Imperial College London
Physics Department

Stochastic Dark Sector

Arad Nasiri

June 2025

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Physics of Imperial College London and the Diploma of Imperial College London.

Abstract

Cosmology probes the largest scales of the Universe, and quantum gravity the smallest. In this thesis, within the framework of quantum gravity phenomenology, we explore how stochastic effects in the dynamics of the dark sector can bridge those regimes.

We first study covariant Brownian motion as an effective signature of spacetime discreteness. We derive the process from first principles via Fokker-Planck and Langevin perspectives, and show its uniqueness when minimally coupled to curvature. Applied to dark matter particles, the resulting stochastic dark matter is governed by a single diffusion rate. We solve for the linear perturbations semi-analytically and numerically, showing that the model suppresses the matter power spectrum at small scales and offers a resolution to the S_8 discrepancy between CMB and low-redshift probes.

We then extend the framework to massless particles. We solve the modified Boltzmann equation analytically in an FRW spacetime, showing that diffusion impacts long-wavelength modes most strongly. We show that a future detection of a stochastic gravitational-wave background with LISA can improve current CMB bounds on the diffusion and drift parameters by over twelve orders of magnitude.

Finally, we examine Everpresent Λ, a stochastic model of dark energy. A key feature of this model is that Λ fluctuates and tracks the dominant energy density in order to resolve the cosmological-constant puzzle without fine-tuning. Accounting for initial conditions, we find hints that the nonlocality in the model is capable of producing inflationary expansion for many e-folds. A subset of Everpresent-Λ realisations fit SN Ia data better than ΛCDM, but the model struggles at describing CMB. We also study the allowed fluctuations of dark-energy density in a more model-independent manner, and find that some variation, especially prior to recombination, can lead to improvements over ΛCDM; but the favoured fluctuations are smaller than is typical in Everpresent Λ.

Declaration of Originality

This is the author's own work, and the contributions and work of others have been appropriately referenced. A full list of references is given in the bibliography. The four main chapters have significant overlap with the following four papers, respectively:

[1] Emma Albertini, Arad Nasiri, and Emanuele Panella. "Stochastic dark matter: Covariant Brownian motion from Planckian discreteness". In: *Phys. Rev. D* 111.2 (2025), p. 023514. DOI: 10.1103/PhysRevD.111.023514. arXiv: 2409.02188 [gr-qc].

[2] Giorgio Mentasti and Arad Nasiri. "Covariant diffusion and drift of the stochastic GW background with LISA". In: (May 2025). arXiv: 2505.24620 [gr-qc].

[3] Santanu Das, Arad Nasiri, and Yasaman K. Yazdi. "Aspects of Everpresent Λ. Part I. A fluctuating cosmological constant from spacetime discreteness". In: *JCAP* 10 (2023), p. 047. DOI: 10.1088/1475-7516/2023/10/047. arXiv: 2304.03819 [gr-qc].

[4] Santanu Das, Arad Nasiri, and Yasaman K. Yazdi. "Aspects of Everpresent Λ. Part II. Cosmological tests of current models". In: *JCAP* 10 (2024), p. 076. DOI: 10.1088/1475-7516/2024/10/076. arXiv: 2307.13743 [astro-ph.CO].

In particular, the Fisher analysis in section 2.4 is done by Giorgio Mentasti and the MCMC analysis in sections 4.2 and 4.3 is performed by Santanu Das.

Copyright

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

"... For else, the whole of heart and soul amounts to nought."

— Hafez

Words cannot fully express my gratitude to my supervisor, Fay Dowker, for all her kind support and wisdom in this journey. She has been patient with my eager yet unpolished enthusiasm, helping mould it into a more mature and reflective attitude toward physics. She encouraged me to explore freely, while masterfully keeping me from straying too far. For all of this, I am humbly grateful.

I owe special thanks to Hessamaddin Arfaei, Shant Baghram, and Bahram Mashhoon, who helped me find the courage to pursue a PhD. I learned far more than they can imagine from their approach to science and way of life.

I am immensely thankful to Yasaman Yazdi for her unwavering presence and guidance, always encouraging me to step out of my comfort zone, and teaching me a great deal through her balanced approach to research and collaboration. Throughout my PhD, I have had the chance to work with some truly amazing people. I am grateful to Santanu Das for his patience with all my cosmology and coding questions. I also thank Stuart Ferguson, Caio Nascimento, Emanuele Panella, Murali Saravanan, Petros Wallden, and Stav Zalel for their ingenuity and brilliance.

I thank Toby Wiseman for his positive feedback early in my PhD. I cherish the company of Bruno Alexandre, Haya AlMuhanna, Guillermo Arias, Stephanie Baines, Carlos Barredo, Paolo Bassani, Sebastian Cespedes, Ben Elder, Kinga Gawrych, Morteza Hosseini, Raymond Isichei, Sumer Jaitly, Sadra Jazayeri, Theo Keseman, Chunhao Li, Jieming Lin, Andreas Mantziris, Karapet Mkrtchyan, Daniel Platt, and Fridrich Valach in the theory group at Imperial.

I deeply thank Niayesh Afshordi for making my one-year experience at PI as smooth and enriching as possible, and for being such a great source of knowledge and inspiration. I am thankful for having had the chance to engage in long, stimulating discussions with Rafael Sorkin. His unflagging curiosity is the ultimate ideal for me to strive toward.

I thank Ghazal Geshnizjani for her kind support during my stay at PI. I enjoyed insightful discussions with Bianca Dittrich, John Moffat, and Dongxue Qu. I feel fortunate to have crossed paths with Hamza Ahmed, Cristian Boghiu, Johanna Borissova, Amir Dehghani, Hassan Khalvati, Amin Moharramipour, and Robin Oberfrank.

I appreciate President's PhD Scholarship at Imperial College, the graduate student exchange grant from the McDonald Canadian Astroparticle Institute, the hospitality I received as a visitor at the Perimeter Institute, and STFC grant ST/W006537/1. This support has provided unique opportunities for me to grow as a physicist.

I am grateful to Saman, Sasan, Leila, Saeideh, Shayan, Mohammad Hassan, and Arshia for their generous and heartwarming support during my early days in London. It made the transition to a new country much smoother.

PhD years come with their ups and downs. The persistent compassion of my cohorts was a light in the gloomier days in London. To Emma, for her never-ending source of energy; to Jan, for our unforgettable hikes; to Giorgio, for his wild riffs and kind heart; and to Pulkit, for bearing my ramblings and complaints.

Over the years, I have grown alongside some lifelong friends, tracing this path shoulder to shoulder, sharing both tears and laughter. Owing to their warm presence, I am a better version of myself. To Amir, for always hearing me out with compassion; to AmirHossein, for our provably identical thought processes; to Yasamin, for seeing through and cheering the best in me; to Ramtin, for his reliable brotherly care; to Mashakh, for our meandering rambles that border on absurdity and brighten even the hardest days; to Yazdan, for his inspiring commitment to friendship; to Hasanz, for our chats on the most uncharted subjects; and to Daei, Navid, and Farbod, for all our intellectual exchanges on poems, maths, and beyond. Thanks are also due to Mamad, Shakiba, Muhammad, Reza, Negin, and Keivan.

Last but not least, I am forever thankful to my parents, without whose constant encouraging support and loving faith, this journey would have never begun.

For those I hold dearest
In memory of Aras Shirani

Contents

Introduction **17**

1 Stochastic dark matter **20**
1.1 Foundation . . . 23
1.1.1 From Fokker-Planck to covariant diffusion equation . . . 23
1.1.2 Stochastic geodesic equation: Langevin equations . . . 29
1.1.3 Swerves models . . . 34
1.2 Cosmological tests . . . 35
1.2.1 Generalised Boltzmann hierarchy . . . 38
1.2.2 Semi-analytic solutions . . . 39
1.2.3 Dark energy model . . . 44
1.2.4 Effect on observables . . . 46
1.2.5 Observational Constraints . . . 47
1.2.6 Discussion . . . 50
1.3 Summary and outlook . . . 52

2 Stochastic gravitational-wave background: constraints on massless diffusion **54**
2.1 The stochastic GW background . . . 56
2.1.1 Wigner function of gravitational waves . . . 57
2.2 Modified transport equation for the distribution function . . . 58
2.2.1 Stochastic massless geodesics . . . 58
2.2.2 Massless diffusion equation . . . 60
2.3 Analytic Solution . . . 60
2.3.1 Perturbative solution for small $\kappa_{1,2}$ . . . 62
2.4 Forecast of $\kappa_{1,2}$ estimation with LISA . . . 63
2.4.1 Diffusion on first order phase transitions and PBH spectra . . . 64
2.5 Summary and outlook . . . 65

3 Stochastic dark energy: introduction to Everpresent Λ **68**
3.1 Background theory . . . 69
3.1.1 Causal set theory . . . 69
3.1.2 Unimodular gravity . . . 71
3.2 Everpresent Λ . . . 73

3.2.1 A phenomenological implementation—Model 1 . . . 75
3.3 Methodology . . . 77
3.3.1 An isotropic and homogeneous universe . . . 78
3.3.2 Numerical recipe . . . 78
3.3.3 $H^2 < 0$ crashed runs . . . 80
3.3.4 ρ_m as a scaling parameter at late times . . . 81
3.4 Results and further investigations . . . 82
3.4.1 Statistical distributions of simulation outcomes . . . 82
3.4.2 Frequency of completed runs . . . 83
3.4.3 Initial conditions . . . 85
3.4.4 Understanding the auto-correlation time . . . 88
3.5 Discussion . . . 91

4 Stochastic dark energy: cosmological tests of Everpresent Λ 93
4.1 Fitting Model 1 to supernova data . . . 93
4.1.1 MCMC analysis . . . 94
4.1.2 Diagnosing good seeds by a single parameter . . . 98
4.1.3 Understanding the results with mock data . . . 100
4.2 Fitting Model 1 to CMB data . . . 102
4.2.1 Methodology . . . 102
4.2.2 Results . . . 103
4.3 Constraining dark energy using CMB . . . 104
4.3.1 During recombination . . . 106
4.3.2 After recombination . . . 108
4.3.3 Before recombination . . . 110
4.4 Summary and outlook . . . 112

Epilogue 114

Bibliography 115

A Technical supplement to Chapter 1 131
A.1 General covariance of Fokker-Planck equation . . . 131
A.1.1 Detailed derivation of covariant diffusion equation . . . 132
A.1.2 Covariance of Liouville operator . . . 133
A.2 Unravelling the covariant diffusion equation . . . 135
A.2.1 Equivalence between Fokker-Planck equation and Itô SDEs . . . 135
A.2.2 From covariant diffusion equation to stochastic trajectories . . . 136
A.2.3 Relation between spacetime covariance and phase space covariance . . . 137
A.3 Generalised Boltzmann hierarchy equations for stochastic dark matter . . . 140
A.3.1 Superhorizon solutions for stochastic dark matter . . . 143
A.4 Dark energy equations when coupled to stochastic dark matter . . . 143
A.5 Derivation of semi-analytic solutions to stochastic dark matter perturbations in FRW . . . 144
A.6 Full posteriors for stochastic dark matter . . . 149

B Technical supplement to Chapter 2 **153**
B.1 Analytic solution to background massless diffusion equation in FRW 153
B.2 Fisher analysis 156

C Technical supplement to Everpresent Λ **158**
C.1 On the equations of motion of global unimodular gravity 158
C.1.1 An alternative derivation of the equations of motion in unimodular gravity . . . 158
C.1.2 An Alternative derivation of (C.5) 159
C.2 On the fluid interpretation of Everpresent Λ 160
C.3 MCMC analysis with a seed number 161
C.4 Full posteriors for fluctuating dark energy 164

List of Figures

1.1 Five realisations of the covariant Brownian motion 32
1.2 Unit disks on the hyperbolic mass shell 33
1.3 Background energy density and equation of state of stochastic dark matter 37
1.4 Evolution of the stochastic dark matter overdensity versus scale factor 38
1.5 Stochastic dark matter and dark energy overdensity and transfer function 44
1.6 The CMB TT power spectrum, $\mathcal{D}_\ell$, for increasing values of the diffusion rate 46
1.7 Suppression of the matter power spectrum in stochastic dark matter scenario 47
1.8 Dependence of σ_8 on the diffusion rate, and the gravitational potential versus scale factor 48
1.9 Summarised corner plots, including S_8 and the diffusion rate, from various combinations of datasets 49

2.1 The physical heat kernel K_- of the massless diffusion equation for different values of $\kappa_{1,2}$ 62
2.2 GW energy density $\Omega_{\rm GW}(\nu)$ of a stochastic background from a first-order phase transition and from primordial black holes for different values of $\kappa_{1,2}$ 66
2.3 Corner plot of the Fisher forecast of the full posterior distribution $P(\Omega_*, \nu_b, \kappa_1, \kappa_2)$ assuming a flat prior and an effective total time of observation with LISA of $T_{\rm obs} = 5$ years. 67

3.1 The evolution of dark energy, matter, and radiation in Everpresent Λ scenario 80
3.2 The statistical distribution of H_0, Ω^0_Λ, and $H_0^4 V$ for Everpresent Λ Model 1 83
3.3 Crashing fraction of Model 1 versus time step 85
3.4 Dark energy density and H^2 versus scale factor for different initial actions S_0 in Model 1 86
3.5 Volume versus scale factor and versus conformal time for different initial actions S_0 in Model 1 87
3.6 The normalised auto-correlation function of dark energy fraction versus $\Delta \log a$ 89
3.7 Dark energy fraction versus redshift for realisations of Model 1 and Model 2 90

4.1 Some statistics from the fit of Model 1 to SN Ia 95
4.2 The corner plot of the seed with lowest χ^2 for SN Ia 97
4.3 Dark energy density comparison between good and bad seeds at low redshift 99
4.4 Om diagnostic versus redshift for good and bad seeds 100
4.5 The Hubble parameter versus redshift for four of the best and four of the worst seeds . 101
4.6 Posterior histograms of Model 1 fit to CMB, with and without dark energy suppression at last scattering surface 104
4.7 CMB power spectrum for best-fit of ΛCDM and Model 1 to Planck data 105

4.8 Early dark energy effect on ionization fraction, visibility function, and CMB power spectrum . . . 107
4.9 A few arbitrary expansion histories from the 200 best-fit realisations of post-recombination fluctuating dark energy to the CMB . . . 109
4.10 A few arbitrary expansion histories from the 200 best-fit realisations of post-recombination fluctuating dark energy for a different set of redshifts . . . 109
4.11 Effect of positive and negative pre-recombination early dark energy on C_ℓ's . . . 110
4.12 A few arbitrary realisations of Ω_Λ from the 200 best-fit realisations of pre-recombination fluctuating dark energy to the CMB . . . 111

A.1 Comparison of the integrand of the time-integral term with its saddle-point approximation for stochastic dark matter overdensity and velocity field . . . 148
A.2 The extended corner plot for stochastic dark matter and ΛCDM from Planck . . . 150
A.3 The extended corner plot for stochastic dark matter and ΛCDM from Planck and low-z Baseline . . . 151
A.4 The extended corner plot for stochastic dark matter and ΛCDM from Baseline and Baseline+DES . . . 152

C.1 Dark energy equation of state versus redshift for one realisation of Everpresent Λ . . . 161
C.2 The corner plot of post-recombination fluctuating dark energy by the CMB . . . 164
C.3 The corner plot of pre-recombination fluctuating dark energy by the CMB . . . 165

List of Tables

1.1 Best-fit χ^2 of stochastic dark matter and ΛCDM from different datasets. 50
1.2 Set of best-fit parameters and posterior averages and standard deviations from the analysis of Baseline and Baseline+DES for stochastic dark matter and ΛCDM. 50

4.1 Best-fit parameters and χ^2 of three variations of Model 1 with CMB 105
4.2 Average and the best-fit parameters of the phenomenological fluctuating dark energy models with CMB 112

List of Notations and Abbreviations

$\hbar,\ c$: The reduced Planck constant and the speed of light. Unless otherwise specified, we set $\hbar = c = 1$.

C_ℓ : Angular power spectrum multipole

f: Phase space distribution function

f: Function encoding the model-independent variations in dark energy, Chapter 4.

G : Newton's gravitational constant

Γ : Diffusion rate for massive particles

H : Hubble parameter

$\mathcal{H}$: Comoving Hubble parameter $= aH$

h : Reduced Hubble constant $H_0/(100\text{km/s/Mpc})$ in Chapters 3 and 4

h^{ij} : Inverse Riemannian metric on the hyperbolic mass shell in Chapter 1

$\boldsymbol{h}_{\mu\nu}$: Gravitational waves in Chapter 2

η : Conformal time

k : Wavevector

k_B : Boltzmann constant

κ : Diffusion constant for massive particles

$\kappa_{1,2}$: Diffusion constant for massless particles

ℓ_p, t_p: Planck length and Planck time

Λ : Cosmological constant

p: Spacetime momentum

P: Pressure

$\mathbb{P}$: Probability

ρ: Energy density

ϱ: Density of causal set elements, presumably $\sim 1/\ell_p^4$

$S_8 \equiv \sigma_8 \sqrt{\frac{\Omega_m}{0.3}}$

σ_8 : Standard deviation of the matter overdensity field, integrated over a spherical top-hat window of comoving radius $8h^{-1}$Mpc

T : Temperature

$\mathcal{T}$: Rescaled cosmological time $d\mathcal{T} = a^3 d\eta$ for the diffusion of massive particles

T : Rescaled cosmological time $\mathsf{T} = a^2 d\eta$ for the diffusion of massless particles

τ : Proper time of massive particles

τ_{opt}, τ_{reio} : Optical depth, Reionization optical depth

x : Spacetime coordinate

z : Cosmological redshift

CMB: Cosmic microwave background

DES: Dark Energy Survey

FRW: Friedmann–Robertson–Walker metric

GR: General Relativity

KiDS: Kilo-Degree Survey

LIGO: Laser Interferometer Gravitational-Wave Observatory

LISA: Laser Interferometer Space Antenna

ΛCDM: Benchmark cosmological model with a cosmological constant (Λ) and Cold Dark Matter

MCMC: Markov Chain Monte Carlo

SDE: Stochastic Differential Equation

SGWB: Stochastic Gravitational-Wave Background

SW, ISW: Sachs-Wolfe effect, Integrated Sachs-Wolfe effect

UV/IR: Ultraviolet / Infrared

Introduction

"Then to the rolling Heav'n itself I cried
Asking, what lamp had Destiny to guide
Her little children stumbling in the dark
And—A blind understanding! Heav'n replied."

— *Omar Khayyam*

The search for a quantum theory of gravity hinges on a deeper understanding of the nature of spacetime on the one hand and uncovering its quantum dynamics on the other. Regarding kinematics, a key concept in several quantum gravity proposals is the idea of spacetime discreteness [1]. In these theories, the smooth spacetime of general relativity is viewed as an effective, coarse-grained approximation of a fundamental discrete structure. Several theoretically compelling reasons support this view. Apart from the immediate reason of the mere existence of Planck length as a distinguished physical scale, discreteness provides a natural explanation for the entropy of black holes based on either counting the microscopic discrete states [2, 3] or computing the entanglement entropy [4]. Moreover, in sum-over-histories quantization, a countable summation over kinematically discrete degrees of freedom is well-defined mathematically and more tractable than continuum path integrals [5, 6]. A discrete spacetime also naturally imposes a UV cutoff in field theories, turning perturbation series into finite sums [7, 8].

However, a major challenge in the search for quantum gravity in general and discrete spacetime in particular is the lack of direct empirical evidence. The Planck energy scale is by many orders of magnitude beyond the reach of current collider experiments. Yet, small quantum gravity effects could accumulate and amplify over cosmic distances or across the age of the universe leaving observable traces in the behaviour or properties of both massive and massless particles. In the era of precision cosmology, this opens a unique window for quantum gravity phenomenology [9]. The interplay between UV and IR physics is also at the heart of several proposed models for dark energy [10–16]. This is in contrast to many of our conventional tools for describing nature, such as renormalization group theory in quantum field theory, which rely on the independence of effects at different scales. This mixing of UV and IR effects also ties in with the expectation that quantum gravity will involve nonlocality in a crucial way. Here, by nonlocality we mean causal influence from within the past light cone [17–19].

On the cosmology side, after the surprising discovery that the expansion of the Universe is accelerating [20, 21], one of the leading candidates for dark energy has been the cosmological constant. Along

with cold dark matter, it has been combined into the very successful ΛCDM model for describing the evolution of the Universe. It is not only effective at the background level but also robust at linear perturbations, giving accurate predictions for the matter power spectrum and for the micro-Kelvin level perturbations in the photon temperature observed in the cosmic microwave background (CMB).

Despite these successes of the standard ΛCDM model, cosmology faces several challenges and open theoretical problems. Although the cosmological constant is the simplest dark energy candidate, its ad hoc value, $\Lambda \sim 10^{-121}$ in Planck units, is an important theoretical motivation for searching for alternative dark energy models. In other words, a fundamental explanation should address the coincidence problem, which asks why Λ is fine-tuned in such a way to only become relevant at the present epoch in cosmic expansion. In addition to such open theoretical problems about the nature of dark matter and dark energy, current cosmology faces tensions between different measurements of H_0 and S_8 within the standard ΛCDM model [22–30]. The former, known as the Hubble tension, is the discrepancy in H_0 inferred from CMB measurements assuming the ΛCDM model and its direct measured value from supernovae [22, 31–38]. Here, the S_8 tension will be of particular focus—a discrepancy that arises when comparing the variance of matter overdensities at the scale of $8h^{-1}$Mpc as measured from CMB [39] with measurements at lower redshifts from probes such as weak lensing [25–27], galaxy power spectrum [40–42], galaxy cluster counts [43], or redshift-space distortion [44, 45]. If these discrepancies are signs of new physics, they provide clear observable targets for well-motivated phenomenological models arising from quantum gravity to provide alternative explanations.

Historically, a common theme in many of the studies on quantum gravity phenomenology has been violating Lorentz invariance, see [46–48] and references therein. Current bounds on Lorentz violations from LIGO events remain weak [49]; however, studies on short gamma-ray bursts [50–53] and flares from active galactic nuclei [54, 55] have placed stringent, in cases trans-Planckian, constraints on the energy scale of such violations in the relativistic dispersion equations. Given these constraints, it is important to invest on Lorentz-invariant quantum gravity phenomenology. Assuming Lorentz invariance is exact at all energies, we would no longer have the option of adding Lorentz violating operators in the standard model Lagrangian (as is systematically done in [56, 57]) or modifying the massive or massless dispersion relation. Instead, quantum gravity could manifest as subtle stochastic effects in the propagation of particles or in the behaviour of dark energy.

Why should quantum gravity have a stochastic observable signature? Stochasticity could be either a manifestation of the discrete nature of spacetime preventing particles from moving exactly on geodesics; so the multitude of Planck-scale degrees of freedom appear to us as stochastic noise in the motion of free particles. This is similar in nature and motivated by the Brownian motion in fluids. Alternatively, stochasticity could result from an uncertainty in the total spacetime volume given a definite fundamentally discrete configuration of spacetime. This uncertainty would in turn drive a stochastic behaviour in the variable conjugate to spacetime volume, i.e. cosmological constant. These are the two main ideas pursued in this thesis.

In the first part of the thesis, we study the covariant Brownian motion. The idea is based on the Lorentz invariant diffusion equation developed in [58] and based on the coordinate transformation properties of the drift and diffusion coefficient in the Fokker-Planck equation found by Sorkin [59]. Here,

we generalise the setup to curved spacetime and find the corresponding stochastic geodesic equation. As in any other random walk process, the variance of the displacement on the mass shell grows with time. Therefore, for the effect to become observable, we need cosmological times. We therefore apply this framework to dark matter particles, leading to a model of stochastic dark matter. We study its implications for cosmological observables such as the CMB angular power spectrum, matter power spectrum, and the S_8 tension.

As another example, we study the massless diffusion equation developed by Dowker et al. [60] and its generally covariant version. We find the corresponding stochastic massless geodesic equation which is phenomenologically described by just two diffusion and drift parameters. These are already constrained by measurements of the CMB blackbody spectrum [60]. We show that, counterintuitively, the effect is dominant at lower frequencies. Therefore, a stochastic gravitational-wave background (SGWB) in the mHz frequency band is ideal for detecting such an effect. In particular, we find that detecting a GW background sourced by realistic models in the primordial universe can constrain the parameters by over 12 orders of magnitude better compared to the bounds from the CMB.

In the second part of the thesis, we study Everpresent Λ [61–63]. This is the oldest example of UV-IR mixing in a cosmological context, and one that epitomises the idea. Everpresent Λ relies only on some basic properties of causal sets, a kind of discrete spacetime proposed to be more fundamental than the conventional continuum spacetimes we are used to. A key feature of this model is that the cosmological constant *fluctuates* over the course of the evolution of the universe. Everpresent Λ dates back to 1987, when Sorkin presented it as a heuristic argument for the magnitude of the cosmological constant [61, 63], well before supernova observations confirmed a value of $\Lambda \sim 10^{-121}$ in Planck units [20, 21]. The heuristic argument has since been fleshed out in subsequent works [64–66]. In this thesis, after reviewing the core argument, we study several statistical properties of a model that encapsulates the idea in an FRW universe. We then study the best-fit realizations of the model to supernovae type Ia (SN Ia) data. We discuss the challenge in fitting a stochastic model like Everpresent Λ to CMB data and develop an understanding on the magnitude of dark energy fluctuations allowed by CMB.

The thesis is structured as follows:

- Chapter 1 introduces the concept of the covariant Brownian motion and stochastic geodesic equation. The idea is applied to dark matter particles, developing a model of stochastic dark matter, and ends with constraints from CMB and $f\sigma_8$ measurements.
- Chapter 2 investigates the massless geodesic equation and makes forecast for its effect on stochastic gravitational-wave background as will be observed by LISA.
- Chapter 3 introduces Everpresent Λ as a fluctuating model of dark energy and expands on its theoretical aspects and statistical properties in the late universe.
- Chapter 4 expands on the cosmological consequences of Everpresent Λ and confronts it with SN Ia and CMB data. Moreover, two generic models of dark energy with sign-agnostic fluctuations are introduced and constrained in order to explore the tracking behaviour of Everpresent Λ in a model-independent way.

Chapter 1

Stochastic dark matter

In this chapter, we explore how quantum gravity could manifest effectively as a Lorentz-invariant stochastic noise in the motion of point particles. To see how spacetime discreteness is linked to the stochastic motion of particles, consider the analogy with the atomic theory of matter. Brownian motion, linked to kinetic theory and the diffusion equation by Einstein in 1905, was key evidence for the atomic nature of matter. The stochastic motion of test particles, whether understood through the behaviour of a single one or the diffusion of many, can be effectively described by just two free parameters: the diffusion coefficient and the friction coefficient. No detailed knowledge of the microscopic dynamics of the background fluid is needed. The inherent stochasticity of the Brownian motion reflects our ignorance of the detailed microscopic configuration of the fluid's atoms or molecules.

