependent bias. This section reviews the necessary background to understand the signal. Readers who are familiar with the scale-dependent bias due to general forms of primordial non-Gaussianity (i.e. beyond the local type) may proceed to Section 3.4.

3.3.1 Galaxy Power Spectrum

The formation of structure in the universe is driven by the growth of density fluctuations in the dark matter [151]. The evolution of the density contrast Fourier mode in redshift space, $\delta_m(\vec{k},z) \equiv \delta\rho_m(\vec{k},z)/\bar{\rho}_m(z)$, can be solved at linear order to give

$$\delta_m(\vec{k},z) = \frac{2k^2 T(k) D(z)}{3\Omega_m H_0^2}\Phi(\vec{k}) \equiv k^2 \mathscr{T}(k,z)\Phi(\vec{k})\,, \tag{3.2}$$

where $k = |\vec{k}|$, $T(k)$ is the transfer function defined such that $T(k \to 0) \to 1$, and $D(z)$ is the linear growth factor normalized as $D(z) = 1/(1+z)$ during matter domination. The primordial Newtonian potential $\Phi(\vec{k})$ encodes the primordial density fluctuations generated during inflation which are concretely of the form

$$\langle \Phi(\vec{k})\Phi(\vec{k}')\rangle = \frac{9}{25}\frac{A_{\rm s}}{k^3}k^{n_{\rm s}-1}(2\pi)^3\delta_D(\vec{k}+\vec{k}')\,, \tag{3.3}$$

where δ_D is the Dirac delta function and we used $\Phi = -\frac{3}{5}\zeta$, with curvature fluctuation ζ, for modes that re-entered the horizon during matter domination. The linear matter power spectrum is therefore given by

$$P_{\text{lin}}(k) = k^4 \mathscr{T}(k,z)^2 \frac{A_\text{s}}{k^{3-(n_\text{s}-1)}} = \frac{4}{25} \frac{A_\text{s} D(z)^2}{\Omega_m^2 H_0^4} T(k)^2 k^{n_\text{s}} \,. \tag{3.4}$$

The nonlinear matter power spectrum $P_m(k)$ follows the linear power spectrum on large scales, $P_m(k) \approx P_{\text{lin}}(k)$ for $k \ll k_{\text{NL}}$, for some scale $k_{\text{NL}} \approx 0.1\, h\,\text{Mpc}^{-1}$ at $z = 0$. Gravitational evolution corrects this behavior on smaller scales, $k \gtrsim k_{\text{NL}}$.

We will generally work with galaxies in the regime $k < k_{\text{NL}}$. The nonlinear effects that are relevant in this regime can be described solely in terms of the bias expansion [152, 153]. In its most general form, the bias expansion is simply the assumption that galaxy formation is a local process:[1]

$$\delta_g(\vec{x}) = \sum_i b_{\mathscr{O}_i} \mathscr{O}_i(\vec{x}) \,, \tag{3.5}$$

where the operators $\mathscr{O}_i(\vec{x})$ are any locally measurable quantities. Since we typically assume that the matter controls the formation of galaxies, these are usually local products of $\delta_m(\vec{x})$ or the tidal tensor [154], $s_{ij} = \left(\frac{\nabla_i \nabla_j}{\nabla^2} - \frac{1}{3}\delta_{ij}\right) \delta_m(\vec{x})$,

$$\mathscr{O} \in \{\delta_m, \nabla^2\delta_m, \delta_m^2, s_{ij}^2, \delta_m^3, \delta s_{ij}^2, \text{Tr}[(\Pi^{[1]})^3]|^{(3)}, \text{Tr}[\Pi^{[1]}\Pi^{[2]}]|^{(3)}, \cdots\} \,, \tag{3.6}$$

where the tensors $\Pi_{ij}^{[n]}$ are defined in [152].[2] In principle, there are an infinite number of operators to consider. At any fixed accuracy, we however understand this as an expansion in powers of the small density contrast, $\delta_m^n \ll 1$, and gradients $R^2\nabla^2 \ll 1$, for some fixed scale R. To model the galaxy power spectrum at one-loop order, we consider the following set of operators (up to third

---

[1] Technically, biasing is best understood as a local process in Lagrangian space. The non-locality in Eulerian space is captured by the Zel'dovich approximation [153].

[2] The superscript (3) indicates that the operator includes only terms up to third order in perturbation theory.

order in δ_m) in the bias expansion:

$$\delta_g(\vec{x}) = b_1 \delta_m(\vec{x}) + b_{\nabla^2} R^2 \nabla^2 \delta_m(\vec{x}) + b_{\delta^2}(\delta_m^2(\vec{x}) - \sigma^2) + b_{s^2}(s^2(\vec{x}) - \langle s^2 \rangle) \\ + b_{\Pi\Pi^{[2]}} \mathrm{Tr}[\Pi^{[1]}\Pi^{[2]}]|^{(3)}(\vec{x}) + \cdots , \tag{3.7}$$

with $\sigma^2 = \langle \delta_m^2 \rangle$. The coefficients $b_{\mathcal{O}_i} = b_1, b_{\nabla^2}, b_{\delta^2}, b_{s^2}, b_{\Pi\Pi^{[2]}}, \cdots$, which we refer to as bias parameters, are constants that are determined by the details of galaxy formation and evolution. Put differently, we parameterize the complex physics of galaxies by these coefficients.

On large scales and for initial conditions set by single-field inflation, this is the complete list of operators for practical purposes on large scales. Because the fluctuations are necessarily adiabatic throughout cosmic history, the evolution of the universe is controlled by a single statistical quantity, which we have chosen to represent in terms of Φ and δ_m. On smaller scales, the evolution of different components of the universe can however lead to some new types of biasing terms.

The redshift evolution of the bias parameters, particularly $b_1(z)$, are important for projecting the sensitivity of future surveys. If (proto-)galaxies form at a high redshift and primarily evolve with the expansion of the universe, then we expect that the large-scale comoving galaxy power spectrum remains constant, $P_{gg}(k,z) \approx P_{gg}(k,z')$. In other words, the power spectrum is due to the inhomogeneities in the density field at the time they are formed, as one would find in Lagrangian biasing. This holds if

$$b_1(z) = D(z)^{-1} b_1(z=0) , \tag{3.8}$$

which is a common simplifying assumption in many forecasts. Changes to the galaxy sample through mergers, for example, can however lead to behavior that differs significantly from Lagrangian biasing. When possible, we will therefore use the biases for surveys that were defined by target selection, but the assumption of this type of time evolution is also important in the context of multi-tracer analyses discussed in Section 3.6.

Given the bias expansion for the galaxy overdensity, we can write down our model for the galaxy power spectrum after including a few additional effects. First, we work in Fourier space by transforming from real-space coordinates $\vec{x}$ to (k, μ), where μ is the cosine between the wavevector $\vec{k}$ and the line-of-sight direction. (Note that the composite operators $\mathcal{O}_i(\vec{x})$ are simply related to $\tilde{\mathcal{O}}_i(k, \mu)$ by a single Fourier integral.) Second, there are redshift-space distortions which arise from the peculiar velocities of galaxies. We account for the Kaiser effect by adding $\Delta b(\mu, z) = f(z)\mu^2$, with the linear growth rate $f \equiv \mathrm{d}\log D/\mathrm{d}a$, to the linear bias b_1 [155]. Collecting all terms in the bias expansion at linear order in $\delta_m(\vec{k})$ therefore results in

$$b(k, \mu, z) = b_1(z) + f(z)\mu^2 + \sum_{n>0} b_{k^{2n}} (kR_*)^{2n}\,, \tag{3.9}$$

where the last term is the gradient biases in Fourier space with the comoving Lagrangian radius R_* of the halos of interest. Third, we describe gravitational nonlinearities in the very mildly nonlinear regime via the one-loop terms P_{22} and P_{13} in standard Eulerian perturbation theory. Finally, the leading-order stochastic contribution is given by P_{ε_0}. Putting it all together, we therefore model the theoretical anisotropic power spectrum for two galaxy samples A and B as follows:

$$\begin{aligned}
P^{\mathrm{th}}_{AB}(k, \mu, z) = {} & b^A(k) b^B(k) [P_{\mathrm{lin}}(k) + P_{22}(k) + P_{13}(k)] + P_{\varepsilon_0} \delta_{AB} \\
& + \left[b^A(k) b^B_{\delta^2} + b^A_{\delta^2} b^B(k) \right] P_{\delta^2}(k) + \left[b^A(k) b^B_{s^2} + b^A_{s^2} b^B(k) \right] P_{s^2}(k) \\
& + b^A_{\delta^2} b^B_{\delta^2} \left(P_{\delta^2\delta^2}(k) - 2\sigma^4 \right) + \left(b^A_{\delta^2} b^B_{s^2} + b^A_{s^2} b^B_{\delta^2} \right) \left(P_{\delta^2 s^2}(k) - \tfrac{4}{3}\sigma^4 \right) \\
& + b^A_{s^2} b^B_{s^2} \left(P_{s^2 s^2}(k) - \tfrac{8}{9}\sigma^4 \right) + \tfrac{3}{2} \left[b^A(k) b^B_{\Pi\Pi^{[2]}} + b^A_{\Pi\Pi^{[2]}} b^B(k) \right] P_{\Pi\Pi^{[2]}}(k)\,,
\end{aligned} \tag{3.10}$$

where we have suppressed the dependence on μ and z on the right-hand side, explicitly define all loop and bias contributions in Appendix B.1, and δ_{AB} is the Kronecker delta. To illustrate these

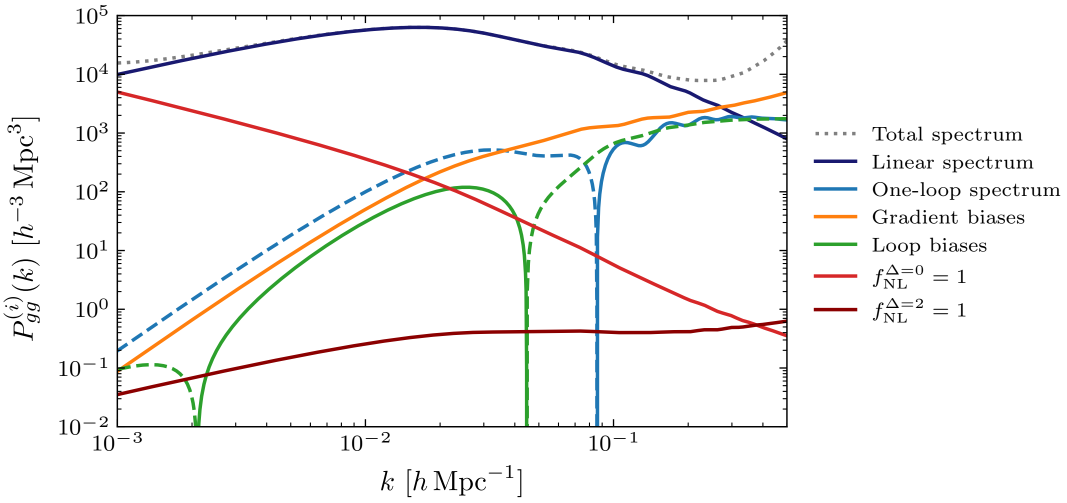

Figure 3.2. Contributions to the galaxy power spectrum $P_{gg}^{(i)}(k)$ at $z = 0$ from the nonlinear, bias and non-Gaussian terms. The linear bias is chosen to be $b_1 = 1.6$ and the gradient biases $b_{k^{2n}}$ (with $n \leq 2$) are taken to be unity. The bias parameters b_{s^2} and $b_{\Pi\Pi^{[2]}}$ are calculated in the Lagrangian local-in-matter-density model [152], and b_{δ^2} is calculated using the halo simulation fit of [156]. For comparison, we already include the scale-dependent bias from general PNG, which we introduce in §3.3.2. The non-Gaussian parameter f_{NL}^{Δ} is set to unity for both cases of $\Delta = 0$ and $\Delta = 2$. The dashed lines indicate negative contributions.

terms, we display the individual contributions in Fig. 3.2.

In practice, we compute the linear matter power spectrum $P_{\mathrm{lin}}(k)$, which underlies all contributions of the theoretical galaxy power spectrum $P^{\mathrm{th}}(k)$, with the Boltzmann solver CAMB [157] using a fiducial ΛCDM cosmological model based on the Planck 2018 TT, TE, EE + lowE + lensing + BAO best-fit cosmology [41]. In Table 3.1, we list the fiducial values of the ΛCDM, non-Gaussian and bias parameters. We set the fiducial values of the three loop biases b_{δ^2}, b_{s^2} and $b_{\Pi\Pi^{[2]}}$ to 0 instead of using the Lagrangian local-in-matter-density model [152] and fit results from halo simulations [156], as we did in Fig. 3.2, since we cover a larger redshift range than the fitted bias relations were derived from. This choice should have minimal impact on the results of this paper because the forecasted constraints on the loop biases are large enough for both choices of fiducial values to be consistent with each other.

Finally, to relate the theoretical galaxy power spectrum $P_{AB}^{\mathrm{th}}(k,\mu,z)$ to the observed galaxy power spectrum $P_{AB}^{\mathrm{obs}}(k,\mu,z)$, we account for two observational effects that are present in any

Table 3.1. Parameters of the fiducial ΛCDM model, based on the Planck 2018 best-fit cosmology [41], with the sum of neutrino masses $\sum m_\nu = 0.06\,\mathrm{eV}$, and the biasing model employed in our forecasts. Except for the optical depth, we vary these parameters in our forecasts unless stated otherwise.

Parameter	Description	Fiducial value
ω_b	Physical density of baryons $\omega_b \equiv \Omega_b h^2$	0.02242
ω_c	Physical density of cold dark matter $\omega_c \equiv \Omega_c h^2$	0.11933
$100\theta_s$	Angular size of the sound horizon at recombination	1.04119
τ	Optical depth due to reionization	0.0561
$\ln(10^{10}A_\mathrm{s})$	Logarithm of the primordial scalar amplitude	3.047
n_s	Scalar spectral index	0.9665
f_NL^Δ	Non-Gaussian amplitude	0
b_1	Linear bias	$b_1(z)$
$b_{k^{2n}}$	Gradient biases ($n \leq 2$)	0
b_{δ^2}	Quadratic bias	0
b_{s^2}	Tidal bias	0
$b_{\Pi\Pi^{[2]}}$	Evolution bias	0

real survey: redshift errors and the Alcock-Paczynski effect. Redshift errors are observational uncertainties in the measurement of galaxy redshifts, $\sigma_{z0}(1+z)$, which we model as an exponential suppression of the power spectrum along the line of sight,

$$f_{\sigma_z}^A(k,\mu,z) = \exp\left\{-k^2\mu^2[\sigma_{z0}^A(1+z)]^2/H^2(z)\right\}\,, \tag{3.11}$$

where σ_{z0}^A is the root-mean-square redshift error of sample A at $z=0$ [158, 159]. (We neglect this effect for spectroscopic surveys by setting $\sigma_{z0}^A \equiv 0$, i.e. $f_{\sigma_z}^A = 1$.) The Alcock-Paczynski effect arises when the true cosmology differs from the fiducial cosmology that is used to convert the measured angular positions and redshifts of the galaxies to comoving wavenumbers in Fourier space. The effect of this mapping on the density contrast is captured by [160],

$$f_{AP}(z) = \frac{D_A^\mathrm{fid}(z)^2}{D_A(z)^2}\frac{H(z)}{H_\mathrm{fid}(z)}\,, \tag{3.12}$$

where $D_A(z)$ is the angular diameter distance to redshift z. To summarize, our model for the observed galaxy power spectrum in redshift space therefore is

$$P_{AB}^{\mathrm{obs}}(k,\mu,z) = f_{AP}(z)\sqrt{f_{\sigma_z}^{A}(k,\mu,z)f_{\sigma_z}^{B}(k,\mu,z)}\,P_{AB}^{\mathrm{th}}(k,\mu,z)\,, \tag{3.13}$$

where $P_{AB}^{\mathrm{th}}(k,\mu,z)$ is given by (3.10).

3.3.2 Scale-Dependent Bias

The conventional description of biasing applies when the initial conditions are Gaussian. Gaussian initial conditions do not correlate modes of different scales, i.e. the collapse of small-scale density fluctuations to form halos in different regions is determined by locally observable quantities, such as the matter density, that vary over cosmological distances due to the long-wavelength fluctuations. On the other hand, mode coupling in the initial conditions due to primordial non-Gaussianity can introduce long-range correlations that are purely related to the statistics of the small-scale fluctuations themselves. When interpreted through the bias expansion, this mode coupling may appear to be non-local (scale-dependent) bias [129, 161, 162, 163] and/or long-range stochastic bias [164]. We will focus on the former since the latter is typically a subdominant contribution to the signal.

The coupling of long and short modes is described by the squeezed or collapsed limits of a non-Gaussian correlation function. In single-field inflation, the bispectrum is constrained by the single-field consistency conditions to take the following form in the squeezed limit [165, 119]:

$$\lim_{\vec{k}_1\to 0}\left\langle \zeta(\vec{k}_1)\zeta(\vec{k}_2)\zeta(\vec{k}_3)\right\rangle = -\left[(n_{\mathrm{s}}-1)+O(k_1^2/k_3^2)\right]P_\zeta(k_1)P_\zeta(k_3)\,(2\pi)^3\delta_D\Big(\textstyle\sum_i \vec{k}_i\Big)\,. \tag{3.14}$$

Furthermore, the leading term, which corresponds to $f_{\mathrm{NL}}^{\mathrm{loc}} = -\frac{5}{12}(n_{\mathrm{s}}-1)$, is unphysical [166, 167]. The leading physical term is suppressed by $O(k_1^2/k_3^2)$ and is due to the coupling of the short modes at horizon crossing to the curvature of the universe due to long-wavelength modes [168], which is typically proportional to the amplitude of equilateral non-Gaussianity, $f_{\mathrm{NL}}^{\mathrm{eq}}$.

In the presence of additional fields, the short-scale power can depend on the long-wavelength values of these fields and not just derivatives of the metric fluctuations [169]. For an additional massless field, the late-time Newtonian potential Φ may be nonlinearly related to a light (isocurvature) field χ during inflation, $\Phi(\vec{x}) = \chi(\vec{x}) - f_{\rm NL}^{\rm loc}\chi^2(\vec{x})$. For a massive field, we have in the superhorizon limit

$$\lim_{\vec{k}_1\to 0} \chi(\vec{k}_1, t) \approx c\frac{H}{k_1^{3/2}}\left(\frac{k_1}{aH}\right)^{\Delta}, \tag{3.15}$$

with a constant c. If χ alters the power spectrum of the short modes at horizon crossing, $k_2 = aH$, then we have a contribution to the bispectrum of the form

$$\lim_{\vec{k}_1\to 0} \left\langle \zeta(\vec{k}_1)\zeta(\vec{k}_2)\zeta(\vec{k}_3)\right\rangle' = f_{\rm NL}^{\Delta}\left(\frac{k_1}{k_2}\right)^{\Delta} P(k_1)P(k_2), \tag{3.16}$$

where we defined $\langle\ldots\rangle = (2\pi)^3\delta_D(\sum_i \vec{k}_i)\,\langle\ldots\rangle'$. To have such a contribution to the bispectrum, we also need χ to mix with ζ. Even without a mixing term, this same effect however appears in the collapsed limit of the trispectrum. For light fields, $\Delta = 3/2 - \sqrt{9/4 - m^2/H^2}$ such that $m = 0$ leads to local PNG ($\Delta = 0$) as expected. For $m/H > 3/2$, $\Delta = \frac{3}{2} \pm i\nu$, with $\nu = \sqrt{m^2/H^2 - 9/4}$, is complex in which case the complex conjugate[3] also contributes which results in the bispectrum being real valued, as required. More generally, it is possible to construct models with a wider range of real and complex values of Δ than those generated by a single massive field [170, 171, 172, 173].

In single-field inflation, physical mode coupling must involve derivatives of the metric fluctuations [166], whose leading behavior in the soft limit corresponds to $\Delta = 2$ [174]. One might naturally be surprised to find that the $m^2 \to \infty$ limit of Δ for a massive field does not yield this value of single-field inflation. It was nicely explained in [143] that $\Delta \approx 3/2 + im/H$

---

[3] For heavy fields, there are two solutions of the long-wavelength behavior which scale as Δ and its complex conjugate $\Delta^\star$. We note that there is also a second solution $\Delta' = (3 - \Delta)$ for light fields which is however suppressed relative to Δ.

is precisely what is expected from the single-particle wavefunction of a massive particle in an expanding universe in the high-mass limit. This contribution to the squeezed limit arises from the physical production of these particles. This particle production is Boltzmann suppressed in some simple models [175], $f_{\rm NL}^{\Delta} \propto {\rm e}^{-\pi m/H}$, but can be enhanced depending on the nature of the interaction [176]. In addition, very massive fields will also produce a sub-leading contribution to the squeezed limit with $\Delta = 2$. This contribution arises from the virtual exchange of a massive field and is equivalent to a purely local interaction (i.e. to integrating out the massive field). It may be surprising that a local term will produce non-local scale-dependent bias. When $T(k) \to 1$, $\Delta = 2$ is indeed a local term and it was long thought that $f_{\rm NL}^{\rm eq}$ could not be measured via scale-dependent bias for this reason (see e.g. [163]). However, because of the evolution of matter after horizon entry ($T(k) \neq 1$), local interactions during inflation are distinguishable from local processes in structure formation at late-times [148], i.e. equilateral PNG can be measured in the power spectrum via scale-dependent bias.

In the conventional picture of galaxy biasing (see §3.3.1), the density contrast of galaxies is determined by a long list of composite operators in the initial density field. The origin of the unique signal of PNG in galaxies is that the behavior of these composite operators becomes a proxy for the light fields during inflation. For example, consider the impact of structure formation from the local amplitudes of fluctuations on a scale R,

$$\sigma_R(\vec{x}) = \int {\rm d}^3 k\, {\rm e}^{-{\rm i}\vec{k}\cdot\vec{x}} \int \frac{{\rm d}^3 p}{(2\pi)^3} \zeta(\vec{p})\zeta(\vec{k}-\vec{p})F(pR)\,, \tag{3.17}$$

for some filter function $F(x)$. This is the local variance of the primordial fluctuations which varies from place to place. Importantly, this variance is correlated with the long-wavelength metric perturbations as follows:

$$\langle \sigma_R(\vec{k})\zeta(\vec{k}')\rangle' \propto f_{\rm NL}^{\Delta}(kR)^{\Delta}P_{\zeta}(k) \quad \Rightarrow \quad \langle \sigma(\vec{k})\delta_m(\vec{k}')\rangle' \propto f_{\rm NL}^{\Delta}\frac{(kR)^{\Delta}}{k^2 T(k)}P_m(k)\,. \tag{3.18}$$

Since the number density of galaxies at a location $\vec{x}$ is related to $\sigma_R(\vec{x})$, this introduces a long-distance correlation between $\delta_g(\vec{x})$ and δ_m that is non-local. The resulting non-local modification of the bias (3.9) is given by [164][4]

$$b(k,\mu,z) = b_1(z) + f(z)\mu^2 + \sum_{n>0} b_{k^{2n}}(kR_*)^{2n} + A f_{\rm NL} \frac{b_\phi(z)}{k^2 \mathscr{T}(k,z)}(kR_*)^\Delta\,, \tag{3.19}$$

where b_ϕ is the non-Gaussian bias parameter and we took $R = R_*$. The constant A normalizes the non-Gaussian parameter $f_{\rm NL}^\Delta$ in the conventional way for PNG in the squeezed limit of the bispectrum,[5] so that the last term in (3.19) is [162, 177, 178, 179, 180]

$$b_{\rm NG}^{\rm loc}(k,z) = f_{\rm NL}^{\rm loc} \frac{b_\phi(z)}{k^2 \mathscr{T}(k,z)} \qquad \text{(local)}, \tag{3.20}$$

$$b_{\rm NG}^{\rm eq}(k,z) = 3 f_{\rm NL}^{\rm eq} \frac{b_\phi(z)}{k^2 \mathscr{T}(k,z)}(kR_*)^2 \qquad \text{(equilateral)}, \tag{3.21}$$

$$b_{\rm NG}^{\Delta}(k,z) = 3 f_{\rm NL}^{\Delta} \frac{b_\phi(z)}{k^2 \mathscr{T}(k,z)}(kR_*)^\Delta \qquad \text{(general exponent } \Delta \in [0,2]\text{)}. \tag{3.22}$$

Note that we do not include the physical orthogonal shape, as it is produced in single-field inflation and, therefore, has the same squeezed limit as the equilateral shape, i.e. corresponds to $\Delta = 2$. Previous work on the scale-dependent bias used an orthogonal template which has $\Delta = 1$ (e.g. [147]), which is however unphysical and has not been used in recent LSS analyses.

Without any additional theoretical input, $b_\phi(z)$ and $f_{\rm NL}^\Delta$ are both unknown quantities and are completely degenerate. If that was the case, we could only measure the combination $f_{\rm NL}^\Delta b_\phi(z)$ in each redshift bin which would make it challenging to infer limits on $f_{\rm NL}^\Delta$ alone. More optimistically, we can fix $b_\phi(z)$ in terms of the linear bias $b_1(z)$ using the universality relations

---

[4] It is important to note that, in principle, the scale-dependent bias is distinguishable from the expansion in $(kR_\star)$ even for $\Delta = 2n$ for any positive integer n. At $k \ll k_{\rm eq}$, $\mathscr{T}(k) \to$ const and the scale-dependent bias scales as k^{2n-2} which is degenerate with b_1 and $b_{k^{2(n-1)}}$. However, since $k_{\rm eq}R_\star \ll 1$, the transfer function introduces a non-trivial scale dependence in the regime $k_{\rm eq} \lesssim k \ll R_\star^{-1}$ which is not captured in the local bias expansion.

[5] Note that for the special bispectrum configurations, when expressed in terms of the general exponent Δ, the equivalent $f_{\rm NL}^\Delta$ definitions are related by $f_{\rm NL}^{\rm loc} = 3 f_{\rm NL}^{\Delta=0}$ and $f_{\rm NL}^{\rm eq} = f_{\rm NL}^{\Delta=2}$.

found in phenomenological models [152],

$$b_\phi(z) = 2\delta_c(b_1(z) - p)\,, \tag{3.23}$$

where $\delta_c = 1.686$ is the critical overdensity for spherical collapse and $p = 1$ if we assume a universal halo mass function. We will adopt this relation in our forecasts, but refer to e.g. [181, 182, 183, 184] for a recent discussion of the validity and risks of this choice.

To conclude, we note that we incorporate the general scale-dependent bias (3.19) in the galaxy power spectrum as described in (3.10) and (3.13), i.e. $b^A = b^A_\Delta(k,\mu,z)$. In the rest of the paper, we follow [148] in assuming a minimum halo mass of $10^{13}M_\odot$, which corresponds to $R_* \approx 2.66\,h^{-1}\,\mathrm{Mpc}$, and truncating the non-local gradient bias expansion at $n = 2$. We illustrate the effect of b^Δ_{NG} on the linear galaxy power spectrum in Fig. 3.1 and compare these non-Gaussian contributions to the nonlinear galaxy power spectrum in Fig. 3.2 at redshift $z = 0$ for linear bias $b_1 = 1.6$. We see from the latter figure that the $\Delta = 0$ contribution does not mimic the behavior of any other bias or nonlinear terms at low k. On the other hand, the $\Delta = 2$ contribution is both much smaller at $k < k_{\mathrm{NL}}$ and similar in behavior to those Gaussian contributions.

3.4 Information Content of the Galaxy Power Spectrum

Our goal in this section is to understand how to design a galaxy survey to best measure f^Δ_{NL} for $\Delta \in [0,2]$. To do so, we will explore the Fisher forecasts for f^Δ_{NL} as a function of Δ for a variety of survey configurations and biasing models.[6] We will compare these results with analytic estimates to understand what aspects of the surveys drive major improvements in sensitivity. While our forecasts will be rooted in specific survey configurations, our goal is to isolate the features of these surveys that impact the forecasted constraining power the most.

---

[6] In Appendix B.2, we also discuss our ability to measure the scaling behavior of the non-Gaussian bias, where we also include Δ as an additional free parameter in the Fisher matrix.

3.4.1 Fisher Information Matrix

In the following, we describe the well-known Fisher matrix forecasting techniques for the galaxy power spectrum that we use. We review the basics here and refer to Appendix B.1 for the experimental specifications of the surveys.

The Fisher matrix for a set of galaxy samples, labeled by $\{A, B, \ldots\}$, is defined by

$$F_{\alpha\beta} = \sum_{z_i} V(z_i) \int_{-1}^{1} \frac{\mathrm{d}\mu}{2} \int_{k_{\mathrm{min}}}^{k_{\mathrm{max}}} \frac{\mathrm{d}k\, k^2}{2\pi^2} \frac{1}{2} \mathrm{Tr}\left[\mathbf{C}_{,\alpha} \mathbf{C}^{-1} \mathbf{C}_{,\beta} \mathbf{C}^{-1}\right], \tag{3.24}$$

where $\theta_{\alpha,\beta}$ are the model parameters, the trace is over the samples and we defined $\mathbf{C}_{,\alpha} \equiv \partial \mathbf{C}/\partial \theta_\alpha$. The integrals are discretized, with the integration over wavenumbers k being limited by the minimum wavenumber k_{min} and the maximum wavenumber k_{max}, and bins of width $\Delta k = k_{\mathrm{min}}$. We note that we vary all parameters listed in Table 3.1 throughout this paper, unless specified otherwise. The elements of the matrix $\mathbf{C}$ are given by

$$\mathbf{C}_{AB}(k, \mu, z) \equiv P_{AB}^{\mathrm{obs}}(k, \mu, z) + \delta_{AB} N_A(k, \mu, z), \tag{3.25}$$

with the shot noise term $N_A(k, \mu, z) = \bar{n}_A(z)^{-1}$, where $\bar{n}_A(z)$ is the average comoving number density of tracers at redshift z. Shot noise arises because we observe a finite number of discrete objects in galaxy surveys. (Note that we absorbed the stochastic term $P_{\varepsilon 0}$ of (3.10) into N_A.) We treat the redshift bins as spherical shells so that the volume of the i-th redshift bin is

$$V(z_i) = \frac{4\pi}{3} f_{\mathrm{sky}} \left[d_c(z_i + \Delta z_i/2)^3 - d_c(z_i - \Delta z_i/2)^3\right], \tag{3.26}$$

where f_{sky} is the sky fraction observed by the survey, $d_c(z)$ is the comoving distance to redshift z and Δz_i is the width of the i^{th} redshift bin.

To ensure that only modes under perturbative control in Eulerian perturbation theory and the bias expansion enter our Fisher matrix calculation, and to exclude long-wavelength modes inaccessible to a given survey due to its finite volume, we restrict the integral in (3.24) to a maximum wavenumber $k_{\max}$ and a minimum wavenumber $k_{\min}$, respectively. We make a conservative choice for $k_{\max}$ by taking the smaller of the two scales indicating the range of validity of perturbation theory and the bias expansion. To be precise, we take it to be the minimum of the nonlinear scale[7] $k_{\rm NL}$ and the wavenumber $k_{\rm halo} \approx 0.19\,h\,{\rm Mpc}^{-1}$ associated with the Lagrangian scale corresponding to the minimum halo mass of the surveyed population which we assume to be $R_* \approx 2.66\,h^{-1}\,{\rm Mpc}$ (cf. §3.3.2), i.e. $k_{\max}(z) = \min\{k_{\rm NL}(z), k_{\rm halo}\}$. For $k_{\min}$, we assume a spherical survey geometry per redshift bin which results in a conservative minimum wavenumber $k_{\min,i} = 2\pi\,[3V(z_i)/(4\pi)]^{1/3}$.

For most of the forecasts in Sections 3.4 and 3.6, we use a futuristic spectroscopic survey with a billion objects divided into 10 redshift bins with $z \leq 5$, which we will hereafter refer to as the *billion-object survey*. We split the galaxy sample of this survey into two populations based on their linear bias at $z = 0$, chosen as $b_1^{(1)}(z=0) = 2.0$ and $b_1^{(2)}(z=0) = 1.2$. We model their redshift evolution according to (3.8) as previously discussed in §3.3.1. In addition, we assume a constant number density across all redshift bins and both galaxy samples, $\bar{n}_g^{(1)}(z) = \bar{n}_g^{(2)}(z) =$ const. The details of this and other surveys used in our forecasts, i.e. the Baryon Oscillation Spectroscopic Survey (BOSS) [149], the Dark Energy Spectroscopic Instrument (DESI) [185], Euclid [186, 187], Vera Rubin Observatory's LSST [188], the Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx) [141], MegaMapper [189, 112] and the billion-object survey are provided in Appendix B.1. Except for SPHEREx and the billion-object survey, we treat all surveys as single-tracer surveys by combining different groups of tracers (if available) into a single effective number density and bias. The impact of this should not significantly impact our forecasts as we explain in Appendix B.1.

Unless mentioned otherwise, we marginalize over all bias parameters and include CMB information on the ΛCDM parameters (but not on PNG) from Planck throughout the paper.[8] We

---

[7] We use $k_{\rm NL} = \frac{\pi}{2R_{\rm NL}}$, with the radius $R_{\rm NL}$ at which the variance of linear fluctuations is $\sigma_R = 1/2$ at a given z.

[8] When marginalizing over the full biasing model for the billion-object survey, for instance, the difference

achieve this by computing the ΛCDM Fisher matrix directly following the Fisher methodology and the experimental specifications of [44] (see also [191] for the latter), marginalized over the optical depth τ. We then combine this marginalized Planck Fisher matrix with the respective LSS Fisher matrix for the five ΛCDM parameters, the non-Gaussian parameter $f_{\rm NL}^{\Delta}$ and the galaxy bias parameters $b_{\mathcal{O}_i}$. We treat each redshift bin with mean redshift $\bar{z}_i$ as an independent survey and independently marginalize over the biases within each redshift bin.

3.4.2 Information and Survey Design

Using Fisher forecasts, our aim is to understand the possible strategies for optimizing the measurement of $f_{\rm NL}^{\Delta}$ for shapes including but not limited to local PNG. The parameters that we most directly control in the design of a survey are the sky fraction ($f_{\rm sky}$), the redshift range ($z \in [z_{\rm min}, z_{\rm max}]$) and the number density of tracers [$\bar{n}_g(z)$]. However, the actual statistical power of the survey is controlled by the smallest ($k_{\rm min}$) and largest ($k_{\rm max}$) wavenumbers that can be reliably measured, the linear bias of the sample [$b_1(z)$], etc. These factors are influenced by the details of the survey, but can also be affected by systematics and astrophysics.

We show the overall landscape of future power spectrum measurements of $f_{\rm NL}^{\Delta}$ for $\Delta \in [0,2]$ in Fig. 3.3. Qualitatively, the forecasts show the overall behavior that we might expect: (i) constraints at $\Delta = 0$ are generally much stronger than at $\Delta = 2$, (ii) higher-order biasing affects larger Δ more than smaller Δ, (iii) larger surveys have more constraining power, and (iv) future LSS surveys will be able to improve over Planck bispectrum constraints from the galaxy power spectrum for $\Delta = 0$, but not for $\Delta = 2$ without additional information on bias parameters. In detail, the forecasts however have elements that are harder to understand without further investigation. First, the forecasts show clear features at values of $\Delta \sim 1.4$ that depend on the marginalization over biasing parameters (see also Fig. 3.6 below). This suggests a qualitative change in where the constraining power is coming from as we vary Δ. Second,

---

between using the ΛCDM covariance from Planck or CMB-S4 [52, 190] for $\Delta \lesssim 1.5$ is less than two percent. In fact, this CMB information decreases $\sigma(f_{\rm NL}^{\Delta})$ only by a few percent for $\Delta \lesssim 1.3$, and by a maximum of 20% (15%) with Planck (CMB-S4) at $\Delta = 2$.

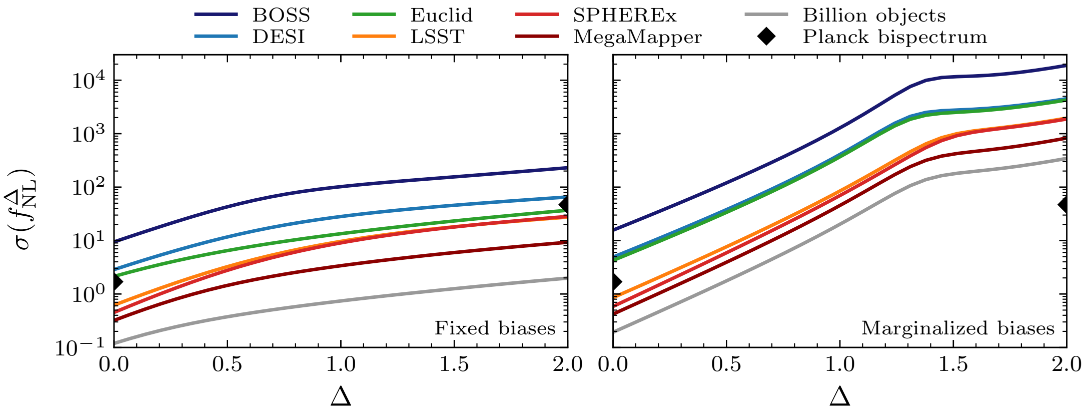

Figure 3.3. Forecasted constraints on $f_{\rm NL}^{\Delta}$ for general scaling $\Delta \in [0, 2]$ of the scale-dependent bias, for current and future surveys, compared to the Planck bispectrum constraints for local and equilateral non-Gaussianity. We also marginalize over the five ΛCDM parameters, use Planck results as priors for the ΛCDM parameters, and either fix the bias parameters (*left*) or marginalize over them (*right*). Note that the SPHEREx forecast uses the five-tracer sample and the billion-object survey forecast uses the double-tracer sample as provided in Appendix B.1. For all the other surveys, we use their combined single-tracer sample.

we see that SPHEREx and LSST produce very similar forecasts despite having very different strengths and weaknesses. Furthermore, both significantly exceed the constraining power of Euclid and DESI which have the advantage of being three-dimensional (spectroscopic) surveys. In order to make sense of these forecasts, we will break them down according to the number density (noise), scales ($k_{\rm min}/k_{\rm max}$), the fiducial bias $[b_1(z)]$, the survey geometry ($f_{\rm sky}$, $z_{\rm max}$) and redshift errors (σ_{z0}) in the following.

Dependence on Number Density

The most basic parameter in any experiment is the signal to noise. In a galaxy survey, shot noise is the dominant (irreducible) noise source. At fixed volume, increasing the number of objects increases the range of wavenumbers k where $P_{gg}(k)$ is measured with at least signal to noise of order one, $P_{gg}(k)/N(k) \gtrsim O(1)$.

For most cosmological parameters, sample variance presents an additional irreducible source of noise. However, for scale-dependent bias, the parameter $f_{\rm NL}^{\Delta}$ can, in principle, be measured without sample variance when the signal to noise is large for multiple tracers [192].

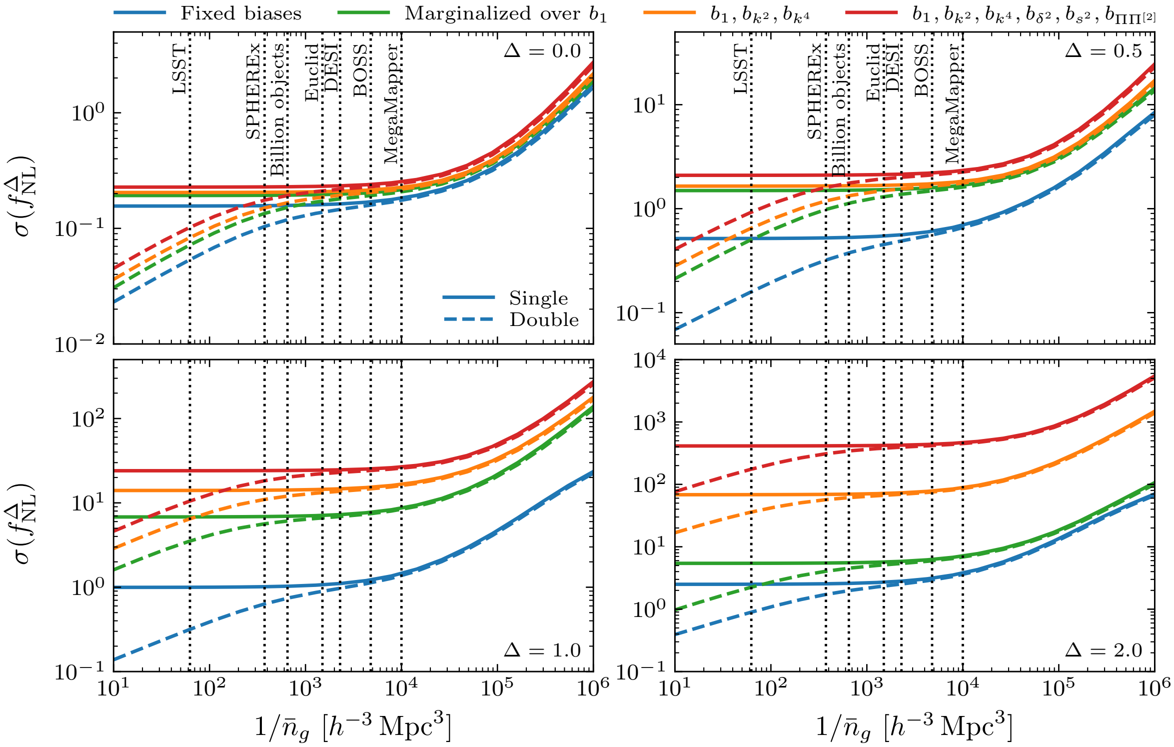

Figure 3.4. Dependence of $\sigma(f_{\mathrm{NL}}^{\Delta})$ on the shot noise, i.e. the inverse survey number density $\bar{n}_g^{-1}$, for the billion-object survey for four representative values of Δ. The solid lines show the single-tracer forecasts and the dashed lines display the double-tracer results. The different colors indicate which bias parameters are marginalized when calculating $\sigma(f_{\mathrm{NL}}^{\Delta})$. The vertical lines mark the averaged shot noise of various surveys which are calculated by taking the average of the number densities of the middle 50% of redshift bins of each survey, i.e. they should be taken as an approximate illustration of the noise level in these surveys.

Concretely, the relationship between $\delta_g(\vec{k})$ and $\delta_m(\vec{k})$ in (3.7) is completely deterministic and is therefore not limited by sample variance. However, this only works if we can measure the same Fourier mode with high signal to noise for tracers with different values of b_ϕ. The noise introduced by sample variance is proportional to $P(k)$, while shot noise is independent of k. As a result, sample-variance cancellation dramatically alters how different scales contribute to the measurement of f_{NL}^{Δ} and, therefore, the qualitative understanding of our forecasts. This is why it is important to distinguish from the outset to what degree multi-tracer sample-variance cancellation is relevant in our forecasts.

The impact of shot noise in both single- and multi-tracer scenarios is illustrated in Fig. 3.4.

Two points should be very clear from these curves: (1) sample-variance cancellation has the potential to dramatically improve the sensitivity to $f_{\rm NL}^{\Delta}$ for all values of Δ, and (2) realistic current and future surveys are very far from being in the regime where sample-variance cancellation has a large impact. In fact, most near future surveys are, at best, at the boundary between single and multiple tracers.[9] As a result, sample-variance cancellation is important for highly accurate forecasts, but the qualitative behavior is consistent with single-tracer forecasts. We therefore use a single-tracer emphasis for the rest of this section. We will revisit the advantages offered by multi-tracer techniques and how astrophysics affects the constraints on $f_{\rm NL}^{\Delta}$ in Section 3.6.

Dependence on Scales

The defining feature of scale-dependent bias is that it is enhanced at small wavenumbers/large distances. For $\Delta < 2$ and $f_{\rm NL}^{\Delta} \neq 0$, we have $P_{gg}(k) \gg b_1^2 P_m(k)$ as $k \to 0$. Since the noise is also a function of k, it is however not a given that the information resides at small k for all $\Delta < 2$. We now investigate this analytically in the Fisher matrix and numerically in our forecasts.

For a single-tracer analysis where we know the bias parameters exactly, i.e. $b_{\mathscr{O}_i}$ are held fixed, the Fisher information for a fiducial $f_{\rm NL}^{\Delta} = 0$ is given by

$$F_{f_{\rm NL}^{\Delta} f_{\rm NL}^{\Delta}} = \sum_{z_i} V(z_i) \int \frac{{\rm d}^3 k}{(2\pi)^3} \frac{[6 b_1 b_\phi(z_i)(kR_*)^{\Delta}]^2 P_m(k,\mu,z_i)^2}{2k^4 \mathscr{T}(k,z_i)^2\,[b_1^2 P_m(k,\mu,z_i) + N(k,\mu,z_i)]^2}\,. \tag{3.27}$$

In the high signal-to-noise regime, $b_1^2 P_m(k) \gg N$, this becomes

$$F_{f_{\rm NL}^{\Delta} f_{\rm NL}^{\Delta}} \approx \frac{9}{\pi^2} \sum_{z_i} \frac{b_\phi^2(z_i)}{b_1^2(z_i)} R_*^{2\Delta} V(z_i) \int {\rm d}k\, \frac{k^{2\Delta-2}}{\mathscr{T}(k,z_i)^2}\,. \tag{3.28}$$

For $2\Delta < 1$, the integral over k is dominated by $k_{\rm min}$ and the Fisher information arises from the

---

[9] Note that the curves in Fig. 3.4 are computed using the spectroscopic billion-object survey (cf. Table B.6). This means that the photometric LSST is not a direct comparison and the line labeled as LSST only actually indicates a spectroscopic follow-up of LSST. The potential sensitivity in a multi-tracer analysis of LSST itself will be discussed in §3.6.2.

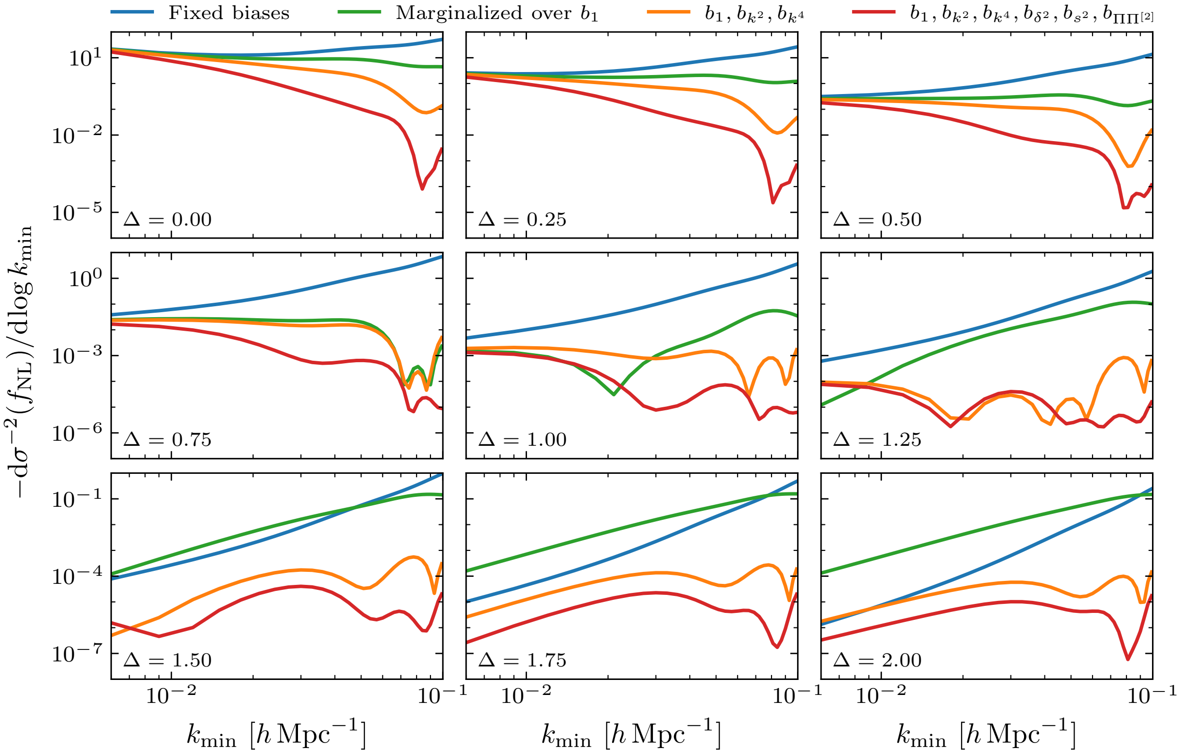

Figure 3.5. The information density with respect to $\log k_{\mathrm{min}}$ in the billion-object survey, with fixed ΛCDM cosmology, and the same k_{min}, $\Delta k = 0.003\,h\,\mathrm{Mpc}^{-1}$ and $k_{\mathrm{max}} = 0.2\,h\,\mathrm{Mpc}^{-1}$ in each redshift bin. The different colors indicate which bias parameters are marginalized when calculating $\sigma(f_{\mathrm{NL}}^{\Delta})$ using the double-tracer version of the survey. A negative (positive) slope of the curves indicates that there is more (less) information on smaller wavenumbers k. We observe that the information density is dominated by small (large) wavenumbers for small (large) scaling exponents.

smallest scales. However, due to the limits placed by noise and volume, realistic surveys are far from the asymptotic $k \to 0$ regime. For modes with the largest signal to noise, $k \approx k_{\mathrm{eq}} \approx 10^{-2}\,h\,\mathrm{Mpc}^{-1}$, the matter power spectrum is flat, $P_m(k \sim k_{\mathrm{eq}}) \propto k^0$ and, therefore, $\mathscr{T}(k,z)^2 \sim k^{-1}$. In this regime, repeating our analysis of (3.28) shows that the integral is dominated by large k for $\Delta \gtrsim 0$. In this regard, we expect that the transition from low to high k is gradual as we change Δ.

This behavior is seen in forecasts by varying k_{min} as illustrated in Fig. 3.5. Since increasing the minimum wavenumber k_{min} removes information, the constraints $\sigma(f_{\mathrm{NL}}^{\Delta})$ always decrease. We therefore show $\mathrm{d}\sigma^{-2}(f_{\mathrm{NL}}^{\Delta})/\mathrm{d}\log k_{\mathrm{min}}$ because this quantity provides a more quantitative illustration of how much information is being lost as we change k_{min}.[10] For $\Delta = 0$, the largest impact is for the smallest values of k_{min}, indicating that the lowest wavenumber has the most information. As we increase $\Delta > 0$, we see that this trend changes: the curves are mostly flat at small k_{min} for $\Delta \lesssim 0.5$. This is a reflection of the fact that there is no sharp transition in Δ due to the impact of $\mathscr{T}(k,z)$ in (3.28).

As we move from large to small scales (or small to large k), we are increasingly sensitive to the marginalization over the bias parameters. The information density in Fig. 3.5 at large wavenumbers becomes highly suppressed as we marginalize so that the true constraining power remains at small wavenumbers. The oscillatory behavior in these figures highlights that the baryon acoustic oscillations contained in the transfer function are imprinted in the non-Gaussian signal and are not absorbed into the bias expansion.

At the same time, this short-distance information remains however crucial for breaking degeneracies between bias and cosmological parameters. This is shown in Fig. 3.6, which illustrates how the maximum wavenumber k_{max} impacts the constraining power. We see that the constraint on f_{NL}^{Δ} improves by roughly a factor of five to six (depending on the biasing model) for $\Delta = 2$ when we choose the less conservative $k_{\mathrm{max}} = k_{\mathrm{NL}}$. This improvement is consistent with the number of additional modes available between k_{halo} and k_{NL}, and the fact that our signal depends on wavenumbers. This difference becomes even more pronounced at higher redshifts since $k_{\mathrm{NL}}(z)$ increases significantly, while k_{halo} is independent of redshifts. For all of our other forecasts, we however use the more conservative choice of $k_{\mathrm{max}}(z) = \min\{k_{\mathrm{halo}}, k_{\mathrm{NL}}(z)\}$ as noted in §3.4.1.

When marginalizing over b_1, we notice a feature around $\Delta \approx 1$ in Fig. 3.6 and not around $\Delta = 0.5$ as we might expect from (3.28). This is a consequence of the information on f_{NL}^{Δ} moving to larger k where $T(k) \neq 1$. Specifically, if we assume $P_m(k) \propto k^n$ for large k, then

---

[10] We consider the logarithmic derivative in order to numerically extract the relevant scaling that we analytically discussed in the previous paragraph. We additionally note that $\log k_{\mathrm{min}}$ is implicitly normalized by $1\,h\,\mathrm{Mpc}^{-1}$.

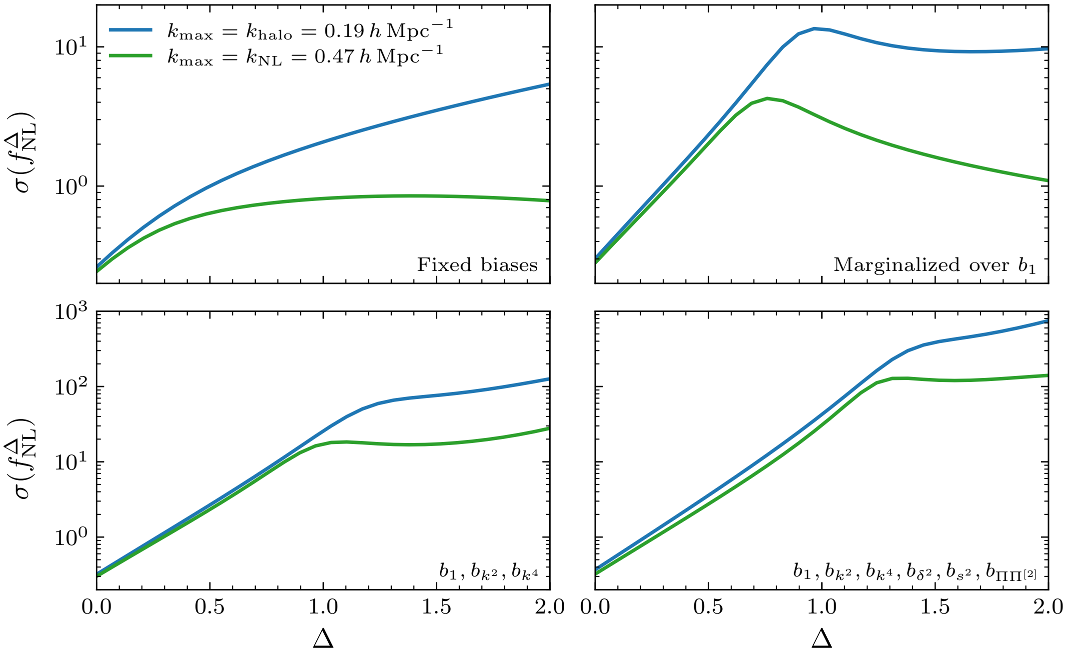

Figure 3.6. Dependence of $\sigma(f_{\rm NL}^{\Delta})$ on the maximum wavenumber where we take $k_{\rm max} = k_{\rm halo}$ and $k_{\rm max} = k_{\rm NL}(z)$, with the latter being less conservative at higher redshifts. Here, we combine the middle two redshift bins ($2 \leq z \leq 3$) of the double-tracer billion-object survey into a single comoving box with a total of about 1.3×10^8 objects. We show this dependence of the constraining power for four biasing models. The choice of $k_{\rm max}$ barely impacts $\sigma(f_{\rm NL}^{\Delta})$ at small Δ, but becomes significant for large Δ. We explain the origin of the feature around $\Delta = 1$ in the main text.

$k^2 T(k) \propto k^{n/2+3/2}$, which implies

$$P_{gg}(k) \approx \left(b_1 + 2b_\phi f_{\rm NL}^{\Delta} C(kR_*)^{\Delta-\frac{n}{2}-\frac{3}{2}} + \ldots\right)^2 P_m(k)\,, \tag{3.29}$$

for some constant C. Using $n \approx -1$ for $k \sim 0.1\,h\,\mathrm{Mpc}^{-1}$ [193], we notice that $\Delta - \frac{3}{2} - \frac{n}{2} \approx \Delta - 1$. At large wavenumbers, the “scale-dependent bias” for $\Delta \approx 1$ is therefore degenerate with the linear bias b_1, resulting in the observed feature.

Dependence on Fiducial Biases

All else being equal, a single-tracer survey with the largest possible (absolute) value of $b_\phi \propto (b_1 - 1)$ will yield the most sensitive measurement of $f_{\rm NL}$. It may therefore seem self-evident that selecting highly biased targets is a central tool in the search for PNG. Having said that, in practice, target selection involves numerous factors which lie beyond the scope of these simple forecasts. Yet, when it comes to understanding the performance of a given survey, the biases of the objects in their sample will strongly influence the overall sensitivity. As a result, it is important to separately understand the role that the fiducial biases play from the aspects of the survey that we can control more directly.

For a single-tracer analysis, we only measure $b_\phi f_{\rm NL}^{\Delta}$. At fixed redshift and assuming $b_\phi \propto (b_1 - 1)$ from (3.23), we therefore have

$$\sigma(f_{\rm NL}^{\Delta})_{\rm single}^{\text{fixed-}z} \propto \frac{1}{b_1 - 1}\,. \tag{3.30}$$

As a result, we can expect a large improvement in a single-tracer analysis from choosing a sample with $b_1 \gtrsim 2$. Figure 3.7 shows how the value of $b_1(z=0)$ influences the overall constraining power of our billion-object survey in both single- and multi-tracer analyses. We see that $\sigma(f_{\rm NL})$ increases sharply when the bias approaches one. It does not diverge because of the redshift dependence of $b_1(z)$ which ensures that some of the redshift bins have $b_\phi \neq 0$. On the other

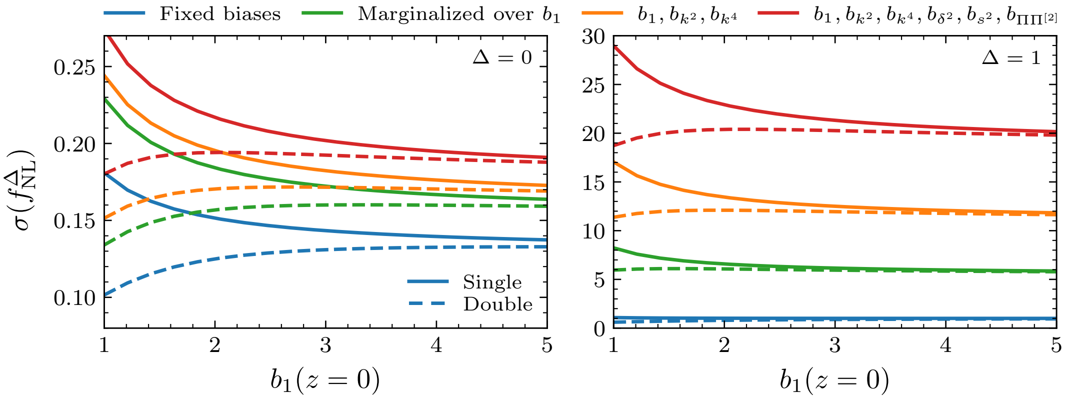

Figure 3.7. Dependence of $\sigma(f_{\rm NL}^{\Delta})$ on the fiducial value of the linear bias for both the single-tracer (solid) and double-tracer (dashed) versions of the billion-object survey for $\Delta = 0$ (*left*) and $\Delta = 1$ (*right*). We show this dependence as a function of the linear bias $b_1(z=0)$ of the single-tracer survey in which case we see the improvement in the constraining power of the survey for larger b_1. In the double-tracer scenario, for which the galaxy sample is split into two populations of equal number density with $b_1^{(A,B)}(z=0) = b_1(z=0) \pm 0.4$, the effects of sample-variance cancellation can be observed. (Note that the y-axis is on a linear and not on a logarithmic scale like in the previous figures.)

hand, sample-variance cancellation in the multi-tracer configuration results in a much weaker dependence of $\sigma(f_{\rm NL}^{\Delta})$ on the linear bias and, in fact, benefits from having one of the tracers with $b_1^B(z=0) < 1$. At the same time, we however see that taking larger values of $b_1(z)$ in a single-tracer analysis closes the gap to the double-tracer case for $b_1(z=0) \gtrsim 3$.

Technically speaking, we cannot choose the linear bias $b_1(z)$ and the number densities $\bar{n}_g(z)$ independently, in particular since more highly biased objects are more rare. In practice, traditional spectroscopic surveys are typically limited by the number of spectra that can be measured and not the number of objects in the universe. However, other types of surveying techniques, such as line intensity mapping (see e.g. [194, 195, 196, 197, 67, 198] in the context of PNG), may get closer to these fundamental limits. We refer to [148] for a discussion of the sensitivities of surveys in the regime where the galaxy power spectrum analysis is limited by the number of available objects.

Dependence on the Volume and Redshift

Given a single redshift bin with volume V and a fixed choice of targets, it is easy to see the impact of this choice on k_{min} and k_{max}. However, when we design a survey, the choice of redshift range, volume and targets is not fixed, but is part of the survey design. The redshift range and volume affect k_{min} and k_{max}, but they also change the biases of the targets and the number densities of objects in each redshift bin. For the purpose of designing a survey, we therefore want to understand the optimal choices for the sky fraction f_{sky} and the redshift range as given by z_{max}.

It is important to compare survey strategies assuming “constant effort”. In practice, it is the time for acquiring spectra that limits the sensitivity of a spectroscopic survey. This is why we vary f_{sky} for a given redshift range while holding the total number of objects fixed. As a result, there is a trade-off between increasing the volume and increasing the shot noise, which needs to be optimized. When it comes to varying the redshift range, the time required to observe each object changes with redshift which means that the total number of objects in the sample decreases with larger z_{max}. To hold the observational effort fixed, we use a simplified model where it takes twice the amount of time to obtain spectra for galaxies at $z > 3$ than for galaxies at $z < 3$ (cf. [111]). We further assume that it is desirable to maintain a constant level of shot noise over the survey volume or that we obtain spectra evenly across the sky and along the radial direction.