Similarly, one can propose the Brownian motion of massive particles in the vacuum of general relativity as a manifestation of underlying spacetime discreteness. In a smooth spacetime manifold, free massive particles follow timelike geodesics. However, if this manifold is merely the continuum limit of a fundamentally discrete spacetime, the geodesic equation would be modified, causing the particles to 'swerve'. Pictorially, as the particles jump from one spacetime atom to another, they will almost surely not find a spacetime atom exactly along the continuum timelike geodesic and will therefore land on the closest available point. This leads to diffusion around the continuum-limit timelike geodesic. In 2003, Dowker et al. proposed a concrete model in the framework of causal set theory, a Lorentz-preserving discrete approach to quantum gravity, to explore this idea [58]. Later in 2008, Dowker et al. added two intrinsic models on a causal set that encompassed the idea [60]. They all feature a nonlocal scale, called the forgetting time τ_f. This is the amount of proper time that a massive particle "remembers" from its past motion, so as to approximately maintain its motion on a geodesic. It is distinct from other typical nonlocal scales in the sense that the geodesic motion is recovered in the $\tau_f \to \infty$ limit. These models, however, should be treated with caution, as the authors themselves noted that a point particle is not an adequate approximation when quantum gravity effects are significant. Rather, they serve as a first approximation to the more complicated problem of the stochastic dynamics of a quantum field due to spacetime discreteness. These models are collectively called the *swerves* models. However, regardless of the exact fundamental discrete degrees of freedom or their interaction with quantum matter, the effective description of a single particle's motion or the distribution of many particles in the continuum approximation should depend only on a few parameters.

One of the earliest works on relativistic Brownian motion was by Dudley in 1965 [67], where it was shown that no Lorentz-invariant stochastic process exists in Minkowski spacetime. One way to understand this is to note that the metric $\eta_{\mu\nu}$ that appears in the flat Laplacian is not positive semi-definite, meaning that it cannot be used as the diffusion tensor for a stochastic process. The resolution is to include the momentum variables in the state of the particle and define the diffusion process on the mass shell in momentum space instead. This defined a diffusion equation on the mass shell valid for a spatially uniform distribution of particles. In [58], Dowker et al. used Sorkin's 1985 formalism of stochastic evolution on the manifold of states [59] to derive a Lorentz invariant diffusion equation for relativistic particles on the mass shell, with only one free parameter: the diffusion constant κ. This generalised Dudley's equation to spatially inhomogeneous distributions. They highlighted the role of proper time elapsed along the worldline of the particle as the evolution parameter for the distribution function. They also placed an upper bound on the diffusion parameter from cold Hydrogen gas in laboratory settings. In parallel, coming from a mathematical perspective, Franchi and Le Jan [68] generalised Dudley's diffusion process to curved Lorentzian manifolds, and studied its asymptotic behaviour in Schwarzschild spacetime. In [69], Kaloper and Mattingly found that the observations of cold molecular clouds in the Milky Way, Bose-Einstein condensate experiments with 87-Rubidium, and the stability of old heavy nuclei put similar bounds on the diffusion parameter. Later in [60], Dowker et al. elaborated on why global Lorentz symmetry is so restrictive, essentially making the Lorentz invariant diffusion equation unique. The authors also considered massless particles, where key differences arise compared to massive particles when Lorentz invariance is enforced. For massless particles, the diffusion does not happen on the whole light cone. Instead, the distribution broadens and drifts in one dimension along the original momentum, without changing direction. The two effects are controlled by two independent free parameters. To test this, they applied the massless diffusion equation to CMB photons, adding a Hubble friction term by hand to extend it to FRW spacetime. Using the blackbody spectrum of CMB photons, they placed upper bounds on the two free parameters. Later in [70], Contaldi et al. used the CMB polarization power spectrum to set upper bounds on a Lorentz invariant diffusion model of photon polarization.

On a related but distinct line of enquiry, there is the problem of finding the relativistic version of the ordinary Brownian motion in a fluid. For this problem, the diffusion no longer needs to be Lorentz invariant because the background fluid selects a preferred frame. Instead, the issue is that the simplest diffusion equation theoretically allows the particles to move outside the light cone. This is also shared by the Ornstein-Uhlenbeck process, which models the diffusion in momentum space rather than position space. To remedy this, Debbasch et al. [71] introduced the relativistic version of the Ornstein-Uhlenbeck process that preserves the particle momentum on the mass shell. So relativistic causality is preserved, but the existence of the fluid rest frame explicitly breaks Lorentz invariance and allows for a version of Einstein relation. Later, Dunkel and Hänggi [72] proposed another relativistic stochastic process with a Lorentz invariant covariance matrix for the noise. Nevertheless, they also included a Lorentz-violating friction term that distinguishes the fluid rest frame and allows for equilibrium solutions. They derived three different Fokker-Planck equations depending on the choice of the stochastic integral (Itô, Stratonovich, or backward Itô). Alcantara Felix and Calogero [73] clarified that only one of the choices leads to a Lorentz invariant diffusion equation in the absence of friction. The equilibrium solution, the coupling to electromagnetic field, and the relevance in plasma physics were also discussed by these

authors, as well as by Haba [74–76]. See [77, 78] for two reviews of the frame-dependent relativistic Brownian motion. Hermann [79] generalised a particular frame-dependent Brownian motion to curved spacetime, and recently, Cai et al. [80, 81] formulated such processes with general diffusion and friction tensors in a generally covariant manner.

In this chapter, we delve into both foundational and cosmological aspects of quantum-gravity-motivated stochastic motion. In the first part, we generalise the Lorentz-invariant diffusion of Dowker et al. [58] to curved spacetime from both the distribution and single-particle perspectives. The resulting stochastic process reproduces that defined by Franchi and Le Jan [68]. Our derivation, however, is based on Sorkin's formulation [59] of the coordinate transformation properties of the terms appearing in a general Fokker-Planck equation. To this end, we formulate the notion of general covariance and admissible coordinate transformations on the phase space of massive particles in curved spacetime. Requiring the Fokker-Planck equation to be covariant under admissible coordinate transformations, we derive a *covariant, frame-independent diffusion equation* in phase space. It can also be seen as the generalisation of the Boltzmann equation to incorporate diffusion on the mass shell. Our derivation is detailed in Section 1.1.1. It emphasizes the geometry of the phase space and gives physical intuition underlying the origin and the form of the diffusion and drift terms, making it accessible to a physics audience.

An alternative perspective on diffusion processes is through the stochastic motion of individual particles where the equations of motion are no longer deterministic but are governed by Langevin or stochastic differential equations (SDEs). In Section 1.1.2, we begin with a general stochastic correction to the geodesic equation. Guided by first principles, including general covariance, mass conservation, and the physical requirement that particles do not travel back in time, we derive two geometric conditions on the diffusion tensor and bias vector. Assuming no background structures that could introduce a preferred frame, and restricting to the zeroth order of the curvature expansion of the diffusion tensor, we derive a unique *stochastic geodesic equation*. In Appendix A.2, we show that these two equations—the covariant diffusion equation and the stochastic geodesic equation—are equivalent and describe the same underlying stochastic process. Note that the final equation has appeared in studies of Franchi and Le Jan's stochastic process [68, 82]. Nevertheless, our formulation explicitly clarifies the extent to which this process is unique. While Lorentz invariance ensures uniqueness in flat spacetime, Brownian motion could couple to curvature. Our constraint equations allow us to write down all such processes. Nevertheless, under minimal coupling, the equations retain their uniqueness, governed solely by one diffusion constant κ.

We highlight that although our motivation stems from considerations in causal set theory, the resulting covariant diffusion equation is independent of causal sets. Similar to the Brownian motion in fluids, the covariant Brownian motion is an effective theory that is agnostic to the UV-complete theory of quantum gravity. Due to its uniqueness, the minimally coupled covariant diffusion can describe any Lorentz-preserving stochastic correction to the motion of point particles, regardless of the underlying microscopic description—be it a covariant effect from loop quantum gravity, quantum foam, spacetime uncertainty, or spontaneous collapse models. This versatility makes it a powerful phenomenological model, serving as the effective description of a range of possible quantum gravity effects.

In the second part of the chapter, we examine the cosmological implications. For the covariant Brownian

motion to have an observable effect, the particles must experience the longest possible proper time. Therefore, we aim to apply the covariant diffusion equation to massive particles in cosmology. While in principle, all massive particles are subject to the covariant Brownian motion, each with its own diffusion constant, our focus is on dark matter particles, the most dominant matter species. This leads to a model of stochastic dark matter. Since the covariant Brownian motion spontaneously heats up the particles, we anticipate stochastic dark matter to have implications for the S_8 tension.

The covariant diffusion equation generalises the Boltzmann equation, modifying both the background and perturbation equations for the stochastic dark matter. In Section 1.2.1, we outline the use of the generalised Boltzmann hierarchy for numerically solving these equations. Intuitively, we expect the diffusion of dark matter particles to counteract clustering and reduce late-time structure growth at small scales. Using the semi-analytic solutions to the distribution function in Section 1.2.2, we confirm this intuition. We show that for certain values of the diffusion parameter, the effect remains negligible in the background, yet small-scale perturbations can be significantly damped. Consequently, the matter power spectrum is suppressed at small scales in a redshift-dependent manner, with stronger suppression at lower redshifts.

A question left unresolved in previous works is the source of energy-momentum non-conservation in swerves models and their effective Lorentz invariant diffusion equation. To address this, we take a pragmatic approach by introducing an imperfect dark energy fluid with heat flux to maintain consistency within Einstein equations, as detailed in Section 1.2.3.

Finally, in Section 1.2.5, we show that the specificity of the diffusion to small scales and low redshifts enables stochastic dark matter to satisfy CMB constraints as well as constraints on the expansion and growth histories from BAO, Ly-α, and uncalibrated SN Ia while reducing and broadening the S_8 posterior from CMB to make it consistent with weak lensing measurements. Interestingly, using an S_8 prior from the Dark Energy Survey (DES) introduces a 1.7σ preference for a non-zero diffusion rate. We end with a discussion on possible implications and future directions.

1.1 Foundation

1.1.1 From Fokker-Planck to covariant diffusion equation

As mentioned earlier, our motivation is the stochastic motion of point particles due to spacetime discreteness. However, all our treatment is an effective description in the continuum, independent of the microscopic details of any discrete theory of quantum gravity. Our key physical principle is spacetime coordinate independence, and local Lorentz invariance on the tangent space at each spacetime point. This is unlike many Lorentz-invariance-violating models in quantum gravity phenomenology [9]. In addition, as in swerves models [60], the quantum gravity description of the stochastic motion most probably involves a nonlocality scale τ_f. In our effective description, however, we assume sufficient coarse-graining, so that we are working at scales where the dynamics is Markovian or local. Our discussion parallels the arguments in [60] but is generalised to arbitrary curved spacetimes.

Drift pseudo-vector and diffusion tensor

Stochastic processes can be described in terms of the evolution law for their distribution in phase space. Let the coordinates on the phase space $\mathcal{S}$ be Z^A and $f_\tau(Z,\tau)$ be the distribution density of particles in this phase space, where τ is the time evolution parameter of f_τ, which in this work we take to be the proper time of the particles. The most general evolution of the density f_τ, assuming continuity in phase space and Markovianity, is governed by the Fokker-Planck equation

$$\partial_\tau f_\tau = -\partial_A\Big(v^A f_\tau\Big) + \partial_A\partial_B\Big(K^{AB} f_\tau\Big)\,, \tag{1.1}$$

where the first term on the right accounts for the transport of the particles according to some least-action principle, and the second term accounts for diffusion around their least-action trajectories. $v^A(Z)$ and $K^{AB}(Z)$ are one-index and two-index objects living in the phase space that have to be determined for any particular dynamics. It is a standard result of stochastic processes that $v^A(Z)$ and $K^{AB}(Z)$ are related to the mean and covariance of the motion of individual particles [83]:

$$\begin{aligned} v^A &= \lim_{\delta\tau\to 0}\frac{1}{\delta\tau}\Big\langle \delta Z^A\Big\rangle \\ K^{AB} &= \lim_{\delta\tau\to 0}\frac{1}{2\delta\tau}\Big\langle \delta Z^A \delta Z^B\Big\rangle. \end{aligned} \tag{1.2}$$

See an alternative derivation based on [59] in Appendix A.2.1.

The total particle-number conservation is implicitly assumed in (1.1) and follows from the conservation of probability for the forward Kolmogorov equation since, when normalised by the total number of particles, f_τ equivalently acts as a probability distribution on phase space. Taking $v^A = 0$ and the diffusion tensor to be the identity matrix, (1.1) is just the familiar heat equation. A derivation of the above Fokker-Planck equation based on first physical principles is outlined in Appendix A.1.

Looking at the equation (1.1), a critical question that comes into mind is what are the transformation rules for v^A and K^{AB} under coordinate transformations in phase space $Z^A \to Z'^M(Z)$. This can be answered by noting that ρ is a scalar density, as discussed in Appendix A.1. As noted by Graham [84] and Sorkin [59] this implies that v^A is a pseudo-vector. When transformed, it mixes with K^{AB}:

$$v'^M = \frac{\partial Z'^M}{\partial Z^A} v^A + \frac{\partial^2 Z'^M}{\partial Z^A \partial Z^B} K^{AB}\,. \tag{1.3}$$

On the other hand, surprisingly, K^{AB} turns out to be a tensor on the phase space:

$$K'^{MN} = \frac{\partial Z'^M}{\partial Z^A}\frac{\partial Z'^N}{\partial Z^B} K^{AB}\,. \tag{1.4}$$

The transformation properties of the drift and diffusion tensor provide a key guiding principle when searching for a covariant description of diffusion in spacetime. Indeed, the derivation reduces to identifying the phase space of massive particles on a curved spacetime and then finding the most general pseudo-vector v^A and tensor K^{AB} that live in this phase space.

Before tackling this problem, it is helpful to replace v^A with an actual vector on the phase space.

The extra term in (1.3) has similarities with the transformation rule of the partial derivatives of a metric tensor. Indeed, it turns out that if we assume $\hat{g}_{AB}$ to be a given metric on the phase space with determinant $\hat{g}$, then

$$u^A = v^A - \frac{1}{\sqrt{|\hat{g}|}}\partial_B\left(K^{AB}\sqrt{|\hat{g}|}\right) \tag{1.5}$$

is an actual vector on phase space. In other words, under general coordinate transformations of phase space, it transforms as

$$u'^M = \frac{\partial Z'^M}{\partial Z^A}u^A\,. \tag{1.6}$$

In [59], Sorkin uses the entropy density s instead of $\log\sqrt{|\hat{g}|}$ to define u. To find the most general vector u^A and tensor K^{AB}, we first specify the phase space for particles of mass m in curved spacetime. In classical physics, phase space consists of three spatial coordinates and three momentum coordinates, with time as the evolution parameter. Here, however, we want to preserve general covariance, so following [58, 60], we use the proper time of the particles τ as the evolution parameter and take all the spacetime coordinates x^μ of a spacetime $\mathcal{M}$ as part of the phase space. Momentum coordinates can be taken to be either vectors or 1-forms living on spacetime. For convenience, in this section, we follow [85], and we choose the latter option with lower indices. Therefore, the phase space can be parametrised by (x^μ, p_ν), where p_ν is a one-form at point x^μ. The set of all such points forms the cotangent bundle of the spacetime. However, we are only interested in the momenta that satisfy the on-shell condition; so we restrict ourselves to a 7-dimensional subspace of the cotangent bundle [85]:

$$\Gamma^+_m = \left\{(x,p)\in T^*\mathcal{M}\;\middle|\; g^{\mu\nu}(x)p_\mu p_\nu = -m^2,\ p^0>0\right\}. \tag{1.7}$$

This is our desired phase space $\mathcal{S}$.

Admissible coordinate transformations

To find vectors and tensors on this phase space, we need to first discuss coordinate transformations. The transformation rules (1.6) and (1.4) were written for arbitrary coordinate transformations of the phase space. However, a coordinate transformation on phase space that mixes the spacetime and momentum labels of the cotangent bundle arbitrarily is not physically meaningful. By construction, the above phase space has spacetime as its base manifold, separate from the momentum coordinates. We cannot use momentum coordinates to label spacetime points in a new coordinate system. Therefore, we do not require our equations to be form-invariant under every possible transformation.

First, we require an *admissible* coordinate transformation to map a spacetime point x^μ to $x'^\alpha(x)$, without any mixing with momentum coordinates. This induces a natural transformation on 1-forms via the push-forward on tangent spaces. On top of that, we can also allow for reparametrisations of the tangent (and cotangent) spaces of the spacetime. The only structure that has to be preserved is the linearity of the tangent space as a vector space. This means that we should include x-dependent Lorentz transformations of the tangent spaces. To formulate this idea, let us use another parametrisation of the phase space (1.7) in terms of *physical momentum* q_a rather than spacetime momentum p_μ. Remember that any physical measurement of the momentum of a particle is performed using an orthonormal tetrad field $e_a{}^\mu(x)$, $a=0,...,4$ that plays the role of a lab frame. The measured momentum is the projection

of the spacetime momentum p on this tetrad, given by $q_a = e_a{}^\mu(x)p_\mu = p.e_a$. After restricting to the mass shell, the physical momentum would live on the 3-dimensional hyperbolic space:

$$\mathbb{H}_3 = \left\{ q \;\middle|\; \eta^{ab} q_a q_b = -m^2, q^0 > 0 \right\}. \tag{1.8}$$

Note that $\mathbb{H}_3$ implicitly depends on m, but for the sake of simplicity, we do not include it in the notation. One can use (x^μ, q_a), $a = 0, ..., 3$ as a coordinate system on Γ_m^+, but then it should be remembered that one of the momentum variables is redundant due to the mass-shell constraint. Instead, we simply restrict to their spatial components and use (x^μ, q_i), $i = 1, 2, 3$, as the coordinates of the phase space.

Now, a local (x-dependent) Lorentz transformation can be simply stated as a change of the orthonormal tetrad field $e_a \to e'_r$, $r = 0, ..., 3$. The new momentum coordinates q'_r would be related to the old ones by a Lorentz transformation at each spacetime point. And once again, we restrict to the spatial components of the momentum q_I, $I = 1, 2, 3$. In summary, the most general admissible coordinate transformation of the phase space is

$$\left(x^\mu, q_i\right) \to \left(x'^\alpha(x),\ q'_I(x, q) = e'_I(x).e^a(x)\ q_a \right). \tag{1.9}$$

It is easy to check that these transformations form a subgroup of the diffeomorphism group of the phase space.

Restricting u^A and K^{AB}

In the spirit of general relativity, we require u^A and K^{AB} to transform like vectors and tensors, respectively, and to be form-invariant only under the set of admissible coordinate transformations. Consider first the diffusive part. Equation (1.2) can be rewritten as

$$K^{\mu B} = \lim_{\delta\tau \to 0} \frac{1}{2\delta\tau} \left\langle \delta x^\mu \delta Z^B \right\rangle. \tag{1.10}$$

Here B can be any of the spacetime or momentum indices. Now dx^μ is of order $d\tau$, and in the limit $d\tau \to 0$, the differential increment dZ^B goes to zero. Therefore $K^{\mu B} = 0$. This is essentially the same argument as in [60]. One might wonder why the same argument does not conclude that any K^{AB} has to be zero. The reason is that there are noise terms present in dq_i that are not of the order $d\tau$. In other words, we are assuming that the stochastic noise terms do not directly act on dx^μ. This assumption will be justified in the next section by demanding the particles not travel back in time.

The only non-zero components of K^{AB} are therefore K^{ij}, where $i, j = 1, 2, 3$ are the momentum indices. For these components, (1.4) reduces to [1]

$$K'^{IJ} = \frac{\partial q'^I}{\partial q^i} \frac{\partial q'^J}{\partial q^j} K^{ij}. \tag{1.11}$$

This is just the statement that K^{ij} should be a tensor on $\mathbb{H}_3$ under Lorentz transformations. As argued in [60], one obvious choice is the hyperbolic inverse metric on $\mathbb{H}_3$, h^{ij}, that is induced from the 4-dimensional Minkowski space where $\mathbb{H}_3$ is embedded. $\mathbb{H}_3$ is a maximally symmetric Riemannian

[1] Note that since q_a lives in the Minkowski spacetime, $q_i = q^i$ and $-q_0 = q^0 = E$.

manifold, so all other naturally defined geometric 2-tensors on it, like the Ricci or Einstein tensors, are proportional to its metric tensor. Therefore, K_{ij} has to be proportional to h_{ij}. More explicitly, for our choice of parametrisation, the metric tensor and its inverse are

$$\mathsf{h}_{ij} = \delta_{ij} - \frac{q_i q_j}{E^2}\,, \tag{1.12}$$

$$\mathsf{h}^{ij} = \delta^{ij} + \frac{q^i q^j}{m^2}\,, \tag{1.13}$$

and we have $K^{ij} = \kappa \mathsf{h}^{ij}$, where κ is a free parameter independent of momentum. So far, we have concluded that the diffusion tensor has to have the following form:

$$K = \begin{pmatrix} 0 & 0 \\ 0 & K^{ij} = \kappa\Big(\delta^{ij} + \frac{1}{m^2} q^i q^j\Big) \end{pmatrix}\,. \tag{1.14}$$

For the true vector u^A, the choice is even more constrained, leaving no room for any additional free parameters. First, we look at the pseudo-vector v^A. Equation (1.2) gives

$$v^\mu = \lim_{\delta\tau\to 0} \frac{1}{\delta\tau}\Big\langle \delta x^\mu \Big\rangle\,. \tag{1.15}$$

By definition, therefore, $v^\mu = p^\mu/m$. Being there no contribution from the diffusion tensor, $K^{\mu B} = 0$, one finds $u^\mu = p^\mu/m$ for the spacetime part of u. This is also how Dowker et al. argued in [60]. They continued by arguing that global Lorentz invariance enforces u^i, the momentum part of u, to be identically zero. This is what limits their derivation to only flat spacetime. Here, we shall see that u^i is non-zero for curved spacetime.

As mentioned earlier, the vector u encodes the dynamics of massive particles in the absence of diffusion processes, according to some least-action principle. For free particles, the action is proportional to the total proper time, and the dynamics is given by the geodesic equation. Alternatively, one can use the Polyakov version of the action of a point particle, which has a corresponding Hamiltonian that after subtraction of a constant term, is $\frac{1}{2m} p_\mu p_\nu g^{\mu\nu}$. The corresponding Hamiltonian vector field is nothing but L/m, where L is the well-known Liouville operator:

$$L = p^\mu \frac{\partial}{\partial x^\mu} + \Gamma^\sigma_{\mu\nu} p_\sigma p^\mu \frac{\partial}{\partial p_\nu}\,. \tag{1.16}$$

Note that this is written in (x^μ, p_ν) parametrisation of the cotangent bundle; see Appendix A.1.2 for its representation in (x^μ, q_i) coordinates of phase space. The integral curves of L/m in the phase space are the trajectories of massive particles moving on the geodesics of spacetime. This is suggesting us to set

$$u = \frac{1}{m} L\,. \tag{1.17}$$

Note that spacetime components of both sides already agree, each being equal to p^μ/m. Not surprisingly, one can prove that L is indeed a vector under the transformations (1.9), as detailed in Appendix A.1.2. But is L/m the only way to extend $\frac{1}{m} p^\mu \partial/\partial x^\mu$ to a true vector on the phase space? Imagine there is a second vector u' on the phase space. Then $u'' = u - u'$ is also a vector field but with spacetime

components being zero. Similar to K^{ij}, now $u''^{\,i}$ must be a Lorentz invariant vector field on $\mathbb{H}_3$. If non-zero, each such vector would introduce preferred local directions in $\mathbb{H}_3$, thereby breaking its maximal symmetry. This is only possible in the presence of spacetime curvature, such that $u''^{\,i}$ would depend on curvature tensors. Our focus, however, is on the *minimal coupling* of the diffusion to the curvature, forcing $u''^{\,i}$ to be identically zero. We will provide an example for non-minimal coupling after we introduce the Langevin equations.

The covariant diffusion equation

We have found the explicit covariant expressions of the terms in the Fokker-Planck equation (1.1). All we have to do is to substitute u and K from (1.17) and (1.14). So far, the proper time τ has served as the evolution parameter, so that f_τ is a function of (τ, x^μ, p_ν). In order to better compare with the usual Boltzmann equation and following [60], we integrate out the unobservable proper time and use a distribution function that only depends on (x^μ, p_ν). This is done via

$$\sqrt{|\hat{g}|} f(x,p) = \int_{-\infty}^{\infty} d\tau \; f_\tau \,. \tag{1.18}$$

Recall that $\hat{g}$ is the determinant of the metric of the phase space. Note that this definition makes f a scalar on the phase space. After adopting the Sasaki metric as the natural choice for $\hat{g}_{AB}$ and performing a few steps on the Fokker-Planck equation, including the use of Liouville's theorem, one derives the following equation as detailed in Appendix A.1.1:

$$\boxed{L[f] = p^\mu \frac{\partial}{\partial x^\mu} f + \Gamma^\sigma_{\mu\nu} p_\sigma p^\mu \frac{\partial}{\partial p_\nu} f = m\kappa \nabla^2_{\mathbb{H}_3} f} \,. \tag{1.19}$$

This is the *covariant diffusion equation* in curved spacetime. It is the unique generally covariant extension to the Boltzmann equation for free, massive, spin 0 particles. It describes the distribution of Brownian particles that experience a frame-independent diffusion effect in curved spacetime. $\nabla^2_{\mathbb{H}_3}$ is the Laplacian operator on the hyperbolic space $\mathbb{H}_3$. κ can be a spacetime scalar in its most general form, but at the zeroth order in a curvature expansion, it is a constant. It has units of M^2/L, therefore being a diffusion constant on the mass shell.

Dudley's distribution function f in [67] differs from ours by a factor of m^3. Apart from re-scaling the parameters, his diffusion equation is the special case of the covariant diffusion equation (1.19) on flat spacetime and for a spatially homogeneous distribution f. Also, the cosmic-time Lorentz invariant diffusion equation in [60], is the special case of (1.19) in flat spacetime once we identify their $\rho_t/(\gamma\sqrt{g})$ and k with our f and κ, respectively.

A reader familiar with the Klein-Kramers equation might see some resemblance with the covariant diffusion equation. The Klein-Kramers equation describes the diffusion of particles in a classical non-relativistic phase space. However, it cannot be derived from the covariant diffusion equation in a non-relativistic limit. The Klein-Kramers has two free parameters—friction and diffusion constants—that can be related to each other via the temperature of the background medium in Einstein relation. There exists the distinguished rest frame of the background fluid. As a result, in the corresponding Langevin equation [86] that describes an Ornstein–Uhlenbeck process, there is a drag term that prevents Brownian

particles from attaining large velocities in that frame, even in relativistic generalisations [77, 78]. The covariant diffusion equation, however, has only a single additional free parameter. There is no preferred frame, and hence the Brownian particles do not feel any drag force. As a matter of fact, the effect acts in the opposite direction: The drift term causes the particles to run away from their rest frame at any given moment. To see this more clearly, we need to find the stochastic dynamics of individual particles. This is the subject of the next section.

1.1.2 Stochastic geodesic equation: Langevin equations

Any diffusion process at the level of the distribution of particles uniquely corresponds to a set of Langevin or stochastic differential equations that describe the Brownian motion of individual particles. In this section, starting from first principles, we derive the most general conditions a stochastic extension to the timelike geodesic equation must satisfy. In the end, we specialize to the case of minimal coupling to curvature and derive the unique covariant Brownian motion in curved spacetime. Similar equations have also been studied in the context of relativistic statistical mechanics, [80, 81], with the difference that the systems of interest were in thermal equilibrium and a form of relativistic Einstein relation was therefore imposed.