Higher-redshifts objects are generally expected to yield larger values of $b_1(z)$ [see the discussion around (3.8)]. Since $b_\phi(z) \propto (b_1(z) - 1)$ from the universality relation (3.23), the signal can be enhanced by a large amount by extending the survey to redshifts where $b_1(z) > 2$ instead of $b_1(z) \approx 1$. Assuming the bias evolution model of (3.8), we see in Fig. 3.8 that we maximize our sensitivity to f_{NL}^{Δ} even at constant effort by maximizing z_{max}. While we only show the results for $\Delta = 1$, the qualitative features in the figure are independent of the scaling exponent Δ. This implies that increasing the redshift range at fixed observational effort increases

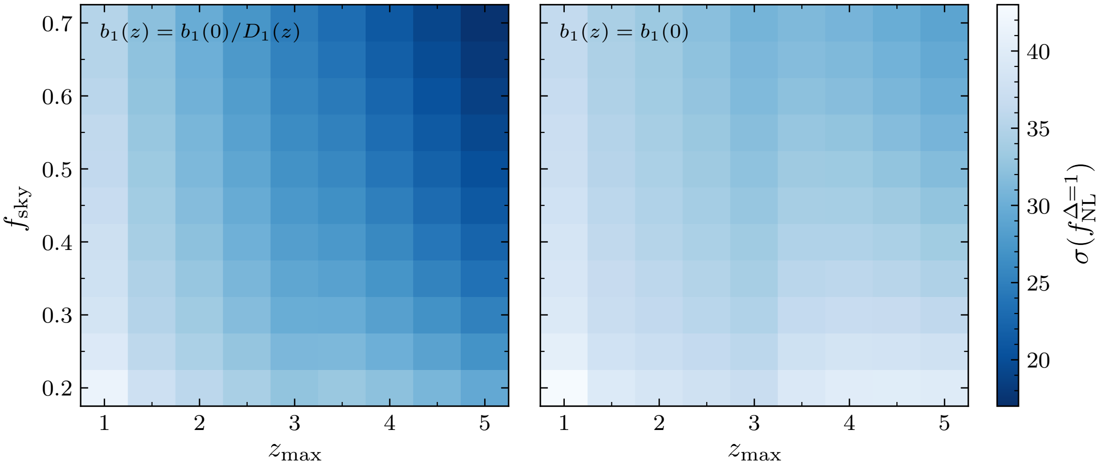

Figure 3.8. Dependence of $\sigma(f_{\rm NL}^{\Delta=1})$ on the sky fraction $f_{\rm sky}$ and the maximum redshift $z_{\rm max}$ at fixed observational effort for the double-tracer billion-object survey. The left panel shows the fiducial case, $b_1(z) \propto D(z)^{-1}$ [cf. (3.8)] and the right panel displays the constant bias case, $b_1(z) = $ const. While we only show the results for $\Delta = 1$, the displayed trends are similar for all scaling exponents Δ.

the constraining power for the entire range of Δ.

Increasing the redshift range also increases the number of modes. If we however hold $b_1(z) = b_1$ fixed and increase $z_{\rm max}$, we see in the right panel of Fig. 3.8 that there is a more complex relationship between $z_{\rm max}$ and the sensitivity to $f_{\rm NL}^{\Delta}$. For instance, keeping $z_{\rm max} = 3$ and increasing $f_{\rm sky}$ also increases the number of modes without the added observing time per object needed at high redshifts. Having said that, the benefit of a large sky fraction is much weaker than the impact of larger b_1 at high z which can be seen by the lower overall sensitivity to $f_{\rm NL}^{\Delta}$ in the right panel.

By comparing the two panels in Fig. 3.8, we can see that the optimal configuration is still at the largest maximum redshift and sky fraction, $z_{\rm max} = 5$ and $f_{\rm sky} = 0.7$, for either biasing model. For smaller $f_{\rm sky}$, the large advantage offered by larger $z_{\rm max}$ however disappears when going from evolving to fixed bias due to the lower overall number density of galaxies at higher redshifts. In this case, more emphasis is placed on high-redshift targets without the corresponding boost in the signal. (The drop-off in $\sigma(f_{\rm NL})$ between $z_{\rm max} = 3.0$ and 3.5 is the result of our simplified choice of dividing the observational effort at $z = 3$.)

To summarize, at fixed observational effort, increasing the redshift range via a larger $z_{\rm max}$ has the largest influence on the sensitivity of a survey. The reason is that $z_{\rm max}$ increases both the size of the signal and the number of modes. The former is particularly important since it offsets the increased noise associated with a large-volume survey. We however note that this conclusion does depend on the redshift dependence of the bias and the observing time needed for acquiring spectra of these high-redshift objects.

Photometric versus Spectroscopic Surveys

The benefit of a three-dimensional survey can simply be estimated by counting modes for many cosmological parameters, including $f_{\rm NL}$ from the bispectrum. The number of modes in a three-dimensional survey scales as $k_{\rm max}^3$, while it scales as $k_{\rm max}^2$ for a two-dimensional survey. The benefits of a three-dimensional survey are however less clear for scale-dependent bias since the information for this signal manifests itself at low k, in particular for $\Delta \lesssim 0.5$. In reality, the trade-off is of course not between a two- and a three-dimensional survey, but between surveys which measure redshifts spectroscopically and photometrically. While photometric redshifts are less precise, they may be good enough to provide three-dimensional information on the scales needed for scale-dependent bias for a significant range of Δ which we investigate now.[11]

The impact of redshift errors on the constraining power of the billion-object survey is shown in Fig. 3.9. We see, somewhat surprisingly, that spectroscopic surveys do not provide a significant improvement to the power spectrum measurement of $f_{\rm NL}^{\Delta}$ for $\Delta \lesssim 1-1.3$. (The specific value of this transition in Δ depends on the employed biasing model.) In practice, photometric redshifts therefore appear to be good enough to access a lot of the available information about the scale-dependent bias. This conclusion is supported by the similarity of the forecasts for SPHEREx and LSST. While the two surveys pursue very different strategies for obtaining redshifts, the resulting forecasts are quite similar. At first sight, this similarity is still surprising because it does

---

[11] Spectroscopic surveys will have additional benefits for controlling systematics, such as projecting out large-scale effects, but typically at the cost of larger shot noise. Quantifying the importance of spectra in this context is a delicate issue which is beyond the scope of this work.

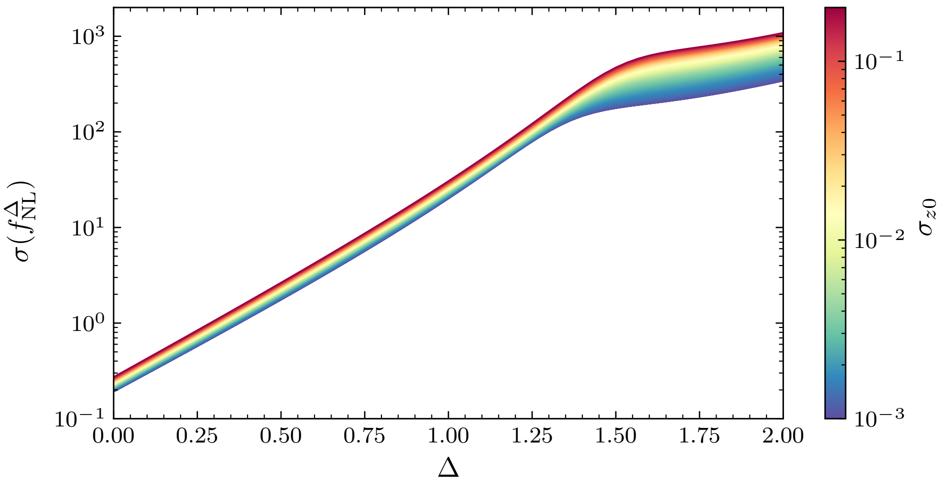

Figure 3.9. Comparison of $\sigma(f_{\rm NL}^{\Delta})$ for a range of photometric redshift errors $\sigma_{z0} \in [0.001, 0.2]$ for the billion-object survey. We observe that photometric redshift errors only mildly degrade the constraining power of a survey for $\Delta \lesssim 1.3$. Photometric surveys therefore seem to be able to access most of the information about the scale-dependent bias with the additional benefit of generally lower noise levels in the galaxy power spectrum.

not reflect the impact of the difference in number density shown in Fig. 3.4. The reason for this, however, is just that multi-tracer analyses are not implemented in conventional LSST forecasts. This naturally raises the question whether a multi-tracer analysis would lead to an advantage of LSST over some spectroscopic surveys given its very high number density (cf. Fig. 3.3). We will explore this in detail in Section 3.6.

Importantly, we have not discussed the impact of systematic effects [199] on photometric and spectroscopic surveys. Given that the signals are the largest on large angular scales (especially for $\Delta \lesssim 0.5$), a variety of atmospheric, detector and astrophysical effects can contaminate our signal. Having said that, we expect that many of the conclusions drawn from these forecasts will be robust over a large parameter range since the scales on which the signals dominate vary with the scaling exponent Δ.

3.5 Constraints and Data Analysis for General Scaling Exponents

In the previous section, we focused on forecasting current and future constraints on $f_{\rm NL}^{\Delta}$, and on how to optimize a survey to measure PNG from the galaxy power spectrum for $\Delta \in [0,2]$. In this section, we will apply this knowledge to constrain $f_{\rm NL}^{\Delta}$ over this same range of scaling exponents Δ using the BOSS DR12 galaxy power spectra. We will first derive and propose an effective and convenient search strategy. We will then perform this measurement and compare it to a direct analysis.

3.5.1 Search Strategy

The scaling exponent Δ is fixed in some inflationary models, while it is a free parameter in others. For the standard local and equilateral shapes, the value of Δ are fixed by their templates to $\Delta = 0$ and $\Delta = 2$, respectively. The physics of these cases is clear since they arise from multi- and single-field dynamics. In quasi-single-field inflation, we have $\Delta = 3/2 - \sqrt{9/4 - m^2/H^2}$ which implies that it is natural to scan over values of $\Delta \in [0,3/2]$ in an analysis of those models. The natural question we face in a practical analysis of current data therefore is whether it is prudent to vary $f_{\rm NL}^{\Delta}$ and $\Delta \in [0,2]$ as free parameters, or whether we should hold the scaling exponent fixed to some value as suggested by a given model, $\Delta = \Delta'$, and only vary $f_{\rm NL}^{\Delta'}$.

The strategy we will follow is to hold Δ fixed to a few discrete values, e.g. $\Delta = 0,2$ or $\Delta = 0.0, 0.1, \ldots, 1.9, 2.0$, and infer the associated constraints on $f_{\rm NL}^{\Delta}$ at each respective Δ. The reason why we prefer this strategy over scanning over both parameters simultaneously is that small changes to Δ have no impact on the signal to noise and, therefore, represent a degenerate direction in the analysis. This becomes evident in the Fisher matrix formalism, where we have for a fiducial value of $f_{\rm NL}^{\Delta} = 0$:

$$\frac{\partial}{\partial \Delta} P_{gg}(k)\bigg|_{f_{\rm NL}^{\Delta}=0} = 6 f_{\rm NL}^{\Delta} b_{\phi}(k)\, k^{-2} \mathscr{T}(k)^{-1} \log(kR_\star)(kR_\star)^{\Delta} b(k) P_{\rm lin}(k)\bigg|_{f_{\rm NL}^{\Delta}=0} = 0\,. \tag{3.31}$$

We therefore have no reason to expect marginalizing over Δ would lead to more accurate results until our data are comfortably excluding $f_{\rm NL}^{\Delta}=0$.

We can quantify to what degree the data can distinguish two values of the scaling exponent, Δ_1 and Δ_2, through the "cosine" between the signals,

$$\cos\left(f_{\rm NL}^{\Delta_1},f_{\rm NL}^{\Delta_2}\right)=\frac{F_{f_{\rm NL}^{\Delta_1}f_{\rm NL}^{\Delta_2}}}{\sqrt{F_{f_{\rm NL}^{\Delta_1}f_{\rm NL}^{\Delta_1}}F_{f_{\rm NL}^{\Delta_2}f_{\rm NL}^{\Delta_2}}}}\equiv\frac{F_{12}}{\sqrt{F_{11}F_{22}}}\,, \tag{3.32}$$

where we defined $F_{ij}\equiv F_{f_{\rm NL}^{\Delta_i},f_{\rm NL}^{\Delta_j}}$. This is precisely the same definition of the cosine used to define the PNG shapes in a bispectrum analysis [200]. We can also generalize this to the Fisher matrix after marginalizing over the bias parameters by inverting the Fisher matrix, truncating to $f_{\rm NL}^{\Delta_1}$ and $f_{\rm NL}^{\Delta_2}$, and inverting the truncated matrix to get the effective Fisher matrix for these two parameters. We therefore take F_{ij} to be the marginalized, two-dimensional Fisher matrix for the parameters $f_{\rm NL}^{\Delta_1}$ and $f_{\rm NL}^{\Delta_2}$ from now on.

The correlation (3.32) between the measurements of $f_{\rm NL}^{\Delta_1}$ and $f_{\rm NL}^{\Delta_2}$ is shown in Fig. 3.10 using the experimental specifications for the BOSS survey of Table B.1 which have yielded forecasted results consistent with performed data analyses in the past (see e.g. [114]).[12] The cosine computed for other surveys is similar (see Appendix B.2). The key take-away from this correlation is that we can effectively constrain almost the entire range $\Delta\in[0,2]$ using measurements only at $\Delta=0$ and 2. Specifically, using the Cauchy-Schwarz inequality and the Cramér-Rao (CR) bound, we have

$$(F_{22})^{-1}\leq\frac{F_{11}}{F_{12}^2}\leq\frac{\sigma^2(f_{\rm NL}^{\Delta_1})}{\cos^2(f_{\rm NL}^{\Delta_1},f_{\rm NL}^{\Delta_2})}\frac{F_{11}}{F_{22}}\,, \tag{3.33}$$

---

[12] Note that for the calculation of the correlation matrix for BOSS, we assumed fiducial values of b_{s^2} and $b_{\Pi\Pi^{[2]}}$ based on the Lagrangian local-in-matter-density biasing model [152] and b_{δ^2} based on halo simulation fit in [156]. This choice (rather than setting them to 0 as in most other cases in this paper) leads to minimal differences in the correlation coefficient. In addition, we impose the same Gaussian priors on the loop biases as quoted in [201] which barely affect the forecasted constraints for small Δ and up to a factor of four at large Δ, as expected. Finally, we employ a slightly larger minimum wavenumber $k_{\rm min}=0.01\,h{\rm Mpc}^{-1}$. All these choices are guided by the data analysis performed in §3.5.2.

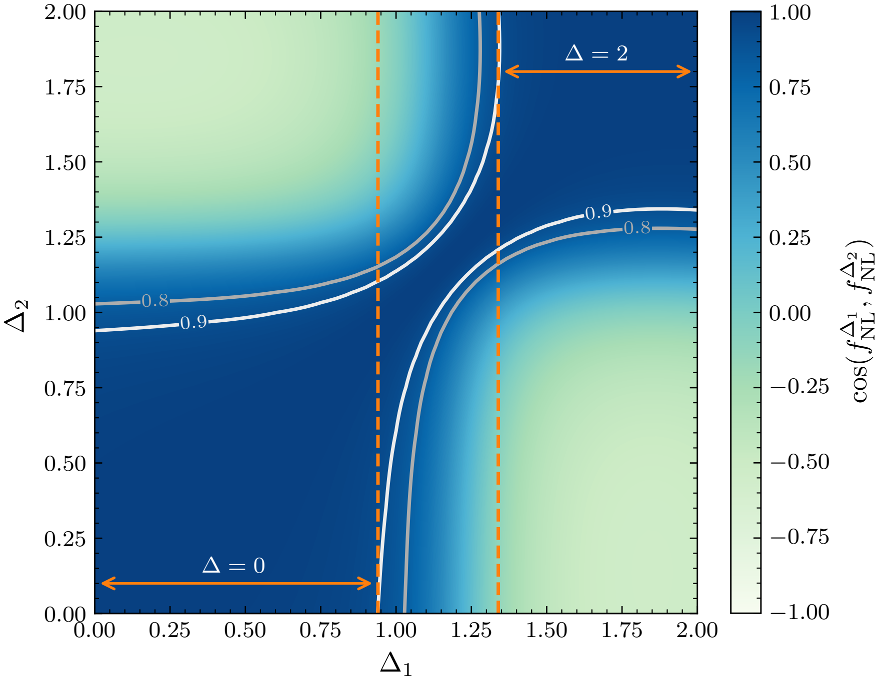

Figure 3.10. Correlation matrix for measurements of galaxy power spectra with different values of the non-Gaussian exponent Δ as defined in (3.32). This was computed for BOSS with fiducial values of $f_{\rm NL}^{\Delta_i} = 0$, fixed ΛCDM parameters and marginalized over all biases. The orange arrows indicate the coverage of the local ($\Delta = 0$) and equilateral ($\Delta = 2$) templates with a correlation coefficient of larger than 0.9.

where $\sigma^2(f_{\rm NL}^{\Delta_1})$ is the variance of the measurement of $f_{\rm NL}^{\Delta_1}$ while holding $f_{\rm NL}^{\Delta_2}$ fixed. The best possible constraint we could place on $f_{\rm NL}^{\Delta_2}$, $\sigma^2(f_{\rm NL}^{\Delta_2})_{\rm CR} = F_{22}^{-1}$, is therefore bounded by the measurement of $f_{\rm NL}^{\Delta_1}$ via

$$\sigma^2(f_{\rm NL}^{\Delta_2})_{\rm CR} \leq \frac{\sigma^2(f_{\rm NL}^{\Delta_1})}{\cos^2(f_{\rm NL}^{\Delta_1}, f_{\rm NL}^{\Delta_2})}\frac{F_{11}}{F_{22}}. \tag{3.34}$$

As a result, we can place an approximate constraint on $f_{\rm NL}^{\Delta_2}$ by rescaling the measurement of $f_{\rm NL}^{\Delta_1}$ using the elements of the marginalized Fisher matrix F_{ij}.

It is important to highlight that this strategy is only effective if our parameter inference is consistent with $f_{\rm NL}^{\Delta} = 0$. To be more precise, all values of Δ look similar at low signal to noise for $f_{\rm NL}^{\Delta}$. On the other hand, the data is no longer independent of Δ if one or more of your $f_{\rm NL}^{\Delta}$ measurements at specific Δ show significant evidence for $f_{\rm NL}^{\Delta} \neq 0$. In that case, a measurement at a different value Δ' could produce a much larger or much smaller signal to noise. We refer to Appendix B.2 for more details on this scenario.

3.5.2 BOSS DR12 Analysis

Based on the strategy laid out above, we inferred constraints on $f_{\rm NL}^{\Delta}$ from the BOSS DR12 dataset following the general setup of the analysis for local, orthogonal and equilateral PNG performed in [202, 136] (see also [135]) with a few different (generally more conservative) choices. In the following, we provide a brief overview of our data analysis, describe these differences and report the limits on PNG that we inferred directly and via the correlation matrix (3.32).

We analyzed the power spectrum from galaxy clustering data of the twelfth and final release of BOSS, referred to as DR12 [149]. This dataset contains the positions of about 1.2×10^6 galaxies between redshifts 0.2 and 0.75 in a cosmic volume of roughly $5.8\,h^{-3}\,{\rm Gpc}^3$ divided into four subsets: the Northern and Southern galactic caps of the two non-overlapping redshift bins $[0.2, 0.5]$ and $[0.5, 0.75]$ with effective mean redshifts $z = 0.38$ and 0.61. We employ the galaxy power spectrum multipoles $P_{gg}^{(\ell)}(k)$, with $\ell = 0, 2, 4$, measured using a quadratic window-function-free estimator [203]. The corresponding covariance matrices were computed from MultiDark-Patchy mock catalogs [204]. This follows the analyses in [202, 136] restricted to the power spectrum multipoles. We model the nonlinear galaxy power spectrum in redshift space as implemented in `CLASS-PT` [205] (based on `CLASS` [206]) which follows the analyses of [207, 208] based on the effective field theory of large-scale structure (EFTofLSS) [209, 210, 211, 212] (see [117, 213] for recent reviews) and galaxy bias expansion [152], including the Alcock-Paczynski effect. We additionally include the non-Gaussian contributions from the scale-dependent bias (3.22) due to $f_{\mathrm{NL}}^{\Delta} \neq 0$ similar to the analyses in [202, 136].[13]

We performed a Markov chain Monte Carlo (MCMC) analysis of the data using the Metropolis-Hastings sampler of `MontePython` [214, 215], varying the theoretical galaxy power spectrum multipoles in each step. We fix the ΛCDM parameters to the Planck 2018 best-fit values [41] (with the sum of neutrino masses $\sum m_\nu = 0.06\,\mathrm{eV}$) instead of including Planck information on ΛCDM since we saw essentially no change in the forecasted constraints on f_{NL}^{Δ} of Section 3.4 when marginalizing over all of our galaxy bias parameters. Per redshift bin and galactic cap, we however separately vary the bias parameters $b_1^{(i)}$, $b_2^{(i)}$ and $b_{\mathcal{G}_2}^{(i)}$ as defined in [216, 217],[14] and analytically marginalize over the other eight biases and EFTofLSS counterterms which appear linearly in the theoretical galaxy power spectrum [218]. We impose flat priors on the linear biases $b_1^{(i)} \in [1,4]$ and the non-Gaussian amplitude f_{NL}^{Δ} (infinitely wide), and Gaussian priors on the other nuisance parameters as described in [201]. Following [202], we also marginalize over the non-Gaussian bias $b_\phi \to N_{b_\phi} b_\phi$ separately for each data subset and impose a wide Gaussian prior on its normalization, $N_{b_\phi} \sim \mathcal{N}(1,5)$, which is motivated by the peak-background split model [177].[15] On the other hand, the non-Gaussian parameter f_{NL}^{Δ} is

---

[13] In this paper, we have not considered the imprints of light inflationary fields in the bispectrum, but focused on the scale-dependent bias in the power spectrum. To remain model agnostic, we also did not include the contribution of a non-zero primordial bispectrum to the galaxy power spectrum.

[14] The bias expansions used in our forecasts and data analysis are equivalent.

[15] It might appear surprising that we marginalize over b_ϕ with a Gaussian prior given the degeneracy with f_{NL}^{Δ}. Note that the mean value of the Gaussian is still set by the universality relation (3.23) with $p = 1$, but its variance is very wide. In addition, we marginalize over b_ϕ in each of the four data subsets separately, as we do with the Gaussian bias parameters. This procedure alleviates degeneracies between f_{NL}^{Δ} and b_1 which can be severe if we assume that the universality relation is exact, especially for $\Delta \sim 1$. This procedure is an imperfect solution to address

commonly varied for all four data subsets at fixed exponent Δ. We (more) conservatively limit the range of wavenumbers to $k \in [0.01, 0.13]\,h\,\mathrm{Mpc}^{-1}$ and $[0.01, 0.16]\,h\,\mathrm{Mpc}^{-1}$ for the low- and high-redshift bin, respectively, after initial tests. We therefore employ a smaller maximum wavenumber than in [136, 202], who tested the validity of the analysis for $\Delta = 0$ and 2 on mock catalogs, and a slightly higher (same) minimum (maximum) wavenumber than the forecasts shown in Fig. 3.3. Since $\Delta k = 0.005\,h\,\mathrm{Mpc}^{-1}$, this choice implies that we employ 24 and 30 k-bins for the respective redshift bins, each power spectrum multipole and both galactic caps. All chains converged with a Gelman-Rubin criterion of $R - 1 < 0.01$ (usually much smaller) for each parameter.

The results of our separate MCMC data analyses for 21 fixed values of the non-Gaussian exponent Δ are shown in Fig. 3.11. Since the posterior distributions for f_{NL}^{Δ} are slightly non-Gaussian, we display the mean values in gray pentagons, and the upper and lower 2σ constraints in filled and unfilled black diamonds, respectively. We do not find any evidence for any PNG shape with $\Delta \in [0, 2]$ and observe the same characteristic functional dependence of $2\sigma(f_{\mathrm{NL}}^{\Delta})$ on Δ as we forecasted in Fig. 3.3.[16] We note that the difference in forecasted constraining power of Fig. 3.3 and the data analysis of Fig. 3.11 can be mainly attributed to the use of Gaussian priors on the higher-order galaxy bias parameters in Section 3.5. In fact, we find reasonable consistency if we include the priors on the Gaussian biases in our Fisher matrices as discussed in §3.5.1. Comparing to [202, 136, 135], our constraints are weaker, as expected due to our more conservative approach and use of less data, but consistent. This is particularly evident for the case of the equilateral shape ($\Delta = 2$) that dramatically benefits from bispectrum information in various ways which is absent in our analysis of the power spectrum alone. On the other hand, our limits on f_{NL}^{Δ} are comparable to the constraints on $f_{\mathrm{NL}}^{\mathrm{eq}}$ when including bispectrum

---

the uncertainty in b_ϕ which we plan to address as part of [219].

[16] The non-Gaussian amplitude f_{NL}^{Δ} is defined relative to the scalar amplitude, $A_{\mathrm{s}}^{1/2} \approx 5 \times 10^{-5}$ [41], such that weak PNG corresponds to $f_{\mathrm{NL}} A_{\mathrm{s}}^{1/2} \ll 1$. This means that our bounds $2\sigma(f_{\mathrm{NL}}^{\Delta}) \sim 10^4$ are approaching the limits of the range of validity of our implicit expansion in linear order in f_{NL}^{Δ}. This concern is however alleviated even with a very conservative inclusion of bispectrum data [202, 135].

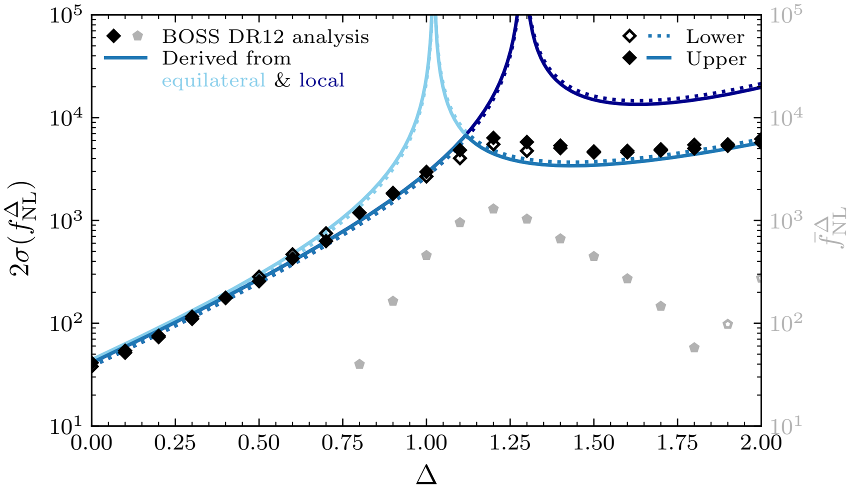

Figure 3.11. Comparison of the measured (diamonds) and derived (lines) values for twice the standard deviation of $f_{\rm NL}^{\Delta}$, $2\sigma(f_{\rm NL}^{\Delta})$, for the BOSS DR12 dataset. The light (dark) blue lines show the approximate value of $2\sigma(f_{\rm NL}^{\Delta})$ derived from the data with $\Delta = 2$ (0), i.e. the equilateral (local) PNG templates using (3.35), with the minimum of these being shown in blue. The diamonds indicate the constraints directly inferred from BOSS DR12 for fixed values of Δ, while the gray pentagons are the corresponding mean values. Filled (unfilled) symbols and solid (dotted) lines represent positive or upper (negative or lower) values. We do not find evidence for $f_{\rm NL}^{\Delta} \neq 0$ for any $\Delta \in [0,2]$, and generally good agreement between the measured and derived limits on general PNG. This demonstrates that our proposed search strategy works on data.

information [202] for $\Delta \lesssim 0.75$. While the bispectrum will always add information, we might however imagine there is a more limited gain in this regime.

We can also compare our bounds on $f_{\rm NL}^{\Delta}$ to those obtained in previous analyses of CMB and LSS data. As already anticipated and explained in Section 3.4, the CMB limits on local and equilateral PNG from the Planck 2018 bispectra are significantly better [106]. This is also the case for the constraints for $\Delta \in [0,3/2]$ derived from Planck 2013 data [220], $|f_{\rm NL}^{\Delta}| \lesssim 50$, which is dominated by the equilateral bispectrum while our bounds arise from the squeezed limit of the quasi-single-field shapes. The only previous analysis constraining $f_{\rm NL}^{\Delta}$ from LSS data was performed in [150] on SDSS-III DR8 data with more limited modeling and range of scales. It is therefore unsurprising that our bounds are considerably more constraining.

In addition to directly inferring constraints on $f_{\rm NL}^{\Delta}$ from the data, there exists a second path to inferring the bounds at a given value of Δ based on the correlation of different values of Δ in the marginalized galaxy power spectrum as discussed in §3.5.1. We take $\Delta = \Delta_1 = 0$ and 2 as the fiducial values for this inference since these scaling exponents correspond to the signal from local and equilateral PNG, respectively, which are usually considered in data analyses. Based on (3.34), we obtain these derived bounds as follows:

$$\left[2\sigma(f_{\rm NL}^{\Delta})\right]_{\rm MCMC} = \left[\frac{1}{\cos(f_{\rm NL}^{\Delta}, f_{\rm NL}^{\Delta=0,2})}\frac{\sigma(f_{\rm NL}^{\Delta})}{\sigma(f_{\rm NL}^{\Delta=0,2})}\right]_{\rm Fisher}\left[2\sigma(f_{\rm NL}^{\Delta=0,2})\right]_{\rm MCMC}, \tag{3.35}$$

where the subscript 'Fisher' ('MCMC') indicates that these quantities are calculated based on the Fisher matrix (from the MCMC analysis). The correlation term $\cos(f_{\rm NL}^{\Delta}, f_{\rm NL}^{\Delta=0,2})$ is displayed in Fig. 3.10, while the standard deviations $\sigma(f_{\rm NL}^{\Delta})$ are computed from the marginalized one-dimensional Fisher matrix as in Section 3.4, but including the same Gaussian priors as in the correlation term and the data analysis. The constraints derived from the measurement at $\Delta = 0$ and 2 are displayed in Fig. 3.11 as the dark and light blue curves. We see that the constraint derived from $2\sigma(f_{\rm NL}^{\Delta=0})$ dominates for $\Delta \lesssim 1.1$, while the constraint derived from $2\sigma(f_{\rm NL}^{\Delta=2})$ dominates for larger Δ, which is consistent with our expectations from Fig. 3.10.

Overall, we find generally reasonable agreement between the directly inferred limits on $f_{\rm NL}^{\Delta}$, and those jointly derived from the local and equilateral measurements. This is especially true for those values of Δ with large correlation and negligible mean values. This implies that we can use (3.35), with the correlation term calculated in a Fisher matrix approach for a given survey, to derive $2\sigma(f_{\rm NL}^{\Delta\neq 0,2})$ in the absence of a detection of $f_{\rm NL}^{\rm loc}, f_{\rm NL}^{\rm eq} \neq 0$. In other words, we have demonstrated that the search strategy that we proposed in §3.5.1 works on BOSS data and allows to dramatically reduce the number of analyses needed to constrain general forms of primordial non-Gaussianity and light inflationary fields as parametrized by Δ from measurements of local and equilateral non-Gaussianity. The agreement between our analysis and forecasts, together with the consistency with prior analyses, is an additional sign that searches for light fields with the galaxy power spectrum will be competitive with the CMB for near-term galaxy surveys.

3.6 Multiple Tracers and the Dependence on Astrophysics

Sample-variance cancellation plays an important part in improving the measurement of primordial non-Gaussianity from the galaxy power spectrum. Applying this technique however requires knowledge of the details of the specific galaxy samples used in the multi-tracer forecast or analysis. In this section, we will first investigate how an optimal survey design can maximize the scientific return of multiple tracers in spectroscopic surveys. Then, we will consider the untapped potential of a multi-tracer configuration of LSST and compare it with to SPHEREx.

3.6.1 Optimizing Spectroscopic Multi-Tracer Analyses

We now explore the potential of multi-tracer analyses and how astrophysical details of the galaxy samples affect the constraints on $f_{\rm NL}^{\Delta}$. Given our biasing model (3.7) with (3.19), this comes down to the description of the non-Gaussian bias parameter $b_\phi(z)$, which depends on three parameters when using the universality relation (3.23): δ_c, $b_1(z)$ and p. Since the parameter δ_c is completely degenerate with $f_{\rm NL}^{\Delta}$, we will focus on exploring what combination of $b_1(z)$ and p

leads to the biggest improvement on $\sigma(f_{\rm NL}^{\Delta})$.

For galaxy power spectrum forecasts to recover all the information in the primordial statistics, we need to apply the multi-tracer technique and its ability for sample-variance cancellation at very high number density. For example, we know the information in the primordial bispectrum for $\Delta > 0$ is dominated by $k \sim k_{\rm max}$ [200]. On the other hand, the non-Gaussian information in the single-tracer power-spectrum forecasts are limited by sample variance at small wavenumbers for $\Delta \lesssim 0.5$. Sample-variance cancellation eliminates this artificial dependence on the noise at low k, limiting a measurement instead by the shot noise associated with galaxy formation [221], which corresponds to modes at $k_{\rm max} \sim k_{\rm halo}$.

Realistic spectroscopic surveys are unfortunately far from being in the fully multi-tracer regime as we have seen in Fig. 3.4. As a result, it is harder to get reliable intuition for how to optimize the measurement of $f_{\rm NL}^{\Delta}$. We have previously explored how to maximize the science of a spectroscopic survey by changing the experimental configuration in terms of $f_{\rm sky}$, $z_{\rm max}$, etc. Here, we will assume a fixed survey geometry and investigate what additional information can be extracted by taking advantage of sample-variance cancellation.

In the fully multi-tracer regime, the inverse variance depends on the linear bias $b_1^{(i)}$ and non-Gaussian bias $b_\phi^{(i)}$ as [222]

$$\sigma(f_{\rm NL}^{\Delta})^{-2} \propto \left| b_1^{(1)} b_\phi^{(2)} - b_1^{(2)} b_\phi^{(1)} \right| . \tag{3.36}$$

On the other hand, we simply maximize $b_1 b_\phi$ in the single-tracer regime. The question therefore is how we can choose the samples to minimize $\sigma(f_{\rm NL}^{\Delta})$ in realistic surveys. In Figure 3.12,[17] we show the relative comparison of $\sigma(f_{\rm NL}^{\Delta})$ between a double- and single-tracer forecast of the

---

[17] Figure 3.12 only shows the forecasted constraints and improvements for the scaling exponent $\Delta = 0$, i.e. local PNG. When we marginalize over all biases, the scaling behavior of the scale-dependent bias however has little effect on the improvement factor and the marginalized biasing model only changes the overall constraining power as indicated by the color bar in the upper panels. In other words, the overall qualitative behavior conveyed in Fig. 3.12 remains intact for different biasing models and scaling exponents.

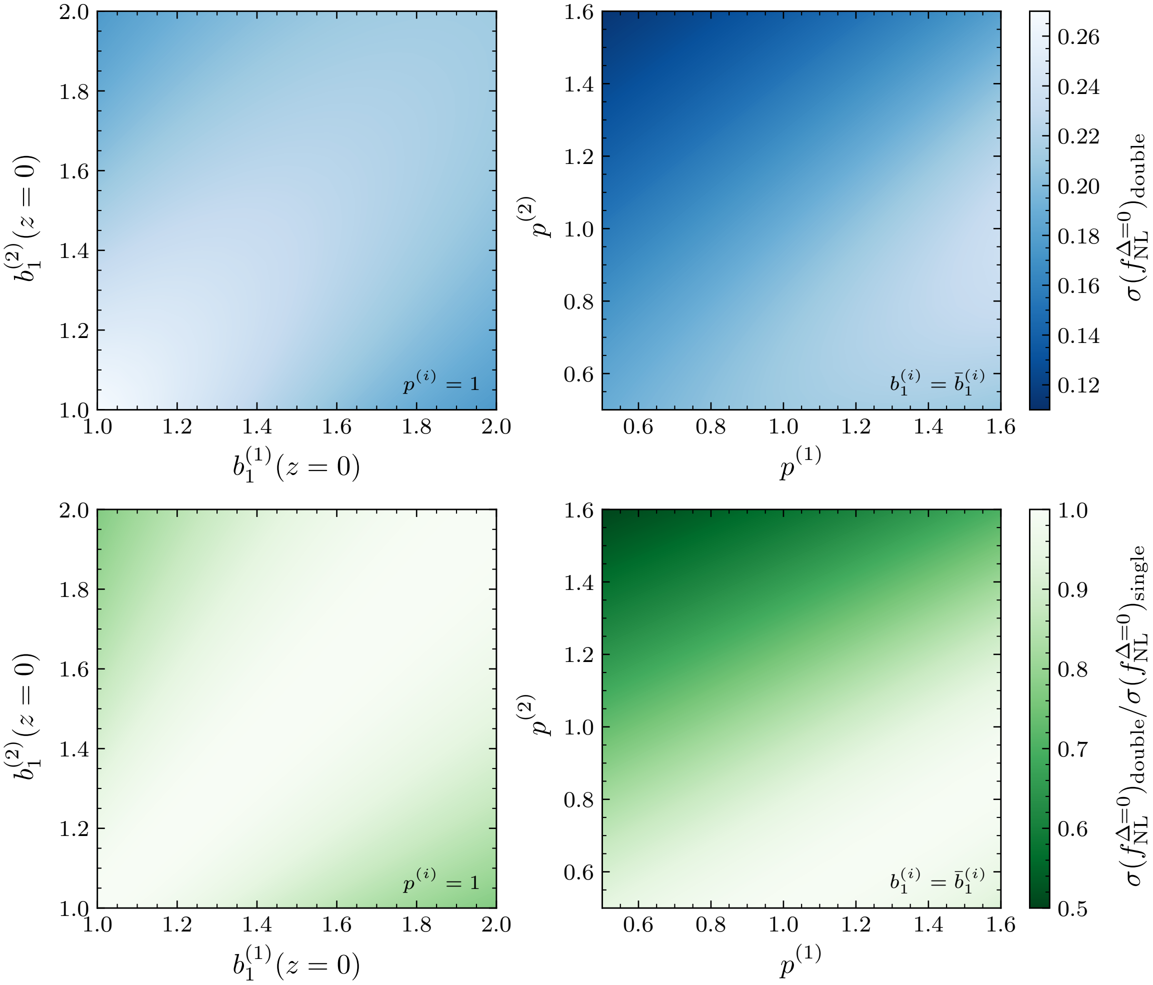

Figure 3.12. Double-tracer forecast for $\sigma(f_{\rm NL}^{\Delta})$ with $\Delta=0$ for the billion-object survey (*top*) and relative comparison to a single-tracer forecast, $\sigma(f_{\rm NL}^{\Delta})_{\rm double}/\sigma(f_{\rm NL}^{\Delta})_{\rm single}$ (*bottom*), with fixed ΛCDM parameters. In the left column, we vary the linear biases $b_1^{(i)}$ of the two samples at fixed $p^{(i)}=1$, i.e. similar to our forecasts in Section 3.4. In the right column, we use the fiducial biases of the samples [cf. (3.8)], with $\bar{b}_1^{(1)}(z=0)=2.0>1.2=\bar{b}_1^{(2)}(z=0)$, and vary the parameter $p^{(i)}$ in the universality relation for $b_\phi^{(i)}$ [cf. (3.23)]. While we only show these results for $\Delta=0$, the qualitative features of these panels remain the same for other values of the scaling exponent.

billion-object survey while holding the ΛCDM parameters fixed.[18] We show this for $\Delta = 0$ (see also e.g. [223]) since we expect this case to benefit the most from a multi-tracer analysis, but explicitly verified that the forecasts with $\Delta > 0$ show at most comparable levels of improvement.

We vary the linear biases of the two subsamples of the survey in the left column of Fig. 3.12, while fixing $p = 1$, as in all of our forecasts in Section 3.4. Note that we only vary the bias at $z = 0$, as labeled in the figure, with the biases of the actual galaxy samples then being calculated based on the bias evolution model of (3.8). The values along the diagonal $b_1^{(1)} = b_1^{(2)}$ correspond to a single-tracer forecast, in which $\sigma(f_{\rm NL}^{\Delta})$ simply decreases with increasing $b_1^{(i)}$. We observe that the multi-tracer improvement is most pronounced here when the difference between $b_1^{(1)}$ and $b_1^{(2)}$ is maximized. We additionally see that the improvement of the double- over the single-tracer approach is limited to about 30%.

So far, we have assumed a universal halo mass function which corresponds to $p = 1$ in the universality relation (3.23). Even though this has been widely adopted in the literature, this assumption is not a perfect description of halos in simulations and will likely not be generally true for actual galaxy samples, cf. e.g. [224, 225, 226, 227, 228, 229, 230, 231, 232, 233]. We therefore now explore the possibility of improving the constraints on $f_{\rm NL}^{\Delta}$ by selecting tracers within the parametrization of b_ϕ through p (assuming such knowledge is already reliably known from simulation results).

We therefore vary the parameters $p^{(i)}$ of the two galaxy samples while using their fiducial biases in the right column of Fig. 3.12. Note that the fact that the linear biases for the two tracers are different, with $b_1^{(1)}(z = 0) = 2.0$ and $b_1^{(2)}(z = 0) = 1.2$, is the reason for the asymmetry of these panels. We observe the largest improvement for small $p^{(1)}$ and large $p^{(2)}$ which is exactly when the combination shown in (3.36) is maximized. In this case, the double-tracer approach can improve the constraints by as much as a factor of two over a single-tracer analysis which

---

[18] When marginalizing over the biases, the difference between the forecasts for fixed and marginalized ΛCDM parameters (with a Planck prior) is at most a few percent. We therefore fix the ΛCDM parameters to their fiducial values for all forecasts in this section.

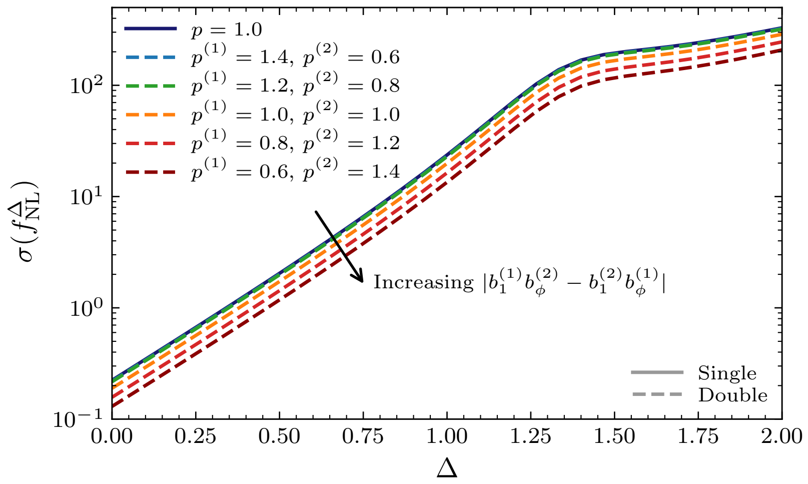

Figure 3.13. Forecasted constraints on f_{NL}^{Δ} for different values of the parameter $p^{(i)}$ in the universality relation (3.23) of b_ϕ for the double-tracer configuration of the billion-object survey and fixed ΛCDM parameters. We observe that $\sigma(f_{\mathrm{NL}}^{\Delta})$ improves for all Δ when $|b_1^{(1)} b_\phi^{(2)} - b_1^{(2)} b_\phi^{(1)}|$ increases which is the direction indicated by the arrow.

is much larger than what we saw for fixed $p^{(i)}$ and varying $b_1^{(i)}$ in the left panels. Figure 3.13 displays the constraining power as a function of the scaling exponent Δ for a few pairs of tracers characterized by different values of $p^{(i)}$.[19] This is further confirmation of our findings and we explicitly see that maximizing the combination of biases given in (3.36) leads to the same relative improvement in $\sigma(f_{\mathrm{NL}}^{\Delta})$ for all Δ.

The results shown in the previous figures suggest that it may be possible to improve $\sigma(f_{\mathrm{NL}}^{\Delta})$ by more than a factor of two if we had knowledge about the parameter $p^{(i)}$ of the tracers. As previously discussed, however, most forecasts for current or near-future galaxy surveys assume the scale-dependent bias takes the form predicted by a universal halo mass function. In particular, these surveys do not select targets based on the value of p (or b_ϕ) which implies that there is some level of uncertainty in the actual value of p for the employed galaxy samples. Since the parameter p is degenerate with f_{NL}^{Δ}, this uncertainty in p could complicate the measurement of f_{NL}^{Δ}. While complete ignorance of p (i.e. marginalizing over it as a free

---

[19] We also tested the possibility of using galaxy samples for which the parameter $p^{(i)}$ is redshift dependent. The results shown here are robust to these changes.

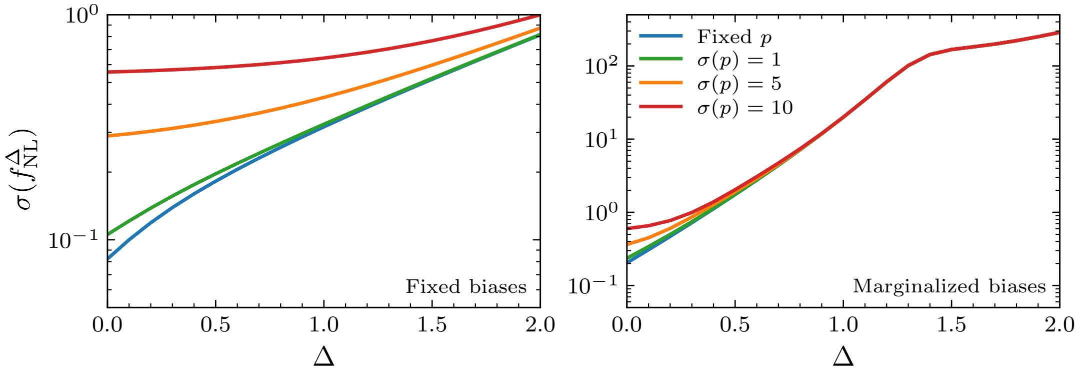

Figure 3.14. Forecast for $\sigma(f_{\rm NL}^{\Delta})$ for different Gaussian priors $\sigma(p)$ on the parameter p in the universality relation (3.23) for the billion-object survey (at fixed ΛCDM parameters and with fiducial value $\bar{f}_{\rm NL}^{\Delta} = 1$ so that p can be independently varied). While marginalizing over p with wide priors significantly impacts the constraints for small Δ in the case of fixed bias parameters (*left*), the realistic case of marginalizing over the biases (*right*) shows a relatively small degradation even for very wide priors on p.

parameter in each redshift bin) would make constraining $f_{\rm NL}^{\Delta}$ impossible, it would be helpful to know the required level of precision so that the universality assumption does not significantly affect the inference of $f_{\rm NL}^{\Delta}$.

We investigate this question by imposing different Gaussian priors on the parameter p for a fiducial value of $p = 1$. Figure 3.14 shows that an uncertainty in p mostly affects $\Delta \lesssim 0.5$ when marginalization over the biases. In addition, the degradation in the constraining power on $f_{\rm NL}^{\Delta}$ is limited to within a factor of less than three (two) for an extremely conservative prior with width $\sigma(p) = 10\,(5)$. This implies that a lack of knowledge of the precise value of p for the given galaxy samples does not significantly affect the forecasts. On the other hand, a careful consideration of the values of p in target selection could potentially lead to important improvements in $\sigma(f_{\rm NL}^{\Delta})$.

3.6.2 Multi-Tracer LSST and Comparison to SPHEREx

So far, we have seen that the multi-tracer forecasts offer somewhat limited potential for improvements in the context of both current and future, realistic spectroscopic surveys. Acquiring spectra takes more observational time than imaging alone which makes it harder for spectroscopic surveys to reach the high number densities needed for sample-variance cancellation to be most effective. In principle, photometric surveys, such as LSST, are therefore not limited in the same way, as indicated in Fig. 3.4, and have the potential to benefit significantly from a multi-tracer analysis. Following the LSST science book [188], our LSST forecasts however use a single tracer as shown in Table B.7 and, consequently, do not capture the full potential of this survey which we will now investigate.

In order to explore the potential for an LSST multi-tracer analysis, we consider a simplified case in which we split the LSST sample of Table B.7 into two subsamples of equal number density but different linear biases $b_1^{(i)}$.[20] In Figure 3.15, we compare the constraints $\sigma(f_{\rm NL}^{\Delta})$ from single- and double-tracer configurations of LSST and the billion-object survey. As for all of our spectroscopic surveys with their given number densities, there is a relatively small change in the constraining power for the billion-object survey.[21] On the other hand, LSST shows a remarkable improvement which is consistent with the factor of roughly three seen in Fig. 3.4.

In this regard, LSST has much more potential for testing inflationary physics than is typically seen in forecasts. Of course, spectroscopic information may prove essential in eliminating systematics or reliably splitting the LSST sample. The improvements that we see in Fig. 3.15 however strongly encourage a future investigation that accounts for systematic effects. In addition, we have not included the information from cosmic-shear observations, which are sensitive to primordial non-Gaussianity in their own right (see e.g. [235]) and may be useful in conjunction with the galaxy power spectra considered here. Cross-correlations with other observables, such as CMB lensing and the kinetic Sunyaev-Zel'dovich effect, would additionally provide complementary constraining power [236, 237, 238].

---

[20] For simplicity, we split the LSST sample in the same way as the sample for the billion-object survey, i.e. we take $b_1^{(1,2)}(z=0) = b_1^{\rm single}(z=0) \pm 0.4$ with equal number densities, $\bar{n}_g^{(1)} = \bar{n}_g^{(2)}$, to retain the single-tracer LSST sample when combined. We refer to [234] for similar forecasts for an LSST sample split into "blue" and "red" galaxies.

[21] In this context, it is important to remember the large maximum redshift for this survey ($z_{\rm max} = 5$) which results in a number density more or less comparable to DESI/Euclid for many redshift bins. The same number of objects in a smaller volume, as would for example be the case for a spectroscopic follow-up survey to LSST with $z_{\rm max} = 3$, would not have the same limitation, but its effectiveness for most other science targets would be reduced [102].

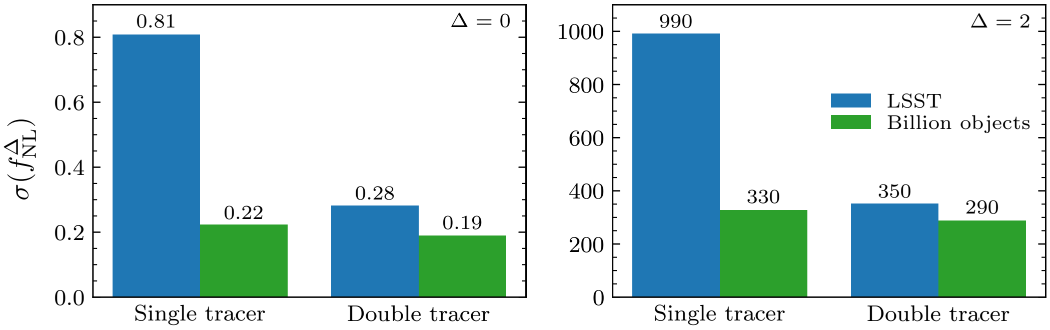

Figure 3.15. Comparison of the constraining power of single- and double-tracer versions of LSST and the billion-object survey (with fixed ΛCDM parameters) for $\Delta = 0$ (*left*) and $\Delta = 2$ (*right*). We can clearly see the significant improvement that a double-tracer analysis of LSST could provide due to its large number density for the entire range of scaling exponents Δ.

Given the potential of LSST as a multi-tracer survey, we should also revisit how it compares to other near-term surveys. We noted in §3.4.2 that LSST and SPHEREx result in similar forecasts despite having different strengths and weaknesses in their designs. SPHEREx is a spectro-photometric survey with a larger sky coverage, while LSST is a photometric survey with a larger number density. The comparison in §3.4.2 is additionally complicated by the fact that we treated LSST as a single-tracer survey following [188] while multi-tracer forecasts would maximize the benefit of the higher number density.

We might have naively imagined that the slight advantage of SPHEREx was due to the smaller photometric redshift error of most of its samples. (Specifically, the redshift uncertainty of LSST is assumed to be $\sigma_{z0} = 0.05$, while the SPHEREx samples are binned with maximum errors of $\sigma_{z0} = \{0.003, 0.01, 0.03, 0.1, 0.2\}$.) This is however not what is found in our forecasts. In Figure 3.9, we see that the potential impact of photometric redshift errors on $f_{\rm NL}^{\Delta}$ for $\Delta \lesssim 1.2$ is at most 50%. At the same time, we observed in Fig. 3.15 that the potential improvements in LSST from splitting the single sample into two is potentially a factor of three or more. This suggests that the number density is a larger effect than the quality of the photometric redshifts.[22] At the same time, there are however many other factors that differ between these surveys and

---

[22] While the observed volume of SPHEREx is in principle more than twice as large (see Appendix B.1), its number densities are much smaller, especially at high redshifts, which results in an effective volume that is smaller by a factor of almost four compared to LSST.

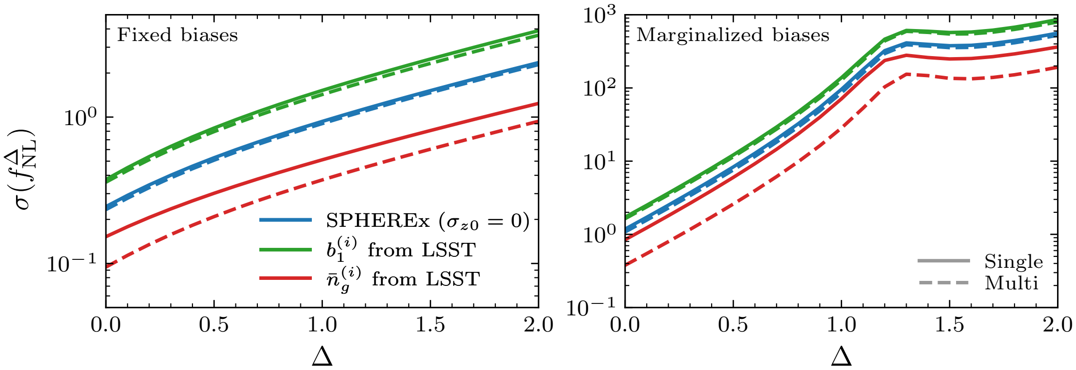

Figure 3.16. Constraining power of SPHEREx with its fiducial biases $b_1^{(i)}$ and number densities $\bar{n}_g^{(i)}$ compared to a survey with either of these observational design factors replaced by LSST-like numbers for fixed (*left*) and marginalized (*right*) Gaussian bias parameters. The solid/dashed lines show the forecasts for a single-/multi-tracer version of these surveys (five samples for SPHEREx and two samples when considering LSST-like design factors) for which we ignore the photometric redshift errors in order to directly infer the sizable impact of these design choices.

we would like to identify the most important individual design factors that set the two surveys apart (at least for the purpose of constraining $f_{\rm NL}^{\Delta}$ from the scale-dependent bias in the power spectrum).

In order to isolate the differences and control the individual factors in survey design, we compare LSST and SPHEREx by their biases and number densities in Fig. 3.16, while ignoring the photometric redshift errors for both surveys, $\sigma_{z0}=0$. To be precise, we take SPHEREx as the fiducial survey and substitute either the linear biases or the number densities from a survey equivalent to LSST (see Appendix B.1 for details). When we keep the number densities of SPHEREx and use the LSST-like biases, i.e. the same biases as in our earlier LSST forecasts just over the larger redshift range of SPHEREx with the same scaling proportional to $D(z)^{-1}$, SPHEREx still significantly outperforms. In this sense, SPHEREx has a built-in advantage from the biases used in our forecasts. If we instead keep the biases of SPHEREx and use the number densities of LSST, i.e. we include objects with LSST-like number densities according to (B.13) beyond its $z_{\rm max}$, the latter becomes more sensitive. In fact, this advantage is quite substantial when we additionally allow for a double-tracer configuration of LSST, which takes full advantage of the difference in number densities. As a result, we should understand that the slightly better sensitivity of SPHEREx indicated in Fig. 3.3 is primarily due to the single-tracer nature of the LSST forecast and the larger biases of the SPHEREx samples.

Before concluding, we have to however state that the largest challenge for both LSST and SPHEREx will of course be identifying and mitigating the impact of large-angle (and other potential) systematics. In this regard, the survey which will be most sensitive to $f_{\rm NL}^{\Delta}$ cannot be easily anticipated ahead of time. Either way, both surveys have the potential to in principle place constraints on general PNG from the galaxy power spectrum alone (i.e. without the galaxy bispectrum) that exceed those inferred from the Planck bispectrum for $\Delta \lesssim 1$.

3.7 Conclusions and Outlook

Primordial non-Gaussianity from scale-dependent bias offers one of the best avenues to test the physics of inflation in the coming decade and beyond. The signal imprinted by light particles coupled to the inflaton is localized on relatively large scales where analyses are less influenced by gravitational nonlinearities, baryonic physics and modeling errors. Equally importantly, the scaling behavior of this signal encodes the mass of these light particles, offering a unique window into the spectrum and interactions of particles during the inflationary epoch.

Pragmatically, the scale-dependent bias offers a method to test for new physics within the power spectra measured in large-scale structure. One might expect that the best measurements of primordial physics should ultimately come from higher-N-point statistics or even from analyses performed directly at the level of the observed maps.[23] Technical obstacles have however slowed down progress and obscured the ultimate sensitivity of LSS surveys. In fact, analyses of the bispectrum only recently became possible for single-field inflation (equilateral PNG) and remain far from the demonstrated sensitivity of the CMB. The established theoretical and observational

---

[23] In some cases, it has already been shown that the signals of extra fields are best understood as map-level features [239, 240, 241, 242, 243, 244] that may carry more signal to noise than low-order statistics.

control of LSS power spectra therefore offer a reliable approach for near-term surveys to impact our understanding of inflation.

Previous work has largely focused on local PNG for which the signal to noise is dominated by only the largest angular scales accessible in a given survey. These measurements are exciting since they have the potential to exceed the sensitivity of the CMB and explore the interesting theoretical threshold of $f_{\rm NL}^{\rm loc} \sim O(1)$. The much richer phenomenology of inflation however raises the question if optimizing a survey for $f_{\rm NL}^{\rm loc}$ is the same as optimizing for understanding inflation more broadly. For example, degeneracies [245, 246, 247] and large-angle (systematic) effects [199, 248, 249, 250, 251, 252] that limit the local PNG measurement may not impact other changes to the statistics. Similarly, some variations in the observational strategy or target selection may not significantly affect the constraining power on $f_{\rm NL}^{\rm loc}$, but could yield a vast increase in the insights that we can gain on other inflationary observables.

In order to extend the reach of current and future surveys to these compelling inflationary targets, we studied primordial non-Gaussianity associated with light fields during inflation over a wide range of their masses. We found that surveys optimized for $f_{\rm NL}^{\rm loc}$ tend to perform equally well over a broader mass range beyond the massless (local) limit. Furthermore, we placed new constraints on these types of PNG by analyzing the BOSS DR12 dataset in two different and consistent ways. These results are consistent with previous analyses in the case of massless fields. In addition, they extend to heavier fields and other types of non-Gaussianity. They also show that bounds on these general non-Gaussian shapes can be simply inferred from scale-dependent bias measurements of $f_{\rm NL}^{\rm loc}$ and $f_{\rm NL}^{\rm eq}$ if they are consistent with zero. Our measured constraints from BOSS data are also consistent with our BOSS forecasts which suggests that our understanding of the constraining power of these surveys, which we described in detail, is reflected in real data. These insights are likely relevant to the analysis strategies for current or near-term observations and the design of future surveys. For instance, we demonstrated the significant improvements in sensitivity that may be available from a multi-tracer analysis of LSST data.

More broadly speaking, additional fields present during inflation are interesting in their own right, far beyond their potential impact on the galaxy power spectrum. While we focused on models leading to power-law corrections to the power spectrum, oscillatory contributions to inflationary spectra can arise from (more) massive fields [175, 143, 253] or chemical potentials [254, 255], for instance. This vast range of phenomenology therefore indicates that there are many opportunities in the analysis of galaxy surveys [256, 234] that remain under-explored and, as we also show in this paper, minor additional work could result in much broader and deeper insights into the primordial universe.

Acknowledgments

We are grateful to Alexandre Barreira, Daniel Baumann, Matteo Biagetti, Kyle Dawson, Olivier Doré, Simone Ferraro, Marilena LoVerde, Azadeh Moradinezhad Dizgah, and Charuhas Shiveshwarkar for helpful discussions. The authors were supported by the US Department of Energy under Grant DE-SC0009919. B. W. also acknowledges support from the Swedish Research Council (Contract No. 638-2013-8993). This work is based on observations obtained by the Sloan Digital Sky Survey III (SDSS-III, http://www.sdss3.org/). Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the US Department of Energy Office of Science. We acknowledge the use of `CAMB` [157], `CLASS` [206], `CLASS-PT` [205], `FAST-PT` [257], `IPython` [96] and `MontePython` [214, 215], and the Python packages `Matplotlib` [97], `NumPy` [99] and `SciPy` [100].

Chapter 3, in full, is a reprint of the material as it appears in Light Fields during Inflation from BOSS and Future Galaxy Surveys, Green, Daniel; Guo, Yi; Han, Jiashu; Wallisch, Benjamin, arXiv:2311.04882. The dissertation author was one of the primary investigators and authors of this paper.

Appendix A

Appendix for Chapter 2

A.1 Computational Details

In this appendix, we provide details underlying our derivation of the axion production rate (2.12) from (2.10). While the final integral (2.12) is eventually evaluated numerically, the analytic reduction of the integral described in Appendix A.1.1 is important for making the numerical evaluation tractable. In Appendix A.1.2, we present additional information on our subsequent calculation of the contributions to $\Delta N_{\rm eff}$ using the Boltzmann equation.