For a massive particle, the least-action trajectory in curved spacetime is given by:

$$dx^\mu = \frac{p^\mu}{m} d\tau \,, \tag{1.20}$$

$$dp^\mu = -\frac{1}{m}\Gamma^\mu_{\alpha\beta} p^\alpha p^\beta d\tau \,, \tag{1.21}$$

which is completely equivalent to the geodesic equation in a spacetime with metric $g_{\mu\nu}$ and associated Christoffel symbols $\Gamma^\mu_{\alpha\beta}$. This is supplied by the mass-shell constraint:

$$g_{\mu\nu} p^\mu p^\nu = -m^2 \,, \tag{1.22}$$

which is trivially conserved by the geodesic motion. Introducing stochasticity into the equations of motion involves adding a random noise term to at least one of the equations. Under the condition of having almost surely continuous stochastic trajectories in spacetime and Markovian (i.e. memoryless) stochastic dynamics, the stochastic differential increment dW^μ is a Gaussian random variable with statistics:

$$\left\langle dW^\mu \right\rangle = 0 \,, \qquad \left\langle dW^\mu dW^\nu \right\rangle = 2D^{\mu\nu} d\tau \,, \tag{1.23}$$

with $D^{\mu\nu}$ being the positive semi-definite covariance matrix of the noise. The mean can in general be non-zero, so in what follows, we explicitly subtract the mean from the noise to have the above statistics. Note that continuity and Markovianity are the effects of coarse-graining. The specific microscopic model might have some discrete jumps and forgetting time τ_f. But when considering appropriate scales, the mentioned assumptions will be valid [60]. As we review in detail in Appendix A.2.1, dW^μ is of order $\mathcal{O}(\sqrt{d\tau})$, a statement made precise by Itô's lemma.

The form of stochastic modification of the equations of motion is strongly constrained by the requirement of covariance and preservation of the particle on the mass shell. First, let us consider (1.20). Adding noise to the evolution equation for x^μ inevitably leads to acausal trajectories in spacetime. To see this,

consider the modified set of equations:

$$dx^\mu = \frac{p^\mu}{m} d\tau + dW^\mu \,, \tag{1.24}$$

$$dp^\mu = -\frac{1}{m} \Gamma^\mu_{\alpha\beta} p^\alpha p^\beta d\tau \,. \tag{1.25}$$

If we perform a boost into a frame where the particle is at rest, then $p^0 = m$, meaning that the equation of motion for the time coordinate reads:

$$dx^0 = d\tau + dW^0 \,. \tag{1.26}$$

No matter the size of the diffusion constant κ, which effectively controls the variance in dW^μ, the particle will eventually experience $dx^0 < 0$ with probability 1, meaning that it will travel back in time. Further discussions on such stochastic processes can be found in [87, 88], where it is shown that such causality violations are suppressed for time scales larger than the particle's Compton wavelength.

To avoid causality violations entirely, the only option left is to add the stochastic term to (1.21):

$$dx^\mu = \frac{p^\mu}{m} d\tau \,, \tag{1.27}$$

$$Dp^\mu = b^\mu d\tau + dW^\mu \,. \tag{1.28}$$

Here, Dp^μ is the covariant differential, which involves the Christoffel symbols. We have explicitly separated the mean of the noise as a drift 4-vector b^μ. The physical origin of this extra drift seems mysterious at first sight. However, note that as the stochasticity is removed for $\kappa \to 0$, both the drift and the noise term should vanish due to dW^μ having zero mean and variance, and we should recover the usual geodesic equation. The necessity of adding the drift term becomes, nonetheless, evident when considering the conservation of mass shell. Using Itô's lemma (i.e. recalling that $dW^\mu dW_\mu$ is of order $d\tau$):

$$\begin{aligned} -d(m^2) = d(p^\mu p_\mu) &= 2p_\mu Dp^\mu + dp^\mu dp_\mu \\ &= 2(p_\mu b^\mu + D^\mu{}_\mu) d\tau + 2p_\mu dW^\mu \,, \end{aligned} \tag{1.29}$$

which has to be zero for the stochastic term not to drive the system off the mass shell. In order for (1.29) to vanish, however, both the $\mathcal{O}(d\tau)$ and the $\mathcal{O}(dW) \sim \mathcal{O}(d\tau^{1/2})$ terms need to vanish independently. Looking at the stochastic term first, this implies:

$$p_\mu dW^\mu = 0 \,, \tag{1.30}$$

meaning that the noise term is orthogonal to the momentum 4-vector, i.e. tangent to the mass shell. We should also make sure that the second order contribution $dW^\mu dW_\mu$ allowed by Itô's lemma preserves the mass shell condition. This means that also the other term present in (1.29) has to be zero on the mass shell:

$$\boxed{p_\mu b^\mu + D^\mu{}_\mu = 0} \,. \tag{1.31}$$

Finally, note that the orthogonality condition (1.30) reduces to a condition on the covariance matrix by

requiring the correlations between $p_\mu dW^\mu$ and dW^ν to vanish on the mass shell:

$$\boxed{p_\mu D^{\mu\nu} = 0}\,. \tag{1.32}$$

In short, equations (1.31) and (1.32) are the most general conditions that the drift and covariance matrix have to satisfy on the mass shell for the stochastic differential equation (1.28) to preserve the mass shell condition. For the equations to be generally covariant, we only need b^μ and $D^{\mu\nu}$ to be a spacetime vector and tensor, respectively. And finally, to have a proper covariance matrix, $D^{\mu\nu}$ must be positive semi-definite.

Given these constraints, what is the most general form of b^μ and $D^{\mu\nu}$? In the absence of any other fields or additional geometric structures, the existing covariant objects that we have are the momentum p^μ, the metric, the covariant derivative, and all of the curvature tensors. We can combine them in any way to produce a vector and a tensor, and as long as we satisfy (1.31) and (1.32), we have an admissible drift and covariance matrix. As highlighted earlier, for the present analysis, we are only interested in the zeroth order in a curvature expansion. A Brownian motion caused by quantum gravity is, in principle, also present in flat spacetime. In curved spacetime, many curvature terms become viable that introduce a non-minimal coupling of the Brownian motion to the curvature. We say more about this possibility in the next subsection. Otherwise, for the rest of the chapter, we advocate for the case of minimally-coupled Brownian motion.

With this assumption, the only covariant objects available are p^μ, $g^{\mu\nu}$, and $p^\mu p^\nu$. Therefore, the drift vector has to be proportional to p^μ, and the covariance matrix has to be a linear combination of the metric and $p^\mu p^\nu$. After plugging in (1.31) and (1.32), the three constants can all be written in terms of one of them, which we call κ. The result is

$$\begin{aligned} b^\mu &= \frac{3\kappa}{m^2} p^\mu \,, \\ D^{\mu\nu} &= \kappa\Big(g^{\mu\nu} + \frac{1}{m^2} p^\mu p^\nu\Big) \,. \end{aligned} \tag{1.33}$$

We have derived that the most general form of the *covariant Brownian motion* of massive particles minimally coupled to gravity and constrained to remain on the mass shell:

$$\boxed{\begin{aligned} dx^\mu &= \frac{p^\mu}{m} d\tau \,, \\ dp^\mu + \frac{1}{m}\Gamma^\mu_{\alpha\beta} p^\alpha p^\beta d\tau &= \frac{3\kappa}{m^2} p^\mu d\tau + dW^\mu \,, \qquad (1.34) \\ \Big\langle dW^\mu \Big\rangle = 0 \,, \quad \Big\langle dW^\mu dW^\nu \Big\rangle &= 2\kappa\Big(g^{\mu\nu} + \frac{1}{m^2} p^\mu p^\nu\Big) d\tau \,. \qquad (1.35) \end{aligned}}$$

It can be trivially extended to spacetimes of dimension other than 4, with the coefficient 3 in the drift term changing to $d-1$. Five different realisations of these equations in 3-dimensional flat spacetime are given in Figure 1.1. The particles perform a Brownian motion on the mass shell, and as a result of the change in their momentum, the spacetime trajectories deviate from a straight line.

It turns out that the parameter κ is the same as the diffusion parameter we encountered in the previous section, and that these stochastic differential equations are completely equivalent to the covariant

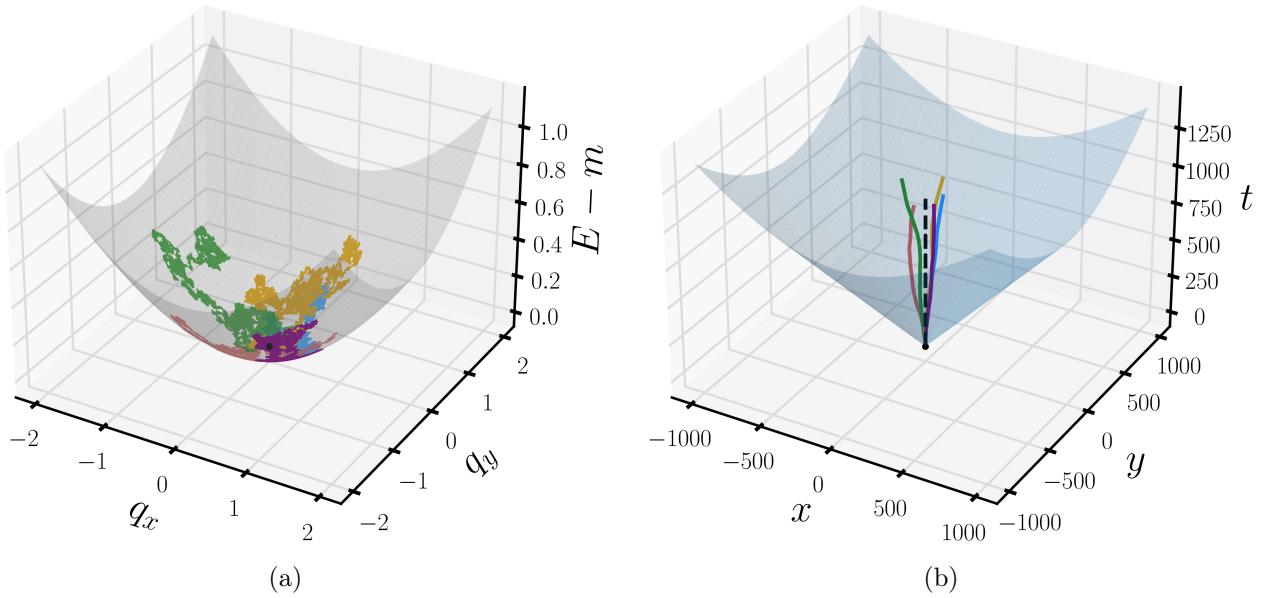


Figure 1.1: Five realisations of the covariant Brownian motion (1.34) in 3d flat spacetime, with $m = 1$, $\kappa = 10^{-3}$, and total proper time $\tau = 10^3$. Figure 1.1a shows the trajectories of the particles in the momentum space and on the mass shell. The particles are undergoing a random walk, and so the trajectories are not differentiable. Figure 1.1b shows the resulting trajectories in spacetime, where the shaded surface is the light cone. The particles have differentiable worldlines that deviate from geodesics.

diffusion equation. Although we used seemingly different physical principles to arrive at the respective equations, there are parallels between the two arguments. First, with the identification

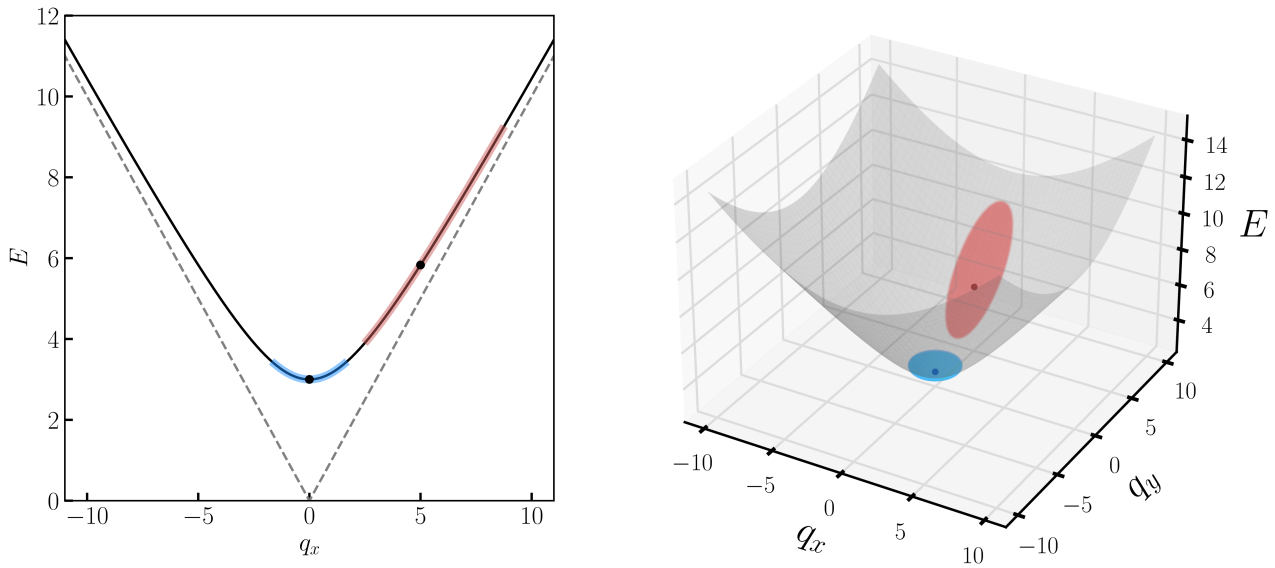
$$K^{ij} = e^i{}_\mu e^j{}_\nu D^{\mu\nu} \,, \tag{1.36}$$

one can prove that the tensor transformation condition of K on the phase space is equivalent to $D^{\mu\nu}$ being orthogonal to p on the mass shell, equation (1.32). Moreover, with the identification

$$v^i = e^i{}_\mu b^\mu \,, \tag{1.37}$$

the vector transformation property of $u = (0, u^i)$ on the phase space is equivalent to the condition (1.31). In Appendix A.2.3, we prove these two statements and show that the covariant diffusion equation and the covariant Brownian motion are equivalent.

As the final note, we highlight that since $D^{\mu\nu}$ has to be positive semi-definite as a covariance matrix, κ must be positive. Therefore, the drift term in (1.34) acts like a force in the direction of the particle's momentum. We have seen that the existence of this drift term is necessary for covariance and the conservation of mass. Here we give a more intuitive account of that. Consider the hyperboloid of mass m, (1.8). As shown in Figure 1.2, in the rest frame of a given particle, the unit sphere around the origin $q^i = 0$ is perfectly symmetric in all directions of the spatial momentum. When a particle starts a random walk from the origin, it has equal chances of going to any spatial momentum direction. However, for a particle with non-zero spatial momentum $|\vec{q}|^2 = q^i q^i$, the unit sphere is skewed toward larger $|\vec{q}|^2$. So when doing a random walk on the mass shell, the particle has a bias toward larger energies. This is the geometric origin of the drift term.



(a) (b)

Figure 1.2: Balls of radius 1/2 in hyperbolic space. Figure 1.2a shows the ball of radius 1/2 centred on $q_x = 0$ in blue and in red the region centred on $q_x = 5$. The centre of the respective regions is represented by a black dot. Note that the latter is asymmetric around its centre, and in particular it extends to far higher momenta than naively expected. The same situation is shown in Figure 1.2b in 2+1 dimension, with the blue region centred on $q_i = 0$ and the red one on $(q_x = 0, q_y = 5)$.

Brownian motion coupled to curvature: an example

Before moving on to cosmology, for completeness, we give an example of a stochastic process coupled to curvature, where the diffusion tensor $D^{\mu\nu}$ is explicitly constructed from curvature tensors. This further demonstrates the strength of the formulation we have developed.

The diffusion tensor derived in the previous section is in fact a projection tensor on the mass shell. It can be used to construct diffusion tensors of higher order in curvature that satisfy the constraint (1.32). So let us define

$$\mathcal{P}^{\mu\nu} = g^{\mu\nu} + \frac{1}{m^2} p^\mu p^\nu . \tag{1.38}$$

The following expression is a valid positive semi-definite diffusion tensor:

$$D^{\mu\nu} = \kappa' \, \mathcal{P}^{\mu\alpha} R_{;\alpha} \mathcal{P}^{\nu\beta} R_{;\beta} \, , \tag{1.39}$$

where R is the Ricci scalar, and κ' is another diffusion parameter. To satisfy the constraint (1.31), it is sufficient to choose $b^\mu = \frac{1}{m^2} p^\mu D^\nu{}_\nu$.

In flat spacetime, the above stochastic process gives no correction to the geodesic motion. Nevertheless, from an effective point of view, all such curvature corrections could in principle appear. It would have been difficult to construct this stochastic process starting directly from the vector and tensor conditions on u and K. The Langevin equation approach we presented provides a systematic framework for constructing all possible covariant Brownian motions coupled to the curvature. It is an interesting problem for future work to investigate the asymptotic behaviour and possible equilibrium solutions of

the above process in FRW spacetime.

1.1.3 Swerves models

Here, we briefly review the swerves models in the context of causal set theory as a possible microscopic origin for the covariant Brownian motion. We emphasize that the covariant Brownian motion, as an effective theory, stands on its own and is not tied to any particular quantum gravity framework. Nevertheless, swerves models provided an important motivation for Dowker et al. [58, 60] for formulating the Lorentz invariant diffusion equation. For us, they provide a toy model with which we can compare the cosmological constraints we obtain in the next section. We will provide another brief overview of causal sets in Chapter 3 with an emphasis on the tools relevant for that context.

Causal set theory provides an approach to quantum gravity where discreteness and manifest Lorentz invariance are central [89–95]. Starting with a given spacetime manifold, one selects a random number of spacetime points with a fixed density ϱ according to the Poisson distribution. This provides a discretization of spacetime. Unlike regular lattices that explicitly distinguish a preferred frame, a random process that samples points based solely on the invariant volume of a spacetime region almost surely does not introduce a preferred frame, and hence preserves Lorentz invariance [96]. Now the discrete set of points has to be complemented with some fundamental degrees of freedom in order to recover the geometry and topology of the spacetime manifold. One option is to encode geometric data like the length of edges between the set of points in the spirit of Regge calculus [97, 98] or causal dynamical triangulation [99, 100]. Instead, in causal set theory, the fundamental degrees of freedom are the causal relations between the discrete elements. This is inspired by a series of papers in the 60's and 70's [101–104] showing that the causal structure of a Lorentzian manifold is enough to recover its topology and metric up to a conformal factor. Therefore, mathematically, causal sets are defined as finite partial orders, with the partial order representing causal ordering.

The dynamics of matter degrees of freedom on causal sets is modelled either using a discrete analogue of the d'Alembertian operator [19, 105, 106] or using the formalism of quantum field theory (QFT) on causal sets [8, 107–110]. Since the developments in these two directions have mostly been formal, it is still early to extract clear phenomenological signatures from these theories. Meanwhile, Dowker et al. [58, 60] constructed simple toy models of the motion of an idealistic point particle on a classical causal set. These models realize the heuristic idea that when moving on a discrete background, particles cannot maintain perfect geodesic motion and experience stochastic fluctuations around the longest chain of events–the discrete analogue of geodesics [111–113].

The simplest swerves model makes use of the embedding information of a causal set in an approximating Minkowski spacetime [58]. In particular, assume that a particle has previously been on causal set elements $(e_1, ..., e_n)$. Its next spacetime point e_{n+1} is chosen such that it satisfies three conditions:

- e_{n+1} is in the future light cone of e_n.
- e_{n+1} is at most within a proper time of τ_f away from e_n. τ_f is called the forgetting time.
- In a frame defined by the momentum 4-vector p_n, the spatial part of $p_{n+1} - p_n$ has the minimum possible magnitude.

Each momentum p_n is defined to be proportional to $e_n - e_{n-1}$ while satisfying the mass-shell condition.

The above model only works on flat spacetime and makes use of continuum information. In [60], Dowker et al. suggested two intrinsic models that only make use of the combinatorial properties of a causal set. The first model chooses the point e_{n+1} such that

- $d(e_{n-1}, e_{n+1}) \leq 2n_f$ and $d(e_n, e_{n+1}) \leq n_f$.
- $d(e_n, e_{n+1})$ is maximized.

Here, $d(a, b)$ denotes the length of the longest chain of events between two causal set elements a and b. n_f is the forgetting number, the discrete analogue of forgetting time. The other intrinsic model shares a similar idea and can be found in [60, 114].

The notion of a geodesic on a causal set as the longest chain of events between two events is highly nonlocal. All these models try to reduce the scale of this nonlocality to τ_f. By doing so, they introduce effective stochasticity in the motion of the particle with a variance that grows with propagation. Philpott showed that the intrinsic models have a faster diffusion given the same forgetting scale. She also simulated the first of the swerves models in two-dimensional Minkowski spacetime. Comparing with the solution to the Lorentz-invariant diffusion equation, she found good agreement with the outcome distribution of the swerves model, confirming that as the swerves models are Lorentz-invariant, they indeed give a microscopic description of the Lorentz-invariant diffusion process. Comparing the resulting distributions with each other and using dimensional analysis, she then related the diffusion constant to the forgetting time. In two dimensions, this is given by

$$\kappa \approx 2m^2 \frac{t_p^4}{\tau_f^5}\,, \tag{1.40}$$

where t_p is the Planck time (or the discreteness time scale to be more precise).

We note that the existence of a nonlocality scale τ_f in the swerves models implies that they are not Markovian processes. It is only on scales large enough compared to the forgetting time that we can approximate the behaviour of swerves models with the covariant diffusion equation. In [115], Dunkel et al. gave a proposal for a relativistic non-Markovian stochastic process. It is an interesting problem for future work to see whether such models provide a better description of swerves.

1.2 Cosmological tests

In [58], Dowker et al. noted that a heating rate as small as a millionth of a degree per second in cold hydrogen gas would have been experimentally detectable in laboratory, leading to the constraint $\kappa \lesssim 10^{-56}\mathrm{kg}^2\mathrm{m}^2\mathrm{s}^{-3}$ for baryons. Later, as noted by Kaloper and Mattingly [69], one might expect much tighter constraints from astrophysical observations of cold molecular clouds in the Milky Way. Their temperatures of about $10 - 20$K, combined with ages comparable to that of the Milky Way itself, would naively improve the previous bound on κ by seven orders of magnitude. However, if baryons undergo covariant Brownian motion and gain energy, part of that energy will subsequently be lost through thermal radiation. This effect weakens the constraint from cold molecular clouds to

$\kappa \lesssim 10^{-53}\mathrm{kg}^2\mathrm{m}^2\mathrm{s}^{-3}$. The same authors also noted that the measured heating rate of Bose-Einstein condensates at microkelvin temperatures gives a similar bound $\kappa \lesssim 10^{-54}\mathrm{kg}^2\mathrm{m}^2\mathrm{s}^{-3}$.

On the other hand, dark matter particles—having very weak or no interaction with standard model particles—are better probes for covariant Brownian motion. Their coldness determines how much structure is formed in the late universe. So we expect tight constraints on the diffusion parameter from cosmological observations.

When applied to dark matter particles, covariant Brownian motion gives rise to a model of stochastic dark matter. Beyond this, we remain agnostic about the particle nature of dark matter, such as its mass or field-theoretic description. Its primary governing equation is the covariant diffusion equation (1.19), written in an FRW universe with metric

$$ds^2 = a(\eta)^2\Big(-\big(1+2\psi\big)d\eta^2 + \big(1-2\phi\big)dx^i dx^i\Big)\,, \tag{1.41}$$

where a is the scale factor, η is the conformal time, and ϕ and ψ are gravitational potentials in Newtonian gauge. To connect the covariant diffusion equation with the background and perturbation equations for stochastic dark matter, we integrate out the momentum variables. The energy-momentum tensor of dark matter is related to its phase space distribution via

$$T_{\mu\nu} = \int \frac{d^3q}{E}\, p_\mu p_\nu f\,. \tag{1.42}$$

Note that the integration measure is the well-known Lorentz-invariant volume element on the mass shell in terms of the physical momentum q^i, while p_μ in the integrand is the spacetime momentum [85]. The background and perturbed energy density are governed by the following equations

$$\bar\rho'_{\rm dm} + 3\mathcal{H}\big(\bar\rho_{\rm dm} + \bar P_{\rm dm}\big) = a\Gamma m\bar n_{\rm dm}\,, \tag{1.43}$$

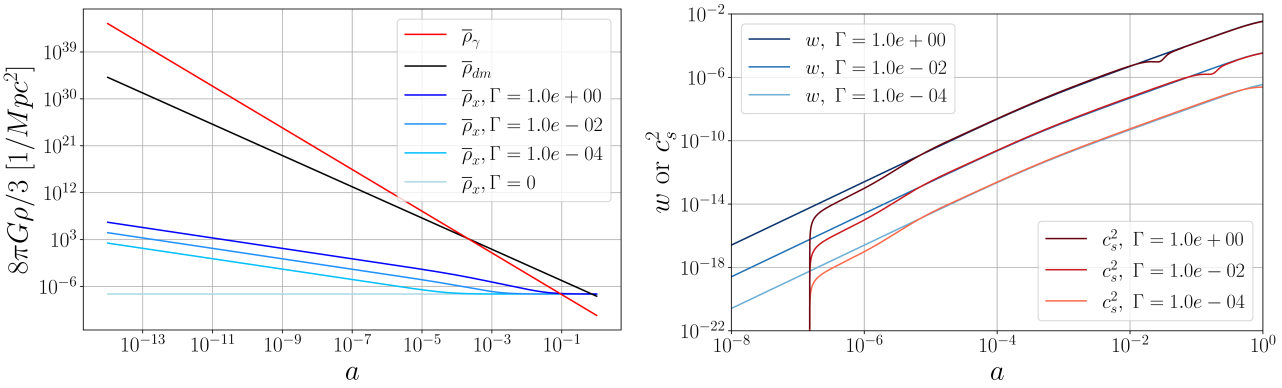
$$\delta\rho' + 3\mathcal{H}(\delta\rho+\delta P) + \big(\bar\rho+\bar P\big)\big(\theta - 3\phi'\big) = a\Gamma\Big(m\delta n + m\bar n\psi\Big)\,, \tag{1.44}$$

where m is the mass of dark matter, the prime denotes a derivative with respect to η, the barred variables are at background, and $P_{\rm dm}$ and $n_{\rm dm}$ are pressure and number density, respectively. θ is the divergence of the velocity field of dark matter. For simplicity, we drop the label dm when there is no risk of confusion. Covariant diffusion has introduced an energy injection term into the continuity equations proportional to the number density. Also note that here we have introduced another constant, Γ, which has the units of 1/time. It is related to the diffusion constant κ and the mass of dark matter particles via

$$\Gamma = \frac{3\kappa}{m^2}\,. \tag{1.45}$$

One can think of Γ as the *diffusion rate* of dark matter particles. As far as background equations like (1.43) are concerned, it has to be compared to H_0. If $\Gamma \ll H_0$, the effect of covariant Brownian motion on the background energy density would be negligible, and stochastic dark matter particles always remain non-relativistic. As we will see later, this condition is indeed necessary for sufficient structure formation.

When solving (1.43), things would simplify greatly if we could set the dark matter pressure density



(a) (b)

Figure 1.3: Figure 1.3a shows the log-log plot of the background energy density of photons, dark matter, and dark energy for various choices of Γ versus scale factor. The units of Γ are km/s/Mpc. One can see that even for values of Γ as large as 1km/s/Mpc, the deviation of dark energy density from a constant Λ is substantial only when its contribution to the total energy density is negligible. Figure 1.3b shows the evolution of the equation of state and sound speed squared of stochastic dark matter at $k = 1/\text{Mpc}$. w_{dm} follows a power law behaviour at all times according to (1.51). The numerical integration of the perturbation equations has started later than the background with $c_s^2 = 0$ as the initial condition. c_s^2 quickly catches up with its superhorizon attractor solution (A.86).

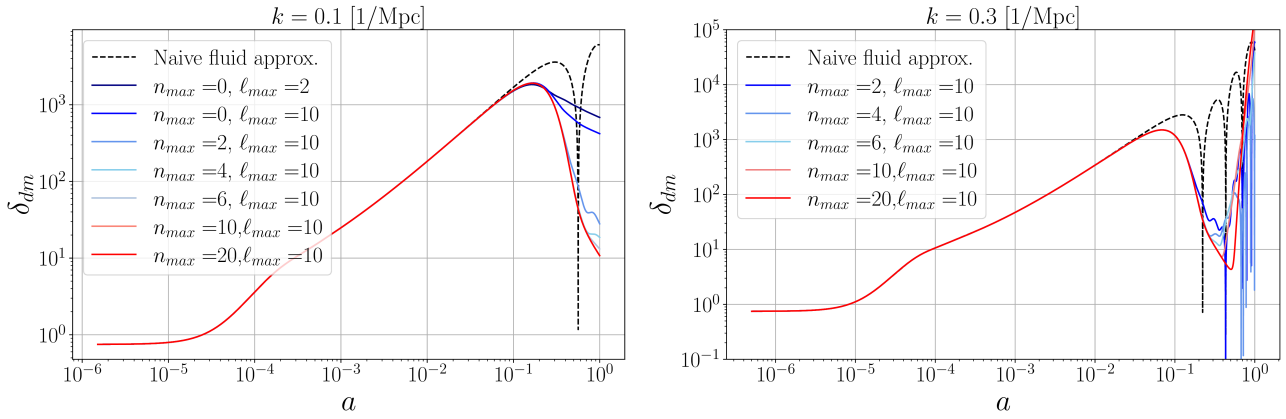
to zero, as for cold dark matter. This is related to whether we can identify $m\bar{n}_{\text{dm}}$ with $\bar{\rho}_{\text{dm}}$ to have a closed system of equations without referencing the number density or mass of dark matter particles. However, we expect heating to be important in general. Indeed, a non-relativistic expansion of E in powers of q/m shows that

$$m\bar{n} = \bar{\rho} - \frac{3}{2}\bar{P}\,, \tag{1.46}$$

The number density is conserved for stochastic dark matter, while the energy density is not. Therefore, pressure density cannot be conserved. Even starting with a perfectly cold distribution, covariant Brownian motion causes dark matter particles to heat up over time, leading to a growing equation of state $w = \bar{P}/\bar{\rho}$ and sound speed $c_s^2 = \delta P/\delta\rho$ as is evident in the numerical solutions presented in Figure 1.3b.

Another possible simplification would be to use a fluid approximation to model stochastic dark matter. While tempting, this approach—where higher moments of velocity are ignored when integrating out the momentum variables—leads to spurious growing acoustic oscillations in the dark matter overdensity $\delta = \delta\rho/\bar{\rho}$, see the dashed curve in Figure 1.4a. In the absence of collisions between stochastic dark matter particles, we do not expect any acoustic oscillations. In essence, the fluid quantities alone cannot capture the irreversibility inherent in Brownian motion.

One way forward is to numerically solve the covariant diffusion equation using Monte Carlo methods [116, 117]. At this stage, however, implementing this in a Boltzmann code remains challenging. Instead, the Boltzmann hierarchy is often employed to solve the perturbation equations [118]. However, the traditional hierarchy is not well-suited to stochastic dark matter. It depends explicitly on the particle mass m, about which we make no assumptions. Also in Boltzmann codes, the background distribution function is assumed to be of the form $f_0(\eta, q) = f_0(aq)$ indicating that particle momentum merely redshifts. In contrast, as it will be discussed in Section 1.2.2, the background distribution function for stochastic dark matter behaves like a heat kernel in the non-relativistic approximation, leading not



(a) (b)

Figure 1.4: Evolution of the stochastic dark matter overdensity versus scale factor for $\Gamma = 0.1$km/s/Mpc and two different wavenumbers k, normalised by the initial photons overdensity. A naive fluid approximation that only keeps the fluid quantities when integrating the momentum variables gives spurious growing oscillations. It is only after taking enough angular and velocity moments of the generalised Boltzmann hierarchy that the solution converges to a decaying power law. Even then, the system of equations becomes unstable at large ky, when the free streaming length of the particles becomes large. This is depicted in Figure 1.4b. Semi-analytic solutions are needed to tame these numerical instabilities.

only to redshifting but also to broadening over time in momentum space. This is why we need the "generalised Boltzmann hierarchy" that introduces both multipole expansions and velocity expansions of the distribution function.

1.2.1 Generalised Boltzmann hierarchy

To numerically solve for the background and perturbation equations of stochastic dark matter, we modify `CLASS` [119, 120] to implement the generalised Boltzmann hierarchy. This was introduced by Nascimento in [121] as an alternative approach to solving for massive neutrinos in Boltzmann codes. It has the advantage of being much easier to implement in `CLASS`, and while Nascimento skipped the full implementation in his paper, he found that the generalised Boltzmann hierarchy gives comparable results to the `ncdm` component in `CLASS` in the high-precision setting with sub-percent accuracy. For us, it also has the advantage of being explicitly independent of the mass of dark matter particles so that we can only work with the diffusion rate Γ.

The main idea is that to have a systematic non-relativistic expansion of the usual Boltzmann hierarchy, one has to multiply the distribution function by $P_\ell(\mu)(q/E)^{2n+\ell}$ where μ is the cosine of the angular component of the momentum, with P_ℓ the Legendre polynomial, and then integrate the momentum variables. The resulting mode is the (n, ℓ) moment of the distribution function which we show by $f_{n,\ell}$. Its evolution equation is found by suitably integrating the covariant diffusion equation, which couples it to the adjacent moments. See Appendix A.3 for the full set of equations and Appendix A.3.1 for a superhorizon solution. The lowest moments in n and l represent fluid variables, which are coupled to higher moments. In practice for numerical solutions, this generalised Boltzmann hierarchy has to be terminated at some finite $n_{\max}$ and $l_{\max}$, in such a way that including even higher moments would not affect the fluid quantities within numerical precision. With trial and error, we find that for our case, $(n_{\max}, l_{\max}) = (20, 10)$ is enough to achieve convergence for fluid quantities, as shown in Figure 1.4a.

The generalised Boltzmann hierarchy provides a straightforward and efficient way of numerically solving for the background and perturbation quantities. However, as already noted in [121], it has the downside of breaking down at large $ky(\eta)$, where k is the wavenumber of perturbations and y is the typical free streaming length at conformal time η. This can be intuitively understood as follows: When the free streaming length of particles becomes comparable to a given length scale, then the perturbations at that length scale become increasingly sensitive to the high-velocity tail of the distribution function, and higher moments in n are needed to achieve convergence. In Figure 1.4b, one can see that choosing a higher $n_{\max}$ integrates further in time, but eventually, the moments start to diverge. Therefore, our numerical solutions using the generalised Boltzmann hierarchy have to be supplemented by a semi-analytic understanding of the behaviour of each of the distribution moments at large $ky(\eta)$. This will be the purpose of the next section.

1.2.2 Semi-analytic solutions

The generalised Boltzmann hierarchy is an infinite system of coupled ODEs and is challenging to solve analytically per se. However, all the equations originate from the single covariant diffusion equation (1.19). In the context of the Boltzmann equation for photons, the line-of-sight integration provides an integral solution for the CMB temperature perturbations. The solution includes the Sachs-Wolfe (SW) term, which accounts for the initial temperature perturbations and gravitational potential. There is also the integrated Sachs-Wolfe (ISW) term, which incorporates a time integral of the varying gravitational potential. Refer to Section 4.3 for more details. The ISW term specifically captures the gravitational redshift of the photons as they climb out of potential wells from the last scattering surface to the present time.

We will follow a similar approach in order to solve for the dark matter perturbations. Analogous to CMB anisotropies, the covariant diffusion equation admits an integral solution comprising SW and ISW terms. There are additional heat kernel integrals to account for diffusion. In general, the time integral cannot be performed explicitly. Nonetheless, in the regime of strong diffusion, when the generalised Boltzmann hierarchy becomes unstable, we can use saddle-point approximations to infer the power law behaviour of all the moments of the distribution. Before that, let us discuss the background distribution first.

Background

To start, let us look at the background covariant diffusion equation in FRW spacetime:

$$\frac{\partial f_0}{\partial \eta} - q\mathcal{H}\frac{\partial f_0}{\partial q} = a\frac{\Gamma m}{3}\frac{1}{q^2}\frac{\partial}{\partial q}\left(Eq^2\partial_q f_0\right). \tag{1.47}$$

Finding a fully relativistic solution to this equation is a difficult task. If there were an extra factor of E on the right, then this equation could have been put exactly in the form of a heat equation on the hyperbolic space, whose heat kernel is known analytically [122]. In fact, this can be done for the diffusion equation in the proper time of the particles, rather than the cosmic or conformal time [58]. However, as we mentioned earlier, the proper time of diffusive particles is not observable, so we need approximations for solving (1.47). It will be seen later that we are only interested in the case where the

dark matter particles remain non-relativistic up to the present time. Therefore, we assume that $E \simeq m$, which makes the right-hand side proportional to the flat Laplacian. Switching to comoving momenta $\tilde{q} = aq$, the equation becomes a simple diffusion equation with the new time parameter

$$d\mathcal{T} = a^3 d\eta \,. \tag{1.48}$$

Therefore, in the non-relativistic approximation, the solution can be written as

$$f_0(\mathcal{T}, \tilde{q}) = \frac{n_0}{\left(\frac{4}{3}\pi\Gamma m^2 \mathcal{T}\right)^{3/2}} e^{-\frac{\tilde{q}^2}{\frac{4}{3}\Gamma m^2 \mathcal{T}}} \,, \tag{1.49}$$

where n_0 is the present number density. We have assumed that the initial distribution function is a delta function at zero momentum. In other words, throughout the chapter, we assume the initial condition for stochastic dark matter is that of cold dark matter. This is of course, an idealization; for any specific particle model of dark matter, the production mechanism gives an initial distribution function [123–125] that should be convolved with the above heat kernel to give the correct distribution at later times. For now, we remain agnostic about the detailed particle model of dark matter and use CDM as the initial condition.

As noted in [58], the distribution (1.49) can be seen as a Maxwell-Boltzmann distribution with a time-dependent temperature

$$k_B T = \frac{2}{3}\Gamma m \frac{\mathcal{T}}{a^2} \,. \tag{1.50}$$

This is one way of seeing that covariant Brownian motion spontaneously warms up the stochastic dark matter particles. Note that in both radiation domination and matter domination, $\mathcal{T} \sim a^3\eta \sim a^2 t$, where t is the cosmic time. Therefore, the temperature grows linearly with the cosmic time and is independent of the choice of a_0, as expected. Also, remember that this equation is only valid as long as the particles remain non-relativistic, so $k_B T \ll m$. This translates to $\Gamma\mathcal{T}/a^2 \ll 1$, or $a\Gamma\eta \ll 1$. So we can only trust these non-relativistic approximations if $\Gamma \ll H_0$.

Having the solution to the background distribution at hand, we can directly infer the behaviour of several background quantities:

$$\begin{aligned} \bar{P}_{\rm dm} &= \frac{2}{3}\frac{\Gamma\mathcal{T}}{a^2} m\bar{n}_{\rm dm} \,, \\ \bar{\rho}_{\rm dm} &= \left(1 + \frac{\Gamma\mathcal{T}}{a^2}\right) m\bar{n}_{\rm dm} \,, \\ w_{\rm dm} &= \frac{2}{3}\frac{\frac{\Gamma\mathcal{T}}{a^2}}{1 + \frac{\Gamma\mathcal{T}}{a^2}} \simeq \frac{2}{3}\frac{\Gamma\mathcal{T}}{a^2} \,. \end{aligned} \tag{1.51}$$

The number density is conserved and scales as $1/a^3$. So for $\Gamma \ll H_0$, we conclude that the deviation of $\bar{\rho}_{\rm dm}$ from $1/a^3$ scaling must be very small, and the equation of state, although growing with time, remains negligible up to the present time. One might be worried that the condition $\Gamma \ll H_0$ is too restrictive, not allowing for any observable signatures. This is only true at the background level. In the next section, we will see that the picture is very different for perturbations. While background is only sensitive to the overall transfer of energy to dark matter particles, perturbations are sensitive to the

diffusion length. From equation (1.49), the typical velocity squared of stochastic dark matter particles at some time $\mathcal{T}$ is

$$v^2 \sim \frac{1}{a^2}\Gamma\mathcal{T} \sim a\Gamma\eta\,, \tag{1.52}$$

which is just the variance of a random walk on the mass shell. The typical diffusion length is given by vt, and so the comoving diffusion scale is

$$\lambda_D \sim vt/a \sim v\eta\,. \tag{1.53}$$

The free streaming length y is just the comoving diffusion length, i.e. $y = \lambda_D$. A typical stochastic dark matter particle could travel a distance of about 1Mpc during the age of the universe, with an equation of state consistently less than $w_{\rm dm} \sim 10^{-6}$. Such a large diffusion length has a significant impact on the small-scale perturbations, as we will see in the subsequent sections.

Perturbations

We expand the distribution function up to first order as

$$f = f_0 + \delta f\,. \tag{1.54}$$

We are mostly interested in solving for δf when the diffusion becomes strong. For the length scales of interest and reasonable values of the diffusion rate, this happens well after radiation-matter equality. We checked numerically that even in the presence of diffusion, the total anisotropic stress is negligible and we can use the approximation $\phi = \psi$. Moreover, to be able to give an analytic solution, we have to go the non-relativistic limit, where $q \ll m$ and $E \simeq m$. This is justified since, as we will see later, we will get a bound of $w \lesssim 4. \times 10^{-6}$ from CMB, which allows velocities of order $v \lesssim 2. \times 10^{-3} c$ (consistent with other reports [126, 127]). In the non-relativistic approximation, the hyperbolic Laplacian in the covariant diffusion term becomes a simple flat Laplacian in momentum space. Under these approximations, if we transform to the comoving momentum variable $\tilde{q} = aq$ again, and use $\mathcal{T}$ defined in (1.48) as the time parameter, the evolution equation for δf becomes

$$\frac{\partial}{\partial\mathcal{T}}\delta f + \frac{1}{ma^4} ik\tilde{q}\mu\,\delta f - \frac{\partial f_0}{\partial\tilde{q}}\frac{m}{a^2} ik\mu\phi = \frac{\Gamma m^2}{3}\delta_{ij}\frac{\partial^2}{\partial\tilde{q}_i\partial\tilde{q}_j}\delta f\,, \tag{1.55}$$

where $\mu = \hat{k}.\hat{q}$. Without the diffusion term on the right-hand side, one could solve the Boltzmann equation using a method similar to line-of-sight integration. This would lead to a term accounting for the initial condition and an integral term accounting for the effect of the gravitational source along the trajectory of particles. Although the diffusion term makes things more complicated, we can still solve this equation. The crucial observation is that (1.55) resembles a diffusion equation with a time-dependent linear potential and a non-homogeneous gravitational source term. Alternatively, by analytically continuing the time variable $\mathcal{T}$ to the imaginary axis, this becomes a Schrödinger equation with a time-dependent linear potential and a non-homogeneous source term. Forgetting about the non-homogeneous term for a moment, this quantum-mechanical problem has been studied and solved in [128]. Following the same steps, we can solve our equation (1.55) without the source term ϕ, and then use Duhamel's principle [129] to solve the full equation including the non-homogeneous source

term. We leave the details for Appendix A.5. The final solution is

$$\delta f(\mathcal{T},k,Q)=\int d^3Q'\left(e^{-\frac{1}{3}\Gamma k^2\mathcal{G}(\mathcal{T},\mathcal{T}_i)}\frac{e^{-\frac{|Q-Q'|^2}{\frac{4}{3}\Gamma m^2(\mathcal{T}-\mathcal{T}_i)}}}{\left(\frac{4}{3}\pi\Gamma m^2(\mathcal{T}-\mathcal{T}_i)\right)^{3/2}}e^{-i\vec{k}.\vec{y}(Q',\mathcal{T}_i)+i\vec{k}.\vec{y}(Q,\mathcal{T})}\delta f(\mathcal{T}_i,k,Q')\right)$$
$$+\int_{\mathcal{T}_i}^{\mathcal{T}}d\mathcal{T}'\int d^3Q'\left(e^{-\frac{1}{3}\Gamma k^2\mathcal{G}(\mathcal{T},\mathcal{T}')}\frac{e^{-\frac{|Q-Q'|^2}{\frac{4}{3}\Gamma m^2(\mathcal{T}-\mathcal{T}')}}}{\left(\frac{4}{3}\pi\Gamma m^2(\mathcal{T}-\mathcal{T}')\right)^{3/2}}e^{-i\vec{k}.\vec{y}(Q',\mathcal{T}')+i\vec{k}.\vec{y}(Q,\mathcal{T})}\frac{m}{a(\mathcal{T}')^2}i\vec{k}.\hat{q}'\frac{\partial f_0}{\partial\tilde{q}}\phi(k,\mathcal{T}')\right). \tag{1.56}$$

$\mathcal{T}_i$ is the initial time when the initial conditions are set, y is the free streaming length, Q is some redefined momentum parameter, and $\mathcal{G}(\mathcal{T},\mathcal{T}')$ is some form of a time difference between $\mathcal{T}'$ and $\mathcal{T}$ (see Appendix A.5 for details):

$$\mathcal{G}(\mathcal{T},\mathcal{T}_i)\sim\frac{\mathcal{T}^3}{a^8}-\frac{\mathcal{T}_i^3}{a_i^8}\sim a\eta^3-a_i\eta_i^3\,. \tag{1.57}$$

The important factor that we have to pay attention to is the first exponential in the two terms. Assuming that $\mathcal{T}_i\ll\mathcal{T}$, the initial-condition term in (1.56) gets exponentially suppressed by $\exp\left(-\frac{1}{3}\Gamma k^2\mathcal{G}(\mathcal{T},\mathcal{T}_i)\right)$ whenever

$$a\Gamma k^2\eta^3\gg 1\,. \tag{1.58}$$

As for the time-integral term, only the late-time part of the integral ($\mathcal{T}'\sim\mathcal{T}$) would give a non-negligible contribution, and the early-time contribution would be exponentially suppressed by the same exponential factor. Therefore, the strength of the diffusion at some scale k is determined by the factor $a\Gamma k^2\eta^3$. We say that diffusion has become relevant at this scale if this number is comparable to 1.

There is yet another more intuitive way of understanding the same result. As we found out in the previous section, at the background level, the strength of the diffusion is controlled by Γ/H. If this is small, then the effect of diffusion can be ignored in the background. However, perturbations are sensitive to the diffusion length. If the growing diffusion length becomes comparable to the length scale of interest, the clustering of dark matter particles starts to get washed away. From the previous section, we know that the typical comoving diffusion length is $\lambda_D\sim v\eta$. To see its relevance, this should be compared to the comoving length scale of interest, $\lambda=1/k$. We have

$$(ky)^2=\left(\frac{\lambda_D}{\lambda}\right)^2\sim a\Gamma k^2\eta^3\,. \tag{1.59}$$

Therefore, $a\Gamma k^2\eta^3$, the parameter that controls the exponential suppression of the initial-condition term and the early contribution of the time integral, is the square of the diffusion length divided by the length scale of interest. When this number becomes very large, only the late-time behaviour of the gravitational potential determines the distribution of stochastic dark matter particles. This aligns with our understanding of diffusion. For stochastic dark matter particles alone, entropy increases, and the dynamics is irreversible, effectively erasing the initial data.

Relation (1.59) indicates that the regime $a\Gamma k^2\eta^3\gg 1$ is also the regime where the generalised Boltzmann hierarchy becomes unstable, according to Section 1.2.1. Therefore, if we could find the analytic behaviour of (1.56) in this regime, we would have everything we need for the full solution to dark

matter perturbations. Fortunately, this is feasible. In this strong diffusion regime, we can ignore the initial-condition term. For the time-integral term, notice that the integrand consists only of either Gaussian or plane wave functions of Q'. So, the momentum integral can be carried out analytically. For the remaining time integral, we can use the steepest descent method or numerical techniques to capture only the very late-time contribution of the integral; see Appendix A.5. For example, the dark matter density contrast is given by the following integral

$$\delta(\eta,k) = -k^2 \int_{\eta_i}^{\eta} d\eta' \; \eta'\big(1-\frac{\eta'}{\eta}\big)\phi(\eta',k) e^{-\frac{1}{315}a\Gamma k^2\eta^3\left(1-6(\frac{\eta'}{\eta})^5+5(\frac{\eta'}{\eta})^6\right)} \,. \tag{1.60}$$

This equation shows that, more precisely, we need

$$a\Gamma k^2\eta^3 \gg 315 \tag{1.61}$$

to be in the strong diffusion regime. The numerical factor comes from specific time integrals in the matter domination era. Now, the integrand is exponentially larger at the right-end of the integration interval. The saddle-point approximation gives

$$\delta(\eta,k) \approx \frac{-21}{2a\Gamma\eta}\,\phi(\eta,k)\,. \tag{1.62}$$

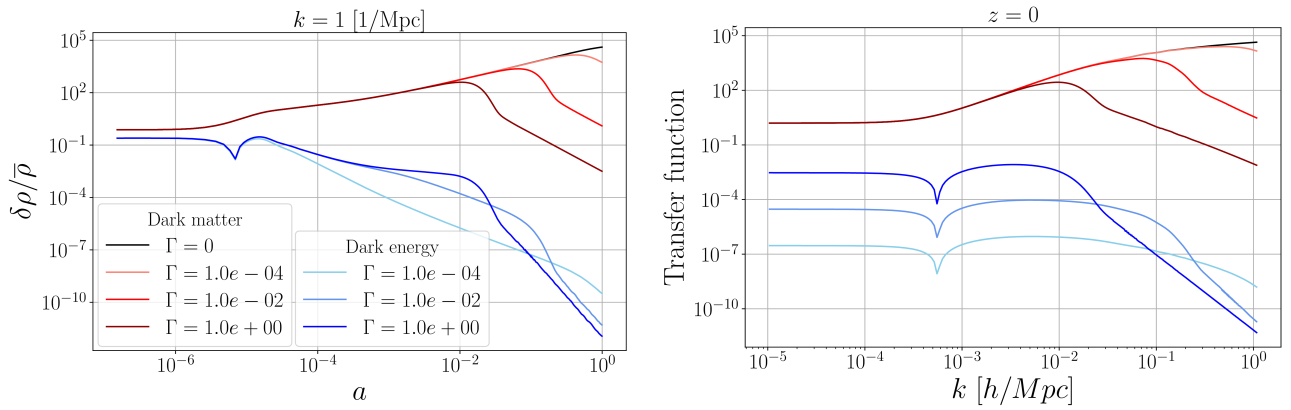
A more accurate numerical approximation gives 23 in the numerator instead; see Appendix A.5. Notice how δ depends only on the present gravitational potential and not on its past history. More importantly, this shows that when diffusion takes over, the density contrast of stochastic dark matter begins to decrease following a power law. In other words, the clustering becomes washed away, and the dark matter distribution tends toward homogeneity. This asymptotic solution matches well with the numerical results from the generalised Boltzmann hierarchy just before the onset of numerical instability. Thus, by combining the two, we get the full solution for the dark matter overdensity and transfer function, as depicted in Figure 1.5. The transfer function further shows the bottom-up smoothing of structures, with more suppression at smaller scales.

Similar expressions to (1.62) can be found for the pressure density, velocity, and anisotropic stress of stochastic dark matter:

$$\frac{\delta P}{\bar\rho} \approx -\phi\,, \quad (1+w)\,\theta \approx -48.6\,\frac{1}{a\Gamma\eta^2}\phi\,, \quad (1+w)\,\sigma \approx -\frac{12.6}{\sqrt{a\Gamma k^2\eta^3}}\,\phi\,. \tag{1.63}$$

Once again, a more accurate numerical integration gives $\delta P/\rho \approx -1.06\,\phi$, see Appendix A.5.

Finally, let us recall the approximations we used to derive these asymptotic solutions. We assumed that the particles are non-relativistic, and we ignored the total anisotropic stress to write $\phi = \psi$. And last, we ignored the effect of dark energy when writing the functions $\mathcal{T}$ and $\mathcal{G}$ in terms of the conformal time η. Since we only switch from numerical integration to the asymptotic solutions at very late times and small length scales, and also since the allowed diffusion rates by the data are small, $\Gamma/H_0 < 10^{-4}$, we expect the error introduced by this approximation in the cosmological observables to be negligible.



(a) (b)

Figure 1.5: Figure 1.5a shows the stochastic dark matter and dark energy overdensity versus scale factor at $k = 1$/Mpc, normalised by the initial photons overdensity. These are numerical solutions obtained by the generalised Boltzmann hierarchy supplemented with the asymptotic solution (1.62) at large $a\Gamma k^2\eta^3$. The dark energy overdensity is at all times much smaller than stochastic dark matter overdensity, except at very early times when the background dark energy density is negligible. Figure 1.5b shows the transfer function of the stochastic dark matter and dark energy versus k at the present time. The perturbations of dark energy are negligible compared to the stochastic dark matter across all scales. The units of Γ are km/s/Mpc and the two plots share the same legend.

1.2.3 Dark energy model

When we add the covariant diffusion term to the Boltzmann equation for stochastic dark matter, the energy-momentum tensor is no longer conserved since [85]:

$$\nabla_\mu T^\mu_{\ \nu} = \int \frac{d^3q}{E}\, p_\nu L[f]\,. \tag{1.64}$$

$L[f]$ being equal to $m\kappa\nabla^2_{\mathbb{H}_3} f$, the right-hand side is no longer zero. However, the Einstein equations require the total energy-momentum tensor to be conserved. Thus, it is necessary to identify the source of the energy and momentum transfer to the stochastic dark matter.

From a fundamental point of view, covariant Brownian motion happens because the very notion of geometry at the Planck scale is ambiguous. The stochastic term in the geodesic equation accounts for our ignorance of the effect of such a discrete or fuzzy geometry on the motion of a point particle. Therefore, the back-reaction of this stochastic motion of a particle on the dynamics of spacetime is of equal importance. Ideally, this should be effectively described by a stochastic term in the Einstein equations. For a single particle moving in spacetime, the energy-momentum tensor has a stochastic contribution from p^μ (1.34). To balance this, there must be a stochastic term on the geometric side of the Einstein equations. Finding such a correction is beyond the scope of this thesis and is an important avenue for future research.

Here, we adopt a phenomenological approach to this problem. Although a single particle exhibits stochastic behaviour, their collective has a deterministic behaviour. Hence, the correction terms to the geometry behave deterministically as well. We propose that this can be effectively described by an imperfect dark energy fluid. In this sense, dark energy should be thought of as part of the geometry rather than a matter species. In other words, an imperfect dark energy fluid represents the mechanism responsible for the back-reaction of covariant Brownian motion on geometry. It transfers

energy and momentum to stochastic dark matter particles, conserving the total energy-momentum tensor. Consequently, it has to obey

$$\nabla_\mu T_x^{\mu\nu} = -\nabla_\mu T^{\mu\nu} \,, \tag{1.65}$$

where $T_x^{\mu\nu}$ is the energy-momentum tensor of dark energy.

At this point, we have to choose a kinematic ansatz for $T_x^{\mu\nu}$. The simplest choice would be to promote the cosmological constant to a spacetime scalar $\Lambda(x)$. In fact, in unimodular gravity, the diffeomorphism invariance is relaxed in such a way that it naturally allows for non-conservation in the matter energy-momentum tensor such that $\nabla_\nu \Lambda = 8\pi G \nabla_\mu T^\mu_\nu$ [130]. Such theories have even been motivated by discreteness at the Planck scale [131, 132]; see [133–139] for cosmological applications and inflationary scenarios. However, this does not apply to our case because (1.64) cannot be written as the gradient of a scalar.