A.1.1 Production Rate Calculation

We want to directly compute the production rate of axions via Compton-like scattering and fermion annihilation using the quantum distribution functions (2.11) for the incoming and outgoing particles which we label by $1,2$ and $3,4$, respectively. The general interaction rate via such two-to-two processes is given by (2.10):

$$\Gamma_\phi = \frac{1}{n_\phi^{\rm eq}} \int {\rm d}\tilde{\Gamma}\, f_1(p_1)\, f_2(p_2)\, [1 \pm f_3(p_3)]\, [1 \pm f_4(p_4)] \sum |\mathscr{M}|^2 \,, \tag{A.1}$$

where we introduced the measure

$${\rm d}\tilde{\Gamma} = \prod_{i=1}^{4} \frac{{\rm d}^3 p_i}{(2\pi)^3 2E_i}\, (2\pi)^4 \delta^{(3)}(\vec{p}_1 + \vec{p}_2 - \vec{p}_3 - \vec{p}_4)\, \delta(E_1 + E_2 - E_3 - E_4)\,. \tag{A.2}$$

In the following, we will change the variables of this 12-dimensional integral, rewrite the integration measure and scattering amplitudes (2.7) and (2.8), and finally reduce the respective production rates to four-dimensional integrals that we will solve numerically.

Parametrization

While there are 12 degrees of freedom in the integration variables $\vec{p}_i$, $i = 1, \ldots, 4$, the four-dimensional energy-momentum conservation enforced by the Dirac delta functions reduces the number of independent degrees of freedom to eight. To parametrize these, we employ the following variables: the absolute values of the total energy, total momentum, one incoming momentum and one outgoing momentum,

$$E = E_1 + E_2 \,, \qquad p = |\vec{p}| = |\vec{p}_1 + \vec{p}_2| \,, \qquad p_1 = |\vec{p}_1| \,, \qquad p_3 = |\vec{p}_3| \,, \tag{A.3}$$

the polar angle ϕ_1 of $\vec{p}_1$, the polar angle ϕ_3 of $\vec{p}_3$, and the remaining two directions of $\vec{p}$. Due to symmetry, we can set $\vec{p} \equiv p\hat{z}$, $\phi_1 \equiv 0$ and $\phi_3 \equiv \phi$. Given this parametrization, we have $p_2 = |\vec{p}_2| = |\vec{p} - \vec{p}_1|$ and $p_4 = |\vec{p}_4| = |\vec{p} - \vec{p}_3|$.

It is useful to define the angle between $\vec{p}$ and $\vec{p}_i$, $\theta_i = \angle(\vec{p}, \vec{p}_i)$, to transform the three-momenta of the particles from Cartesian to spherical coordinates:

$$\begin{aligned} \vec{p}_1 &= p_1 \left(\sin\theta_1, 0, \cos\theta_1\right), & \vec{p}_2 &= p_2 \left(-\sin\theta_2, 0, \cos\theta_2\right), \\ \vec{p}_3 &= p_3 \left(\sin\theta_3 \cos\phi, \sin\theta_3 \sin\phi, \cos\theta_3\right), & \vec{p}_4 &= p_4 (-\sin\theta_4 \cos\phi, -\sin\theta_4 \sin\phi, \cos\theta_4) \,. \end{aligned} \tag{A.4}$$

We can express these angles in terms of the absolute values of the momenta for $i = 1, 3$ and $j = 2, 4$ as follows:

$$\cos\theta_i = \frac{p^2 + p_i^2 - p_{i+1}^2}{2pp_i} \,, \quad \sin\theta_i = \frac{g(p, p_i, p_{i+1})}{2pp_i} \,,$$

$$\cos\theta_j = \frac{p^2 + p_j^2 - p_{j-1}^2}{2pp_j}\,, \quad \sin\theta_j = \frac{g(p, p_{j-1}, p_j)}{2pp_j}\,,$$

where we introduced $g(p,p_i,p_j) = \sqrt{(p+p_i+p_j)(p+p_i-p_j)(p-p_i+p_j)(-p+p_i+p_j)}$. Finally, it is also helpful to introduce the angles $\theta_{13} = \angle(\vec{p}_1,\vec{p}_3)$ and $\theta_{14} = \angle(\vec{p}_1,\vec{p}_4)$, which can be parametrized as

$$\begin{aligned} \cos\theta_{13} &= \sin\theta_1\sin\theta_3\cos\phi + \cos\theta_1\cos\theta_3 \equiv a\cos\phi + b\,, \\ \cos\theta_{14} &= -\sin\theta_1\sin\theta_4\cos\phi + \cos\theta_1\cos\theta_4 \equiv -c\cos\phi + d\,, \end{aligned} \tag{A.5}$$

where we defined a short-hand notation in terms of a, b, c and d in the last equalities, respectively.

Integral Measure

We now turn to the measure of the production rate integral introduced in (A.2) and express it in terms of the new coordinates. By inserting $1 = \int \mathrm{d}^3p\,\delta^{(3)}(\vec{p} - \vec{p}_1 - \vec{p}_2)$, we get

$$\begin{aligned} \int \mathrm{d}\tilde{\Gamma} &= \int \mathrm{d}^3p_1 \int \mathrm{d}^3p_2 \int \mathrm{d}^3p_3 \int \mathrm{d}^3p_4 \, \frac{\delta^{(3)}(\vec{p}_1 + \vec{p}_2 - \vec{p}_3 - \vec{p}_4)\,\delta(E_1 + E_2 - E_3 - E_4)}{(2\pi)^8 16 E_1E_2E_3E_4} \\ &= \int \mathrm{d}^3p_1 \int \mathrm{d}^3p_2 \left[\int \mathrm{d}^3p\,\delta^{(3)}(\vec{p} - \vec{p}_1 - \vec{p}_2)\right] \\ &\qquad \times \int \mathrm{d}^3p_3 \int \mathrm{d}^3p_4 \, \frac{\delta^{(3)}(\vec{p}_1 + \vec{p}_2 - \vec{p}_3 - \vec{p}_4)\,\delta(E_1 + E_2 - E_3 - E_4)}{(2\pi)^8 16 E_1E_2E_3E_4} \\ &= \int \mathrm{d}^3p \int \mathrm{d}^3p_1 \int \mathrm{d}^3p_3 \, \frac{\delta(E_1 + E_2 - E_3 - E_4)}{(2\pi)^8 16 E_1E_2E_3E_4}\bigg|_{\vec{p}_2 = \vec{p} - \vec{p}_1,\, \vec{p}_4 = \vec{p} - \vec{p}_3}\,, \end{aligned} \tag{A.6}$$

where we introduced a short-hand notation for imposing $\vec{p}_j = \vec{p} - \vec{p}_{j-1}$ in the last line which we will further abbreviate below. As previously mentioned, we have the freedom to take $\vec{p} = p\hat{z}$ for fixed p since the direction of $\vec{p}$ does not affect the production rate. Similarly, we can rotate all the particle momenta $\vec{p}_i$ around the total momentum $\vec{p}$ because only their relative angle matters. This motivates choosing the polar angle between the planes spanned by $(\vec{p}_1, \vec{p}_2)$ and $(\vec{p}_3, \vec{p}_4)$, $\phi = \phi_3 - \phi_1$, as an integration variable (which is equivalent to the choice mentioned above). With these choices, the measure becomes

$$\begin{aligned}\int \mathrm{d}\tilde{\Gamma} &= \int \mathrm{d}p\, p^2 \int \mathrm{d}p_1\, p_1^2\, \mathrm{d}\phi_1\, \mathrm{d}\cos\theta_1 \int \mathrm{d}p_3\, p_3^2\, \mathrm{d}\phi_3\, \mathrm{d}\cos\theta_3\, \frac{\delta(E_1+E_2-E_3-E_4)}{(2\pi)^7 8E_1E_2E_3E_4}\bigg|_{\vec{p}_2,\vec{p}_4} \\ &= \int \mathrm{d}E \int \mathrm{d}p\, p^2 \int \mathrm{d}\phi \left[\prod_{i=1,3} \int \mathrm{d}p_i\, p_i^2 \int \mathrm{d}\cos\theta_i\, \delta(E-E_i-E_{i+1})\right] \frac{1}{512\pi^6 E_1E_2E_3E_4}\bigg|_{\vec{p}_2,\vec{p}_4},\end{aligned} \tag{A.7}$$

where we inserted $1 = \int \mathrm{d}E\, \delta(E-E_1-E_2)$ to introduce the total energy E. To evaluate the integrals over the azimuthal angles θ_i, we use $E_{i+1} = E - E_i = (m_{i+1}^2 + p^2 + p_i^2 - 2pp_i \cos\theta_i)^{1/2}$ and $\delta(g(p)) = \sum_k \delta(p-\bar{p}_k)/|g'(\bar{p}_k)|$, with the simple zeros $\bar{p}_k$ of the function $g(p)$, to reparametrize the Dirac delta functions:

$$\int_{-1}^{1} \mathrm{d}\cos\theta_i\, \delta(E-E_i-E_{i+1}) = \frac{E_{i+1}}{pp_i} \int_{-1}^{1} \mathrm{d}\cos\theta_i\, \delta\left(\cos\theta_i - \frac{m_{i+1}^2 + p^2 + p_i^2 - E_{i+1}^2}{2pp_i}\right), \tag{A.8}$$

which we can directly evaluate subject to the finite integration limits.

Finally, we have to appropriately treat the integration limits which we have neglected so far. The integral in (A.8) is non-vanishing only if

$$\left(\frac{m_{i+1}^2 + p^2 + p_i^2 - E_{i+1}^2}{2pp_i}\right)^2 \leq 1\,. \tag{A.9}$$

After defining $A = s + m_i^2 - m_{i+1}^2$ and $B = A^2 + 4E^2 m_i^2$, and introducing the Mandelstam variable $s = E^2 - p^2$, we can expand this inequality to

$$16s^2 \left(p_i^2\right)^2 + \left(8sB - 16E^2A^2\right) p_i^2 + B^2 - 16E^2A^2m_i^2 \leq 0\,, \tag{A.10}$$

which determines the integration limits. To satisfy this quadratic inequality in p_i^2, its determinant must be positive, which requires $s \geq (m_i + m_{i+1})^2$. Since the total momentum p is positive, this implies for the total momentum and energy:

$$p \leq \sqrt{E^2 - (m_i + m_{i+1})^2}\,, \qquad E \geq m_i + m_{i+1}\,. \tag{A.11}$$

Because these requirements have to be simultaneously satisfied for $i = 1, 3$, we get

$$E_{\min} = \max_{i=1,3}\{m_i + m_{i+1}\}\,, \qquad p_{\max} = \min_{i=1,3}\left\{\sqrt{E^2 - (m_i + m_{i+1})^2}\right\}, \tag{A.12}$$

which can also be directly inferred from kinematic considerations. The two solutions to the quadratic inequality (A.10) for p_i^2 then imply the integration limits for p_i, $i = 1, 3$, to be

$$p_i^{\min,\max} = \frac{1}{2s}\left| E\sqrt{[s - (m_i - m_{i+1})^2][s - (m_i + m_{i+1})^2]} \mp (s + m_i^2 - m_{i+1}^2)p \right|, \tag{A.13}$$

with the minus (plus) sign being associated with the lower (upper) limit. To put it all together, we therefore arrive at the following result for the measure:

$$\int \mathrm{d}\tilde{\Gamma} = \int_{E_{\min}}^{\infty} \mathrm{d}E \int_0^{p_{\max}} \mathrm{d}p \int_{p_1^{\min}}^{p_1^{\max}} \mathrm{d}p_1 \int_{p_3^{\min}}^{p_3^{\max}} \mathrm{d}p_3 \int_0^{2\pi} \mathrm{d}\phi \left.\frac{p_1 p_3}{512\pi^6 E_1 E_3}\right|_{\vec{p}_2, \vec{p}_4}, \tag{A.14}$$

with the integration limits given by (A.12) and (A.13).

Scattering Amplitudes

Before turning to the entire production rate calculation, we first rewrite the amplitudes of Compton-like scattering and fermion annihilation, which we provided in (2.7) and (2.8), as a function of the invariant Mandelstam variables s, t and u. In terms of our integration variables, these invariants are given by

$$\begin{aligned} s &= (E_1+E_2)^2-(\vec{p}_1+\vec{p}_2)^2 = E^2-p^2\,, \\ t &= (E_1-E_3)^2-(\vec{p}_1-\vec{p}_3)^2 = m_1^2+m_3^2-2E_1E_3+2p_1p_3\cos\theta_{13}\,, \\ u &= (E_1-E_4)^2-(\vec{p}_1-\vec{p}_4)^2 = m_1^2+m_4^2-2E_1E_4+2p_1p_4\cos\theta_{14}\,. \end{aligned} \tag{A.15}$$

This implies that the scattering amplitude of the Compton-like process can be rewritten as

$$\sum|\mathscr{M}|^2_{(a)} = 16\pi A_\psi\,|\tilde{\varepsilon}_\psi|^2 \frac{2p_1p_3^2(1-\cos\theta_{13})^2}{(E^2-p^2-m_\psi^2)\,(E_4-p_4\cos\theta_{14})}\,, \tag{A.16}$$

while the scattering amplitude of the annihilation process in these coordinates is

$$\sum|\mathscr{M}|^2_{(b)} = 16\pi A_\psi\,|\tilde{\varepsilon}_\psi|^2 \frac{(E^2-p^2)^2}{4p_3p_4(E_1-p_1\cos\theta_{13})(E_1-p_1\cos\theta_{14})}\,. \tag{A.17}$$

We will now separately insert these expressions into the integral (A.1) to obtain the final expressions for the respective production rates.

Compton-Like Scattering Rate

For Compton-like scattering, $\{\gamma,g\}+\psi\to\phi+\psi$, we take the momenta of the massless bosons to be p_1 and p_3, while the fermion ψ has incoming momentum p_2, outgoing momentum p_4 and mass m_ψ. This means that the energies and momenta can be expressed as

$$E_1=p_1\,,\qquad E_3=p_3\,,\qquad E_j=E-p_{j-1}\,,\qquad p_j=\sqrt{(E-p_{j-1})^2-m_\psi^2}\,,$$

for $j=2,4$, while the integration limits for $i=1,3$ are

$$E_{\min}=m_\psi\,,\quad p_{\max}=\sqrt{E^2-m_\psi^2}\,,\quad p_i^{\min}=\frac{E^2-m_\psi^2-p^2}{2(E+p)}\,,\quad p_i^{\max}=\frac{E^2-m_\psi^2-p^2}{2(E-p)}\,.$$

The production rate associated with the scattering process therefore is

$$\begin{aligned}\Gamma_{(a)} &= \frac{1}{512\pi^6 n_\phi^{\mathrm{eq}}}\int \mathrm{d}E \int \mathrm{d}p \int \mathrm{d}p_1 \int \mathrm{d}p_3 \int \mathrm{d}\phi\, f_1(E_1)\, f_2(E_2)\,[1+f_3(E_3)]\,[1-f_4(E_4)])\sum|\mathscr{M}|^2 \\ &= \frac{A_\psi|\tilde{\varepsilon}_\psi|^2}{16\pi^5 n_\phi^{\mathrm{eq}}}\int \mathrm{d}E \int \mathrm{d}p \int \mathrm{d}p_1 \int \mathrm{d}p_3\, \frac{1}{\left(\mathrm{e}^{\frac{p_1}{T}}-1\right)\left(\mathrm{e}^{\frac{E-p_1}{T}}+1\right)\left(1-\mathrm{e}^{-\frac{p_3}{T}}\right)\left(1+\mathrm{e}^{-\frac{E-p_3}{T}}\right)} \\ &\qquad \times \frac{p_1 p_3^2}{E^2-p^2-m_\psi^2}\int_0^{2\pi}\mathrm{d}\phi\, \frac{(1-\cos\theta_{13})^2}{E_4-p_4\cos\theta_{14}}\,.\end{aligned}$$

We can further simplify this expression by analytically performing the angular integral,

$$\frac{1}{2\pi}\int_0^{2\pi}\mathrm{d}\phi\,\frac{p_4(1-\cos\theta_{13})^2}{E_4-p_4\cos\theta_{14}} = \frac{c^2(1-b)^2+[a^2 f+2ac(1-b)](f-\sqrt{f^2-c^2})}{c^2\sqrt{f^2-c^2}}\,,$$

where we employed the short-hand notation introduced in (A.5), with

$$\cos\theta_i = \frac{m_\psi^2-E^2+p^2+2Ep_i}{2pp_i},\quad \cos\theta_4 = \frac{E^2-m_\psi^2+p^2-2Ep_3}{2p\sqrt{(E-p_3)^2-m_\psi^2}},\quad f = \frac{E-p_3}{\sqrt{(E-p_3)^2-m_\psi^2}}-d,$$

for $i=1,3$. Finally, we can use $a/c = p_4/p_3$ to arrive at the final expression for the Compton-like production rate,

$$\begin{aligned}\Gamma_{(a)} = \frac{A_\psi|\tilde{\varepsilon}_\psi|^2}{8\pi^4 n_\phi^{\mathrm{eq}}}\int_{m_\psi}^{\infty}\mathrm{d}E\int_0^{p_{\max}}\mathrm{d}p\int_{p_1^{\min}}^{p_1^{\max}}\mathrm{d}p_1\int_{p_3^{\min}}^{p_3^{\max}}\mathrm{d}p_3\,\frac{1}{\left(\mathrm{e}^{\frac{p_1}{T}}-1\right)\left(\mathrm{e}^{\frac{E-p_1}{T}}+1\right)}\frac{p_1}{E^2-p^2-m_\psi^2} \\ \times\frac{1}{\left(1-\mathrm{e}^{-\frac{p_3}{T}}\right)\left(1+\mathrm{e}^{-\frac{E-p_3}{T}}\right)}\frac{p_3^2(1-b)^2+\left[p_4^2 f+2p_3p_4(1-b)\right]\left(f-\sqrt{f^2-c^2}\right)}{p_4\sqrt{f^2-c^2}}\,,\end{aligned}$$

with $p_4^2\,(f^2-c^2) = (E_4\cos\theta_1 - p_4\cos\theta_4)^2 + m_\psi^2\sin^2\theta_1$.

Fermion Annihilation Rate

For fermion-antifermion annihilation, $\psi+\bar{\psi}\to\{\gamma,g\}+\phi$, we assign the incoming momenta p_1 and p_2 to the fermion and antifermion with mass m_ψ, and the outgoing momenta p_3

and p_4 to the massless vector boson and axion, respectively. In consequence, the energies and momenta associated with this process are

$$E_1=\sqrt{p_1^2+m_\psi^2}\,,\quad E_3=p_3\,,\quad E_j=E-p_{j-1}\,,\quad p_2=\sqrt{(E-E_1)^2-m_\psi^2}\,,\quad p_4=E-p_3\,,$$

for $j=2,4$, while the integration limits are given by

$$E_{\min}=2m_\psi\,,\quad p_{\max}=\sqrt{E^2-4m_\psi^2}\,,\quad p_1^{\min,\max}=\frac{1}{2}\left|\frac{E}{\sqrt{s}}\sqrt{s-4m_\psi^2}\mp p\right|\,,\quad p_3^{\min,\max}=\frac{E\mp p}{2}\,,$$

where the minus (plus) signs are associated with the minimum (maximum) particle momenta. We can therefore write the interaction rate for the annihilation process as

$$\begin{aligned}\Gamma_{(b)}&=\frac{1}{512\pi^6 n_\phi^{\text{eq}}}\int\mathrm{d}E\int\mathrm{d}p\int\mathrm{d}p_1\int\mathrm{d}p_3\int\mathrm{d}\phi\,\frac{p_1}{E_1}f_1(E_1)f_2(E_2)[1+f_3(E_3)][1+f_4(E_4)])\\&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\times\sum|\mathscr{M}|^2\\&=\frac{A_\psi|\tilde{\varepsilon}_\psi|^2}{128\pi^5 n_\phi^{\text{eq}}}\int\mathrm{d}E\int\mathrm{d}p\int\mathrm{d}p_1\int\mathrm{d}p_3\,\frac{1}{\left(\mathrm{e}^{\frac{E_1}{T}}+1\right)\left(\mathrm{e}^{\frac{E-E_1}{T}}+1\right)\left(1-\mathrm{e}^{-\frac{p_3}{T}}\right)\left(1-\mathrm{e}^{-\frac{E-p_3}{T}}\right)}\\&\qquad\qquad\times\frac{(E^2-p^2)^2}{E_1p_3p_4}\int_0^{2\pi}\mathrm{d}\phi\,\frac{p_1}{(E_1-p_1\cos\theta_{13})(E_1-p_1\cos\theta_{14})}.\end{aligned}$$

As for the Compton-like process, we can again compute the angular integral analytically and express it in terms of the short-hand notation introduced in (A.5),

$$\frac{1}{2\pi}\int_0^{2\pi}\mathrm{d}\phi\,\frac{p_1^2}{(E_1-p_1\cos\theta_{13})(E_1-p_1\cos\theta_{14})}=\frac{\frac{a}{\sqrt{(f-b)^2-a^2}}+\frac{c}{\sqrt{(f-d)^2-c^2}}}{a(f-d)+c(f-b)}\,,$$

with $f=E_1/p_1$ and the following underlying expressions:

$$\cos\theta_1=\frac{m_\psi^2-(E-E_1)^2+p^2+p_1^2}{2pp_1},\ \cos\theta_3=\frac{-E^2+p^2+2Ep_3}{2pp_3},\ \cos\theta_4=\frac{E^2+p^2-2Ep_3}{2p(E-p_3)}.$$

After rewriting the result of the angular integral, the final expression for the production rate via fermion annihilation that we implemented numerically is given by

$$\Gamma_{(b)} = \frac{A_\psi |\tilde{\varepsilon}_\psi|^2}{32\pi^4 n_\phi^{\mathrm{eq}}} \int_{2m_\psi}^{\infty} \mathrm{d}E \int_0^{p_{\max}} \mathrm{d}p \int_{p_1^{\min}}^{p_1^{\max}} \mathrm{d}p_1 \int_{p_3^{\min}}^{p_3^{\max}} \mathrm{d}p_3 \, \frac{1}{\left(\mathrm{e}^{\frac{E_1}{T}}+1\right)\left(\mathrm{e}^{\frac{E-E_1}{T}}+1\right)} \frac{(E^2-p^2)p_1}{E_1}$$
$$\times \frac{1}{\left(1-\mathrm{e}^{-\frac{p_3}{T}}\right)\left(1-\mathrm{e}^{-\frac{E-p_3}{T}}\right)} \left(\frac{1}{p_3\, h(\theta_3)} + \frac{1}{p_4\, h(\theta_4)}\right),$$

where we defined $h(\theta_k) = \sqrt{m_\psi^2 \sin^2\theta_1 + (E_1\cos\theta_1 - p_1\cos\theta_k)^2}$.

A.1.2 Boltzmann Equation and $\Delta \mathbf{N}_{\mathbf{eff}}$

Having obtained the total production rate $\Gamma = 2\Gamma_{(a)} + \Gamma_{(b)}$ by numerically computing the four-dimensional integrals stated above, we computed the resulting contribution to the radiation density in the early universe as parameterized by N_{eff}. In the following, we provide additional details of the underlying computational steps.

First, we solve the Boltzmann equation (2.13) to calculate the axion number density $n_\phi(t)$. Instead of directly solving (2.13), we however adopt the conventional change of variables to the dimensionless time variable $x = m/T$ and the dimensionless comoving number density $Y_\phi = n_\phi/s$, with the entropy density $s = 2\pi^2 g_{*s}T^3/45$. Conservation of entropy in the early universe, $a^3 s = \mathrm{const}$, implies $\dot{s}/s = -3\dot{a}/a = -3H$, where the overdot denotes a derivative with respect to time t, and allows to express (2.13) as

$$\dot{Y}_\phi = \Gamma_\phi \left(Y_\phi^{\mathrm{eq}} - Y_\phi\right). \tag{A.18}$$

Rewriting the derivative with respect to t in terms of x leads to

$$Hx\frac{\mathrm{d}Y_\phi}{\mathrm{d}x} = \left(1 - \frac{1}{3}\frac{\mathrm{d}\log g_{*s}}{\mathrm{d}\log x}\right)\Gamma_\phi \left(Y_\phi^{\mathrm{eq}} - Y_\phi\right), \tag{A.19}$$

where we used the definition of the entropy density s.

We numerically solve this equation for Y_ϕ from an initial condition of no axions, $Y_\phi(T_0) = 0$, to the final late-time value of $Y_{\phi,\infty} = Y_\phi(T_\infty)$. We take the initial temperature to be the temperature of the electroweak crossover, $T_0 = T_{\mathrm{EW}} = 159.5\,\mathrm{GeV}$ [258], since the Lagrangian (2.6) and, therefore, the computed production rates are only valid after electroweak symmetry breaking. Due to Boltzmann suppression for $T \ll m_\psi$, it is sufficient to compute the number density for $T_\infty = m_\psi/100$.

Finally, we have to convert the computed value of $Y_{\phi,\infty}$ to a contribution to ΔN_{eff} as defined in (2.3). At late times, the axion energy and number densities, the photon density and the entropy density are given by

$$\rho_\phi = \frac{\pi^2}{30} T_\phi^4\,, \qquad n_\phi = \frac{\zeta(3)}{\pi^2} T_\phi^3\,, \qquad \rho_\gamma = \frac{\pi^2}{15} T_\gamma^4\,, \qquad s = \frac{2\pi^2}{45} g_{*s,\infty} T_\gamma^3\,, \tag{A.20}$$

with the effective number of relativistic degrees of freedom in entropy only counting photons and neutrinos, $g_{*s,\infty} = 43/11$. We therefore arrive at the following expression for the contribution to the effective number of relativistic degrees of freedom:

$$\Delta N_{\mathrm{eff}} = \frac{4}{7}\left(\frac{11\pi^4}{90\zeta(3)} g_{*s,\infty} Y_{\phi,\infty}\right)^{4/3} = \frac{4}{7}\left(\frac{43\pi^4}{90\zeta(3)} Y_{\phi,\infty}\right)^{4/3}\,, \tag{A.21}$$

which approximately evaluates to $\Delta N_{\mathrm{eff}} \approx 74.84\, Y_{\phi,\infty}^{4/3}$.

A.2 Production Rate Comparisons and Uncertainties

In this appendix, we examine a few aspects of the computed production rate of axions and other pNGBs. We first consider the differences between employing the full quantum distribution functions $f^i(p)$, $i = f, b$, instead of the classical Boltzmann distribution function and the relative importance of the Compton-like and fermion-antifermion annihilation processes in the production (Appendix A.2.1). Then, we describe the uncertainties associated with our calculations involving the axion coupling to the bottom and charm quarks due to the QCD phase transition and their potential impact on our predictions for ΔN_{eff} as a function of the coupling constants $\Lambda_{\{b,c\}}$ (Appendix A.2.2).

A.2.1 Quantum Statistics and Production Rates

Commonly, the distribution functions in the production rate are approximated by Boltzmann distributions and the Bose enhancement and Pauli blocking are neglected. We went beyond these approximations and consistently included the quantum nature using the Bose-Einstein and Fermi-Dirac distribution functions, including the effects of Bose enhancement and Pauli blocking.

When neglecting the quantum statistics, it is convenient to performe an integration over the momentum in the center-of-mass frame to obtain the cross section. Since the Bose-enhancement and Pauli-blocking terms depend on the energy of the outgoing particles, they however render the cross-section integral more complicated. In this case, some previously employed approximation schemes break down because the approximate integrand peaks in unphysical regimes. Unlike for freeze-out calculations above the electroweak scale, this is particularly noticeable for axion couplings to matter at lower temperatures, as considered in this work. While the calculation of [32] included Bose enhancement and Pauli blocking in a simplified fashion as $\left[1 \pm f_3\right]\left[1 \pm f_4\right] \rightarrow \frac{1}{2}\left([1 \pm f_3(p_1)][1 \pm f_4(p_2)] + \{p_1 \leftrightarrow p_2\}\right)$, this meant that the outgoing momenta were approximated by the incoming momenta. In the case of an

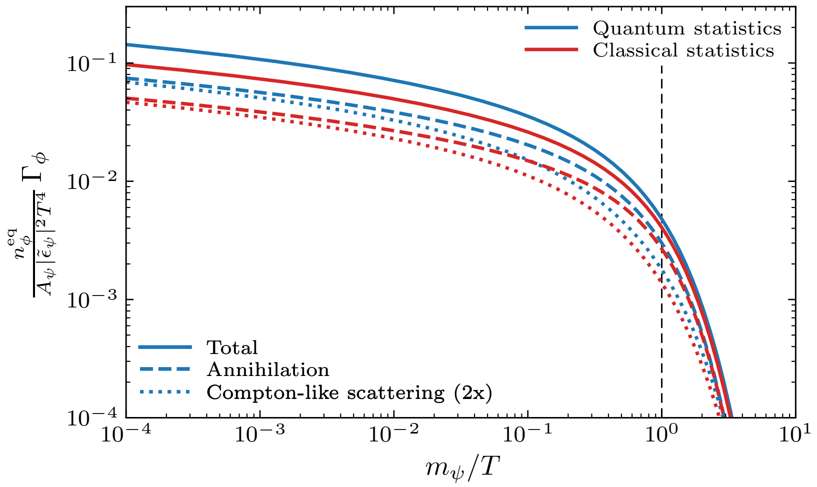

Figure A.1. Dimensionless rescaling of the total interaction rate Γ_ϕ as a function of m_ψ/T together with its contributions from fermion-antifermion annihilation and Compton-like scattering (for both fermions and antifermions). We compare the results of our full calculation with the approximate result neglecting the Bose-Einstein and Fermi-Dirac statistics as well as Bose enhancement and Pauli blocking (see Fig. 2.3). The vertical dashed line indicates $T = m_\psi$ which is approximately the temperature where decoupling occurs. We see that the annihilation rate is larger than the Compton-like scattering rate at all temperatures. We also observe that the difference between using the Boltzmann approximation instead of the full quantum statistics is most pronounced for the annihilation rate at large temperatures $T \gg m_\psi$.

incoming and outgoing boson, e.g. in the Compton-like process, the integrand however diverges at low momenta,

$$\lim_{p_1 \to 0} \frac{\mathrm{d}^3 p_1}{2E_1} f_1^b(p_1) \left[1 + f_3^b(p_1)\right] \sim \frac{\mathrm{d}p_1}{p_1}\,. \tag{A.22}$$

We remedied these shortcomings by going beyond any of these approximations and incorporating the full quantum statistics in our analytic and numerical computation of the production rate as described in Appendix A.1. At the same time, this allows us to compare our full calculation to the results based on the commonly-employed Boltzmann approximation without Bose enhancement and Pauli blocking.

We present the results of this comparison of quantum and classical statistics in Fig. A.1 for the processes relevant to the axion production calculated in this work.[1] We see that the

---

[1] We note that we assumed production to occur after electroweak symmetry breaking in our calculation of the underlying scattering amplitudes, i.e. these results should not be extrapolated to arbitrarily large energies.

difference between the quantum and classical production rates is fairly substantial when $T \gg m$. This considerable difference justifies the concern that the employed statistics can non-trivially impact production rates and, therefore, the bounds on the couplings from $\Delta N_{\rm eff}$ measurements. Having said that, the impact on the production rate relevant to these bounds is actually somewhat small since the dominant source of axions will be produced when $T \approx m$, and the effects of Bose enhancement and Pauli blocking are reduced when the number densities are suppressed at low temperatures.

In addition, we separately break down the impact on the annihilation and Compton-like processes. We can observe that the production rate receives approximately equal contributions from Compton-like scattering and fermion-antifermion annihilation at high temperatures, $T \gg m_\psi$, but is dominated by the former and latter process for $T \ll m_\psi$ and around $T = m_\psi$, respectively. In addition, we see that both processes show the same difference between their classical and quantum evaluation for $T \gtrsim m_\psi/10$, with the quantum annihilation rate approaching its classical counterpart for low temperatures whereas the quantum Compton-like scattering rate remains elevated compared to its classical treatment. While this is not directly related to the failure of the approximation scheme for the Compton-like process in (A.22), both effects are tied to the correct implementation of the Bose enhancement.

A.2.2 Uncertainties in the Axion-Quark Calculation

The masses of the bottom and charm quarks are close to the energy scale of the QCD phase transition. As a result, we expect that decoupling of the axion occurs during the QCD phase transition for its couplings to these quarks because this happens at temperatures $T \approx m_i/10$ for large interaction strengths. Since the production rates involve external gluons, this would in principle require a non-perturbative calculation, such as with lattice QCD, to determine the exact contribution to $\Delta N_{\rm eff}$. This is a particularly critical issue as the number of degrees of freedom changes rapidly with temperature during the transition, which implies that the predictions of $\Delta N_{\rm eff}(\Lambda_{\{b,c\}})$ are extremely sensitive to the temperature of decoupling.

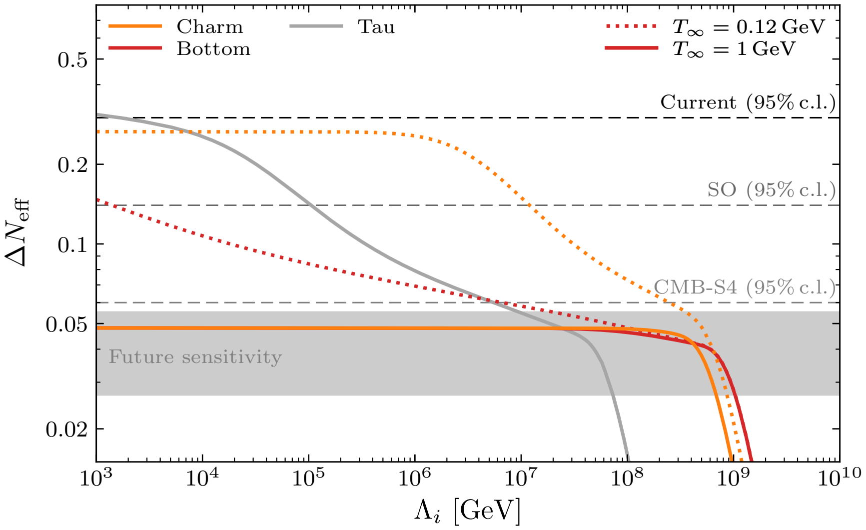

Figure A.2. Contributions to ΔN_{eff} from coupling to the charm (orange) and bottom (red) quarks. For comparison, the gray line indicates the prediction for axion-tau interactions (cf. Figures 2.4 and 2.9). The solid line indicates the conservative assumption made in the main text where axion decoupling is imposed by hand at $T_\infty = 1\,\mathrm{GeV}$. The dashed line indicates our less-conservative estimate where we allow the axions to remain in equilibrium through the QCD phase transition, with $\alpha_s = 1$, and force decoupling at $T_\infty = 120\,\mathrm{MeV}$. We see that the predictions depend sensitively on non-perturbative physics during the QCD phase transition.

In the absence of a non-perturbative calculation of the axion production rate,[2] we report our results in terms of conservative and less-conservative estimates for the processes involving the bottom and charm quarks. In the main text, we only presented the conservative estimates which were computed by cutting off the Boltzmann evolution at a final temperature of $T_\infty = 1\,\mathrm{GeV}$ when the strong coupling constant $\alpha_s \approx 0.5$. In essence, we force the axion decoupling by hand at the onset of the QCD phase transition when perturbation theory starts to break down. This likely underestimates the contributions to ΔN_{eff} at larger couplings (smaller Λ_i) for which the axion is surely still in equilibrium for $T < 1\,\mathrm{GeV}$.

In Figure A.2, we included a less conservative calculation (dashed lines), that stops at the lower end of the QCD phase transition, $T_\infty = 120\,\mathrm{MeV}$, assuming that the strong coupling α_s is fixed to unity in the regime where the naive coupling would exceed one. We note that the results of [53] show the same qualitative behavior. This figure illustrates that there is little

---

[2] Alternatively, one could also attempt to match across the QCD phase transition, along the lines of [72, 73].

difference between the two estimates for weak couplings (large Λ_i) because the axion is no longer in equilibrium during the QCD phase transition. The difference becomes large at stronger interaction strengths (smaller Λ_i) since the axion remains in equilibrium at $T < 1\,\text{GeV}$ in the less-conservative scenario. Larger couplings keep the axion in equilibrium to progressively lower temperatures and the contribution to ΔN_{eff} climbs accordingly. The less conservative estimate suggests that the coupling to the bottom and charm quarks may lie within the sensitivity of CMB-S4 and potentially even the Simons Observatory. This motivates a non-perturbative computation of axion production from bottom and charm couplings during the QCD phase transition.

Appendix B

Appendix for Chapter 3

B.1 Forecasting Details

In this appendix, we collect additional information on our LSS Fisher forecasts. We first provide the full set of definitions underlying our model for the galaxy power spectrum and then describe the experimental specifications of all galaxy surveys employed in the main text.

B.1.1 Galaxy Power Spectrum Model

We model the galaxy power spectrum at one-loop order in standard Eulerian perturbation theory and in the bias expansion where the galaxy overdensity δ_g is a function of terms up to third order in the linear matter density contrast δ_m [cf. (3.7)]. The theoretical galaxy power spectrum for two biased tracers A and B is then given by (3.10). The various loop contributions in that equation are defined as follows:

$$\sigma^4 = \int_{\vec{q}} P_{\rm lin}^2(q)\,, \tag{B.1}$$

$$P_{22}(k) = 2\int_{\vec{q}} \left[F_2^{(s)}(\vec{q},\vec{k}-\vec{q})\right]^2 P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.2}$$

$$P_{13}(k) = 6P_{\rm lin}(k)\int_{\vec{q}} F_3^{(s)}(\vec{k},\vec{q},-\vec{q})P_{\rm lin}(q)\,, \tag{B.3}$$

$$P_{\delta^2}(k) = 2\int_{\vec{q}} F_2^{(s)}(\vec{q},\vec{k}-\vec{q})P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.4}$$

$$P_{s^2}(k) = 2\int_{\vec{q}} F_2^{(s)}(\vec{q},\vec{k}-\vec{q})\left(\mu_-^2-\tfrac{1}{3}\right)P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.5}$$

$$P_{\delta^2\delta^2}(k) = 2\int_{\vec{q}} P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.6}$$

$$P_{\delta^2 s^2}(k) = 2\int_{\vec{q}} \left(\mu_-^2-\tfrac{1}{3}\right)P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.7}$$

$$P_{s^2 s^2}(k) = 2\int_{\vec{q}} \left(\mu_-^2-\tfrac{1}{3}\right)^2 P_{\rm lin}(q)P_{\rm lin}(|\vec{k}-\vec{q}|)\,, \tag{B.8}$$

$$P_{\Pi\Pi^{[2]}}(k) = 2P_{\rm lin}(k)\int_{\vec{q}} \left\{\frac{2}{7}\left[\left(\frac{\vec{k}\cdot\vec{q}}{kq}\right)^2-1\right]\frac{2}{3}\mu_-^2+\frac{8}{63}\right\}P_{\rm lin}(q)\,, \tag{B.9}$$

with $\int_{\vec{q}} \equiv \int \frac{\mathrm{d}q^3}{(2\pi)^3}$ and $\mu_- \equiv \frac{\vec{q}\cdot(\vec{k}-\vec{q})}{q|\vec{k}-\vec{q}|}$. The functions $F_2^{(s)}(\vec{k}_1,\vec{k}_2)$ and $F_3^{(s)}(\vec{k}_1,\vec{k}_2,\vec{k}_3)$ are the symmetric, second- and third-order kernels of δ_m in standard perturbation theory [151],

$$F_2^{(s)}(\vec{k}_1,\vec{k}_2) = \frac{5}{7}+\frac{2}{7}\left(\frac{\vec{k}_1\cdot\vec{k}_2}{k_1k_2}\right)^2+\frac{1}{2}\frac{\vec{k}_1\cdot\vec{k}_2}{k_1k_2}\left(\frac{k_1}{k_2}+\frac{k_2}{k_1}\right)\,, \tag{B.10}$$

$$\begin{aligned} F_3^{(s)}(\vec{k}_1,\vec{k}_2,\vec{k}_3) = &\frac{k^2}{27}\left[\frac{\vec{k}_1\cdot(\vec{k}_2+\vec{k}_3)}{k_1^2(\vec{k}_2+\vec{k}_3)^2}G_2^{(s)}(\vec{k}_2,\vec{k}_3)+2\text{ cyclic}\right] \\ &+\frac{7}{54}\vec{k}\cdot\left[\frac{\vec{k}_1+\vec{k}_2}{(\vec{k}_1+\vec{k}_2)^2}G_2^{(s)}(\vec{k}_1,\vec{k}_2)+2\text{ cyclic}\right] \\ &+\frac{7}{54}\vec{k}\cdot\left[\frac{\vec{k}_1}{k_1^2}F_2^{(s)}(\vec{k}_2,\vec{k}_3)+2\text{ cyclic}\right]\,, \end{aligned} \tag{B.11}$$

where $\vec{k} \equiv \vec{k}_1 + \vec{k}_2 + \vec{k}_3$ and $G_2^{(s)}(\vec{k}_1, \vec{k}_2)$ is the symmetric, second-order kernel of the velocity divergence θ_m,

$$G_2^{(2)}(\vec{k}_1, \vec{k}_2) = \frac{3}{7} + \frac{4}{7}\left(\frac{\vec{k}_1 \cdot \vec{k}_2}{k_1 k_2}\right)^2 + \frac{1}{2}\frac{\vec{k}_1 \cdot \vec{k}_2}{k_1 k_2}\left(\frac{k_1}{k_2} + \frac{k_2}{k_1}\right). \tag{B.12}$$

B.1.2 Survey Specifications

In the following, we provide detailed information about the spectroscopic and photometric redshift surveys that we employ in our Fisher forecasts, including the assumed redshift distribution in bias and number density.

Spectroscopic Surveys

The experimental specifications of the spectroscopic redshift surveys BOSS, DESI, Euclid, SPHEREx, MegaMapper and the billion-object survey are provided in Tables B.1 to B.6. We provide the linear bias b_1 and number density $\bar{n}_g$ for each sample and redshift bin with mean redshift $\bar{z}$ and spherical volume V. For DESI, different types of tracers were combined into one galaxy sample with a single effective number density and bias following [44, 114]. Since the number density of DESI is not in the regime where sample-variance cancellation is most effective (see Fig. 3.4), our results should be relatively insensitive to this choice. For all surveys, we treat each redshift bin as being independent so that the Fisher matrix for the entire survey is the sum of the Fisher matrices associated with each separate redshift bin.

Photometric Surveys

To forecast Vera Rubin Observatory's LSST, which is a photometric redshift survey, we use the "gold" sample defined in their science book [188], which includes more than four billion galaxies over 20000 deg^2. The number of objects in each redshift bin is calculated according to

$$N_g(z_{\rm min}, z_{\rm max}) = \frac{N_{\rm tot}\Omega_{\rm sky}}{2z_0^2}\left(2z_0^2 + 2z_0 z + z^2\right)\mathrm{e}^{-z/z_0}\Big|_{z=z_{\rm min}}^{z=z_{\rm max}}, \tag{B.13}$$

where $N_{\rm tot} = 0.55\,{\rm arcmin}^{-2}$, $\Omega_{\rm sky} = 20\,000\,{\rm deg}^2$ and $z_0 = 0.3$. The bias is taken to be $b_1(z) = 0.95/D(z)$ and we assume a photometric redshift error of $\sigma_{z0} = 0.05$. Table B.7 explicitly provides the derived experimental specifications for this survey with $z_{\rm max} = 3$ and twelve redshift bins with width $\Delta z = 0.25$.

Table B.1. Basic specifications for a BOSS-like survey [114] (inspired by [259] as detailed in [44]), covering a sky area of 10 252 $\deg^2$ with a total of about 1.2×10^6 objects in a volume of roughly $6.3\,h^{-3}\,\mathrm{Gpc}^3$.

$\bar{z}$	0.35	0.625
b_1	1.634	1.877
$10^3 \bar{n}_g\ [h^3\,\mathrm{Mpc}^{-3}]$	0.275	0.142
$V\ [h^{-3}\,\mathrm{Gpc}^3]$	2.18	4.15

Table B.2. Basic specifications for DESI [114] (derived from [185] as explained in [44]), covering a sky area of 14 000 $\deg^2$ with a total of about 2.7×10^7 objects in a volume of roughly $58\,h^{-3}\,\mathrm{Gpc}^3$.

$\bar{z}$	0.05	0.15	0.25	0.35	0.45	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
b_1	1.40	1.48	1.55	1.61	1.67	2.05	1.71	1.71	1.53	1.45	1.48	1.47	1.47	1.69	1.68	2.27	2.45	2.47
$10^3 \bar{n}_g\ [h^3\,\mathrm{Mpc}^{-3}]$	38.8	15.7	3.96	0.883	0.0992	0.591	1.31	0.920	0.779	0.466	0.398	0.387	0.180	0.133	0.110	0.0387	0.0197	0.0208
$V\ [h^{-3}\,\mathrm{Gpc}^3]$	0.0356	0.229	0.560	0.979	1.44	2.39	2.83	3.24	3.61	3.94	4.24	4.49	4.71	4.90	5.05	5.18	5.29	5.37

Table B.3. Basic specifications for the Euclid survey [44] (derived from [260]), covering a sky area of 15 000 $\deg^2$ with a total of about 4.9×10^7 objects in a volume of roughly $71\,h^{-3}\,\mathrm{Gpc}^3$.

$\bar{z}$	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	2.05
b_1	1.06	1.11	1.16	1.21	1.27	1.33	1.38	1.44	1.51	1.54	1.63	1.70	1.85	1.90	1.26
$10^3 \bar{n}_g\ [h^3\,\mathrm{Mpc}^{-3}]$	0.637	1.46	1.63	1.50	1.33	1.14	1.00	0.837	0.652	0.512	0.357	0.246	0.149	0.0904	0.0721
$V\ [h^{-3}\,\mathrm{Gpc}^3]$	2.56	3.03	3.47	3.87	4.23	4.54	4.81	5.05	5.25	5.41	5.55	5.67	5.76	5.83	5.88

Table B.4. Basic specifications for SPHEREx [261], covering a sky fraction $f_{\rm sky} = 0.75$ with a total of about 7.0×10^8 objects in a volume of roughly $450\,h^{-3}\,\mathrm{Gpc}^3$. SPHEREx is a spectro-photometric survey and the observed objects are divided into five samples based on their photometric redshift uncertainty bin, with respective maximum error $\sigma_{z0} = \{0.003, 0.01, 0.03, 0.1, 0.2\}$. The volume V and the galaxy number density $\bar{n}_g$ are given in units of $h^{-3}\,\mathrm{Gpc}^3$ and $h^3\,\mathrm{Mpc}^{-3}$, respectively.

$\bar{z}$	0.1	0.3	0.5	0.7	0.9	1.3	1.9	2.5	3.1	3.7	4.3
V	0.584	3.40	7.43	11.5	15.1	60.4	71.4	73.6	71.8	68.3	64.2
$b_1^{(1)}$	1.3	1.5	1.8	2.3	2.1	2.7	3.6	2.3	3.2	2.7	3.8
$10^5\bar{n}_g^{(1)}$	997	411	50.1	7.05	3.16	1.64	0.359	0.0807	0.184	0.150	0.113
$b_1^{(2)}$	1.2	1.4	1.6	1.9	2.3	2.6	3.4	4.2	4.3	3.7	4.6
$10^5\bar{n}_g^{(2)}$	1230	856	282	93.7	43.0	5.00	0.803	0.383	0.328	0.107	0.0679
$b_1^{(3)}$	1.0	1.3	1.5	1.7	1.9	2.6	3.0	3.2	3.5	4.1	5.0
$10^5\bar{n}_g^{(3)}$	1340	857	362	294	204	21.2	0.697	0.202	0.143	0.193	0.0679
$b_1^{(4)}$	0.98	1.3	1.4	1.5	1.7	2.2	3.6	3.7	2.7	2.9	5.0
$10^5\bar{n}_g^{(4)}$	2290	1290	535	495	415	79.6	7.75	0.787	0.246	0.193	0.136
$b_1^{(5)}$	0.83	1.2	1.3	1.4	1.6	2.1	3.2	4.2	4.1	4.5	5.0
$10^5\bar{n}_g^{(5)}$	1490	752	327	250	183	73.4	25.3	5.41	2.99	0.941	0.204

Table B.5. Basic specifications for MegaMapper [189], covering a sky area of $14\,000\,\deg^2$ with a total of about 6.6×10^7 objects in a volume of roughly $200\,h^{-3}\,\mathrm{Gpc}^3$.

$\bar{z}$	2.0	3.0	4.0	5.0
b_1	2.5	4.0	3.5	5.5
$10^3\bar{n}_g$ $[h^3\,\mathrm{Mpc}^{-3}]$	0.98	0.12	0.10	0.040
V $[h^{-3}\,\mathrm{Gpc}^3]$	54.0	54.4	50.0	44.8

Table B.6. Basic specifications for our (futuristic) spectroscopic billion-object survey, covering a sky fraction $f_{\mathrm{sky}} = 0.5$ with a total of 10^9 objects in a volume of roughly $330\,h^{-3}\,\mathrm{Gpc}^3$. The observed objects are divided into two samples based on their linear bias, $b_1(z=0) = 2.0$ and 1.2.

$\bar{z}$	0.25	0.75	1.25	1.75	2.25	2.75	3.25	3.75	4.25	4.75
V $[h^{-3}\,\mathrm{Gpc}^3]$	4.78	20.6	32.8	38.8	40.8	40.7	39.4	37.7	35.9	34.0
$b_1^{(1)}$	2.28	2.94	3.66	4.41	5.17	5.94	6.72	7.51	8.29	9.07
$10^3\bar{n}_g^{(1)}$ $[h^3\,\mathrm{Mpc}^{-3}]$	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54
$b_1^{(2)}$	1.37	1.76	2.19	2.64	3.10	3.57	4.03	4.50	4.97	5.45
$10^3\bar{n}_g^{(2)}$ $[h^3\,\mathrm{Mpc}^{-3}]$	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.54

Table B.7. Basic specifications for Vera Rubin Observatory's LSST [188], covering a sky area of $20\,000\,\mathrm{deg}^2$ with a total of about 4×10^9 objects in a volume of roughly $170\,h^{-3}\,\mathrm{Gpc}^3$. The photometric redshift error is $\sigma_{z0} = 0.05$.

$\bar{z}$	0.125	0.375	0.625	0.875	1.125	1.375
b_1	1.01	1.16	1.31	1.48	1.65	1.82
$10^3\bar{n}_g$ $[h^3\,\mathrm{Mpc}^{-3}]$	292	183	108	63.6	36.8	21.0
V $[h^{-3}\,\mathrm{Gpc}^3]$	0.709	3.93	8.10	11.9	14.9	17.0
$\bar{z}$	1.625	1.875	2.125	2.375	2.625	2.875
b_1	2.00	2.18	2.36	2.54	2.73	2.91
$10^3\bar{n}_g$ $[h^3\,\mathrm{Mpc}^{-3}]$	11.9	6.46	3.55	1.94	0.969	0.646
V $[h^{-3}\,\mathrm{Gpc}^3]$	18.4	19.2	19.7	19.9	19.8	19.6

B.2 Measuring the Scaling Behavior

We have assumed for all forecasts throughout the paper that the fiducial value of the non-Gaussian amplitude $f_{\rm NL}^{\Delta}$ vanishes. This is of course motivated by the fact that we have only been placing observational bounds on PNG that are consistent with $f_{\rm NL}^{\Delta} = 0$, including in Section 3.5. We additionally always considered the scaling exponent Δ to be fixed. In this appendix, we will consider nonzero fiducial PNG amplitudes and clarify to what degree measurements of the galaxy power spectrum are sensitive to the precise value of Δ.

We defined the correlation between the non-Gaussian signals of two different values of Δ in (3.32) in terms of the respective Fisher matrix elements (after marginalization) for a fiducial value of $\bar{f}_{\rm NL}^{\Delta_{1,2}} = 0$. For BOSS, we can observe in Fig 3.10 that this correlation is close to unity, even for widely separated values of Δ. We see in Fig. B.1 that this statement still holds for the billion-object survey. In fact, the range of Δ which are highly correlated with the local shape further expanded. Naively, this suggests there is little information about the shape of the signal, described by Δ, encoded in these observables.

To first approximation, the fiducial value of $\bar{f}_{\rm NL}^{\Delta} = 0$ is the reason why we see little ability to distinguish different values of Δ. Since there is no signal to measure, it is not surprising that the precise shape is not very important. If we were instead to detect $f_{\rm NL}^{\Delta}$, we would expect that the signal to noise would be a fairly sensitive function of Δ and we could therefore measure the true value of Δ precisely. This is exactly what is found in Fig. B.2, in which we show the forecasted constraints for measuring Δ, $\sigma(\Delta)$, assuming that $f_{\rm NL}^{\Delta}$ is detected at a significance of 5σ in the billion-object survey. The key result is that if we detect $f_{\rm NL}^{\Delta}$ while holding Δ fixed, we expect to have enough sensitivity to measure the scaling exponent with $\sigma(\Delta) \sim 0.05 - 0.5$. In this sense, the strategy of searching for $f_{\rm NL}^{\Delta}$ at fixed Δ is a reliable strategy to search for new physics, but would yield additional information about the origin of the signal even at the threshold of a detection. When we instead marginalize over Δ, we require a significantly larger non-Gaussian amplitude $f_{\rm NL}^{\Delta}$ to be detected. In this case, the corresponding uncertainties on Δ

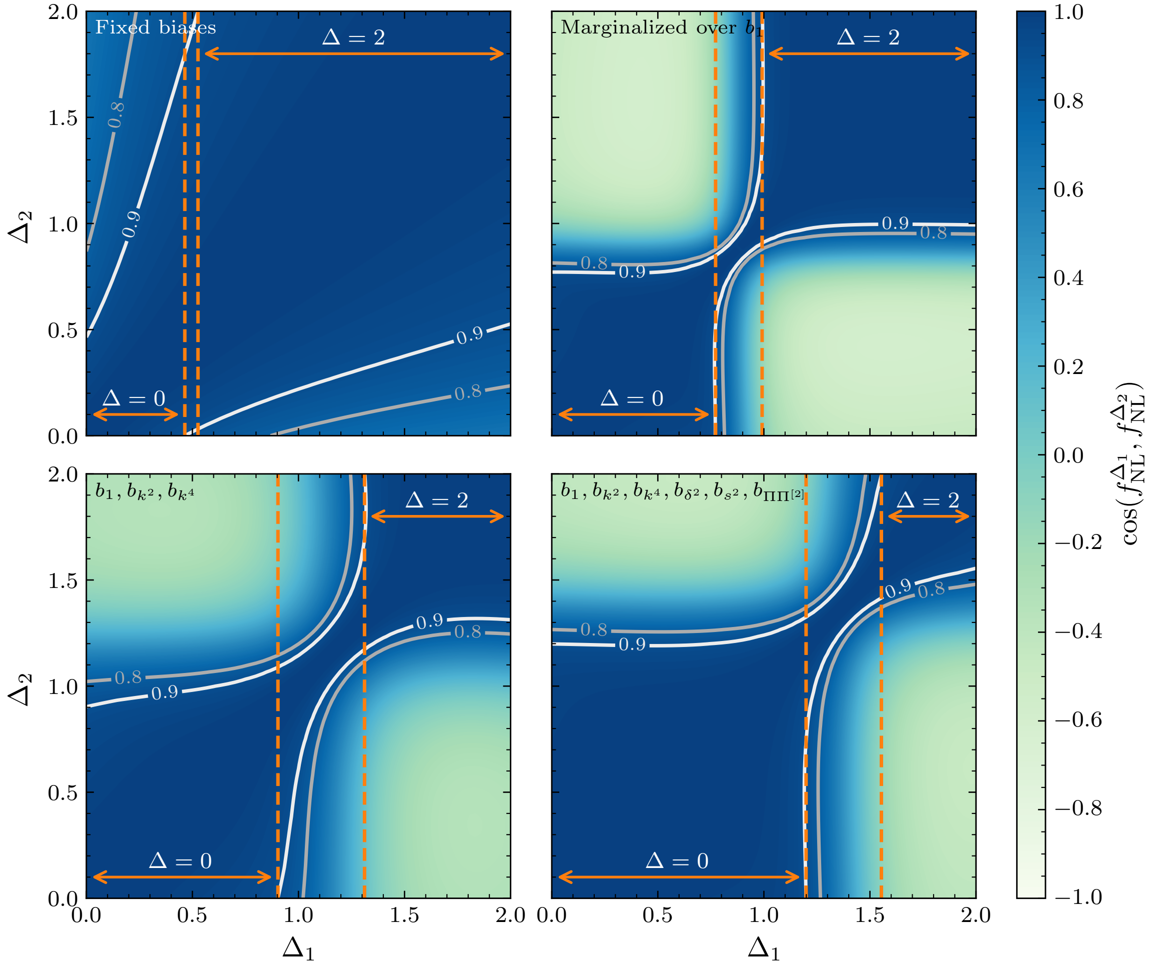

Fixed biases

Marginalized over b_1

b_1, b_{k^2}, b_{k^4}

$b_1, b_{k^2}, b_{k^4}, b_{\delta^2}, b_{s^2}, b_{\Pi\Pi^{[2]}}$

Figure B.1. Correlation matrix for measurements of galaxy power spectra with different values of the non-Gaussian scaling exponent Δ as defined in (3.32) for the billion-object survey and four biasing models. The lower right panel is equivalent to Fig. 3.10, which shows the correlations for BOSS. As in that figure, we hold the ΛCDM parameters fixed, take the fiducial non-Gaussian amplitudes to be zero, $\bar{f}_{\rm NL}^{\Delta_i} = 0$, and indicate the coverage of the local and equilateral templates with a correlation coefficient of at least 0.9 by the orange lines. We can clearly observe the impact of the biasing model and the different experimental specifications compared to BOSS.

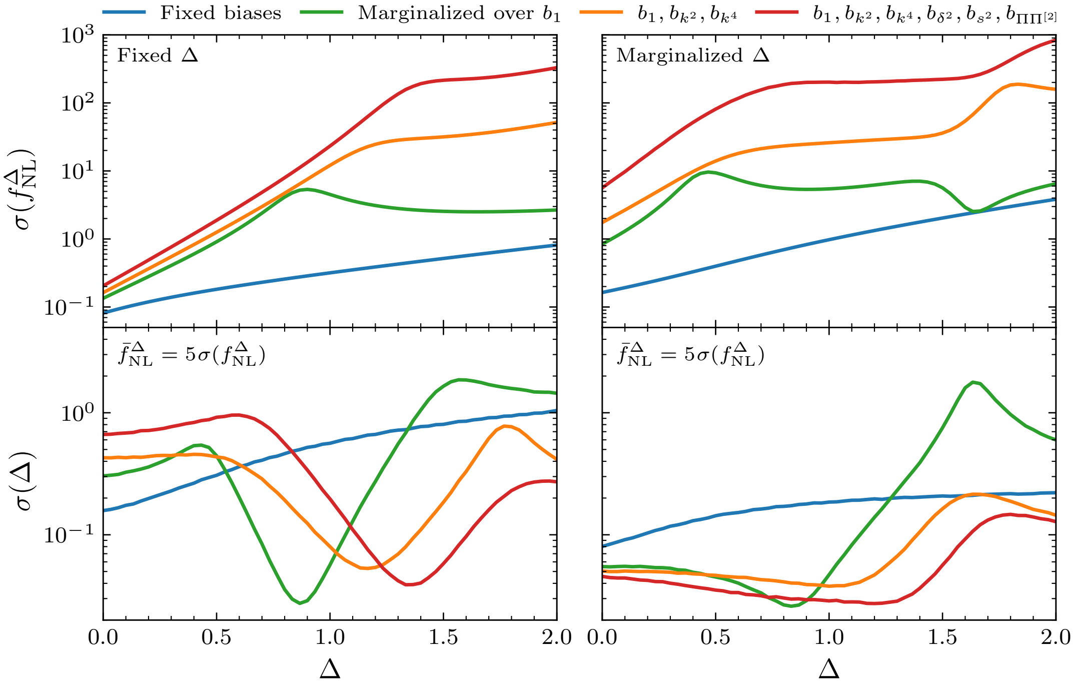

Figure B.2. Forecasts for $\sigma(f_{\rm NL}^{\Delta})$ (*top*) and $\sigma(\Delta)$ (*bottom*) as a function of Δ if the scaling exponent is either fixed (*left*) or varied (*right*) for the billion-object survey. To compute the constraint on Δ, we take the fiducial value of the non-Gaussian amplitude, $\bar{f}_{\rm NL}^{\Delta}$, to be the value necessary to achieve a 5σ detection in the respective forecast. We fix the ΛCDM parameters and marginalize over four different biasing models.

would however be generally smaller, with $\sigma(\Delta) < 0.1$.

At first sight, it might seem surprising that the curves of $\sigma(\Delta)$ in Fig. B.2 have minima at the values of Δ where the curves of $\sigma(f_{\mathrm{NL}}^{\Delta})$ have maxima or downward bends. By comparing the curves for different biasing models, we can however see that these features are the result of marginalizing over the bias parameters. For example, we noted in the main text that the maximum in $\sigma(f_{\mathrm{NL}}^{\Delta})$ near $\Delta = 1$ when marginalizing over b_1 arises because the linear bias b_1 and $f_{\mathrm{NL}}^{\Delta=1}$ are degenerate at high k. The fact that the measurement of f_{NL}^{Δ} is less constraining at this point then also implies that we are very sensitive to the exact value of the scaling exponent Δ. When we hold $f_{\mathrm{NL}}^{\Delta}/\sigma(f_{\mathrm{NL}}^{\Delta})$ fixed, we should therefore expect that $\sigma(\Delta)$ is smaller in regions of Δ where f_{NL}^{Δ} is degenerate with bias parameters.