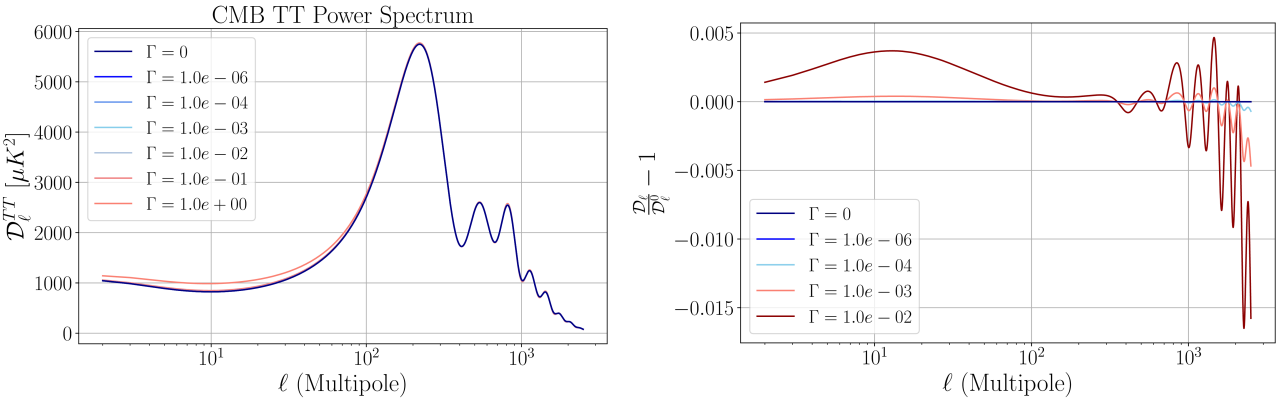
The next simplest ansatz is to consider a perfect fluid dark energy. If one chooses the dark energy equation of state to be $w_x = -1$, then the perturbed velocity divergence of dark energy, θ_x, would not appear in any dynamical equation as it is always multiplied by $1 + w_x$. Then the spatial part of the equation (1.65) enforces a constraint equation on the pressure perturbation of dark energy, which translates into a constraint on its total sound speed squared, c_x^2. Such a constraint equation fixes c_x^2 in terms of the energy density of dark energy and velocity divergence of dark matter, $\theta_{\rm dm}$. As a result, c_x^2 can in general become negative, creating gravitational instability in the perturbation equations. So one has to allow for $w_x \neq -1$. In this case, it has been shown that interacting dark energy-dark matter theories could have early-time instabilities [140]. To avoid the instabilities in the case that energy flows from dark energy to dark matter, one has to ensure $w_x < -1$; i.e. one needs phantom dark energy [141]. To avoid physical problems associated with phantom dark energy models [142], we discard this option as well.

Finally, we get to imperfect fluid models for dark energy. In [143–147] it is shown that many modified gravities and scalar dark energy models can be expressed in GR using an imperfect fluid energy-momentum tensor. In its most general form, it contains heat flux and anisotropic stress terms in addition to the perfect fluid quantities. For our current purpose, since (1.65) has four equations, we only need four dynamical degrees of freedom. Therefore we use the following ansatz

$$T_x^{\mu\nu} = P_x g^{\mu\nu} + (P_x + \rho_x) U^\mu U^\nu + \mathcal{Q}_x^\mu U^\nu + \mathcal{Q}_x^\nu U^\mu \,. \tag{1.66}$$

$\mathcal{Q}_x^\mu$ is the heat flux vector that satisfies $\mathcal{Q}_x^\mu U_\mu = 0$, so it has three independent degrees of freedom. We set $w_x = -1$, and to avoid gravitational instability, we impose $c_x^2 = 1$ manually. Unlike perfect fluids, setting $w_x = -1$ does not enforce any constraint equation on c_x^2 due to the presence of $\mathcal{Q}_x^\mu$. We still have three extra degrees of freedom in U^μ. To fix them, we make the extra assumption that the imperfect dark energy fluid is tightly coupled to the stochastic dark matter so that they share the same velocity vector U^μ even at the perturbations level. The background and perturbation equations derived from (1.65) can be found in Appendix A.4.

We highlight again that this is a purely phenomenological way of dealing with the back-reaction of the covariant Brownian motion on geometry, and the correct formulation will most probably correct our



(a) (b)

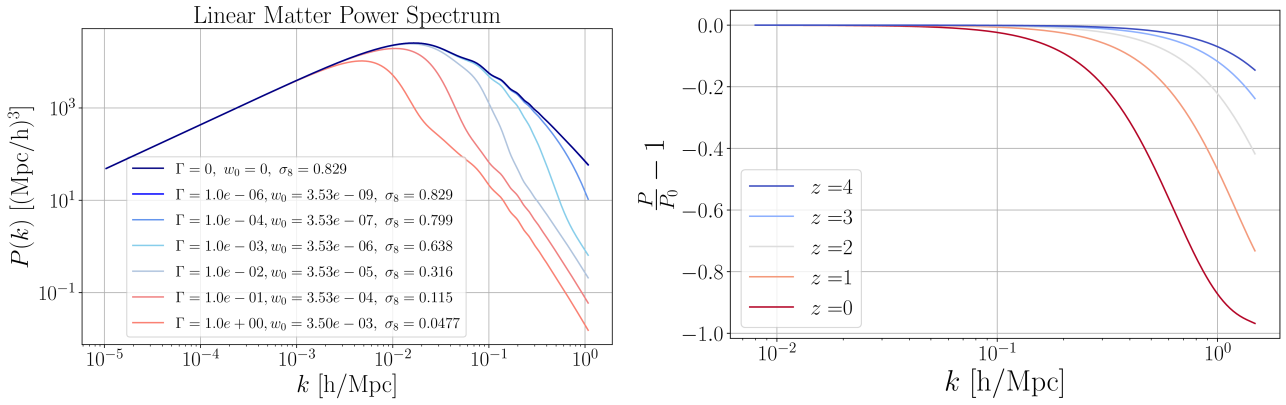
Figure 1.6: The CMB TT power spectrum, $\mathcal{D}_\ell = \frac{\ell(\ell+1)}{2\pi} C_\ell$, for increasing values of Γ in km/s/Mpc. Figure 1.6b shows the fractional change in the power spectrum compared to ΛCDM, $\mathcal{D}_\ell^0$. At low ℓ there is an enhancement due to an increased late ISW contribution. At large ℓ, the diffusion has a de-lensing effect.

assumptions. For the present, we believe this is suitable enough for our purposes since the main physics that we are confident about and want to investigate is the diffusion of dark matter. As the dark energy sector is more speculative, we want it to have minimal deviations from a constant Λ. In other words, we want the background energy density of dark energy to closely resemble a cosmological constant and the clustering of dark energy to be negligible compared to dark matter. In Figure 1.3a, one can see that if $\Gamma \ll H_0$, the deviation from a cosmological constant behaviour is indeed small. Also, Figure 1.5 shows that the dark energy overdensity is small compared to dark matter, therefore having a negligible impact on structure formation.

1.2.4 Effect on observables

Figure 1.6a shows the CMB TT power spectrum for various Γ values. Apart from an enhancement of C_ℓ at $\ell < 100$ for Γ values around 1km/s/Mpc, other changes are less visually apparent. Figure 1.6b better shows the fine details by plotting the fractional change in C_ℓ compared to ΛCDM ($\Gamma = 0$). Here, we can better see the change at small angular scales. This pattern resembles the decaying dark matter model discussed in [148]. As detailed in Appendix A.5 and Figure 1.8b, stochastic dark matter induces a decay in the gravitational potential at small scales even during matter domination. A time-dependent gravitational potential enhances the late ISW effect, hence increasing the power at small ℓ. On the other hand, due to the suppression of structures at small scales, the amount of lensing is reduced, which results in sharper peaks but an overall decrease in the power at large ℓ.

Figure 1.7a shows the effect of increasing values of Γ on the matter power spectrum. We can clearly see the power-law suppression at small scales, which is accompanied by increasing values of the present-day equation of state of dark matter and decreasing values of σ_8—the standard deviation of matter density perturbations at the scale of $8h^{-1}$Mpc—and of the derived parameter $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$. Figure 1.7b shows the relative suppression of the matter power spectrum compared to ΛCDM at different redshifts for a fixed diffusion rate. The suppression of power at small scales is redshift-dependent, with a stronger suppression at lower redshifts. These are the distinguishing features of stochastic dark matter. For example, the interacting dark energy-dark matter model studied in [149], with a net flow of energy



(a) (b)

Figure 1.7: Suppression of the matter power spectrum in stochastic dark matter scenario. Figure 1.7a shows the power-law suppression of the matter power spectrum at $z = 0$ for increasing values of Γ in km/s/Mpc, which is accompanied by increasing values of $w_{\rm dm}$ at present and decreasing values of σ_8. Figure 1.7b shows the fractional change in the matter power spectrum compared to the power spectrum of ΛCDM, P_0, for $\Gamma = 1.44 \times 10^{-4}$km/s/Mpc at various redshifts.

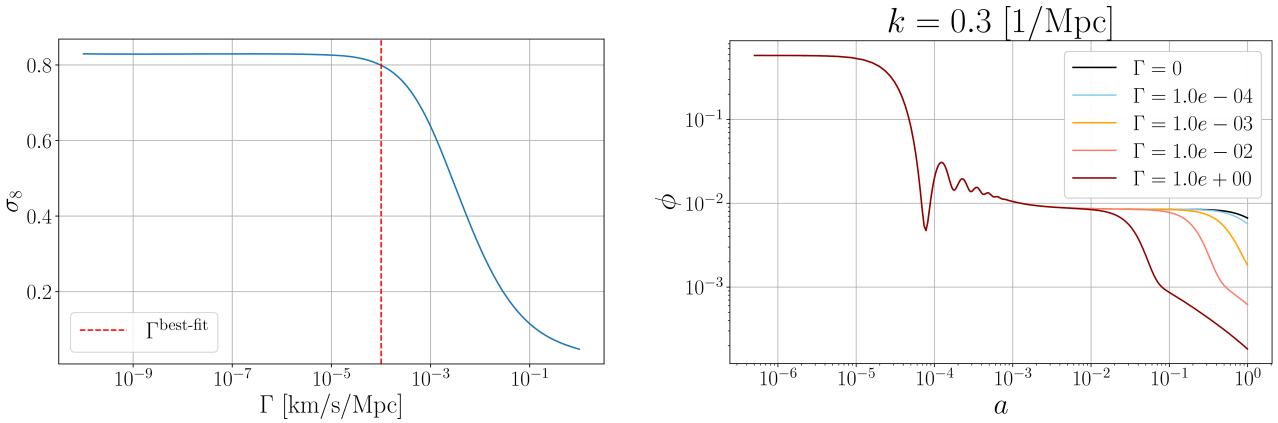
from dark energy to dark matter, was shown to suppress the matter power spectrum and relieve the S_8 tension. However, the power spectrum is modified at all scales k, there is a shift in the location of the peak, and the suppression plateaus at very large k relative to ΛCDM. For massive neutrinos and many warm dark matter or interacting dark matter models, the interesting new physics happens at sufficiently early times so that the relative matter power spectrum is weakly dependent on redshift at late times [125, 150–152]. The closest models in the literature to stochastic dark matter are decaying dark matter models that feature similar redshift-dependent suppression of power at small scales to offer resolution for the S_8 tension [153]. However, they also modify the matter power spectrum at large scales, and the suppression plateaus at very large k [154, 155]. Unlike such scenarios, for stochastic dark matter, the new physics is happening only at small scales and at late times. This feature helps the model to satisfy CMB lensing constraints which are mainly sensitive to matter overdensity at $2 \lesssim z \lesssim 5$, while suppressing the matter power spectrum strong enough at $z \lesssim 1$ to decrease S_8 and make it consistent with the weak lensing measurements.

Figure 1.8a shows the dependence of σ_8 on Γ while keeping the other cosmological parameters fixed. One can see that for $\Gamma < 10^{-5}$km/s/Mpc, σ_8 remains virtually unchanged, while for $\Gamma > 10^{-3}$km/s/Mpc, the decay in structure formation is too strong. This gives a clear picture of the order of magnitude of Γ if it is to satisfy the cosmological constraints.

1.2.5 Observational Constraints

The stochastic dark matter scenario has only a single additional free parameter Γ to the standard six free parameters of ΛCDM: $\{H_0, \Omega_{\rm dm}h^2, \Omega_b h^2, n_s, ln(10^{10}A_s), \tau_{reio}\}$. To constrain these parameters from cosmological measurements, we perform Markov Chain Monte Carlo (MCMC) using `MontePython` [156, 157]. We employ flat priors for all parameters, including a flat prior of $[0, 1.0]$km/s/Mpc for Γ.

As explained in [30], any model that claims a resolution for the S_8 tension should be tested against a variety of different datasets. Our datasets include CMB high-ℓ TT, TE, EE, low-ℓ TT, EE, and lensing power spectra [39, 158], as well as expansion and growth histories data from SDSS DR7 Main



(a) (b)

Figure 1.8: Figure 1.8a shows the dependence of σ_8 on Γ. Increasing the diffusion rate decreases the clustering of matter. To have enough structure formation, we need to constrain the diffusion rate to $\Gamma \lesssim 10^{-3}$km/s/Mpc. Figure 1.8b shows the gravitational potential ϕ, normalised by the initial photons overdensity, versus scale factor at $k = 0.3$/Mpc. Unlike CDM, for which ϕ is constant during matter domination and only starts to decay during dark matter-dark energy transition epoch, here the decay starts earlier.

Galaxy Sample (MGS) [159], BOSS DR12 LOWZ and LRG samples [160], eBOSS DR16 LRG and Quasar (QSO) samples [161–165], eBOSS DR16 Ly-α [166], and uncalibrated Pantheon+ SN Ia catalog [167].[2] Together, we call these our Baseline. We also include weak lensing measurements as the main drivers of the S_8 tension in ΛCDM. The latest cosmic shear measurement by the DES collaboration reports $S_8 = 0.759^{+0.025}_{-0.023}$ [25, 26]. Following similar analyses of the S_8 tension for power-suppressing models [151–153], we only use this measured value of S_8 as a prior for our analysis. This requires model independence of the S_8 measurements. Studies by both the KiDS and DES collaborations [168, 169] indicate robustness of S_8 value for various models, including a massive neutrino model νCDM and a massive sterile neutrino model $N_{\text{eff}} - m_{\text{eff}}$ which show power suppression at small scales. However, to the best of our knowledge, none of these tested models show significant modification to the late-time structure growth like the stochastic dark matter. So when using S_8 measurements as priors, this caveat should be borne in mind. A more complete analysis of stochastic dark matter in light of weak lensing measurements, with accurate modelling of non-linearities or with proper scale cuts is beyond the scope of the present thesis, and we leave it to future work. We have also tested that using the cosmic shear measurement by KiDS DR4 [27], $S_8 = 0.776^{+0.029}_{-0.027}$, does not change our conclusions.

A note of caution is in order. With an updated redshift calibration methodology, the cosmic shear measurement from KiDS DR5 [170] is consistent with the CMB prediction from Planck2018, showing no S_8 tension within ΛCDM. While such a resolution of the tension within ΛCDM is reassuring, the significance of the tension in the past few years has stemmed from the accumulation of several independent measurements. For this reason, and pending forthcoming independent results from LSST and Euclid, we do not incorporate this latest KiDS S_8 into the present analysis. It should be noted that if one were to adopt such a prior for S_8, the following results may turn out weaker.

We want to assess how stochastic dark matter affects the S_8 tension and how it compares to ΛCDM. First, we only keep those datasets that, when combined with Pantheon+ and an S_8 prior, result in

[2] More on supernova data is discussed in Section 4.1.

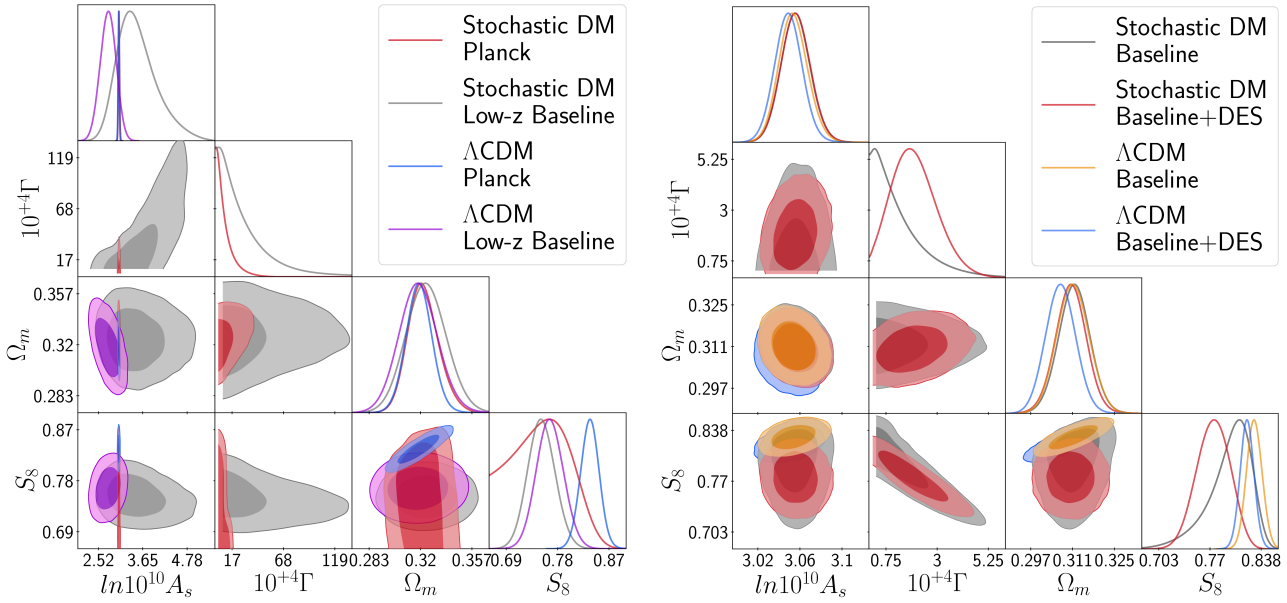


Figure 1.9: 68% and 95% confidence-level posterior distributions of the amplitude of primordial fluctuations, the diffusion rate of stochastic dark matter, Ω_m, and S_8 for stochastic dark matter and ΛCDM from various combinations of datasets. Figure 1.9a emphasizes the discrepancy in S_8 between Planck and low-z measurements in ΛCDM, while showing consistency in the case of stochastic dark matter. Figure 1.9b shows the 1.7σ preference for non-zero Γ when all of the datasets are taken into account together with the S_8 prior from DES.

a posterior average $S_8 < 0.785$ for ΛCDM. These are SDSS DR7 MGS, eBOSS DR16 LRG, eBOSS DR16 Ly-α, uncalibrated Pantheon+ SN Ia, and DES, which we call the "Low-z Baseline" data. Figure 1.9a compares the posterior plots of ΛCDM and stochastic dark matter for Planck2018 with Low-z Baseline. For ΛCDM, there is a 2.4σ tension in the respective S_8 posteriors. In contrast, the stochastic dark matter model effectively resolves this tension. The S_8 posterior from Planck2018 under this model has an extended lower tail, such that it is consistent with the low-z constraint. It is important to note that this is not just through having a large parameter space and enlarged error bars. The diffusion rate Γ specifically targets S_8, while maintaining the posterior precision of other parameters at almost the same level as ΛCDM. This is shown in Appendix A.6 by comparing the full posteriors of ΛCDM and stochastic dark matter for Planck2018. In particular, stochastic dark matter does not affect the H_0 tension. This is not surprising because the allowed range for the diffusion rate satisfies $\Gamma \ll H_0$, and hence the background is unaffected by the model.

We have seen that in the stochastic dark matter model, the Planck data has a very low constraining power on S_8, making it consistent with the Low-z Baseline data. The consistency allows us to combine these datasets. Another approach to evaluate the tension or consistency in S_8 is to compare the best-fit χ^2 values for a given model when combining all datasets, both with and without the S_8 prior. The tension can then be estimated by calculating the square root of the χ^2 difference between these two cases [153]. Table 1.1 gives the minimum χ^2 for various combinations of datasets. After including the S_8 prior, the minimum χ^2 for ΛCDM increases by ~ 6.7, while for stochastic dark matter, it only increases by ~ 1.2. Again, this shows a reduction in the S_8 tension from 2.6σ in ΛCDM to 1.1σ in stochastic dark matter.

Table 1.1: Best-fit χ^2 of stochastic dark matter and ΛCDM from different datasets. $\Delta\chi^2$ denotes the best-fit χ^2 of stochastic dark matter minus that of ΛCDM.

	Planck	**Low-z Baseline**	**Baseline**	**Baseline+DES**
Stochastic DM χ^2	1192.44	1418.16	4215.26	4216.5
ΛCDM χ^2	1191.76	1419.30	4212.98	4219.7
$\Delta\chi^2$	$+0.68$	-1.14	$+2.28$	-3.2

Finally, we use the Baseline and Baseline+DES to get the strongest constraints on Γ and S_8 for stochastic dark matter. Figure 1.9b shows the posterior plots of this analysis. First, in the case of ΛCDM, adding the DES datapoint has the effect of slightly shifting Ω_m to lower values to obtain a smaller S_8. For stochastic dark matter, however, Ω_m remains virtually unchanged after adding the DES datapoint. Instead, by tuning Γ, S_8 becomes consistent with the DES measurement. Although the posteriors are consistent with the diffusion rate being zero, here we see a preference for non-zero Γ that ranges between 1.2σ and 1.7σ, depending on whether we use the S_8 prior from KiDS or DES. We highlight that this non-zero preference is entirely driven by the S_8 prior from weak lensing measurements. While the Baseline still prefers the ΛCDM model, as shown in Table 1.1, stochastic dark matter provides a better best-fit χ^2 with Baseline+DES by about 3.2. Precise model comparison, however, remains beyond the scope of this thesis and is deferred to future studies.

Table 1.2: Set of best-fit parameters and posterior averages and standard deviations from the analysis of Baseline and Baseline+DES. At each entry, the number on top is the best-fit parameter.

Parameter	**Stochastic DM**		**ΛCDM**	
	Baseline	**Baseline+DES**	**Baseline**	**Baseline+DES**
$100\Omega_b h^2$	2.244 $2.239^{+0.013}_{-0.014}$	2.245 $2.240^{+0.013}_{-0.014}$	2.235 $2.240^{+0.013}_{-0.014}$	2.240 $2.245^{+0.013}_{-0.013}$
$\Omega_{\rm dm} h^2$	0.1190 $0.1198^{+0.0009}_{-0.0009}$	0.1195 $0.1196^{+0.0009}_{-0.0009}$	0.1200 $0.1197^{+0.0009}_{-0.0009}$	0.1188 $0.1190^{+0.0008}_{-0.0008}$
$\ln(10^{10} A_s)$	3.060 $3.055^{+0.014}_{-0.015}$	3.057 $3.055^{+0.014}_{-0.015}$	3.052 $3.053^{+0.014}_{-0.015}$	3.056 $3.047^{+0.013}_{-0.015}$
n_s	0.9680 $0.9661^{+0.0037}_{-0.0036}$	0.9676 $0.9664^{+0.0036}_{-0.0036}$	0.9671 $0.9662^{+0.0036}_{-0.0036}$	0.9687 $0.9674^{+0.0036}_{-0.0037}$
τ_{reio}	0.0630 $0.0591^{+0.0067}_{-0.0078}$	0.0606 $0.0594^{+0.0070}_{-0.0077}$	0.0574 $0.0583^{+0.0068}_{-0.0076}$	0.0618 $0.0567^{+0.0066}_{-0.0074}$
$10^{+4}\Gamma$	0.157 $1.29^{+0.24}_{-1.29}$	1.44 $2.00^{+0.84}_{-1.18}$	— —	— —
H_0	67.82 $67.52^{+0.39}_{-0.40}$	67.70 $67.60^{+0.39}_{-0.39}$	67.40 $67.54^{+0.39}_{-0.40}$	67.88 $67.84^{+0.38}_{-0.38}$
Ω_m	0.3076 $0.3119^{+0.0053}_{-0.0054}$	0.3096 $0.3107^{+0.0052}_{-0.0053}$	0.3133 $0.3115^{+0.0052}_{-0.0055}$	0.3064 $0.3074^{+0.0050}_{-0.0050}$
S_8	0.820 $0.795^{+0.036}_{-0.017}$	0.786 $0.774^{+0.023}_{-0.022}$	0.832 $0.829^{+0.010}_{-0.010}$	0.821 $0.819^{+0.009}_{-0.010}$

1.2.6 Discussion

The best-fit diffusion rate from Baseline+DES is

$$\boxed{\Gamma = 1.44 \times 10^{-4} \text{km/s/Mpc}}. \tag{1.67}$$

This is about 6 orders of magnitude smaller than H_0, implying a "diffusion time" Γ^{-1} that is a million times the age of the universe. At first glance, this could suggest a fine-tuning problem from a quantum gravity perspective, potentially making (1.67) too small to be explained even by Planckian effects.

However, Γ^{-1} should not be considered the primary time scale for evaluating the strength of quantum gravity effects. Rather, it is the time scale over which the covariant Brownian motion affects the background energy density on cosmological scales. In principle, different quantum gravity theories would provide distinct mechanisms for diffusion. Within each model, Γ has to be related to some more fundamental parameters. Thus, the interpretation of the above value is highly model-dependent.

As a specific example, we consider the swerves models as introduced in section 1.1.3. Combining equation (1.40) with the definition of Γ in (1.45), we can relate Γ to the forgetting time:

$$\Gamma \approx 6\frac{t_p^4}{\tau_f^5} \quad \rightarrow \quad \tau_f \sim \left(\Gamma^{-1} t_p^4\right)^{1/5} . \tag{1.68}$$

Strictly speaking, this relation is confirmed only in flat two-dimensional spacetime. However, assuming that the ratio of powers is similar in more general cases, (1.67) gives an estimate for the forgetting time of stochastic dark matter particles

$$\tau_f \sim 10^{13} t_p \, . \tag{1.69}$$

This is a time scale much larger than the Planck time, but still four orders of magnitude below typical time scales in the Standard Model of particle physics and significantly smaller than any cosmological scale. Another nonlocal scale that appears in the discussion of d'Alembertian on causal sets has been estimated to be 20 orders of magnitude larger than the Planck length [19]. Had these nonlocal scales been of the order of cosmological scales, we would have had a fine-tuning problem. Further implications of the above forgetting time require more sophisticated models of matter propagation on a causal set, potentially relating the forgetting time to the Compton wavelength of the particle or its wave-packet properties [171].

How does our best-fit Γ compare to previous bounds on the diffusion parameter? Since we are agnostic about the mass of dark matter particles, we cannot translate our result into a bound on κ. Instead, the bounds on κ presented in [58, 69], which are based on cold baryonic systems, can be translated into bounds on their diffusion rate Γ_b. These studies suggest an upper bound within the range

$$\Gamma_b \lesssim [0.1, 10]\text{km/s/Mpc} \, . \tag{1.70}$$

Therefore, in terms of the diffusion rate, our cosmological bounds on dark matter are much tighter. Nevertheless, due to the larger power of the Planck time in (1.68), they yield approximately the same lower bound (1.69) on the forgetting time of baryons. The authors in [69] also estimated an upper bound on the diffusion constant κ of massive neutrinos from existing bounds on hot dark matter. Although their bound on κ is much tighter than that for baryonic systems, the small mass of neutrinos results in a weaker bound on the diffusion rate than (1.70).