Let us finally return to the correlation coefficients between scale-dependent biases with different scaling exponent Δ. In order to be consistent, the impact of marginalization over the bias parameters must also appear in $\cos(f_{\mathrm{NL}}^{\Delta_1}, f_{\mathrm{NL}}^{\Delta_2})$. Specifically, the impact of the marginalization has a large impact on $\sigma(f_{\mathrm{NL}}^{\Delta})$ for larger values of Δ. For the inferred constraints to be consistent with our forecasts, this has to imply that the correlation between large and small Δ decreases by a similar factor. If this was not the case, our constraints inferred from $f_{\mathrm{NL}}^{\mathrm{loc}}$ at large Δ would be stronger than directly measuring f_{NL}^{Δ}.

This expectation is precisely what occurs, as shown in Fig. B.1. The correlation coefficients for larger Δ are affected by the number of bias parameters that are marginalized over. With more marginalized bias parameters, the range over which $\Delta < 2$ and $\Delta = 2$ are degenerate decreases significantly. This is expected since the (equilateral) signal for $\Delta = 2$ gets most of its information from large wavenumbers, where the biases play a significant role (cf. §3.4.2). In contrast, the increase of the range of strong overlap with $\Delta = 0$ is likely due to the increasing difficulty to distinguish similar shapes of Δ with the inclusion of more bias parameters. Finally, it is also worth noticing that the non-Gaussian signal with $\Delta = 0$ (local PNG) appears to be anti-correlated with the scale-dependent bias signal for $\Delta \sim 2$ in the last three panels of Fig. B.1

in which we marginalize over the biasing model.

Bibliography

[1] R. Peccei and H. Quinn, "CP Conservation in the Presence of Instantons," *Phys. Rev. Lett.* **38** (1977) 1440.

[2] S. Weinberg, "A New Light Boson?," *Phys. Rev. Lett.* **40** (1978) 223.

[3] A. Hook, "TASI Lectures on the Strong CP Problem and Axions," *PoS* **TASI2018** (2019) 004, `arXiv:1812.02669 [hep-ph]`.

[4] P. Graham, D. E. Kaplan, and S. Rajendran, "Cosmological Relaxation of the Electroweak Scale," *Phys. Rev. Lett.* **115** (2015) 221801, `arXiv:1504.07551 [hep-ph]`.

[5] J. Preskill, M. Wise, and F. Wilczek, "Cosmology of the Invisible Axion," *Phys. Lett. B* **120** (1983) 127.

[6] L. Abbott and P. Sikivie, "A Cosmological Bound on the Invisible Axion," *Phys. Lett. B* **120** (1983) 133.

[7] M. Dine and W. Fischler, "The Not So Harmless Axion," *Phys. Lett. B* **120** (1983) 137.

[8] A. Davidson and K. Wali, "Minimal Flavor Unification via Multi-Generational Peccei-Quinn Symmetry," *Phys. Rev. Lett.* **48** (1982) 11.

[9] F. Wilczek, "Axions and Family Symmetry Breaking," *Phys. Rev. Lett.* **49** (1982) 1549.

[10] D. Reiss, "Can the Family Group Be a Global Symmetry?," *Phys. Lett. B* **115** (1982) 217.

[11] J. Feng, T. Moroi, H. Murayama, and E. Schnapka, "Third-Generation Familons, B Factories and Neutrino Cosmology," *Phys. Rev. D* **57** (1998) 5875, `arXiv:hep-ph/9709411 [hep-ph]`.

[12] Y. Chikashige, R. Mohapatra, and R. Peccei, "Are There Real Goldstone Bosons Associated with Broken Lepton Number?," *Phys. Lett. B* **98** (1981) 265.

[13] Z. Chacko, L. Hall, T. Okui, and S. Oliver, "CMB Signals of Neutrino Mass Generation," *Phys. Rev. D* **70** (2004) 085008, `arXiv:hep-ph/0312267 [hep-ph]`.

[14] P. Svrcek and E. Witten, "Axions In String Theory," *JHEP* **06** (2006) 051, `arXiv:hep-th/0605206`.

[15] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, "String Axiverse," *Phys. Rev. D* **81** (2010) 123530, `arXiv:0905.4720 [hep-th]`.

[16] D. J. E. Marsh, "Axion Cosmology," *Phys. Rept.* **643** (2016) 1, `arXiv:1510.07633 [astro-ph.CO]`.

[17] N. Arkani-Hamed, T. Cohen, R. D'Agnolo, A. Hook, H. Kim, and D. Pinner, "Solving the Hierarchy Problem at Reheating with a Large Number of Degrees of Freedom," *Phys. Rev. Lett.* **117** (2016) 251801, `arXiv:1607.06821 [hep-ph]`.

[18] Z. Chacko, N. Craig, P. Fox, and R. Harnik, "Cosmology in Mirror Twin Higgs and Neutrino Masses," *JHEP* **07** (2017) 023, `arXiv:1611.07975 [hep-ph]`.

[19] Z. Chacko, D. Curtin, M. Geller, and Y. Tsai, "Cosmological Signatures of a Mirror Twin Higgs," *JHEP* **09** (2018) 163, `arXiv:1803.03263 [hep-ph]`.

[20] J. Jaeckel and A. Ringwald, "The Low-Energy Frontier of Particle Physics," *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405, `arXiv:1002.0329 [hep-ph]`.

[21] R. Essig *et al.*, "Working Group Report: New Light Weakly Coupled Particles," in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013), Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. `arXiv:1311.0029 [hep-ph]`.

[22] P. Graham, I. Irastorza, S. Lamoreaux, A. Lindner, and K. van Bibber, "Experimental Searches for the Axion and Axion-Like Particles," *Ann. Rev. Nucl. Part. Sci.* **65** (2015) 485, `arXiv:1602.00039 [hep-ex]`.

[23] E. Braaten and T. Yuan, "Calculation of Screening in a Hot Plasma," *Phys. Rev. Lett.* **66** (1991) 2183.

[24] M. Bolz, A. Brandenburg, and W. Buchmüller, "Thermal Production of Gravitinos," *Nucl. Phys. B* **606** (2001) 518, `arXiv:hep-ph/0012052 [hep-ph]`.

[25] E. Massó, F. Rota, and G. Zsembinszki, "On Axion Thermalization in the Early Universe," *Phys. Rev. D* **66** (2002) 023004, `arXiv:hep-ph/0203221 [hep-ph]`.

[26] P. Graf and F. Steffen, "Thermal Axion Production in the Primordial Quark-Gluon Plasma," *Phys. Rev. D* **83** (2011) 075011, `arXiv:1008.4528 [hep-ph]`.

[27] D. Cadamuro, S. Hannestad, G. Raffelt, and J. Redondo, "Cosmological Bounds on Sub-MeV Mass Axions," *JCAP* **02** (2011) 003, `arXiv:1011.3694 [hep-ph]`.

[28] D. Cadamuro and J. Redondo, "Cosmological Bounds on Pseudo-Nambu-Goldstone Bosons," *JCAP* **02** (2012) 032, `arXiv:1110.2895 [hep-ph]`.

[29] C. Brust, D. E. Kaplan, and M. Walters, "New Light Species and the CMB," *JHEP* **12** (2013) 058, `arXiv:1303.5379 [hep-ph]`.

[30] S. Weinberg, "Goldstone Bosons as Fractional Cosmic Neutrinos," *Phys. Rev. Lett.* **110** (2013) 241301, `arXiv:1305.1971 [astro-ph.CO]`.

[31] A. Salvio, A. Strumia, and W. Xue, "Thermal Axion Production," *JCAP* **01** (2014) 011, `arXiv:1310.6982 [hep-ph]`.

[32] D. Baumann, D. Green, and B. Wallisch, "New Target for Cosmic Axion Searches," *Phys. Rev. Lett.* **117** (2016) 171301, `arXiv:1604.08614 [astro-ph.CO]`.

[33] R. Cyburt, B. Fields, K. Olive, and T.-H. Yeh, "Big Bang Nucleosynthesis: 2015," *Rev. Mod. Phys.* **88** (2016) 015004, `arXiv:1505.01076 [astro-ph.CO]`.

[34] C. Pitrou, A. Coc, J.-P. Uzan, and E. Vangioni, "Precision Big Bang Nucleosynthesis with Improved Helium-4 Predictions," *Phys. Rept.* **754** (2018) 1, `arXiv:1801.08023 [astro-ph.CO]`.

[35] B. Fields, K. Olive, T.-H. Yeh, and C. Young, "Big Bang Nucleosynthesis after Planck," *JCAP* **03** (2020) 010, `arXiv:1912.01132 [astro-ph.CO]`.

[36] S. Bashinsky and U. Seljak, "Neutrino Perturbations in CMB Anisotropy and Matter Clustering," *Phys. Rev. D* **69** (2004) 083002, `arXiv:astro-ph/0310198 [astro-ph]`.

[37] B. Follin, L. Knox, M. Millea, and Z. Pan, "First Detection of the Acoustic Oscillation Phase Shift Expected from the Cosmic Neutrino Background," *Phys. Rev. Lett.* **115** (2015) 091301, `arXiv:1503.07863 [astro-ph.CO]`.

[38] D. Baumann, D. Green, J. Meyers, and B. Wallisch, "Phases of New Physics in the CMB," *JCAP* **01** (2016) 007, `arXiv:1508.06342 [astro-ph.CO]`.

[39] C. Brust, Y. Cui, and K. Sigurdson, "Cosmological Constraints on Interacting Light Particles," *JCAP* **08** (2017) 020, `arXiv:1703.10732 [astro-ph.CO]`.

[40] G. Choi, C.-T. Chiang, and M. LoVerde, "Probing Decoupling in Dark Sectors with the Cosmic Microwave Background," *JCAP* **06** (2018) 044, `arXiv:1804.10180 [astro-ph.CO]`.

[41] N. Aghanim *et al.* (Planck Collaboration), "Planck 2018 Results. VI. Cosmological Parameters," *Astron. Astrophys.* **641** (2020) A6, `arXiv:1807.06209 [astro-ph.CO]`.

[42] N. Blinov and G. Marques-Tavares, "Interacting Radiation after Planck and Its Implications for the Hubble Tension," *JCAP* **09** (2020) 029, `arXiv:2003.08387 [astro-ph.CO]`.

[43] D. Baumann, D. Green, and M. Zaldarriaga, "Phases of New Physics in the BAO Spectrum," *JCAP* **11** (2017) 007, `arXiv:1703.00894 [astro-ph.CO]`.

[44] D. Baumann, D. Green, and B. Wallisch, "Searching for Light Relics with Large-Scale Structure," *JCAP* **08** (2018) 029, `arXiv:1712.08067 [astro-ph.CO]`.

[45] D. Baumann, F. Beutler, R. Flauger, D. Green, A. Slosar, M. Vargas-Magaña, B. Wallisch, and C. Yèche, "First Constraint on the Neutrino-Induced Phase Shift in the Spectrum of Baryon Acoustic Oscillations," *Nat. Phys.* **15** (2019) 465, `arXiv:1803.10741 [astro-ph.CO]`.

[46] D. Green and A. Ridgway, "The Phase of the BAO on Observable Scales," *JCAP* **12** (2020) 050, `arXiv:2008.05026 [astro-ph.CO]`.

[47] D. Baumann, "Primordial Cosmology," *PoS* **TASI2017** (2018) 009, `arXiv:1807.03098 [hep-th]`.

[48] B. Wallisch, *Cosmological Probes of Light Relics*. PhD thesis, University of Cambridge, 2018. `arXiv:1810.02800 [astro-ph.CO]`.

[49] D. Green *et al.*, "Messengers from the Early Universe: Cosmic Neutrinos and Other Light Relics," *Bull. Am. Astron. Soc.* **51** (2019) 159, `arXiv:1903.04763 [astro-ph.CO]`.

[50] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory*. Addison-Wesley, Reading, MA, USA, 1995.

[51] P. Ade *et al.* (Simons Observatory Collaboration), "The Simons Observatory: Science Goals and Forecasts," *JCAP* **02** (2019) 056, `arXiv:1808.07445 [astro-ph.CO]`.

[52] K. Abazajian *et al.* (CMB-S4 Collaboration), "CMB-S4 Science Case, Reference Design, and Project Plan," `arXiv:1907.04473 [astro-ph.IM]`.

[53] R. Ferreira and A. Notari, "Observable Windows for the QCD Axion Through the Number of Relativistic Species," *Phys. Rev. Lett.* **120** (2018) 191301, `arXiv:1801.06090 [hep-ph]`.

[54] F. D'Eramo, R. Ferreira, A. Notari, and J. Bernal, "Hot Axions and the H_0 Tension," *JCAP* **11** (2018) 014, `arXiv:1808.07430 [hep-ph]`.

[55] F. Arias-Aragón, F. D'Eramo, R. Ferreira, L. Merlo, and A. Notari, "Production of Thermal Axions across the Electroweak Phase Transition," *JCAP* **03** (2021) 090, `arXiv:2012.04736 [hep-ph]`.

[56] D. Ghosh and D. Sachdeva, "Constraints on Axion-Lepton Coupling from Big Bang Nucleosynthesis," *JCAP* **10** (2020) 060, `arXiv:2007.01873 [hep-ph]`.

[57] R. Ferreira, A. Notari, and F. Rompineve, "Dine-Fischler-Srednicki-Zhitnitsky Axion in the CMB," *Phys. Rev. D* **103** (2021) 063524, `arXiv:2012.06566 [hep-ph]`.

[58] J. Dror, H. Murayama, and N. Rodd, "Cosmic Axion Background," *Phys. Rev. D* **103** (2021) 115004, `arXiv:2101.09287 [hep-ph]`.

[59] R. Bollig, W. DeRocco, P. Graham, and H.-T. Janka, "Muons in Supernovae: Implications for the Axion-Muon Coupling," *Phys. Rev. Lett.* **125** (2020) 051104, `arXiv:2005.07141 [hep-ph]`.

[60] D. Croon, G. Elor, R. Leane, and S. McDermott, “Supernova Muons: New Constraints on Z' Bosons, Axions and ALPs,” *JHEP* **01** (2021) 107, `arXiv:2006.13942 [hep-ph]`.

[61] A. Caputo, G. Raffelt, and E. Vitagliano, “Muonic Boson Limits: Supernova Redux,” *Phys. Rev. D* **105** (2022) 035022, `arXiv:2109.03244 [hep-ph]`.

[62] K. Akita and M. Yamaguchi, “A Precision Calculation of Relic Neutrino Decoupling,” *JCAP* **08** (2020) 012, `arXiv:2005.07047 [hep-ph]`.

[63] J. Froustey, C. Pitrou, and M. Volpe, “Neutrino Decoupling Including Flavor Oscillations and Primordial Nucleosynthesis,” *JCAP* **12** (2020) 015, `arXiv:2008.01074 [hep-ph]`.

[64] J. Bennett, G. Buldgen, P. De Salas, M. Drewes, S. Gariazzo, S. Pastor, and Y. Y. Y. Wong, “Towards a Precision Calculation of N_{eff} in the Standard Model II: Neutrino Decoupling in the Presence of Flavor Oscillations and Finite-Temperature QED,” *JCAP* **04** (2021) 073, `arXiv:2012.02726 [hep-ph]`.

[65] N. Sehgal *et al.*, “CMB-HD: An Ultra-Deep, High-Resolution Millimeter-Wave Survey Over Half the Sky,” *Bull. Am. Astron. Soc.* **51** (2019) 6, `arXiv:1906.10134 [astro-ph.CO]`.

[66] D. Schlegel *et al.*, “Astro2020 APC White Paper: The MegaMapper: A $z > 2$ Spectroscopic Instrument for the Study of Inflation and Dark Energy,” *Bull. Am. Astron. Soc.* **51** (2019) 229, `arXiv:1907.11171 [astro-ph.IM]`.

[67] A. Slosar *et al.* (PUMA Collaboration), “Packed Ultra-wideband Mapping Array (PUMA): A Radio Telescope for Cosmology and Transients,” *Bull. Am. Astron. Soc.* **51** (2019) 53, `arXiv:1907.12559 [astro-ph.IM]`.

[68] N. Sailer, E. Castorina, S. Ferraro, and M. White, “Cosmology at High Redshift – A Probe of Fundamental Physics,” *JCAP* **12** (2021) 049, `arXiv:2106.09713 [astro-ph.CO]`.

[69] F. Herren and M. Steinhauser, “Version 3 of RunDec and CRunDec,” *Comput. Phys. Commun.* **224** (2018) 333, `arXiv:1703.03751 [hep-ph]`.

[70] K. Saikawa and S. Shirai, “Primordial Gravitational Waves, Precisely: The Role of Thermodynamics in the Standard Model,” *JCAP* **05** (2018) 035, `arXiv:1803.01038 [hep-ph]`.

[71] S. Borsanyi *et al.*, “Calculation of the Axion Mass Based on High-Temperature Lattice Quantum Chromodynamics,” *Nature* **539** (2016) 69, `arXiv:1606.07494 [hep-lat]`.

[72] F. D’Eramo, F. Hajkarim, and S. Yun, “Thermal Axion Production at Low Temperatures: A Smooth Treatment of the QCD Phase Transition,” `arXiv:2108.04259 [hep-ph]`.

[73] F. D’Eramo, F. Hajkarim, and S. Yun, “Thermal QCD Axions across Thresholds,” *JHEP* **10** (2021) 224, `arXiv:2108.05371 [hep-ph]`.

[74] M. Millea, L. Knox, and B. Fields, "New Bounds for Axions and Axion-Like Particles with keV-GeV Masses," *Phys. Rev. D* **92** (2015) 023010, `arXiv:1501.04097 [astro-ph.CO]`.

[75] N. DePorzio, W. Xu, J. Muñoz, and C. Dvorkin, "Finding eV-Scale Light Relics with Cosmological Observables," *Phys. Rev. D* **103** (2021) 023504, `arXiv:2006.09380 [astro-ph.CO]`.

[76] W. Xu, J. Muñoz, and C. Dvorkin, "Cosmological Constraints on Light (but Massive) Relics," `arXiv:2107.09664 [astro-ph.CO]`.

[77] G. Raffelt, *Stars as Laboratories for Fundamental Physics*. University of Chicago Press, Chicago, IL, USA, 1996.

[78] G. Raffelt, "Astrophysical Axion Bounds," *Lect. Notes Phys.* **741** (2008) 51, `arXiv:hep-ph/0611350 [hep-ph]`.

[79] G. Raffelt, "Neutrinos and the Stars," *Proc. Int. Sch. Phys. Fermi* **182** (2012) 61, `arXiv:1201.1637 [astro-ph.SR]`.

[80] N. Bar, K. Blum, and G. D'Amico, "Is There a Supernova Bound on Axions?," *Phys. Rev. D* **101** (2020) 123025, `arXiv:1907.05020 [hep-ph]`.

[81] S. Andreas, O. Lebedev, S. Ramos-Sanchez, and A. Ringwald, "Constraints on a Very Light CP-Odd Higgs of the NMSSM and Other Axion-Like Particles," *JHEP* **08** (2010) 003, `arXiv:1005.3978 [hep-ph]`.

[82] M. Dolan, F. Kahlhoefer, C. McCabe, and K. Schmidt-Hoberg, "A Taste of Dark Matter: Flavor Constraints on Pseudoscalar Mediators," *JHEP* **03** (2015) 171, `arXiv:1412.5174 [hep-ph]`.

[83] L. Calibbi, D. Redigolo, R. Ziegler, and J. Zupan, "Looking Forward to Lepton-Flavor-Violating ALPs," *JHEP* **09** (2021) 173, `arXiv:2006.04795 [hep-ph]`.

[84] B. Hansen, H. Richer, J. Kalirai, R. Goldsbury, S. Frewen, and J. Heyl, "Constraining Neutrino Cooling using the Hot White Dwarf Luminosity Function in the Globular Cluster 47 Tucanae," *Astrophys. J.* **809** (2015) 141, `arXiv:1507.05665 [astro-ph.SR]`.

[85] T. Battich, A. Córsico, L. Althaus, and M. Miller Bertolami, "First Axion Bounds from a Pulsating Helium-Rich White Dwarf Star," *JCAP* **08** (2016) 062, `arXiv:1605.07668 [astro-ph.SR]`.

[86] E. Armengaud *et al.* (EDELWEISS Collaboration), "Searches for Electron Interactions Induced by New Physics in the EDELWEISS-III Germanium Bolometers," *Phys. Rev. D* **98** (2018) 082004, `arXiv:1808.02340 [hep-ex]`.

[87] F. Capozzi and G. Raffelt, "Axion and Neutrino Bounds Improved with New Calibrations of the Tip of the Red-Giant Branch Using Geometric Distance Determinations," *Phys. Rev. D* **102** (2020) 083007, `arXiv:2007.03694 [astro-ph.SR]`.

[88] G. Lucente and P. Carenza, "Supernova Bound on Axion-Like Particles Coupled with Electrons," *Phys. Rev. D* **104** (2021) 103007, `arXiv:2107.12393 [hep-ph]`.

[89] L. Di Luzio, M. Fedele, M. Giannotti, F. Mescia, and E. Nardi, "Stellar Evolution Confronts Axion Models," *JCAP* **02** (2022) 035, `arXiv:2109.10368 [hep-ph]`.

[90] W. DeRocco, P. Graham, and S. Rajendran, "Exploring the Robustness of Stellar Cooling Constraints on Light Particles," *Phys. Rev. D* **102** (2020) 075015, `arXiv:2006.15112 [hep-ph]`.

[91] R. Essig, R. Harnik, J. Kaplan, and N. Toro, "Discovering New Light States at Neutrino Experiments," *Phys. Rev. D* **82** (2010) 113008, `arXiv:1008.0636 [hep-ph]`.

[92] T. Aoyama *et al.*, "The Anomalous Magnetic Moment of the Muon in the Standard Model," *Phys. Rept.* **887** (2020) 1, `arXiv:2006.04822 [hep-ph]`.

[93] B. Abi *et al.* (Muon g-2 Collaboration), "Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm," *Phys. Rev. Lett.* **126** (2021) 141801, `arXiv:2104.03281 [hep-ex]`.

[94] S. Borsanyi *et al.*, "Leading Hadronic Contribution to the Muon Magnetic Moment from Lattice QCD," *Nature* **593** (2021) 51, `arXiv:2002.12347 [hep-lat]`.

[95] T. Ohl, "Drawing Feynman Diagrams with LaTeX and Metafont," *Comput. Phys. Commun.* **90** (1995) 340, `arXiv:hep-ph/9505351 [hep-ph]`.

[96] F. Pérez and B. Granger, "IPython: A System for Interactive Scientific Computing," *Comput. Sci. Eng.* **9** (2007) 21.

[97] J. Hunter, "Matplotlib: A 2D Graphics Environment," *Comput. Sci. Eng.* **9** (2007) 90.

[98] S. Lam, A. Pitrou, and S. Seibert, "Numba: A LLVM-Based Python JIT Compiler," in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC (LLVM '15), Austin, TX, USA, November 15, 2015.* 2015.

[99] C. Harris *et al.*, "Array Programming with NumPy," *Nature* **585** (2020) 3572, `arXiv:2006.10256 [cs.MS]`.

[100] P. Virtanen *et al.*, "SciPy 1.0 – Fundamental Algorithms for Scientific Computing in Python," *Nat. Methods* **17** (2020) 261, `arXiv:1907.10121 [cs.MS]`.

[101] D. Green *et al.*, "Snowmass Theory Frontier: Astrophysics and Cosmology," `arXiv:2209.06854 [hep-ph]`.

[102] C. Chang *et al.*, "Report of the Topical Group on Cosmic Frontier 5 Dark Energy and Cosmic Acceleration: Cosmic Dawn and Before for Snowmass 2021," `arXiv:2209.08265 [hep-ex]`.

[103] D. Green, "TASI Lectures on Cosmic Signals of Fundamental Physics," `arXiv:2212.08685 [hep-ph]`.

[104] R. Flauger, V. Gorbenko, A. Joyce, L. McAllister, G. Shiu, and E. Silverstein, "Snowmass White Paper: Cosmology at the Theory Frontier," `arXiv:2203.07629 [hep-th]`.

[105] A. Achúcarro *et al.*, "Inflation: Theory and Observations," `arXiv:2203.08128 [astro-ph.CO]`.

[106] Y. Akrami *et al.* (Planck Collaboration), "Planck 2018 Results. IX. Constraints on Primordial Non-Gaussianity," *Astron. Astrophys.* **641** (2020) A9, `arXiv:1905.05697 [astro-ph.CO]`.

[107] P. D. Meerburg *et al.*, "Primordial Non-Gaussianity," *Bull. Am. Astron. Soc.* **51** (2019) 107, `arXiv:1903.04409 [astro-ph.CO]`.

[108] M. Biagetti, "The Hunt for Primordial Interactions in the Large-Scale Structure of the Universe," *Galaxies* **7** (2019) 71, `arXiv:1906.12244 [astro-ph.CO]`.

[109] G. Cabass, E. Pajer, and F. Schmidt, "How Gaussian Can Our Universe Be?," *JCAP* **01** (2017) 003, `arXiv:1612.00033 [hep-th]`.

[110] V. Assassi, D. Baumann, D. Green, and L. McAllister, "Planck-Suppressed Operators," *JCAP* **01** (2014) 033, `arXiv:1304.5226 [hep-th]`.

[111] D. Schlegel *et al.* (DESI Collaboration), "A Spectroscopic Road Map for Cosmic Frontier: DESI, DESI-II, Stage-5," `arXiv:2209.03585 [astro-ph.CO]`.

[112] D. Schlegel *et al.*, "The MegaMapper: A Stage-5 Spectroscopic Instrument Concept for the Study of Inflation and Dark Energy," `arXiv:2209.04322 [astro-ph.IM]`.

[113] J. Annis, J. Newman, and A. Slosar, "Snowmass2021 Cosmic Frontier: Report of the CF04 Topical Group on Dark Energy and Cosmic Acceleration in the Modern Universe," `arXiv:2209.08049 [astro-ph.CO]`.

[114] F. Beutler, M. Biagetti, D. Green, A. Slosar, and B. Wallisch, "Primordial Features from Linear to Nonlinear Scales," *Phys. Rev. Res.* **1** (2019) 033209, `arXiv:1906.08758 [astro-ph.CO]`.

[115] M. Alvarez *et al.*, "Testing Inflation with Large-Scale Structure: Connecting Hopes with Reality," `arXiv:1412.4671 [astro-ph.CO]`.

[116] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, "LSS Constraints with Controlled Theoretical Uncertainties," `arXiv:1602.00674 [astro-ph.CO]`.

[117] G. Cabass, M. Ivanov, M. Lewandowski, M. Mirbabayi, and M. Simonović, "Snowmass White Paper: Effective Field Theories in Cosmology," *Phys. Dark Univ.* **40** (2023) 101193, `arXiv:2203.08232 [astro-ph.CO]`.

[118] G. Cabass, M. Ivanov, O. Philcox, M. Simonović, and M. Zaldarriaga, "Constraining Single-Field Inflation with MegaMapper," *Phys. Lett. B* **841** (2023) 137912, `arXiv:2211.14899 [astro-ph.CO]`.

[119] P. Creminelli and M. Zaldarriaga, "Single-Field Consistency Relation for the Three-Point Function," *JCAP* **10** (2004) 006, `arXiv:astro-ph/0407059`.

[120] P. Creminelli, J. Noreña, M. Simonović, and F. Vernizzi, "Single-Field Consistency Relations of Large-Scale Structure," *JCAP* **12** (2013) 025, `arXiv:1309.3557 [astro-ph.CO]`.

[121] P. Creminelli, J. Gleyzes, M. Simonović, and F. Vernizzi, "Single-Field Consistency Relations of Large-Scale Structure. Part II: Resummation and Redshift Space," *JCAP* **02** (2014) 051, `arXiv:1311.0290 [astro-ph.CO]`.

[122] P. Creminelli, J. Gleyzes, L. Hui, M. Simonović, and F. Vernizzi, "Single-Field Consistency Relations of Large-Scale Structure. Part III: Test of the Equivalence Principle," *JCAP* **06** (2014) 009, `arXiv:1312.6074 [astro-ph.CO]`.

[123] D. Lyth, C. Ungarelli, and D. Wands, "The Primordial Density Perturbation in the Curvaton Scenario," *Phys. Rev. D* **67** (2003) 023503, `arXiv:astro-ph/0208055`.

[124] M. Zaldarriaga, "Non-Gaussianities in Models with a Varying Inflaton Decay Rate," *Phys. Rev. D* **69** (2004) 043508, `arXiv:astro-ph/0306006`.

[125] M. Sasaki, J. Valiviita, and D. Wands, "Non-Gaussianity of the Primordial Perturbation in the Curvaton Model," *Phys. Rev. D* **74** (2006) 103003, `arXiv:astro-ph/0607627`.

[126] X. Chen and Y. Wang, "Large Non-Gaussianities with Intermediate Shapes from Quasi-Single-Field Inflation," *Phys. Rev. D* **81** (2010) 063511, `arXiv:0909.0496 [astro-ph.CO]`.

[127] X. Chen and Y. Wang, "Quasi-Single-Field Inflation and Non-Gaussianities," *JCAP* **04** (2010) 027, `arXiv:0911.3380 [hep-th]`.

[128] D. Baumann and D. Green, "Signatures of Supersymmetry from the Early Universe," *Phys. Rev. D* **85** (2012) 103520, `arXiv:1109.0292 [hep-th]`.

[129] N. Dalal, O. Doré, D. Huterer, and A. Shirokov, "The Imprints of Primordial Non-Gaussianities on Large-Scale Structure: Scale-Dependent Bias and Abundance of Virialized Objects," *Phys. Rev. D* **77** (2008) 123514, `arXiv:0710.4560 [astro-ph]`.

[130] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, "Constraints on Local Primordial Non-Gaussianity from Large-Scale Structure," *JCAP* **08** (2008) 031, `arXiv:0805.3580 [astro-ph]`.

[131] A. Ross *et al.*, "The Clustering of Galaxies in SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: Constraints on Primordial Non-Gaussianity," *Mon. Not. Roy. Astron. Soc.* **428** (2013) 1116, `arXiv:1208.1491 [astro-ph.CO]`.

[132] B. Leistedt, H. Peiris, and N. Roth, "Constraints on Primordial Non-Gaussianity from 800 000 Photometric Quasars," *Phys. Rev. Lett.* **113** (2014) 221301, `arXiv:1405.4315 [astro-ph.CO]`.

[133] E. Castorina *et al.*, "Redshift-Weighted Constraints on Primordial Non-Gaussianity from the Clustering of the eBOSS DR14 Quasars in Fourier Space," *JCAP* **09** (2019) 010, `arXiv:1904.08859 [astro-ph.CO]`.

[134] E.-M. Mueller *et al.*, "Primordial Non-Gaussianity from the Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey II: Measurements in Fourier Space with Optimal Weights," *Mon. Not. Roy. Astron. Soc.* **514** (2022) 3396, `arXiv:2106.13725 [astro-ph.CO]`.

[135] G. D'Amico, M. Lewandowski, L. Senatore, and P. Zhang, "Limits on Primordial Non-Gaussianities from BOSS Galaxy Clustering Data," `arXiv:2201.11518 [astro-ph.CO]`.

[136] G. Cabass, M. Ivanov, O. Philcox, M. Simonović, and M. Zaldarriaga, "Constraints on Multi-Field Inflation from the BOSS Galaxy Survey," *Phys. Rev. D* **106** (2022) 043506, `arXiv:2204.01781 [astro-ph.CO]`.