1.3 Summary and outlook

Inspired by swerves models that describe the deviation of massive particles from geodesics in discrete spacetimes of causal set theory, we built on previous work on the Lorentz-invariant diffusion equation in flat spacetime. We laid the groundwork needed to formulate all possible covariant, frame-independent stochastic processes in curved spacetime. By focusing on the minimal coupling case, we derived the unique covariant Brownian motion.

We applied the covariant Brownian motion framework to dark matter particles, developing a model of stochastic dark matter. We found that the effect of Brownian motion on perturbations becomes significant much earlier than on the background. It counteracts the dark matter clustering and dissipates overdensities caused by gravitational collapse, leading to a bottom-up suppression of the matter power spectrum. This allows stochastic dark matter to meet cosmological constraints on large scales and at higher redshifts while resolving the S_8 tension.

A key question is the behaviour of stochastic dark matter at sub-megaparsec scales, and whether the model can satisfy Milky Way satellite measurements [172, 173] and Ly-α forest flux power spectrum data [174, 175], which impose strong constraints on warm dark matter models. Since matter power suppression is strongest at $z \lesssim 1$, Milky Way satellite measurements are expected to place tighter constraints on stochastic dark matter. Addressing this requires fully exploring the non-linear regime through N-body simulations for stochastic dark matter, an important avenue for future research. This will also provide more accurate theoretical computations to compare with CMB lensing and weak-lensing measurements. Moreover, it would determine whether stochastic dark matter can resolve small-scale problems of CDM [176–180].

We can further test the covariant Brownian motion by applying it to other matter particles, including baryons and electrons. The bounds (1.70) already rule out significant diffusion in baryons in the early universe that could affect CMB spectral distortions. Nonetheless, these bounds allow for spontaneous heating of baryons in the late universe. Our formulation of the covariant diffusion equation in curved spacetime and its semi-analytic treatment in FRW spacetime pave the way for a systematic treatment of baryons. The diffusion rate of baryons can then be constrained through various methods, including its impact on the matter power spectrum, Ly-α bounds on the temperature of the intergalactic medium [181], and the heating of gas-rich dwarf galaxies [182]. We expect the resulting bounds on the diffusion rate of baryons to improve 1.70, making them comparable to those we have found for dark matter.

There are several open questions on the theoretical side. Our motivation for considering covariant Brownian motion stemmed from discrete quantum gravity. Another motivation concerns the effect of vacuum fluctuations of a scalar or electromagnetic field on a point particle coupled to that field [183–185]. As noted in [183], this process should be Lorentz invariant in the vacuum. The effective approach we developed in Section 1.1.2 can also inform all possible covariant stochastic processes in the presence of an additional field. Future work should explore the connection between our effective approach and existing Langevin equations. Another independent motivation for studying particle diffusion could arise from spontaneous collapse models [186]. In relativistic models [187, 188], we expect that a suitable Markovian, point-particle limit of these dynamics will recover the covariant

Brownian motion. Furthermore, attempts to reconcile quantum field theory with general relativity without quantizing the gravitational degrees of freedom [189] lead to collapse within the quantum system and a stochastic modification to Einstein equations. It is plausible, therefore, that alongside effects due to the stochasticity of spacetime itself [190], an accurate description of the phenomenology of the dark sector in such theories necessarily requires taking the covariant Brownian motion of particles into account.

Returning to our original motivation from discrete spacetime, a natural generalisation is to consider particles of higher spin. Lorentz invariant fluctuations of photon polarization have been studied in [70]. Our work prepares the ground for generalising this to curved spacetime and finding the corresponding Langevin equations on the Bloch sphere. Moving on from the point particle approximation, it is crucial to find a field-theoretic basis for covariant Brownian motion through a stochastic correction to the Klein-Gordon equation. A key question is whether Lorentz invariance imposes similar strict restrictions on stochastic corrections within the classical covariant phase space of scalar fields [191]. We are interested in recovering covariant Brownian motion from the particle approximation of such a Langevin Klein-Gordon equation. An example of a causal but nonlocal scalar field theory on a causal set exists, which, in the continuum limit, yields nonlocal corrections to the Klein-Gordon equation [19, 192, 193]. It remains an open question whether this nonlocality effectively manifests as a stochastic effect and whether it can recover covariant Brownian motion.

Finally, there is the question of back-reaction. In this work, for cosmological purposes, we used a phenomenological imperfect dark energy fluid to balance total energy-momentum conservation. Some assumptions were made to ensure the dark energy sector is well-behaved and stable, but these must ultimately be justified by a self-consistent treatment of covariant Brownian motion within the framework of Einstein equations. The Langevin-Einstein equations in [194] offer a potential path forward if we demand the noise tensor to balance energy-momentum non-conservation.

Chapter 2

Stochastic gravitational-wave background: constraints on massless diffusion

With the detection of numerous binary black hole merger events by LIGO and the upcoming plans for the Einstein Telescope and the Laser Interferometer Space Antenna (LISA), there is growing interest in using gravitational waves to test models inspired by quantum gravity that go beyond general relativity. While some studies have focused on modifications to the GW sources [195–198], many have concentrated on the effects of quantum spacetime on the propagation of GWs. Although such effects are expected to be small, they can accumulate and amplify over the cosmological distances traversed by GWs. Some studies are mostly motivated by models that arise in the context of a particular quantum gravity theory or modified gravity [199–201], while others highlight their theory-agnostic approach [202–205].

In the previous chapter, we provided a detailed study on the massive diffusion equation and its cosmological implications for dark matter. However, that equation is not applicable to gravitational waves. The limit $m \to 0$ of the relevant equations is problematic because in this limit, the mass shell becomes the light cone which has a singularity and its induced metric is degenerate. Nevertheless, it is possible to construct a unique Lorentz-invariant second-order diffusion operator on the light cone. Unlike the massive case, it has two free parameters but allows for diffusion and drift in only one dimension along the initial null momentum. This was found by Dowker et al. in [60] and the parameters were constrained using the observation that the CMB blackbody spectrum closely follows the Planck's distribution.

In this chapter, following the derivation of the stochastic geodesic equation for massive particles, we first derive its massless counterpart. We then present its corresponding Fokker-Planck equation, which is just the curved spacetime generalization of the massless diffusion equation in [60]. We emphasize that the massless dispersion relation is preserved. Therefore, any Lorentz-invariant quantum gravity effect on massless particles can be effectively described by this model.

Next, we present an analytic study of the massless diffusion equation in FRW spacetime, showing that the strength of the effect is proportional to the inverse of the frequency. This contrasts sharply with Lorentz-violating effects, which are stronger for large frequencies and are best probed using high-energy gamma rays. However, the variance of the diffusion is proportional to the elapsed affine time. The

affine time, defined with a suitable normalisation, measures the spacetime area swept by one wavelength [60] and is Lorentz invariant. The larger this area, the greater the stochastic Planck-scale effects on the propagating massless particle. Therefore, for a fixed diffusion constant and propagation distance, lower-frequency waves are the most promising for detecting the diffusion of gravitons. Since LISA probes GW frequencies in the mHz range, we must be able to improve the bounds on the diffusion and drift constants by many orders of magnitude compared to those derived from $\sim$ 100GHz CMB photons. Indeed, our forecast shows that, in the presence of a GW background with non-trivial spectral features, LISA could tighten the bounds on the diffusion and drift constants by over 12 orders of magnitude relative to those in [60].

In principle, one can put bounds on the diffusion effect with LIGO, Virgo, KAGRA, or other ground-based interferometers, since they are already operating and probe frequencies just about 4 orders of magnitude greater than LISA's. However, LIGO has not yet detected a GW background, and there are no universally accepted strong candidates for cosmological sources at those frequencies. On the other hand, the background sourced by astrophysical events in the LIGO band is not well described by a stochastic ensemble of gravitons but instead arises from the superposition of many unresolved weak coherent events. In addition, resolved sources such as compact binary coalescences, are not describable as a statistical ensemble of GWs. Since a Lorentz invariant diffusion at the field level for $\boldsymbol{h}_{\mu\nu}$ (which is necessary for a resolved source search) is currently lacking, we cannot constrain the diffusion parameter with LIGO catalog data at this time.

LISA is the first space-based GW interferometer, scheduled for launch in the 2030s. With its unprecedented sensitivity in the millihertz frequency band, LISA will open a new observational window into the early universe, probing stochastic GW backgrounds of cosmological origin that are inaccessible to ground-based detectors. Cosmological sources such as first-order phase transitions, primordial black hole (PBH) formation, and inflationary processes can generate relic GW backgrounds whose spectral features encode fundamental physics at energy scales far beyond the reach of particle colliders. LISA's capability to measure the spectral shape, amplitude, and anisotropies of a stochastic gravitational-wave background makes it a unique tool to test early-universe scenarios and constrain new physics models [206–208]. Strong first-order phase transitions in the early universe arise in many extensions of the Standard Model, notably in scenarios with extended scalar sectors or dark sectors. They serve as a prime example of cosmological signals potentially observable by LISA. During a first-order phase transition, bubbles of the new vacuum nucleate, expand, and collide, sourcing a stochastic GW signal characterised by a broken power-law spectrum with a peak frequency set by the transition temperature and dynamics of the bubble walls. LISA is particularly sensitive to transitions occurring at energy scales from tens of GeV up to tens of TeV, offering a probe of electroweak-scale and beyond–Standard-Model physics [206]. As the second example, primordial black holes, hypothesized to form from large density perturbations in the radiation-dominated era, represent another potential source of cosmological GW backgrounds in the LISA band. Strictly speaking, these are second-order GWs produced as a by-product of large scalar perturbations that are responsible for the formation of PBHs. The resulting spectrum of GWs also features a broken power-law shape, with characteristic frequencies determined by the PBH mass distribution and the underlying primordial power spectrum [208]. These two primordial power spectrum candidates, with grounded theoretical motivation, non-trivial frequency features in their spectrum,

and sourced in the early universe, can carry a clear imprint of the aforementioned graviton covariant diffusion effect during their propagation.

This chapter is organized as follows. In Section 2.1, we introduce the main quantities used to characterise the stochastic GW background and link them to the graviton distribution function in phase space. In Section 2.2, we find the modified massless geodesic equation and Boltzmann equation in the presence of a Lorentz invariant diffusion effect. In Section 2.3, we analytically solve the modified Boltzmann equation and we study the solution in the limit where the diffusion effects are small. This solution is employed in Section 2.4, where we show how to parametrise a realistic GW spectrum in the presence of the diffusion effects and we forecast the constraints posed on both the background model parameters and the diffusion and drift constants with LISA. In Section 2.5, we comment on the results.

2.1 The stochastic GW background

We consider a stochastic GW background as a superposition of plane waves. In the TT gauge, this can be written as

$$\boldsymbol{h}_{ij}(t,\mathbf{x}) = \int_{-\infty}^{\infty} d\nu \int d^2\mathbf{n} \sum_A \boldsymbol{h}^A(\nu,\mathbf{n}) e^A_{ij}(\mathbf{n}) e^{-2\pi i\nu(t-\mathbf{n}\cdot\mathbf{x})}\,, \tag{2.1}$$

where i, j run over spatial coordinates and the index A runs over all orthogonal polarisation states, including $+$ and $\times$. We consider only plane waves with a standard dispersion relation propagating at the speed of light. The unit vector $\mathbf{n}$ is aligned with the momentum of each wave with frequency ν.

The amplitude of $\boldsymbol{h}^A(\nu,\mathbf{n})$ in (2.1) is a stochastic variable at each frequency ν and direction $\mathbf{n}$, while the polarisation tensors are defined with normalisation $\sum_A e^A_{ij}(\mathbf{n}) e^{ij}_A(\mathbf{n}) = 2$. We assume the GW background to be stationary, isotropic, Gaussian, and unpolarised, as it is predicted by the vast majority of the cosmological models:

$$\langle \boldsymbol{h}^A(\nu,\mathbf{n}) \boldsymbol{h}^{B\star}(\nu',\mathbf{n}')\rangle = \delta_{AB}\delta(\nu-\nu')\delta^{(2)}(\mathbf{n}-\mathbf{n}') S_h(\nu)\,, \tag{2.2}$$

where S_h is the frequency power spectrum of the gravitational waves, while $\delta(\nu)$ and $\delta^{(2)}(\mathbf{n})$ are the 1-d and 2-d Dirac deltas, respectively. One can relate the power spectrum in Equation (2.2) to the graviton distribution function in phase space f via the GW energy density

$$\rho_{\rm GW} = \frac{\langle \dot{\boldsymbol{h}}_{ij} \dot{\boldsymbol{h}}^{ij}\rangle}{32\pi G} = \frac{2\pi^2}{G}\int_0^\infty d\nu\, \nu^2 S_h(\nu)\,. \tag{2.3}$$

If we define $f(E)$ to be the phase space distribution of gravitons for homogeneous distributions, with $E = 2\pi\hbar\nu$ being the energy, then the energy density is found simply by integrating $Ef(E)$ over the momentum space:

$$\rho_{\rm GW} = 4\pi \int_0^\infty dE\, E^3 f(E)\,. \tag{2.4}$$

Comparing the last two equations, one obtains the relationship between the distribution function and

power spectrum

$$f(E) = \frac{1}{32\pi^3 G}\frac{S_h(\nu)}{\nu}\,. \tag{2.5}$$

As the definition of the GW spectrum S_h may vary depending on conventions, it is always convenient to introduce the logarithmic GW energy density $\Omega_{\rm GW}(\nu)$

$$\frac{\rho_{\rm GW}}{\rho_c} = \int_{-\infty}^{\infty} d\ln\nu\ \Omega_{\rm GW}(\nu)\,, \tag{2.6}$$

where ρ_c is the present critical density. We can finally relate $\Omega_{\rm GW}$ to f:

$$f(E) = \frac{3H_0^2}{512\pi^6 G}\frac{1}{\nu^4}\Omega_{\rm GW}(\nu)\,, \tag{2.7}$$

where H_0 is the Hubble constant today. What matters most for our later analysis is the frequency scaling in $f \sim \Omega_{\rm GW}(\nu)/\nu^4$.

2.1.1 Wigner function of gravitational waves

There is yet another way to derive the distribution f in terms of $\Omega_{\rm GW}$ which is more appealing from a theoretical perspective. The distribution function is closely related to the Wigner function, and both satisfy the Boltzmann equation. The Wigner function is the Fourier-transformed point-separated two-point function, and for gravitational waves in the TT-gauge in flat spacetime, it is defined at spacetime point x and four-momentum p as

$$\mathcal{W}(x,p) = \frac{1}{32\pi G}\int \frac{d^4x'}{(2\pi)^4}\ e^{-ip.x'}\langle \boldsymbol{h}_{ij}(x+x'/2)\boldsymbol{h}^{ij}(x-x'/2)\rangle\,, \tag{2.8}$$

see [209] for the extension of Wigner functions to curved spacetime. This definition in the TT-gauge has the property that the energy-momentum tensor of the gravitational waves would be given by

$$t_{\mu\nu}(x) = \int d^4p\ p_\mu p_\nu \mathcal{W}(x,p)\,. \tag{2.9}$$

For another discussion of the Wigner function of the linearized gravitational field, see [210]. From the equation of motion of $\boldsymbol{h}_{ij}$, one can show that ignoring quantum fluctuations, $\mathcal{W}$ is non-zero only on the light cone, i.e. when $p^2 = 0$. Therefore, we can define the distribution function f as for scalar fields [211]

$$\mathcal{W}(x,p) = 2\ \delta(p^2) f(x,p)\,. \tag{2.10}$$

One can see that this definition, together with (2.9), gives the expression (2.4) for the energy density. Now since $\mathcal{W}$ is in a sense the Fourier transform of the two-point function, one could anticipate a close relation to the power spectrum S_h. Indeed, using the isotropic expressions (2.1) and (2.2) for $\boldsymbol{h}_{ij}$, one can find the Wigner function in terms of the power spectrum

$$\mathcal{W} = \frac{1}{8\pi^2 G}\frac{1}{p}S_h\,\delta(p^2)\,. \tag{2.11}$$

This, together with the definition of the distribution function (2.10), gives the relation between f and S_h that we found earlier by just comparing the energy density. However, knowing that $\mathcal{W}$ is a Wigner function ensures us that the distribution f satisfies the Boltzmann equation. Therefore, any effective modification to the Boltzmann equation from quantum gravity would show up in the transport equation for f, which is related to $\Omega_{\rm GW}$ via (2.7).

2.2 Modified transport equation for the distribution function

In the absence of interactions and collisions, the phase space distribution of the gravitational waves, f, satisfies a Boltzmann equation

$$L[f] \equiv p^\mu \frac{\partial}{\partial x^\mu} f + \Gamma^\sigma_{\mu\nu} p_\sigma p^\mu \frac{\partial}{\partial p_\nu} f = 0\,. \tag{2.12}$$

The term proportional to the Christoffel symbol is what gives the Hubble friction in FRW spacetime. To find the Lorentz invariant diffusive correction to this equation, this time we start with stochastic corrections to the massless geodesic equation and subsequently find the corresponding Fokker-Planck equation.

2.2.1 Stochastic massless geodesics

As mentioned earlier, many proposals motivated by quantum gravity modify the dispersion relation $p^2 = 0$ and hence break Lorentz invariance. However, motivated by stringent constraints on Lorentz violation [50–55], we maintain the light-cone propagation of gravitons and preserve $p^2 = 0$. Then, at an effective level, any quantum gravity effect must appear in the geodesic equation for p instead, and the modification should be such that $p^2 = 0$ is preserved by the new dynamics. In general, a modification to the geodesic equation as a function of an affine parameter λ includes a drift and a noise term, where the drift has to be a vector and the covariance matrix of the noise has to be a tensor:

$$dx^\mu = p^\mu d\lambda \tag{2.13}$$

$$dp^\mu + \Gamma^\mu_{\alpha\beta} p^\alpha p^\beta d\lambda = b^\mu d\lambda + dW^\mu\,. \tag{2.14}$$

As already pointed out in [60], the normalisation of the affine time is not arbitrary but set by (2.13) and the knowledge of physical momentum. We have separated the mean of the noise as a drift vector $b^\mu(p)$ so that the Gaussian noise dW^μ has zero mean, and its covariance tensor $D^{\mu\nu}$ has to be determined. Preserving the light-cone condition $p^2 = 0$ leads to the same conditions (1.30), (1.31), and (1.32) as the massive case. Once again, we are only interested in the minimal coupling case. However, there is an important distinction with the massive case: The same arguments as the previous section allow for two free parameters due to $p^2 = 0$. The drift vector and covariance tensor must be proportional to p^μ and $p^\mu p^\nu$, respectively, but possibly with two different proportionality constants. In other words, the conditions (1.31) and (1.32) enforce dW^μ to be null and orthogonal to p^μ. As a result, it has to be proportional to the momentum:

$$dW^\mu = p^\mu dW_\lambda\,, \tag{2.15}$$

where dW_λ is the differential increment of a 1-dimensional Wiener process in λ. We conclude that the covariant Brownian motion of massless particles is described by

$$\boxed{\begin{aligned} dx^\mu &= p^\mu d\lambda\,, \qquad (2.16)\\ dp^\mu + \Gamma^\mu_{\alpha\beta} p^\alpha p^\beta d\lambda &= \kappa_2 p^\mu d\lambda + p^\mu dW_\lambda\,, \qquad (2.17)\\ \langle dW_\lambda \rangle = 0\,, \quad \langle dW_\lambda^2 \rangle &= 2\kappa_1 d\lambda\,. \qquad (2.18) \end{aligned}}$$

κ_1 and κ_2 are two free parameters that have units of M/L. The massless geodesic equation corresponds to the limit $\kappa_1, \kappa_2 \to 0$. Note that κ_1 controls the variance of the noise and must be non-negative, while κ_2 can be positive or negative depending on whether the microscopic effect causes a blue or red tilt.

Interestingly, in an FRW spacetime, the momentum equation has a closed form solution. With the initial condition $p^\mu(\lambda = \lambda_i) = p_i^\mu$, this is given by:

$$p^\mu = \frac{a_i^2}{a^2} p_i^\mu \; e^{(\kappa_2 - \kappa_1)\Delta\lambda + W(\Delta\lambda)}\,. \qquad (2.19)$$

$\Delta\lambda$ is the total elapsed affine time. $W(\Delta\lambda)$ is the integral of the noise and is a Wiener process with variance $2\kappa_1\Delta\lambda$. In terms of physical energy $E = ap^0$,

$$E = \frac{a_i}{a} E_i e^{(\kappa_2 - \kappa_1)\Delta\lambda + W_{\Delta\lambda}}\,. \qquad (2.20)$$

The above solution recovers the $1/a$ redshifting of the particles for zero diffusion while clearly showing that the strength of the diffusion and drift is controlled by the dimensionless quantity $\kappa_{1,2}\Delta\lambda$. Note that given a frame, high-energy particles traverse the same distance in a shorter affine time. So low-frequency particles would experience stronger diffusion and drift. At first, this appears counter-intuitive, as we are used to the idea of larger frequencies probing smaller scales. However, frequency or the number of oscillations of a propagating wave is frame-dependent, while the elapsed affine time defined above is Lorentz invariant. The elapsed affine time measures the spacetime area swept by one de-Broglie wavelength of the particle [60]. The longer the wavelength, the larger this area. So a quantum-gravity effect that acts as independent stochastic kicks at Planck-scale volumes would add up with a standard deviation proportional to the square root of the spanned spacetime area of a wave.

From this, it is easy to show that the momentum of the diffusing particle has mean

$$\langle p^\mu \rangle = \frac{a_i^2}{a^2} p_i^\mu e^{\kappa_2 \Delta\lambda}\,, \qquad (2.21)$$

and variance:

$$\mathrm{Var}[p^\mu(\lambda)] = \frac{a_i^4}{a^4} (p_i^\mu)^2 e^{2\kappa_2\Delta\lambda} \left(e^{2\kappa_1\Delta\lambda} - 1 \right)\,. \qquad (2.22)$$

Unlike the massive case, the particle cannot change its direction on the light cone. To see this, note that all the components of the momentum get rescaled equally in time. Moreover, the scaling factor, being an exponential, is always positive. So the particle does not change orientation and never reaches zero frequency.

2.2.2 Massless diffusion equation

Using the projection relations (1.36) and (1.37) or the prescription in Appendix A.2.1, we can write down the covariant diffusion equation for massless particles corresponding the the above SDE:

$$p^\mu \frac{\partial}{\partial x^\mu} f + \Gamma^\sigma_{\mu\nu} p_\sigma p^\mu \frac{\partial}{\partial p_\nu} f = \big(3\kappa_1 - \kappa_2\big) \frac{1}{E} \frac{\partial}{\partial E} \big(E^2 f\big) + \kappa_1 \frac{1}{E} \frac{\partial}{\partial E} \Big(E^3 \frac{\partial}{\partial E} f\Big) \,. \tag{2.23}$$

Geometrically, this reflects the fact that $u^i = \big(\kappa_2 - 3\kappa_1\big) q^i$ is a true vector and $K^{ij} = \kappa_1 q^i q^j$ is a tensor on the phase space of massless particles, as explained by Dowker et al. [60]. As noted by the authors, there is currently no quantum gravity model of massless particles that explains the origin of the two free parameters, and the equation remains purely phenomenological. The above equation generalizes their Lorentz invariant equation to curved spacetime if we identify $k_1 = \kappa_1$, $k_2 = \kappa_2 - 3\kappa_1$, and $\rho_t \propto a^3 E^2 f$.

For a flat FRW universe with scale factor a and conformal time η, we get

$$\frac{1}{a} \frac{\partial}{\partial \eta} f - \frac{1}{a} E \mathcal{H} \frac{\partial}{\partial E} f = \big(3\kappa_1 - \kappa_2\big) \frac{1}{E^2} \frac{\partial}{\partial E} \big(E^2 f\big) + \kappa_1 \frac{1}{E^2} \frac{\partial}{\partial E} \Big(E^3 \frac{\partial}{\partial E} f\Big) \,, \tag{2.24}$$

where $\mathcal{H} = a'/a$. This equation reflects the fact that a peaked distribution of gravitons broadens and drifts over time. In the next section, we show the analytic solutions to this equation for a general expansion history.

2.3 Analytic Solution

Here we present two heat kernel solutions to the covariant diffusion equation and study their behaviour. Introducing the comoving energy as $\epsilon = aE$ and a new time variable T as $dT = a^2 d\eta$, the left side of Equation (2.24) simplifies to $a\partial_{\mathsf{T}} f$. Following the steps detailed in Appendix B.1, one can construct two heat kernels for this equation

$$\boxed{K_\pm(\mathsf{T}, \epsilon; \epsilon_i) = \frac{1}{\epsilon_i} \exp\left(-\frac{\epsilon + \epsilon_i}{\kappa_1(\mathsf{T} - \mathsf{T}_i)}\right) \left(\frac{\epsilon_i}{\epsilon}\right)^{3 - \frac{\kappa_2}{2\kappa_1}} \frac{\sqrt{\epsilon\epsilon_i}}{\kappa_1(\mathsf{T} - \mathsf{T}_i)} I_{\pm(1 - \frac{\kappa_2}{\kappa_1})}\Big(\frac{2\sqrt{\epsilon\epsilon_i}}{\kappa_1(\mathsf{T} - \mathsf{T}_i)}\Big)}\,, \tag{2.25}$$

where I_α is the modified Bessel function. The heat kernels satisfy the initial condition

$$K_\pm(\mathsf{T}_i, \epsilon; \epsilon_i) = \delta(\epsilon - \epsilon_i) \,. \tag{2.26}$$

One can check this using the zeroth-order term in the asymptotic expansion of the Bessel function. When $\kappa_1 \Delta\mathsf{T}$ is small (compared to energy scales ϵ and ϵ_i), $K_\pm$ is indeed a function of ϵ peaked at ϵ_i and integrates to 1. So in the limit $\kappa_1 \Delta\mathsf{T} \to 0$, it becomes the delta function, as expected.

The above heat kernels provide two solutions to Equation (2.24) given an initial condition at some initial time T_i is provided by $f(\mathsf{T}_i, \epsilon_i)$ for all $\epsilon_i > 0$. These can be written as the following integral

$$\boxed{f(\mathsf{T}, \epsilon) = \int_0^\infty d\epsilon_i \; f(\mathsf{T}_i, \epsilon_i) K_\pm(\mathsf{T}, \epsilon; \epsilon_i)}\,. \tag{2.27}$$

This is similar to a convolution but with the difference that the kernel is not a function of $\epsilon - \epsilon_i$ because the diffusion operator is not translationally invariant in ϵ.

However, the fact that we have found two different solutions to our diffusion equation is puzzling at the start. Mathematically, this happens because $\epsilon = 0$ is a singular point in the original PDE (2.24), and hence there is no uniqueness of solutions. Nevertheless, there must be only one physically acceptable solution. We can show that this is indeed the case by demanding the total particle number to be finite and conserved.

From the solution (2.20) to the stochastic geodesic equation, we know that the underlying particle model conserves the particle number; i.e. no particle gets lost at the tip of the light cone or zero frequency. So we expect the distribution of the particles to be conservative as well. In terms of the heat kernels, we demand $\int_0^\infty d\epsilon\ \epsilon^2 K_\pm$ to be equal to ϵ_i^2 at all times.

In Appendix B.1, we show that for the given heat kernels, we have

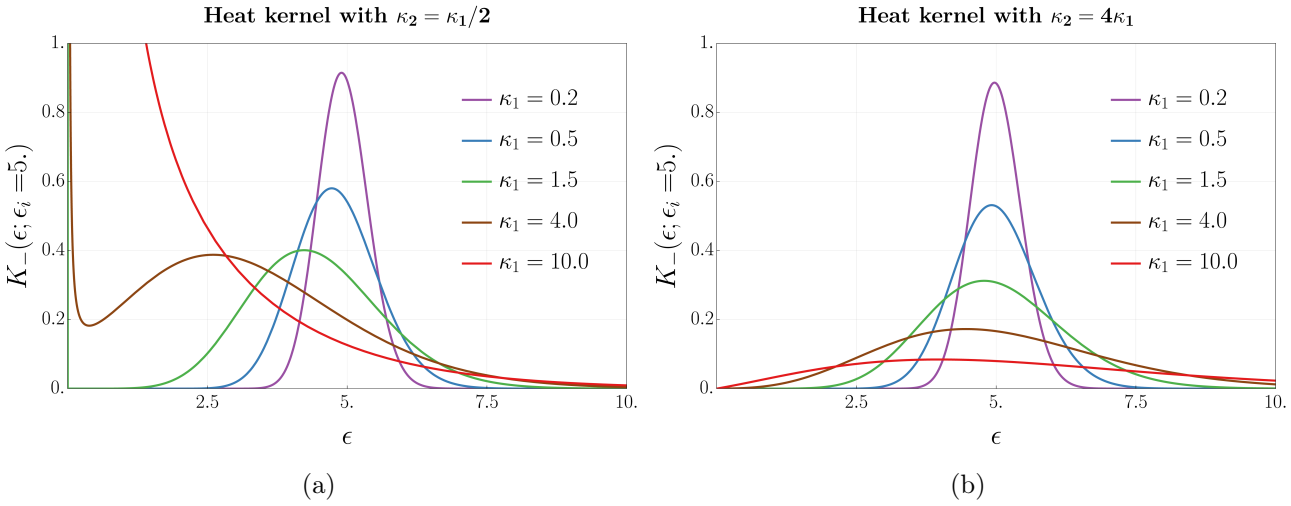
$$\int_0^\infty d\epsilon\ \epsilon^2 K_+ = \begin{cases} \epsilon_i^2\,, & \text{if } \kappa_2/\kappa_1 \in \{1,2,3,...\} \\ \epsilon_i^2\Big(1 - \dfrac{\Gamma(1-\kappa_2/\kappa_1, \frac{\epsilon_i}{\kappa_1 \Delta\mathsf{T}})}{\Gamma(1-\kappa_2/\kappa_1)}\Big)\,, & \text{if } \kappa_2/\kappa_1 \notin \{1,2,3,...\} \end{cases}\,, \tag{2.28}$$

$$\int_0^\infty d\epsilon\ \epsilon^2 K_- = \begin{cases} \epsilon_i^2\,, & \text{if } \kappa_2 > 0 \\ \epsilon_i^2\Big(1 - \dfrac{\Gamma(1-\kappa_2/\kappa_1, \frac{\epsilon_i}{\kappa_1 \Delta\mathsf{T}})}{\Gamma(1-\kappa_2/\kappa_1)}\Big)\,, & \text{if } \kappa_2/\kappa_1 \in \{0,-1,-2,...\} \\ \infty\,, & \text{else} \end{cases}\,. \tag{2.29}$$

So for $\kappa_2 > 0$, K_- gives a physically acceptable heat kernel. For natural number values of κ_2/κ_1, K_+ is also acceptable, but then it actually coincides with K_- because for integer values of α we have $I_{-\alpha}(z) = I_\alpha(z)$. So for $\kappa_2 > 0$, we have found a unique physically acceptable solution.

For $\kappa_2 \leq 0$, however, none of the above heat kernels are physically allowed, as they either gradually lose particles at $\epsilon = 0$ or give a divergent particle number. Note that this does not mean that the massless diffusion model is only valid for $\kappa_2 > 0$, but that we have not been able to find the correct heat kernel with our approach for solving the diffusion equation.

Nevertheless, it is important to note that the heat kernels $K_\pm$ are still valid for $\kappa_2 \leq 0$ if we avoid small frequencies; i.e. small values of $\frac{\epsilon}{\kappa_1 \Delta\mathsf{T}}$. To back up this claim, notice that the difference between $I_{-\alpha}(z)$ and $I_\alpha(z)$ is proportional to the modified Bessel function of the second kind. Therefore, the difference $K_+ - K_-$ is exponentially small for large values of $\frac{\epsilon}{\kappa_1 \Delta\mathsf{T}}$, and becomes noticeable only as we approach $\frac{\epsilon}{\kappa_1 \Delta\mathsf{T}} \to 0$, where the power law behaviours of K_+ and K_- are different. In particular, for $\kappa_2 > 0$, K_- behaves like $\epsilon^{\kappa_2/\kappa_1 - 3}$ at small frequencies. This can be seen in Figure 2.1 for two different rations of κ_2/κ_1. Also note that the curves for increasing values of κ_1 could be equivalently seen as increasing time $\Delta\mathsf{T}$. For explicit expressions for the power-law behaviour of the two kernels close to zero frequency see Appendix B.1.



(a) (b)

Figure 2.1: The physical heat kernel K_- for a time laps of $\Delta\mathsf{T} = 0.1$, and different values of $\kappa_{1,2}$ in arbitrary units. $\epsilon_i = 5$, so all the kernels are evolved from an initial delta function at $\epsilon = 5$. A larger value for κ_2 causes a relative drift toward higher energies.

2.3.1 Perturbative solution for small $\kappa_{1,2}$

Although we have the full analytic solution of the massless diffusion equation given by (2.27), it would be computationally expensive to perform the integral for every value of κ_1, κ_2, and ϵ in a numerical forecast. Moreover, all that matters for the sake of the Fisher forecast (which will be treated in detail in the next section) is the Taylor expansion of the solution for f up to the first order in κ_1 and κ_2.

Therefore, we expand our analytic solution to first order in $\kappa_{1,2}$. More precisely, the small dimensionless parameter over which we perform the Taylor expansion is $\frac{\kappa_1\Delta\mathsf{T}}{\epsilon}$. So this can be interpreted as an approximation on small enough time scales or at large enough frequencies as well. For primordial gravitational waves, $\Delta\mathsf{T} \approx a_0^2\eta_0/5$, where η_0 is the conformal age of the universe. This time-of-flight is dominated by the late-universe matter-dominated evolution. For a mHz experiment like LISA, $\frac{\kappa_1\Delta\mathsf{T}}{\epsilon}$ is of order 1 for $\kappa_1 \sim 10^{-55}\mathrm{kg\,m^2\,s^{-3}}$. So for our forecast, we restrict attention to values of $\kappa_{1,2}$ smaller than this. Otherwise, for values of $\kappa_{1,2}$ comparable to this number or larger, the effect on the spectrum would be very drastic, as can be seen with the red curves in Figure 2.1. Note that the regime of small $\frac{\kappa_1\Delta\mathsf{T}}{\epsilon}$ is also the regime in which $K_\pm$ coincide, and so the choice of the heat kernel becomes irrelevant. Also, the aforementioned issue of invalidity of the heat kernels for $\kappa_2 \leq 0$ is irrelevant for the same reason.

In the regime where $\frac{\epsilon}{\kappa_1\Delta\mathsf{T}} \gg 1$, we need the asymptotic expansion of the modified Bessel function

$$I_\alpha(z) = \frac{e^z}{\sqrt{2\pi z}}\left(1 - \frac{4\alpha^2 - 1}{8z} + \mathcal{O}(1/z^2)\right). \tag{2.30}$$

Since $\alpha = \pm(1 - \frac{\kappa_2}{\kappa_1})$ only appears as α^2, both $K_\pm$ have the same asymptotic expansion as expected. We then plug the above expansion in $K_\pm$, and change the heat kernel integration variable in (2.27) from ϵ_i to y using $\sqrt{\epsilon_i} = \sqrt{\epsilon} + y\sqrt{\kappa_1\Delta\mathsf{T}}$. After Taylor-expanding the integrand around $\kappa_1\Delta\mathsf{T} = 0$, the integrand becomes of the form $\exp(-y^2)$ times a series in y. So in the perturbative regime of small

diffusion, the heat kernel integral (2.27) becomes a simple Gaussian integral, giving

$$f(\mathsf{T},\epsilon) = f_i(\epsilon) + \left[2\left(3-\frac{\kappa_2}{\kappa_1}\right)\frac{1}{\epsilon}f_i(\epsilon) + \left(6-\frac{\kappa_2}{\kappa_1}\right)\partial_\epsilon f_i(\epsilon) + \epsilon\partial_\epsilon^2 f_i(\epsilon)\right]\kappa_1\Delta\mathsf{T} + \mathcal{O}\left(\left(\frac{\kappa_1\Delta\mathsf{T}}{\epsilon}\right)^2\right), \quad (2.31)$$

where $f_i(\epsilon) := f(\mathsf{T}_i,\epsilon)$. Given that we know the initial power spectrum analytically, this is very simple to evaluate numerically.

2.4 Forecast of $\kappa_{1,2}$ estimation with LISA

We aim to forecast the observability of a stochastic GW background by estimating the posterior distribution over the model parameters θ that characterise the power spectrum $\Omega_{\rm GW}(\nu;\theta)$. Specifically, we define $\theta = \{\theta_0,\kappa_1,\kappa_2\}$, where θ_0 is the set of parameters that model the unperturbed spectrum, i.e. in the absence of the diffusion effect, $\Omega_0(\nu;\theta_0)$, and $\kappa_{1,2}$ are the newly introduced diffusion and drift parameters that appear explicitly in Equation (2.31). For forecasting purposes, we specify the fiducial value for the model parameters $\bar\theta = \{\bar\theta_0,\bar\kappa_1,\bar\kappa_2\}$. Provided that $\bar\kappa_1,\bar\kappa_2 \ll 1$ in our units, we express Equation (2.31) as

$$\begin{aligned} f(\epsilon;\theta) &\simeq f_0(\epsilon;\theta_0) + \kappa_1 f_1(\epsilon;\theta_0) + \kappa_2 f_2(\epsilon;\theta_0)\,, \\ f_{1,2}(\epsilon;\theta_0) &= \left.\frac{\partial f}{\partial\kappa_{1,2}}\right|_{\kappa_{1,2}=\bar\kappa_{1,2}}. \end{aligned} \quad (2.32)$$

At the present time where $a_0 = 1$, the frequency is given by $\nu = \frac{\epsilon}{2\pi\hbar}$. Using Equation (2.7), the above equation translates into an expression for the GW energy density:

$$\begin{aligned} \Omega_{\rm GW}(\nu;\theta) &\simeq \Omega_0(\nu;\theta_0) + \kappa_1\Omega_1(\nu;\theta_0) + \kappa_2\Omega_2(\nu;\theta_0)\,, \\ \Omega_i(\nu;\theta_0) &= \left.\frac{\partial\Omega_{\rm GW}(\nu;\theta)}{\partial\kappa_i}\right|_{\kappa_i=\bar\kappa_i}. \end{aligned} \quad (2.33)$$

A stochastic GW background search with LISA exploits the fact that the instrument provides three data streams with (approximately) uncorrelated instrumental noise and a correlated signal in the presence of a GW background. It can be shown that the three correlators obtained between the LISA data streams follow a Gaussian distribution. Since the correlators between the data streams are proportional to the GW power spectrum and therefore the GW energy density, the likelihood function in the LISA correlators can be written as a function of the theoretical parameters θ. We report the explicit form of the likelihood function without reproducing the derivation, which can be found in [212]:

$$\begin{aligned} \mathcal{L}(\theta) &= \frac{1}{\mathcal{N}}e^{-\frac{1}{2}\chi^2(\theta)}\,, \\ \chi^2(\theta) &= T_{\rm obs}\int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\,\frac{4\pi\left(\Omega_{\rm GW}(\nu;\theta)-\Omega_{\rm measured}(\nu)\right)^2}{\Omega_n^2(\nu)}\,, \end{aligned} \quad (2.34)$$

where $\mathcal{N}$ is a normalisation constant and $\Omega_n(\nu)$ is the sensitivity curve for an isotropic stochastic GW background, obtained from the combined LISA data channels and for a total time of observation $T_{\rm obs}$. The frequency bandwidth for LISA is assumed to be between $\nu_{\rm min} = 0.1$mHz to $\nu_{\rm max} = 100$mHz. The

computation of $\Omega_n(\nu)$ follows the same steps of the derivation of the noise curve in [212] for an isotropic GW background that leads to their equation (4.43). The posterior distribution on the parameters θ can be obtained via the Bayes theorem given a prior function $\Pi(\theta)$

$$\mathbb{P}(\theta) = \Pi(\theta)\mathcal{L}(\theta)\,. \tag{2.35}$$

The explicit evaluation of the likelihood function in (2.34) is in general computationally expensive, given the high dimensionality of the space of parameters and the non-trivial dependence of the signal on these parameters in the LISA data. For forecasting purposes, approximate techniques, such as the Fisher method outlined in appendix B.2, provide the same quantitative and qualitative information as simulation-based approaches. They allow us to reliably forecast the degree of detectability and estimate the theoretical parameters. In appendix B.2, we analytically derive the covariance matrix for the theoretical parameters in the two specific models considered in the following subsection. This provides insight into how diffusion effects can correlate with the parameters of the unperturbed GW background model.

2.4.1 Diffusion on first order phase transitions and PBH spectra

We apply the outlined formalism to the specific case of the GW background from a strong first-order phase transition and the background produced in a PBH scenario. In particular, we adopt the broken power-law parametrisation employed in [213] as it effectively describes the spectral shape of GW backgrounds generated by first-order phase transitions:

$$\Omega_0\left(\nu\right) = \Omega_* \left(\frac{\nu}{\nu_b}\right)^{n_1} \left[\frac{1}{2} + \frac{1}{2}\left(\frac{\nu}{\nu_b}\right)^{a_1}\right]^{\frac{n_2-n_1}{a_1}} \equiv \Omega_*\, \psi_0(\nu;\nu_b)\,, \tag{2.36}$$

where we set $n_1 = 2.4$, $n_2 = -2.4$, and $a_1 = 1.2$ as it has been found from the numerical fit to a strong phase transition [213, 214], while we leave the overall energy density amplitude Ω_* and the characteristic frequency ν_b as free parameters. In our notation, this means that the theoretical parameters to be inferred are $\theta = \{\Omega_*, \nu_b, \kappa_1, \kappa_2\}$. We employ the same parametrisation to model the PBH background, which serves as our second toy model, with different values for $n_1 = 3$, $n_2 = -8$, and $a_1 = 2$. The choice of n_1 is motivated by the fact that for a physical GW energy density, a $\Omega_{\rm GW} \sim \nu^3$ scaling is expected at low frequencies [215].[1] The parameters n_2 and a_1 are chosen to mimic the high-frequency behaviour of the GW energy density displayed in Figure 2 of [216]. The perturbed $\Omega_{\rm GW}$ in (2.33) can be rewritten as

$$\Omega_{\rm GW}(\nu;\theta) \simeq \Omega_* \left[\psi_0(\nu;\nu_b) + \kappa_1\psi_1(\nu;\nu_b) + \kappa_2\psi_2(\nu;\nu_b)\right]\,. \tag{2.37}$$

Note that from equations (2.33) and (2.36), the functions $\psi_{1,2}(\nu;\nu_b) = \Omega_{1,2}(\nu;\nu_b)/\Omega_*$ do not depend on Ω_*. The plot of $\Omega_{\rm GW}(\nu;\nu_b)$ is shown in Figure 2.2 for different choices of Ω_*, ν_b, κ_1, and κ_2 for both models under consideration. Qualitatively, the presence of a diffusion term with $\kappa_1 > 0$ results in a blue-shift and broadening of the GW energy density, while $\kappa_2 > 0$ sharpens any peak present in the unperturbed spectrum.

[1] We note that the first-order phase transition energy density has a different low-frequency scaling since the numerical fit in [214] is performed on a simulation that does not take the Hubble expansion into account.

Figure 2.3 shows the full corner plot for the forecast posterior distribution $P(\Omega_*, \nu_b, \kappa_1, \kappa_2)$ for the theoretical parameters in the first-order phase transition and the PBH models. We assume flat priors on the parameters and impose that Ω_*, ν_b, and κ_1 are positive. We find that if the spectrum has a strong enough signal in the LISA frequency band, all the parameters can be constrained, and there are no fundamental degeneracies between them. The effect of the diffusion parameter κ_1 is anti-correlated with the characteristic frequency ν_b, as they both change the peak frequency. Similarly, κ_2 always shows some degree of anti-correlation with the overall GW energy density amplitude Ω_*, since they both enhance the height of the peak as they increase. The posterior distribution for the first-order phase transition case with $\bar{\nu}_b = 1\text{mHz}$ (top-left panel in Figure 2.3) and the posterior distribution for the PBH scenario with $\bar{\nu}_b = 5\text{mHz}$ (bottom-right panel) show stronger correlations between the parameters compared to the same models with $\bar{\nu}_b = 5\text{mHz}$ and $\bar{\nu}_b = 1\text{mHz}$ respectively. The reason for these correlations has to be attributed to the fact that the Fisher matrix is built by non-trivial combinations of $\psi_1(\nu;\nu_b)$, $\psi_2(\nu;\nu_b)$, and $\frac{\partial\psi_0(\nu;\nu_b)}{\partial\nu_b}$, along with the instrumental noise curve. We emphasize that even if the GW background may not have a peak that falls perfectly in LISA's most sensitive frequency band, any deviation from a power-law behaviour in the unperturbed spectrum can produce spectral features that allow the inference to break the degeneracy between the newly introduced diffusion and drift parameters $\kappa_{1,2}$ and the ones used to model the unperturbed GW background, as the perturbations depend on the derivatives in frequency of the unperturbed spectrum.

From Figure 2.3 we find that the parameters $\kappa_{1,2}$ can be constrained at the 1σ level to $\kappa_1 \leq 1.8 \times 10^{-56}\text{kg}\,\text{m}^2\,\text{s}^{-3}$ and $\kappa_2 \leq 3.7 \times 10^{-56}\text{kg}\,\text{m}^2\,\text{s}^{-3}$ under the detection of a first-order phase transition GW background, and to $\kappa_1 \leq 6.5 \times 10^{-60}\text{kg}\,\text{m}^2\,\text{s}^{-3}$ and $\kappa_2 \leq 2.6 \times 10^{-59}\text{kg}\,\text{m}^2\,\text{s}^{-3}$ in the PBH scenario. It is important to note that the first-order phase transitions and PBH backgrounds considered in [213] and [216], respectively, have relatively high amplitudes, while other scenarios may produce a possibly smaller signal. An inspection of the Fisher matrix reveals the rough dependence of the 1σ constraint on $\kappa_{1,2}$ in terms of Ω_*, provided that the peak of the signal stands in the LISA sensitivity band and is not excessively broad. This is given by

$$\boxed{\sigma_{\kappa_1} \simeq \sigma_{\kappa_2} \simeq \left(\frac{10^{-11}}{\Omega_*}\right) 10^{-56}\text{kg}\,\text{m}^2\,\text{s}^{-3}}\,. \tag{2.38}$$

2.5 Summary and outlook

Assuming Lorentz invariance to be an exact symmetry of quantum gravity that is also preserved effectively at low energies, we analyzed the covariant diffusion model of massless particles. The model has already been applied to photons and constrained via CMB blackbody spectrum [60]. Here we solved the modified Boltzmann equation analytically at the background FRW level and showed explicitly that the effect dominates at low frequencies. This contrasts with the intuitive expectation that high-energy particles should be more sensitive to new Planck-scale physics. Nevertheless, this low-frequency dominance arises as a consequence of Lorentz invariance. As a result, quantum gravity phenomenology in this case would manifest at the largest scales.

We then performed a Fisher forecast on the two parameters of the model for LISA. Assuming a bright

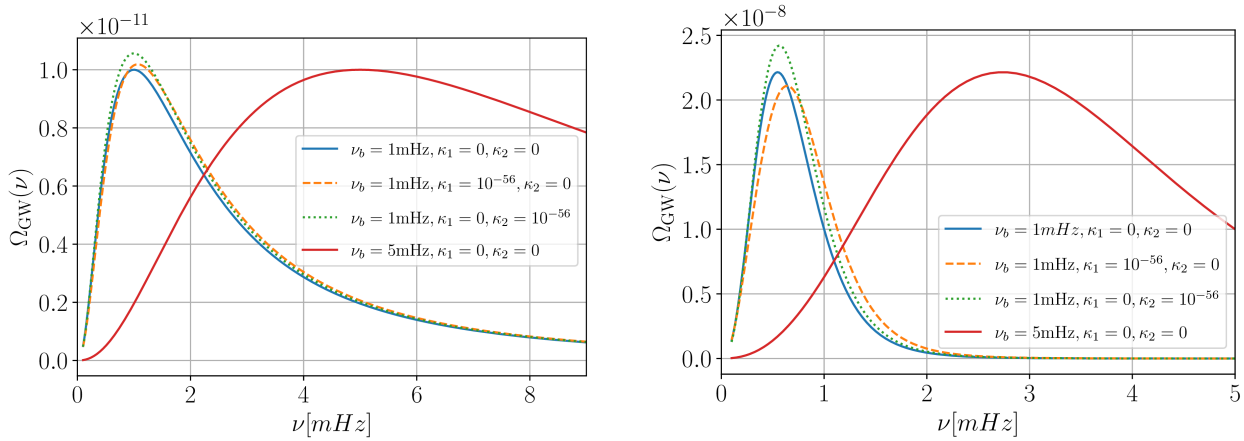


Figure 2.2: Plot of the GW energy density $\Omega_{\rm GW}(\nu)$ of a stochastic background from a first-order phase transition (left) and from primordial black holes (right). We show in dashed lines how the diffusion effect broadens and shifts the unperturbed spectra for different values of $\kappa_{1,2}$ in $\mathrm{kg\,m^2\,s^{-3}}$. We parametrise the unperturbed spectra as broken power laws as in Equation (2.36), with different fiducial values for the parameters.

enough background signal either from first-order phase transitions or second-order gravitational waves from large scalar perturbations, we found that we will be able to put stringent bounds on the diffusion and drift parameters of the model. This opens the possibility of detecting the covariant diffusion of gravitons with diffusion and drift constants beyond the sensitivity of CMB experiments. Note, however, that the diffusion and drift parameters might as well differ between species of massless particles, or even depend on Lorentz-invariant properties of a given waveform [217]. It is therefore an interesting problem for future work to go beyond the point-particle or geometric optics limit and model the covariant diffusion of scalar or higher-spin waves; see [218, 219] for some relevant progress. This will shed light on how $\kappa_{1,2}$ could depend on waveform properties and possibly the microscopic physics that is responsible for the diffusion. Ultimately, the interpretation of a bound like (2.38) will depend on the specific quantum gravity framework that gives rise to the diffusion.

It is important to remark that the present analysis with LISA assumes that all the other sources of resolved and unresolved GW signals are perfectly fitted in the datastream, the GW background is stationary and isotropic, and that the instrumental noise is Gaussian with a well-characterised power spectral density. Even if we expect that some complications will be present in the LISA data analysis (along with several other experimental challenges), we would like to insist on the fact that the effect of graviton diffusion on the GW spectrum can be easily taken into account, as it depends on just the first two derivatives in frequency of the unperturbed graviton distribution function and introduces just two phenomenological parameters. Any GW background with a non-trivial frequency dependence is distorted by these diffusion effects in a very specific way, producing distinct signatures that are very difficult to mimic with some features of the unperturbed model.

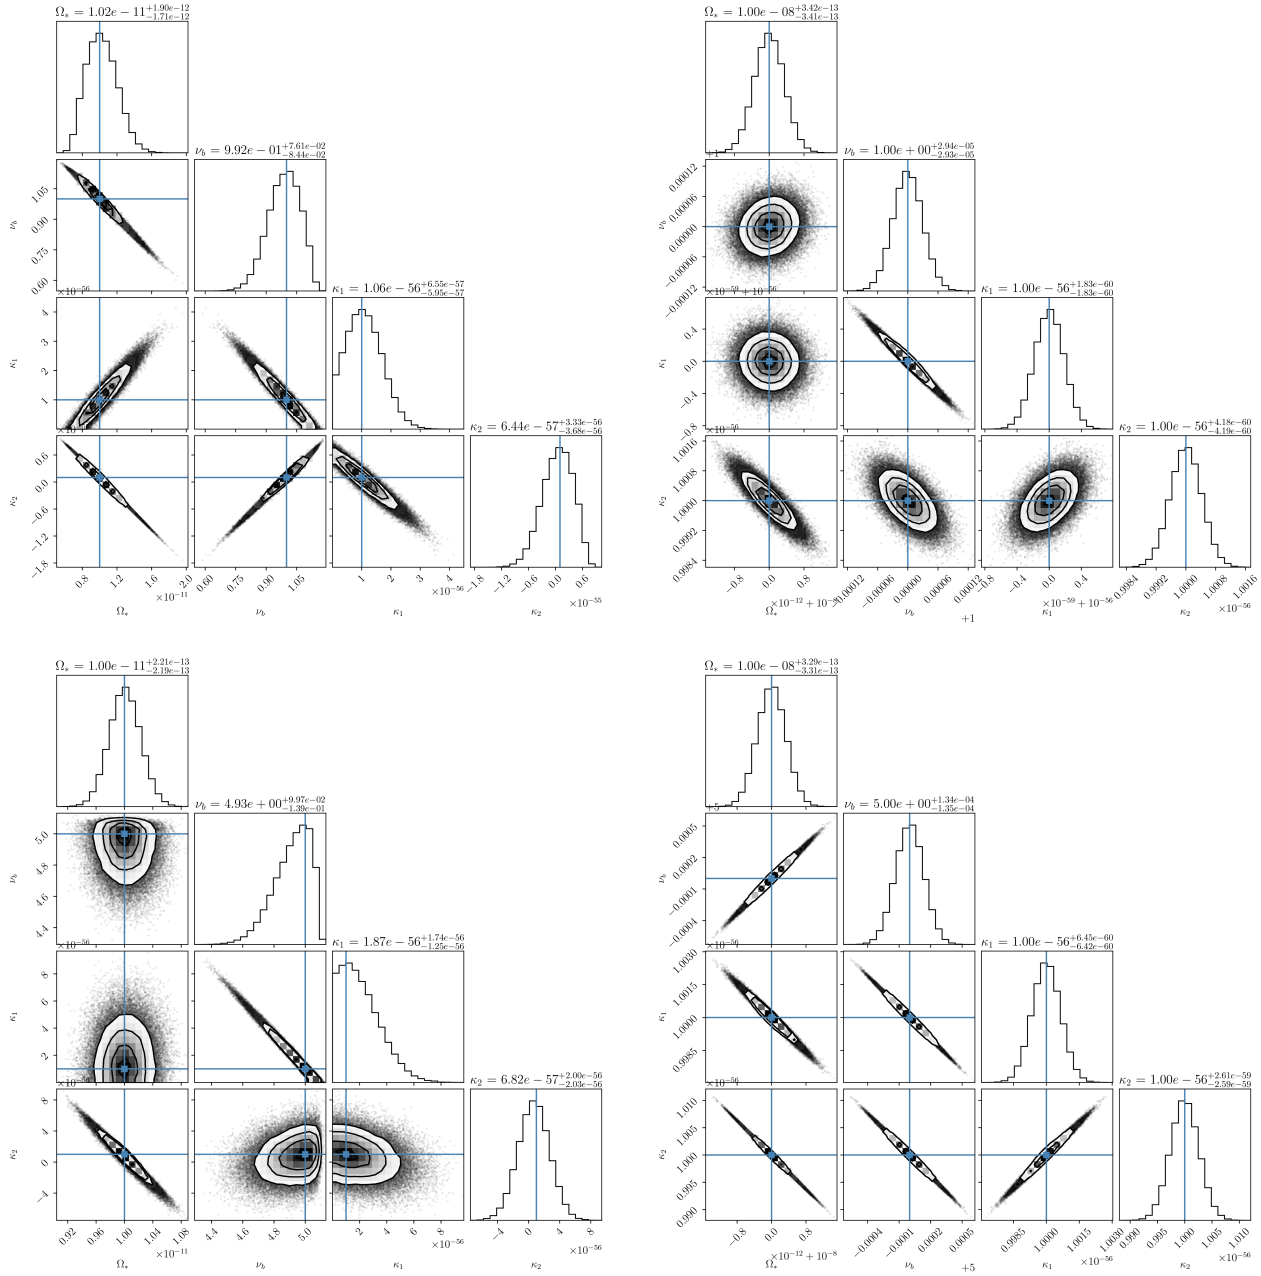


Figure 2.3: Corner plot of the forecast full posterior distribution $P(\Omega_*, \nu_b, \kappa_1, \kappa_2)$ assuming a flat prior and an effective total time of observation with LISA of $T_{\rm obs} = 5$ years. Top left: a first-order phase transition scenario with fiducial values $\bar{\Omega}_* = 10^{-11}$ and $\bar{\nu}_b$ =1mHz. Top right: a PBH scenario with fiducial values $\bar{\Omega}_* = 10^{-8}$ and $\bar{\nu}_b$ =1mHz. Bottom left: a first-order phase transition scenario with fiducial values $\bar{\Omega}_* = 10^{-11}$ and $\bar{\nu}_b$ =5mHz. Bottom right: a PBH scenario with fiducial values $\bar{\Omega}_* = 10^{-8}$ and $\bar{\nu}_b$ =5mHz. Fiducial values of $\bar{\kappa}_1 = 10^{-56}\mathrm{kg\,m^2\,s^{-3}}$ and $\bar{\kappa}_2 = 10^{-56}\mathrm{kg\,m^2\,s^{-3}}$ are assumed in all the four cases.