[137] M. Rezaie *et al.*, "Local Primordial Non-Gaussianity from the Large-Scale Clustering of Photometric DESI Luminous Red Galaxies," `arXiv:2307.01753 [astro-ph.CO]`.

[138] A. Krolewski *et al.* (DESI Collaboration), "Constraining Primordial Non-Gaussianity from DESI Quasar Targets and Planck CMB Lensing," `arXiv:2305.07650 [astro-ph.CO]`.

[139] M. Cagliari, E. Castorina, M. Bonici, and D. Bianchi, "Optimal Constraints on Primordial Non-Gaussianity with the eBOSS DR16 Quasars in Fourier Space," `arXiv:2309.15814 [astro-ph.CO]`.

[140] E. Komatsu and D. Spergel, "Acoustic Signatures in the Primary Microwave Background Bispectrum," *Phys. Rev. D* **63** (2001) 063002, `arXiv:astro-ph/0005036`.

[141] O. Doré *et al.* (SPHEREx Collaboration), "Cosmology with the SPHEREX All-Sky Spectral Survey," `arXiv:1412.4872 [astro-ph.CO]`.

[142] R. de Putter, J. Gleyzes, and O. Doré, "Next Non-Gaussianity Frontier: What Can a Measurement with $\sigma(f_{\rm NL}) \lesssim 1$ Tell Us about Multi-Field Inflation?," *Phys. Rev. D* **95** (2017) 123507, `arXiv:1612.05248 [astro-ph.CO]`.

[143] N. Arkani-Hamed and J. Maldacena, "Cosmological Collider Physics," `arXiv:1503.08043 [hep-th]`.

[144] R. de Putter and O. Doré, "Designing an Inflation Galaxy Survey: How to Measure $\sigma(f_{\rm NL}) \sim 1$ using Scale-Dependent Galaxy Bias," *Phys. Rev. D* **95** (2017) 123513, `arXiv:1412.3854 [astro-ph.CO]`.

[145] E.-M. Mueller, W. Percival, and R. Ruggeri, "Optimizing Primordial Non-Gaussianity Measurements from Galaxy Surveys," *Mon. Not. Roy. Astron. Soc.* **485** (2019) 4160, `arXiv:1702.05088 [astro-ph.CO]`.

[146] E. Sefusatti, J. Fergusson, X. Chen, and E. P. S. Shellard, "Effects and Detectability of Quasi-Single-Field Inflation in the Large-Scale Structure and Cosmic Microwave Background," *JCAP* **08** (2012) 033, `arXiv:1204.6318 [astro-ph.CO]`.

[147] J. Noreña, L. Verde, G. Barenboim, and C. Bosch, "Prospects for Constraining the Shape of Non-Gaussianity with the Scale-Dependent Bias," *JCAP* **08** (2012) 019, `arXiv:1204.6324 [astro-ph.CO]`.

[148] J. Gleyzes, R. de Putter, D. Green, and O. Doré, "Biasing and the Search for Primordial Non-Gaussianity Beyond the Local Type," *JCAP* **04** (2017) 002, `arXiv:1612.06366 [astro-ph.CO]`.

[149] S. Alam *et al.* (BOSS Collaboration), "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** (2017) 2617, `arXiv:1607.03155 [astro-ph.CO]`.

[150] N. Agarwal, S. Ho, and S. Shandera, "Constraining the Initial Conditions of the Universe using Large-Scale Structure," *JCAP* **02** (2014) 038, `arXiv:1311.2606 [astro-ph.CO]`.

[151] F. Bernardeau, S. Colombi, E. Gaztanaga, and R. Scoccimarro, "Large-Scale Structure of the Universe and Cosmological Perturbation Theory," *Phys. Rept.* **367** (2002) 1, `arXiv:astro-ph/0112551`.

[152] V. Desjacques, D. Jeong, and F. Schmidt, "Large-Scale Galaxy Bias," *Phys. Rept.* **733** (2018) 1, `arXiv:1611.09787 [astro-ph.CO]`.

[153] M. Schmittfull, M. Simonović, V. Assassi, and M. Zaldarriaga, "Modeling Biased Tracers at the Field Level," *Phys. Rev. D* **100** (2019) 043514, `arXiv:1811.10640 [astro-ph.CO]`.

[154] P. McDonald and A. Roy, "Clustering of Dark Matter Tracers: Generalizing Bias for the Coming Era of Precision LSS," *JCAP* **08** (2009) 020, `arXiv:0902.0991 [astro-ph.CO]`.

[155] N. Kaiser, "Clustering in Real Space and in Redshift Space," *Mon. Not. Roy. Astron. Soc.* **227** (1987) 1.

[156] T. Lazeyras, C. Wagner, T. Baldauf, and F. Schmidt, "Precision Measurement of the Local Bias of Dark Matter Halos," *JCAP* **02** (2016) 018, `arXiv:1511.01096 [astro-ph.CO]`.

[157] A. Lewis, A. Challinor, and A. Lasenby, "Efficient Computation of CMB Anisotropies in Closed FRW Models," *Astrophys. J.* **538** (2000) 473, `arXiv:astro-ph/9911177`.

[158] H.-J. Seo and D. Eisenstein, "Probing Dark Energy with Baryon Acoustic Oscillations from Future Large Galaxy-Redshift Surveys," *Astrophys. J.* **598** (2003) 720, `arXiv:astro-ph/0307460`.

[159] H. Zhan and L. Knox, "Baryon Oscillations and Consistency Tests for Photometrically Determined Redshifts of Very Faint Galaxies," *Astrophys. J.* **644** (2006) 663, `arXiv:astro-ph/0509260`.

[160] C. Alcock and B. Paczynski, "An Evolution-Free Test for Non-Zero Cosmological Constant," *Nature* **281** (1979) 358.

[161] M. LoVerde, A. Miller, S. Shandera, and L. Verde, "Effects of Scale-Dependent Non-Gaussianity on Cosmological Structures," *JCAP* **04** (2008) 014, `arXiv:0711.4126 [astro-ph]`.

[162] S. Matarrese and L. Verde, "The Effect of Primordial Non-Gaussianity on Halo Bias," *Astrophys. J. Lett.* **677** (2008) L77, `arXiv:0801.4826 [astro-ph]`.

[163] V. Assassi, D. Baumann, and F. Schmidt, "Galaxy Bias and Primordial Non-Gaussianity," *JCAP* **12** (2015) 043, `arXiv:1510.03723 [astro-ph.CO]`.

[164] D. Baumann, S. Ferraro, D. Green, and K. Smith, "Stochastic Bias from Non-Gaussian Initial Conditions," *JCAP* **05** (2013) 001, `arXiv:1209.2173 [astro-ph.CO]`.

[165] J. Maldacena, "Non-Gaussian Features of Primordial Fluctuations in Single-Field Inflationary Models," *JHEP* **05** (2003) 013, `arXiv:astro-ph/0210603`.

[166] E. Pajer, F. Schmidt, and M. Zaldarriaga, "The Observed Squeezed Limit of Cosmological Three-Point Functions," *Phys. Rev. D* **88** (2013) 083502, `arXiv:1305.0824 [astro-ph.CO]`.

[167] R. de Putter, O. Doré, and D. Green, "Is There Scale-Dependent Bias in Single-Field Inflation?," *JCAP* **10** (2015) 024, `arXiv:1504.05935 [astro-ph.CO]`.

[168] P. Creminelli, A. Perko, L. Senatore, M. Simonović, and G. Trevisan, "The Physical Squeezed Limit: Consistency Relations at Order q^2," *JCAP* **11** (2013) 015, `arXiv:1307.0503 [astro-ph.CO]`.

[169] A. Testa and M. Wise, "Impact of Transforming to Conformal Fermi Coordinates on Quasi-Single-Field Non-Gaussianity," *Phys. Rev. D* **102** (2020) 023533, `arXiv:2004.06126 [astro-ph.CO]`.

[170] D. Green, M. Lewandowski, L. Senatore, E. Silverstein, and M. Zaldarriaga, "Anomalous Dimensions and Non-Gaussianity," *JHEP* **10** (2013) 171, `arXiv:1301.2630 [hep-th]`.

[171] H. An, M. Wise, and Z. Zhang, "De Sitter Quantum Loops as the Origin of Primordial Non-Gaussianities," *Phys. Rev. D* **99** (2019) 056007, `arXiv:1806.05194 [hep-ph]`.

[172] M. McAneny and A. Ridgway, "New Shapes of Primordial Non-Gaussianity from Quasi-Single-Field Inflation with Multiple Isocurvatons," *Phys. Rev. D* **100** (2019) 043534, `arXiv:1903.11607 [astro-ph.CO]`.

[173] D. Green, Y. Huang, C.-H. Shen, and D. Baumann, "Positivity from Cosmological Correlators," `arXiv:2310.02490 [hep-th]`.

[174] P. Creminelli, J. Noreña, and M. Simonović, "Conformal Consistency Relations for Single-Field Inflation," *JCAP* **07** (2012) 052, `arXiv:1203.4595 [hep-th]`.

[175] T. Noumi, M. Yamaguchi, and D. Yokoyama, "Effective Field Theory Approach to Quasi-Single-Field Inflation and Effects of Heavy Fields," *JHEP* **06** (2013) 051, `arXiv:1211.1624 [hep-th]`.

[176] R. Flauger, M. Mirbabayi, L. Senatore, and E. Silverstein, "Productive Interactions: Heavy Particles and Non-Gaussianity," *JCAP* **10** (2017) 058, `arXiv:1606.00513 [hep-th]`.

[177] F. Schmidt and M. Kamionkowski, "Halo Clustering with Non-Local Non-Gaussianity," *Phys. Rev. D* **82** (2010) 103002, `arXiv:1008.0638 [astro-ph.CO]`.

[178] V. Desjacques, D. Jeong, and F. Schmidt, "Accurate Predictions for the Scale-Dependent Galaxy Bias from Primordial Non-Gaussianity," *Phys. Rev. D* **84** (2011) 061301, `arXiv:1105.3476 [astro-ph.CO]`.

[179] V. Desjacques, D. Jeong, and F. Schmidt, "Non-Gaussian Halo Bias Re-Examined: Mass-Dependent Amplitude from the Peak-Background Split and Thresholding," *Phys. Rev. D* **84** (2011) 063512, `arXiv:1105.3628 [astro-ph.CO]`.

[180] T. Giannantonio, C. Porciani, J. Carron, A. Amara, and A. Pillepich, "Constraining Primordial Non-Gaussianity with Future Galaxy Surveys," *Mon. Not. Roy. Astron. Soc.* **422** (2012) 2854, `arXiv:1109.0958 [astro-ph.CO]`.

[181] A. Barreira, "On the Impact of Galaxy Bias Uncertainties on Primordial Non-Gaussianity Constraints," *JCAP* **12** (2020) 031, `arXiv:2009.06622 [astro-ph.CO]`.

[182] A. Barreira, "Predictions for Local PNG Bias in the Galaxy Power Spectrum and Bispectrum and the Consequences for $f_{\rm NL}$ Constraints," *JCAP* **01** (2022) 033, `arXiv:2107.06887 [astro-ph.CO]`.

[183] A. Barreira, T. Lazeyras, and F. Schmidt, "Galaxy Bias from Forward Models: Linear and Second-Order Bias of IllustrisTNG Galaxies," *JCAP* **08** (2021) 029, `arXiv:2105.02876 [astro-ph.CO]`.

[184] A. Barreira, “Can We Actually Constrain $f_{\rm NL}$ using the Scale-Dependent Bias Effect? An Illustration of the Impact of Galaxy Bias Uncertainties using the BOSS DR12 Galaxy Power Spectrum,” *JCAP* **11** (2022) 013, `arXiv:2205.05673 [astro-ph.CO]`.

[185] A. Aghamousa *et al.* (DESI Collaboration), “The DESI Experiment Part I: Science, Targeting and Survey Design,” `arXiv:1611.00036 [astro-ph.IM]`.

[186] R. Laureijs *et al.* (Euclid Collaboration), “Euclid Definition Study Report,” `arXiv:1110.3193 [astro-ph.CO]`.

[187] R. Scaramella *et al.* (Euclid Collaboration), “Euclid Preparation – I. The Euclid Wide Survey,” *Astron. Astrophys.* **662** (2022) A112, `arXiv:2108.01201 [astro-ph.CO]`.

[188] P. Abell *et al.* (LSST Science & LSST Project Collaborations), “LSST Science Book, Version 2.0,” `arXiv:0912.0201 [astro-ph.IM]`.

[189] S. Ferraro *et al.*, “Inflation and Dark Energy from Spectroscopy at $z > 2$,” *Bull. Am. Astron. Soc.* **51** (2019) 72, `arXiv:1903.09208 [astro-ph.CO]`.

[190] K. Abazajian *et al.* (CMB-S4 Collaboration), “Snowmass 2021 CMB-S4 White Paper,” `arXiv:2203.08024 [astro-ph.CO]`.

[191] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, “Towards a Cosmological Neutrino Mass Detection,” *Phys. Rev. D* **92** (2015) 123535, `arXiv:1509.07471 [astro-ph.CO]`.

[192] U. Seljak, “Extracting Primordial Non-Gaussianity Without Cosmic Variance,” *Phys. Rev. Lett.* **102** (2009) 021302, `arXiv:0807.1770 [astro-ph]`.

[193] C. Orban and D. Weinberg, “Self-Similar Bumps and Wiggles: Isolating the Evolution of the BAO Peak with Power-Law Initial Conditions,” *Phys. Rev. D* **84** (2011) 063501, `arXiv:1101.1523 [astro-ph.CO]`.

[194] S. Camera, M. Santos, and R. Maartens, “Probing Primordial Non-Gaussianity with SKA Galaxy Redshift Surveys: A Fully Relativistic Analysis,” *Mon. Not. Roy. Astron. Soc.* **448** (2015) 1035, `arXiv:1409.8286 [astro-ph.CO]`.

[195] Y.-C. Li and Y.-Z. Ma, “Constraints on Primordial Non-Gaussianity from Future HI Intensity Mapping Experiments,” *Phys. Rev. D* **96** (2017) 063525, `arXiv:1701.00221 [astro-ph.CO]`.

[196] A. Moradinezhad Dizgah and G. Keating, “Line Intensity Mapping with [CII] and CO(1-0) as Probes of Primordial Non-Gaussianity,” *Astrophys. J.* **872** (2019) 126, `arXiv:1810.02850 [astro-ph.CO]`.

[197] R. Ansari *et al.* (Cosmic Visions 21 cm Collaboration), “Inflation and Early Dark Energy with a Stage-II Hydrogen Intensity Mapping Experiment,” `arXiv:1810.09572 [astro-ph.CO]`.

[198] D. Karagiannis, J. Fonseca, R. Maartens, and S. Camera, “Probing Primordial Non-Gaussianity with the Power Spectrum and Bispectrum of Future 21 cm Intensity Maps,” *Phys. Dark Univ.* **32** (2021) 100821, `arXiv:2010.07034 [astro-ph.CO]`.

[199] A. Pullen and C. Hirata, “Systematic Effects in Large-Scale Angular Power Spectra of Photometric Quasars and Implications for Constraining Primordial Non-Gaussianity,” *Publ. Astron. Soc. Pac.* **125** (2013) 705, `arXiv:1212.4500 [astro-ph.CO]`.

[200] D. Babich, P. Creminelli, and M. Zaldarriaga, “The Shape of Non-Gaussianities,” *JCAP* **08** (2004) 009, `arXiv:astro-ph/0405356`.

[201] O. Philcox and M. Ivanov, “BOSS DR12 Full-Shape Cosmology: ΛCDM Constraints from the Large-Scale Galaxy Power Spectrum and Bispectrum Monopole,” *Phys. Rev. D* **105** (2022) 043517, `arXiv:2112.04515 [astro-ph.CO]`.

[202] G. Cabass, M. Ivanov, O. Philcox, M. Simonović, and M. Zaldarriaga, “Constraints on Single-Field Inflation from the BOSS Galaxy Survey,” *Phys. Rev. Lett.* **129** (2022) 021301, `arXiv:2201.07238 [astro-ph.CO]`.

[203] O. Philcox, “Cosmology Without Window Functions: Quadratic Estimators for the Galaxy Power Spectrum,” *Phys. Rev. D* **103** (2021) 103504, `arXiv:2012.09389 [astro-ph.CO]`.

[204] F.-S. Kitaura *et al.*, “The Clustering of Galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Mock Galaxy Catalogues for the BOSS Final Data Release,” *Mon. Not. Roy. Astron. Soc.* **456** (2016) 4156, `arXiv:1509.06400 [astro-ph.CO]`.

[205] A. Chudaykin, M. Ivanov, O. Philcox, and M. Simonović, “Nonlinear Perturbation Theory Extension of the Boltzmann Code CLASS,” *Phys. Rev. D* **102** (2020) 063533, `arXiv:2004.10607 [astro-ph.CO]`.

[206] D. Blas, J. Lesgourgues, and T. Tram, “The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation Schemes,” *JCAP* **07** (2011) 034, `arXiv:1104.2933 [astro-ph.CO]`.

[207] M. Ivanov, M. Simonović, and M. Zaldarriaga, “Cosmological Parameters from the BOSS Galaxy Power Spectrum,” *JCAP* **05** (2020) 042, `arXiv:1909.05277 [astro-ph.CO]`.

[208] G. D’Amico, J. Gleyzes, N. Kokron, K. Markovic, L. Senatore, P. Zhang, F. Beutler, and H. Gil-Marín, “The Cosmological Analysis of the SDSS/BOSS Data from the Effective Field Theory of Large-Scale Structure,” *JCAP* **05** (2020) 005, `arXiv:1909.05271 [astro-ph.CO]`.

[209] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, “Cosmological Nonlinearities as an Effective Fluid,” *JCAP* **07** (2012) 051, `arXiv:1004.2488 [astro-ph.CO]`.

[210] J. J. Carrasco, M. Hertzberg, and L. Senatore, "The Effective Field Theory of Cosmological Large-Scale Structure," *JHEP* **09** (2012) 082, `arXiv:1206.2926 [astro-ph.CO]`.

[211] E. Pajer and M. Zaldarriaga, "On the Renormalization of the Effective Field Theory of Large-Scale Structure," *JCAP* **08** (2013) 037, `arXiv:1301.7182 [astro-ph.CO]`.

[212] J. J. Carrasco, S. Foreman, D. Green, and L. Senatore, "The Effective Field Theory of Large-Scale Structure at Two Loops," *JCAP* **07** (2014) 057, `arXiv:1310.0464 [astro-ph.CO]`.

[213] M. Ivanov, "Effective Field Theory for Large-Scale Structure," `arXiv:2212.08488 [astro-ph.CO]`.

[214] B. Audren, J. Lesgourgues, K. Benabed, and S. Prunet, "Conservative Constraints on Early Cosmology: An Illustration of the MontePython Cosmological Parameter Inference Code," *JCAP* **02** (2013) 001, `arXiv:1210.7183 [astro-ph.CO]`.

[215] T. Brinckmann and J. Lesgourgues, "MontePython 3: Boosted MCMC Sampler and Other Features," *Phys. Dark Univ.* **24** (2019) 100260, `arXiv:1804.07261 [astro-ph.CO]`.

[216] V. Assassi, D. Baumann, D. Green, and M. Zaldarriaga, "Renormalized Halo Bias," *JCAP* **08** (2014) 056, `arXiv:1402.5916 [astro-ph.CO]`.

[217] L. Senatore, "Bias in the Effective Field Theory of Large-Scale Structure," *JCAP* **11** (2015) 007, `arXiv:1406.7843 [astro-ph.CO]`.

[218] O. Philcox, M. Ivanov, M. Zaldarriaga, M. Simonović, and M. Schmittfull, "Fewer Mocks and Less Noise: Reducing the Dimensionality of Cosmological Observables with Subspace Projections," *Phys. Rev. D* **103** (2021) 043508, `arXiv:2009.03311 [astro-ph.CO]`.

[219] D. Green, J. Han, and B. Wallisch, *work in progress*.

[220] P. A. R. Ade *et al.* (Planck Collaboration), "Planck 2013 Results. XXIV. Constraints on Primordial Non-Gaussianity," *Astron. Astrophys.* **571** (2014) A24, `arXiv:1303.5084 [astro-ph.CO]`.

[221] R. de Putter, "Primordial Physics from Large-Scale Structure Beyond the Power Spectrum," `arXiv:1802.06762 [astro-ph.CO]`.

[222] A. Barreira and E. Krause, "Towards Optimal and Robust f_{NL} Constraints with Multi-Tracer Analyses," *JCAP* **10** (2023) 044, `arXiv:2302.09066 [astro-ph.CO]`.

[223] D. Karagiannis, R. Maartens, J. Fonseca, S. Camera, and C. Clarkson, "Multi-Tracer Power Spectra and Bispectra: Formalism," `arXiv:2305.04028 [astro-ph.CO]`.

[224] V. Desjacques, U. Seljak, and I. Iliev, "Scale-Dependent Bias Induced by Local Non-Gaussianity: A Comparison to *N*-Body Simulations," *Mon. Not. Roy. Astron. Soc.* **396** (2009) 85, `arXiv:0811.2748 [astro-ph]`.

[225] A. Pillepich, C. Porciani, and O. Hahn, "Universal Halo Mass Function and Scale-Dependent Bias from *N*-Body Simulations with Non-Gaussian Initial Conditions," *Mon. Not. Roy. Astron. Soc.* **402** (2010) 191, `arXiv:0811.4176 [astro-ph]`.

[226] M. Grossi, L. Verde, C. Carbone, K. Dolag, E. Branchini, F. Iannuzzi, S. Matarrese, and L. Moscardini, "Large-Scale Non-Gaussian Mass Function and Halo Bias: Tests on *N*-Body Simulations," *Mon. Not. Roy. Astron. Soc.* **398** (2009) 321, `arXiv:0902.2013 [astro-ph.CO]`.

[227] B. Reid, L. Verde, K. Dolag, S. Matarrese, and L. Moscardini, "Non-Gaussian Halo Assembly Bias," *JCAP* **07** (2010) 013, `arXiv:1004.1637 [astro-ph.CO]`.

[228] C. Wagner and L. Verde, "*N*-Body Simulations with Generic Non-Gaussian Initial Conditions II: Halo Bias," *JCAP* **03** (2012) 002, `arXiv:1102.3229 [astro-ph.CO]`.

[229] N. Hamaus, U. Seljak, and V. Desjacques, "Optimal Constraints on Local Primordial Non-Gaussianity from the Two-Point Statistics of Large-Scale Structure," *Phys. Rev. D* **84** (2011) 083509, `arXiv:1104.2321 [astro-ph.CO]`.

[230] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, "Large-Scale Bias and Efficient Generation of Initial Conditions for Non-Local Primordial Non-Gaussianity," *Phys. Rev. D* **85** (2012) 083002, `arXiv:1108.5512 [astro-ph.CO]`.

[231] T. Baldauf, U. Seljak, L. Senatore, and M. Zaldarriaga, "Linear Response to Long-Wavelength Fluctuations using Curvature Simulations," *JCAP* **09** (2016) 007, `arXiv:1511.01465 [astro-ph.CO]`.

[232] M. Biagetti, T. Lazeyras, T. Baldauf, V. Desjacques, and F. Schmidt, "Verifying the Consistency Relation for the Scale-Dependent Bias from Local Primordial Non-Gaussianity," *Mon. Not. Roy. Astron. Soc.* **468** (2017) 3277, `arXiv:1611.04901 [astro-ph.CO]`.

[233] A. Barreira, G. Cabass, F. Schmidt, A. Pillepich, and D. Nelson, "Galaxy Bias and Primordial Non-Gaussianity: Insights from Galaxy Formation Simulations with IllustrisTNG," *JCAP* **12** (2020) 013, `arXiv:2006.09368 [astro-ph.CO]`.

[234] A. Moradinezhad Dizgah and C. Dvorkin, "Scale-Dependent Galaxy Bias from Massive Particles with Spin during Inflation," *JCAP* **01** (2018) 010, `arXiv:1708.06473 [astro-ph.CO]`.

[235] D. Anbajagane, C. Chang, H. Lee, and M. Gatti, "Primordial Non-Gaussianities with Weak Lensing: Information on Nonlinear Scales in the Ulagam Full-Sky Simulations," `arXiv:2310.02349 [astro-ph.CO]`.

[236] M. Schmittfull and U. Seljak, "Parameter Constraints from Cross-Correlation of CMB Lensing with Galaxy Clustering," *Phys. Rev. D* **97** (2018) 123540, `arXiv:1710.09465 [astro-ph.CO]`.

[237] M. Münchmeyer, M. Madhavacheril, S. Ferraro, M. Johnson, and K. Smith, "Constraining Local Non-Gaussianities with Kinetic Sunyaev-Zel'dovich Tomography," *Phys. Rev. D* **100** (2019) 083508, `arXiv:1810.13424 [astro-ph.CO]`.

[238] S.-F. Chen, H. Lee, and C. Dvorkin, "Precise and Accurate Cosmology with CMB×LSS Power Spectra and Bispectra," *JCAP* **05** (2021) 030, `arXiv:2103.01229 [astro-ph.CO]`.

[239] M. Münchmeyer and K. Smith, "Higher N-Point Function Data Analysis Techniques for Heavy Particle Production and WMAP Results," *Phys. Rev. D* **100** (2019) 123511, `arXiv:1910.00596 [astro-ph.CO]`.

[240] M. Biagetti, A. Cole, and G. Shiu, "The Persistence of Large-Scale Structures I: Primordial Non-Gaussianity," *JCAP* **04** (2021) 061, `arXiv:2009.04819 [astro-ph.CO]`.

[241] D. Baumann and D. Green, "The Power of Locality: Primordial Non-Gaussianity at the Map Level," *JCAP* **08** (2022) 061, `arXiv:2112.14645 [astro-ph.CO]`.

[242] M. Biagetti, J. Calles, L. Castiblanco, A. Cole, and J. Noreña, "Fisher Forecasts for Primordial Non-Gaussianity from Persistent Homology," *JCAP* **10** (2022) 002, `arXiv:2203.08262 [astro-ph.CO]`.

[243] A. Andrews, J. Jasche, G. Lavaux, and F. Schmidt, "Bayesian Field-Level Inference of Primordial Non-Gaussianity using Next-Generation Galaxy Surveys," *Mon. Not. Roy. Astron. Soc.* **520** (2023) 5746, `arXiv:2203.08838 [astro-ph.CO]`.

[244] G. Jung, D. Karagiannis, M. Liguori, M. Baldi, W. Coulton, D. Jamieson, L. Verde, F. Villaescusa-Navarro, and B. Wandelt, "Quijote-PNG: Quasi-Maximum Likelihood Estimation of Primordial Non-Gaussianity in the Nonlinear Dark Matter Density Field," *Astrophys. J.* **940** (2022) 71, `arXiv:2206.01624 [astro-ph.CO]`.

[245] M. LoVerde, "Spherical Collapse in $\nu\Lambda$CDM," *Phys. Rev. D* **90** (2014) 083518, `arXiv:1405.4858 [astro-ph.CO]`.

[246] S. Vagnozzi, T. Brinckmann, M. Archidiacono, K. Freese, M. Gerbino, J. Lesgourgues, and T. Sprenger, "Bias Due to Neutrinos Must Not Uncorrect'd Go," *JCAP* **09** (2018) 001, `arXiv:1807.04672 [astro-ph.CO]`.

[247] C. Shiveshwarkar, T. Brinckmann, M. Loverde, and M. McQuinn, "Post-Inflationary Contamination of Local Primordial Non-Gaussianity in Galaxy Power Spectra," `arXiv:2306.07517 [astro-ph.CO]`.

[248] D. Alonso, P. Bull, P. Ferreira, R. Maartens, and M. Santos, "Ultra-Large-Scale Cosmology in Next-Generation Experiments with Single Tracers," *Astrophys. J.* **814** (2015) 145, `arXiv:1505.07596 [astro-ph.CO]`.

[249] E. Castorina and A. Moradinezhad Dizgah, "Local Primordial Non-Gaussianities and Super-Sample Variance," *JCAP* **10** (2020) 007, `arXiv:2005.14677 [astro-ph.CO]`.

[250] E. Castorina and E. di Dio, "The Observed Galaxy Power Spectrum in General Relativity," *JCAP* **01** (2022) 061, `arXiv:2106.08857 [astro-ph.CO]`.

[251] R. Martinez-Carrillo, J. Hidalgo, K. Malik, and A. Pourtsidou, "Contributions from Primordial Non-Gaussianity and General Relativity to the Galaxy Power Spectrum," *JCAP* **12** (2021) 025, `arXiv:2107.10815 [astro-ph.CO]`.

[252] M. Foglieni, M. Pantiri, E. Di Dio, and E. Castorina, "Large-Scale Limit of the Observed Galaxy Power Spectrum," *Phys. Rev. Lett.* **131** (2023) 111201, `arXiv:2303.03142 [astro-ph.CO]`.

[253] H. Lee, D. Baumann, and G. Pimentel, "Non-Gaussianity as a Particle Detector," *JHEP* **12** (2016) 040, `arXiv:1607.03735 [hep-th]`.

[254] S. Behbahani and D. Green, "Collective Symmetry Breaking and Resonant Non-Gaussianity," *JCAP* **11** (2012) 056, `arXiv:1207.2779 [hep-th]`.

[255] A. Bodas, S. Kumar, and R. Sundrum, "The Scalar Chemical Potential in Cosmological Collider Physics," *JHEP* **02** (2021) 079, `arXiv:2010.04727 [hep-ph]`.

[256] P. D. Meerburg, M. Münchmeyer, J. Muñoz, and X. Chen, "Prospects for Cosmological Collider Physics," *JCAP* **03** (2017) 050, `arXiv:1610.06559 [astro-ph.CO]`.

[257] J. McEwen, X. Fang, C. Hirata, and J. Blazek, "FAST-PT: A Novel Algorithm to Calculate Convolution Integrals in Cosmological Perturbation Theory," *JCAP* **09** (2016) 015, `arXiv:1603.04826 [astro-ph.CO]`.

[258] M. D'Onofrio and K. Rummukainen, "Standard Model Cross-Over on the Lattice," *Phys. Rev. D* **93** (2016) 025003, `arXiv:1508.07161 [hep-ph]`.

[259] F. Beutler *et al.* (BOSS Collaboration), "The Clustering of Galaxies in the Completed SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in Fourier Space," *Mon. Not. Roy. Astron. Soc.* **464** (2017) 3409, `arXiv:1607.03149 [astro-ph.CO]`.

[260] A. Font-Ribera, P. McDonald, N. Mostek, B. Reid, H.-J. Seo, and A. Slosar, "DESI and Other Dark Energy Experiments in the Era of Neutrino Mass Measurements," *JCAP* **05** (2014) 023, `arXiv:1308.4164 [astro-ph.CO]`.

[261] O. Doré *et al.* (SPHEREx Collaboration), “SPHEREx Public Products Github Repository: Forecasted Galaxy Number Density and Bias for Fives Samples and Multiple Galaxy Redshift Bins,” 2020. https://github.com/SPHEREx/Public-products/blob/master/galaxy_density_v28_base_cbe.txt.