Chapter 3

Stochastic dark energy: introduction to Everpresent Λ

In recent years, cosmological model building has taken on a new flavour, in part following some perplexing observational findings. In addition to the Hubble tension, baryon acoustic oscillation (BAO) data [220, 221] appears to allow for a negative cosmological constant [222–226] for the relevant epoch. See also the latest DESI results [227, 228] which seem to favour an evolving dark energy over a cosmological constant. As a result, models with a variable cosmological "constant" are being considered more seriously. Some of these examples include [62, 229, 230].

In this chapter we will be considering one of the earliest fluctuating cosmological constant models, known as Everpresent Λ. Everpresent Λ dates back to 1987 and originates from a heuristic argument by Sorkin [61, 63, 231] based on unimodular gravity and causal set theory. This model predicted that the value of the cosmological constant should fluctuate, but with a standard deviation that at any time is of the order of the total energy density in the Universe at that time. This prediction was verified for the current epoch with the supernova observations made in 1998. It was seen that, at present, the contribution of dark energy is of the order of the matter density. Subsequently, more practical phenomenological models have been proposed for producing full expansion histories that have the character of Everpresent Λ. The first such cosmological model based on Sorkin's argument was proposed in [64]. This model has further been studied in [65, 66]. A second model was also proposed in [66], which is computationally more efficient, but has much of the characteristics of the first phenomenological model.

The main inputs to the Everpresent Λ proposal are the assumptions that (1) spacetime volume plays the role of time in a path integral based dynamics, and (2) there is a statistical correspondence between the number of elements in the discrete causal set and the spacetime volume of the continuum spacetime which approximates it. From these assumptions, a simple relation can be drawn between the value of the cosmological constant at each epoch and the volume of the cosmos at that epoch, consistent with the value observed today. The heuristic argument has since been fleshed out in subsequent works. For example, in [64], the first concrete model with the basic features of Everpresent Λ in an FRW universe was constructed and studied via simulations; the model that we study in this chapter is essentially the

same as this original one. In [65] the incorporation of various alternative combinations of the Friedmann equations into the model in [64] was explored; they found that similar results can be obtained under a range of combinations. In [66], a simpler phenomenological model (dubbed "Model 2") of Everpresent Λ based on a Gaussian process, inspired by the first model but easier to simulate, was studied. In [232] and [233], observational constraints were placed on the amplitude of fluctuations, under some assumptions; these studies provide insight into what we should aim to change in future models of Everpresent Λ. Finally, in [234], several avenues for incorporating inhomogeneities were discussed.

Despite progress through the works mentioned above, work on Everpresent Λ has largely stayed close to the original heuristic argument and early models based on it. Causal set theory itself is a minimalistic approach to quantum gravity, with causal structure and fundamental discreteness as its key inputs. Everpresent Λ in turn, as mentioned above, uses basic properties of causal set theory as input. On the one hand, it is rather remarkable that such simple ingredients can be used to naturally explain the otherwise unnatural value of the cosmological constant. On the other hand, the simplicity of the ingredients and theoretical context are what also make it challenging to refine and build more complex models of Everpresent Λ which can confront high precision cosmological data. Nevertheless, phenomenologically at least, we may draw on ideas outside of causal set theory to supplement and refine the Everpresent Λ model. First, however, it is essential to better understand the model in its current state.

The aim of this chapter is to take steps towards placing Everpresent Λ on firmer footing. We specifically do this by expanding on some aspects of the original heuristic argument, and by presenting a detailed study of a phenomenological model based on it. We begin in Section 3.1 with a review of the main two background topics of causal set theory and unimodular gravity. In Section 3.2 we present the Everpresent Λ idea and a phenomenological implementation of it. Further explanation of our methodology is described in Section 3.3. We carry out a detailed study of the phenomenological model presented in Section 3.2 and the results, including the statistical distribution of the present day Hubble parameter and the auto-correlation function of Ω_Λ, are summarised in Section 3.4. We end in Section 3.5 with a discussion of open questions and future directions of this program.

3.1 Background theory

3.1.1 Causal set theory

Causal set theory is an approach to quantum gravity that posits that spacetime is fundamentally discrete [89–91] (see also [92–95]). This introduces a minimum length scale that naturally resolves the UV divergences of quantum field theories. However, at macroscopic scales, this discreteness is small enough that spacetime appears to be continuous.

A causal set $\mathcal{C}$ at its most basic level is a discrete spacetime with the following properties: First, $\mathcal{C}$ is a partially ordered set where some elements of $\mathcal{C}$ may causally precede other elements of $\mathcal{C}$. In particular, this means that if element x precedes element y ($x \prec y$), then y lies in the future light cone of x. Second, the *transitivity* of the partial ordering means that if $x \prec y$ and $y \prec z$, then $x \prec z$. In relativistic language this means that if $x \prec y$, the future light cone of y is a subset of the future light cone of x. Furthermore, the *asymmetry* of the partial ordering means that one cannot have both $x \prec y$ and $y \prec x$.

This, together with the transitivity property, implies that a spacetime cannot have closed causal loops. Finally, the causal set has to be *locally finite*: Between any two elements x and z, there are a finite number of elements y such that $x \prec y \prec z$. This encodes the fundamental discreteness of spacetime in causal set theory.

The above properties describe a discrete spacetime with the basic causal structure we know from relativity. However, the spacetime of general relativity (GR) has the further properties that it is a differentiable manifold and that it has local Lorentz symmetry, i.e., that it looks locally the same in any inertial frame. The vast majority of causal sets with the properties enumerated in the previous paragraph, do not resemble a continuum manifold [235]. Therefore we need additional constraints to ensure that a causal set is manifoldlike and physical. While ideally the fundamental laws of causal set theory governing the dynamics (a work in progress) would generate spacetimes with the desired properties, in the meantime we can generate causal sets that closely approximate a continuum spacetime through a process called *sprinkling*.

Sprinkling is the sampling of points from a volume of a Lorentzian spacetime manifold $\mathcal{M}$. We would like the sampled points to satisfy the following condition: On average, the number of sprinkled points (or elements) N in a region should be proportional to the region's spacetime volume V. This is known as the number-volume correspondence. The proportionality factor is denoted by $\varrho_{cs} = \langle N \rangle / V \equiv 1/\ell_{cs}^d$, where ℓ_{cs} is the discreteness scale which is believed to be close to the Planck length, and d is the spacetime dimension which we will take to be 4 henceforth. In the limit $N \to \infty$ (or similarly $\ell_{cs} \to 0$), we should retrieve the continuum limit. This number-volume correspondence for arbitrary volumes satisfies the requirement that the sprinkled points should preserve the local Lorentz symmetry of the spacetime [236, 237].

A sampling that satisfies the above condition is achieved through a Poisson process. In a (Poisson) sprinkled region with volume V, the probability of having N elements is

$$\mathbb{P}(N) = \frac{1}{N!} e^{-\varrho_{cs} V} (\varrho_{cs} V)^N \,, \tag{3.1}$$

i.e., the number of elements in a volume is given by the Poisson distribution. These elements form a causal set with the relations between its elements induced by the causal structure of the underlying continuum spacetime. As desired, the mean is $\langle N \rangle = \varrho_{cs} V$, which is the statement of the number-volume correspondence, and the standard deviation $\delta N = \sqrt{\varrho_{cs} V}$, which is also the smallest possible deviation [238] from an exact number-volume correspondence.

Conversely, given a causal set $\mathcal{C}$, a Lorentzian manifold is said to be a good continuum approximation of it if $\mathcal{C}$ is a typical outcome of a Poisson sprinkling on the manifold [236]. Lorentzian manifolds that approximate $\mathcal{C}$ (with number of elements N) have a volume $V = \frac{1}{\varrho_{cs}}(N \pm \sqrt{N})$, i.e.,

$$\delta V = \frac{\sqrt{N}}{\varrho_{cs}} \,. \tag{3.2}$$

The uncertainty arises when replacing N by $\varrho_{cs} V$ and vice versa in the expression for the Poisson distribution in (3.1), and then using the known formulas for the mean and standard deviation.

So far, we have only discussed purely kinematic aspects of causal sets. Also relevant to the discussion of Everpresent Λ is the dynamics. We could envisage a causal set growing by new elements being born sequentially. Each new element will have some causal relations to the already existing elements. The total number of causal set elements, which corresponds to the spacetime volume in the continuum, can thus be viewed as an external time parameter appropriate for defining the dynamics (e.g. as a constraint in a path integral). The clock ticks once every time a new element is born. Such models were studied in [239, 240].

On the other hand, given a causal set $\mathcal{C}$, we can define the Benincasa-Dowker-Glaser action $S_{\rm BDG}[\mathcal{C}]$ which sums over the numbers of causal intervals of certain sizes in $\mathcal{C}$. It has been shown that in the continuum limit, i.e. $\lim_{\rho_{cs}\to\infty} S_{\rm BDG}[\mathcal{C}]$, this action tends to the Einstein-Hilbert action $S_{\rm EH}$. $S_{\rm BDG}[\mathcal{C}]$ was first proposed in $3+1$D by [241] and was later extended to other dimensions in [242].

In analogy with gravitational path integrals, we can use $S_{\rm BDG}$ to define a path integral for causal sets. Instead of integrating over different spacetime configurations (e.g. those that extend between two 3-geometries), we sum over all causal sets with a given number of elements [6]

$$\mathcal{Z}(N) = \sum_{|\mathcal{C}|=N} e^{iS_{\rm BDG}[\mathcal{C}]} \,. \tag{3.3}$$

This form of the path integral (or sum over histories) for causal sets will be our motivation to introduce unimodular gravity in the next subsection.

3.1.2 Unimodular gravity

Unimodular gravity was originally proposed as a mild modification of GR with the potential to solve the old cosmological constant problem, namely to address why the vacuum energy of field theories does not seem to gravitate [63, 243–250]. As discussed in the previous subsection, holding the number of elements fixed is the natural constraint to impose on a path integral for causal sets (c.f. (3.3)). Motivated by this, let us consider the continuum gravitational path integral while keeping fixed the total spacetime volume V. In other words, to obtain the path integral with the initial and final spatial hypersurfaces Σ_i and Σ_f, we integrate over all compact 4-manifolds $\mathcal{M}$ that connect Σ_i to Σ_f and satisfy $\mathrm{Vol}(\mathcal{M}) = V$:

$$\mathcal{Z}_{\rm UG}^{global}(V) = \int_{\mathrm{Vol}(\mathcal{M})=V} \frac{\mathcal{D}g_{\mu\nu}}{\mathrm{Vol[Diff]}}\, e^{iS[g]} \,, \tag{3.4}$$

where $S[g]$ is the Einstein-Hilbert action plus matter fields and vacuum energy $\tilde{\Lambda}$. The gauge group consists of diffeomorphisms that keep the total volume of a given $(\mathcal{M}, g_{\mu\nu})$ unaltered. However, the total volume is an invariant under gauge transformations (in terms of a linear gauge transformation $x^\mu \to x^\mu + \xi^\mu$, $g_{\mu\nu} \to g_{\mu\nu} + \nabla_\mu\xi_\nu + \nabla_\nu\xi_\mu$, this is just the statement that $\int d^4x\sqrt{-g}\ \nabla^\mu\xi_\mu = 0$, which is trivial if ξ^μ is zero on the boundary). That is why we have divided by the volume of the diffeomorphism group, Vol[Diff]. The superscript "global" refers to the fact that only the total volume of the spacetime $\mathcal{M}$ is constrained. In contrast to the global version (see [63, 251, 252] for discussions on global unimodular gravity), there is also a local version of the theory where the volume measure

$d^4x\sqrt{-g}$ is constrained to be locally equal to some (nondynamical) background 4-form ω. This is the formulation of unimodular gravity that is usually discussed in the literature. It posits a seemingly stronger constraint on the metric, and reduces one of its degrees of freedom at each spacetime point. At the classical level of local unimodular gravity, it is well known (see e.g. [245, 253, 254]) that the variational principle for the constrained metric only gives the traceless Einstein equations. We shall see that the same is true for global unimodular gravity, and then conclude that at the classical level, local and global unimodular gravity are equivalent (even in their treatment of the cosmological constant).

We can implement the volume constraint in (3.4) as an additional term in the action with a Lagrange multiplier σ, such that the action becomes

$$S[g;\sigma] = S[g] - \sigma\left(\int d^4x\sqrt{-g} - V\right) . \tag{3.5}$$

To take into account this new term, we must integrate over the Lagrange multiplier such that the path integral is expressed as

$$\mathcal{Z}_{\text{UG}}^{global}(V) = \int \frac{\mathcal{D}g_{\mu\nu}}{\text{Vol[Diff]}}\frac{d\sigma}{2\pi}\, e^{iS[g]-i\sigma\left(\int d^4x\sqrt{-g}-V\right)} . \tag{3.6}$$

Note that by performing the σ integral in the path integral above, one obtains a delta function containing the volume constraint, which is how we will treat it in the next section.

Returning our attention to (3.5), we can vary the action by varying the metric and Lagrange multiplier (after a redefinition of σ by a factor of $8\pi G$) as

$$\delta S = \frac{1}{16\pi G}\int d^4x\sqrt{-g}\;\delta g^{\mu\nu}\left(E_{\mu\nu} + \sigma g_{\mu\nu}\right) - \frac{1}{8\pi G}\delta\sigma\left(\int d^4x\sqrt{-g} - V\right) , \tag{3.7}$$

where $E_{\mu\nu} = G_{\mu\nu} - 8\pi G T_{\mu\nu} + \tilde{\Lambda}g_{\mu\nu}$, and the term containing $E^{\mu\nu}$ is the variation of $S[g]$. Setting the variations with $\delta g_{\mu\nu}$ and $\delta\sigma$ to zero independently, we obtain the equations of motion

$$G_{\mu\nu} - 8\pi G T_{\mu\nu} + \tilde{\Lambda}g_{\mu\nu} + \sigma g_{\mu\nu} = 0 , \tag{3.8}$$

and the metric constraint equation

$$\int d^4x\sqrt{-g} = V , \tag{3.9}$$

respectively. Taking the trace in (3.8), we obtain

$$-R - 8\pi G T + 4\tilde{\Lambda} + 4\sigma = 0 . \tag{3.10}$$

Solving for σ and substituting it back into (3.8), we see that the $\tilde{\Lambda}$ term is cancelled, and we are left with

$$G_{\mu\nu} - 8\pi G T_{\mu\nu} + \frac{1}{4}g_{\mu\nu}(R + 8\pi G T) = 0 , \tag{3.11}$$

which is the trace-free form of the Einstein equations. The full Einstein equations can be recovered by taking the covariant derivative of (3.11), as for example shown in [245] for local unimodular gravity. Since σ and $\tilde{\Lambda}$ are both constant, so must $\frac{1}{4}(R + 8\pi G T)$ be, which we can therefore rename to be our

effective cosmological constant Λ. We therefore have

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu} \,. \tag{3.12}$$

This Λ is an integration constant, and is unrelated to any coupling constant in the action or the vacuum energy. (Note that we initially separate any vacuum energy from $T_{\mu\nu}$ and add its contribution to $\tilde{\Lambda}$, so that T does not include vacuum energy). This is how unimodular gravity offers a resolution to the old cosmological constant problem. For an alternative derivation of these equations, see Appendix C.1.

Another way to see the same result, namely that $\tilde{\Lambda}$ disappears from all physical quantities, is to note that the path integral with external source $J^{\mu\nu}$ is written as

$$\mathcal{Z}_{\rm UG}^{global}(V, J^{\mu\nu}) = \int_{{\rm Vol}(\mathcal{M})=V} \frac{\mathcal{D}g_{\mu\nu}}{{\rm Vol[Diff]}} \; e^{iS[g]+i\int d^4x\sqrt{-g}\; g_{\mu\nu}J^{\mu\nu}} \,. \tag{3.13}$$

Now if we separate the vacuum energy term from the Einstein-Hilbert term and the matter action as $S = S_{\rm EH} + S_{\rm m} - \frac{\tilde{\Lambda}}{8\pi G}\int d^4x\sqrt{-g}$ and use the condition $\int d^4x\sqrt{-g} = V$ of the path integral, we can simply take the non-geometric phase factor out of the integral

$$\mathcal{Z}_{\rm UG}^{global}(V, J^{\mu\nu}) = e^{-i\tilde{\Lambda}V/(8\pi G)} \int_{{\rm Vol}(\mathcal{M})=V} \frac{\mathcal{D}g_{\mu\nu}}{{\rm Vol[Diff]}} \; e^{iS_{\rm EH}[g]+iS_{\rm m}[g]+i\int d^4x\sqrt{-g}\; g_{\mu\nu}J^{\mu\nu}} \,. \tag{3.14}$$

Physically relevant quantities, including n-point functions of the metric (if one could make sense of them in quantum gravity), are computed by functionally differentiating $\mathcal{Z}_{\rm UG}^{global}(V, J^{\mu\nu})$ with respect to the external source and then normalising by dividing by $\mathcal{Z}_{\rm UG}^{global}(V, 0)$ [255]. Therefore the factor $e^{-i\tilde{\Lambda}V/(8\pi G)}$ cancels out, and everything is as if inside the path integral we are using the action $S_{\rm EH+m} = S_{\rm EH} + S_{\rm m}$ without any vacuum energy. From now on, therefore, we will assume that the action does not include a vacuum energy term.

As is evident from the above discussion, unimodular gravity does not provide any insight on the new cosmological constant problem or on the coincidence problem. In the context of unimodular gravity, why the observed dark energy density is of the order of the present critical density remains an open problem. Everpresent Λ brings insights from causal set theory, which together with unimodular gravity, fill this gap and provide an explanation for the observed value of the dark energy density.

3.2 Everpresent Λ

We now review how Everpresent Λ follows from aspects of causal set theory and unimodular gravity presented in the previous section. A sketch of our discussion below, as well as several alternative arguments leading to same conclusion, can be found in [256].

The path integral of global unimodular gravity (3.4) can be rewritten to incorporate the volume constraint using a delta function

$$\mathcal{Z}_{\rm UG}^{global}(V) = \int \frac{\mathcal{D}g_{\mu\nu}}{{\rm Vol[Diff]}} \; \delta(V - \int d^4x\sqrt{-g}) \; e^{iS_{\rm EH+m}[g]} \,. \tag{3.15}$$

The above partition function depends explicitly on V, the total volume of spacetime. Now instead of working in "volume space", let us perform a Fourier transform and express the partition function in the conjugate "ϖ space" as

$$\mathcal{Z}(\varpi) = \int dV \; e^{-i\varpi V} \mathcal{Z}_{\rm UG}^{global}(V) \,. \tag{3.16}$$

At this point the symbol "ϖ" is an arbitrary symbol for the conjugate Fourier transform variable. There is no need to worry about the physical meaning of the negative values for V inside the integral (3.16) because (3.15) has a delta function at a positive value. Performing this delta function integration, and redefining ϖ with a factor of $8\pi G$, we find

$$\mathcal{Z}(\varpi) = \int \frac{\mathcal{D}g_{\mu\nu}}{\text{Vol[Diff]}} \; \exp\left(iS_{\rm EH+m}[g] - \frac{2i\varpi}{16\pi G} \int d^4x\sqrt{-g} \right) . \tag{3.17}$$

The Fourier parameter ϖ has turned out to play the role of the cosmological constant Λ, and $\mathcal{Z}(\varpi)$ is just the partition function of GR. The equation of motion coming from $\mathcal{Z}(\varpi)$ is exactly the same as (3.12), so we are justified to identify $\varpi = \Lambda$. We infer that the 4-volume V and the cosmological constant Λ are quantum mechanically conjugate. [1] This is not a rigorous derivation, as there are many subtleties that require a full theory of quantum gravity to appreciate. However, we see it as a well-motivated conjecture to expect from a quantum unimodular gravity. The uncertainty principle for these conjugate variables is

$$\frac{\delta\Lambda}{8\pi G} \,.\, \delta V \geq \frac{\hbar}{2} \,. \tag{3.18}$$

Naturally one would ask what the standard deviations δV and $\delta\Lambda$ mean. Recall from (3.3) that from the point of view of a fundamental causal set path integral, we should hold the total number of elements N fixed. A consequence of this for the effective continuum picture we have been working with in this section is that we must consider spacetime volumes V that are consistent with a causal set with N elements. Therefore V is no longer held strictly fixed as it was before, in our review of conventional unimodular gravity. We saw in the arguments leading to (3.2) that such volumes would be Poisson distributed with mean $\langle V \rangle = N/\varrho_{cs}$ and standard deviation $\delta V = \sqrt{N}/\varrho_{cs}$. We can interpret this as a wave function distribution, where the universe is in a superposition of different volumes with spread δV. Note that in allowing ourselves to work with a continuum wave function, we have implicitly assumed that we are in a regime where the number-volume correspondence makes sense. In particular, this means that we are in a macroscopic regime and not too far away from a causal set approximated by a classical geometry.

For $N, V \gg 1$ (which is indeed the case we are interested in), the Poisson distribution is very well approximated by a normal distribution $\mathcal{N}(N/\varrho_{cs}, \sqrt{N}/\varrho_{cs})$. Moreover, if our universe is in a Gaussian wave function in V-space, it will also be described by a Gaussian wave function in Λ-space (through a Fourier transform). Since Gaussian wave functions saturate the uncertainty principle, we expect the state of the universe to be in a superposition of different Λ values with spread

[1] If by conjugacy one means the existence of a commutation relation like $[\Lambda, V] \sim 1$ in a canonical analysis, then there is currently no proof available. As an example of what has been done in this direction, we refer the reader to [257], in which a BRST quantization of local unimodular gravity was done. In that case, the metric was expanded around some background, and the fluctuations $h_{\mu\nu}$ were quantized. Commutation relations such as $[\varpi(p), (h^{00} + h^{33})(q)] \sim \delta(p-q)$ were obtained. However, this is still some way off from the desired commutation relations between ϖ and a quantized geometric quantity such as V. A better understanding of quantum gravity may be needed to derive more general relations of this kind.

$$\frac{\delta\Lambda}{8\pi G} = \frac{\hbar}{2}\frac{1}{\delta V}\,. \tag{3.19}$$

Finally, using (3.2) and the relation between number and volume, we find

$$\delta\Lambda = 4\pi\left(\frac{\ell_p}{\ell_{cs}}\right)^2 \frac{1}{\sqrt{V}}\,. \tag{3.20}$$

We used $\ell_p = \sqrt{G\hbar}$ as the Planck length and $\ell_{cs} = \varrho_{cs}^{-1/4}$ for the discreteness scale (in 4 dimensions) of causal sets.

The standard deviation of Λ in (3.20) can be estimated for our universe by assuming the discreteness length to be of order of the Planck length, and setting V to the volume of the observable universe.

By dimensional analysis, in an FRW universe this volume should be of order H^{-4} in fundamental units, H being the Hubble parameter. In our background solutions for the Friedmann equations in the next section, we will see that this is indeed the case. Roughly then,

$$\delta\Lambda \sim H^2\,. \tag{3.21}$$

We will assume that the mean value about which Λ fluctuates is zero.[2] This assumption does not stem from any fundamental aspect of causal set theory or unimodular gravity and will require justification in the future. See [258–260] for some possible arguments in favour of this, using ideas from Euclidean quantum gravity. See also footnote 4 in [64] and towards the end of [234] for a Lorentzian guess (involving a non-manifoldlike phase of the universe) at how this may come about. Therefore, $\Lambda = \langle\Lambda\rangle \pm \delta\Lambda \sim H^2$, which is indeed a correct estimate for the current value of Λ. Furthermore, this is telling us that Λ is not temporally constant in our picture. It fluctuates such that at any time in the expansion history of the universe, its standard deviation is of the order of $1/\sqrt{V}$ (in FRW this is comparable to the critical energy density at that time); hence the name *Everpresent* Λ, as it will always be relatable to the physical scale of the universe in this way.

3.2.1 A phenomenological implementation—Model 1

Perhaps the closest analogue to the $\delta\Lambda\delta V$ uncertainty relation is the energy-time uncertainty in quantum mechanics. It has been suggested [63, 261, 262] that in quantum gravity, an increasing spacetime volume V (such as the volume of the past light cone of an observer) could serve as a *time* variable to treat the seemingly timeless nature of GR and the Wheeler-DeWitt equation.

Just as the $\delta E\delta t$ uncertainty in quantum mechanics alone cannot provide a dynamical model for energy non-conservation in short time intervals, the $\delta\Lambda\delta V$ uncertainty relation does not provide a practical model for how to evolve Λ. In [64] the authors suggested a phenomenological model (referred to as

[2] This is a consistent choice since the mean of Λ is arbitrary (or at least undetermined) in unimodular gravity. It remains the case even if one takes into account the vacuum energy $\tilde{\Lambda}$ (from the cosmological constant term in $S_{\rm EH}$). To see this, we can absorb $\tilde{\Lambda}$ into the definition of the Fourier variable ϖ in (3.17). This amounts to a constant shift in the effective cosmological constant $\Lambda = \tilde{\Lambda} + \langle\varpi\rangle + \delta\varpi$. But since $\langle\varpi\rangle$ is arbitrary, we can choose it so that the shifted mean $\tilde{\Lambda} + \langle\varpi\rangle$ is 0. A more complete theory should advise the correct mean for Λ. The fluctuations $\delta\varpi$ remain unaffected by the vacuum energy.

"Model 1" in [66]) for the dynamics of Λ that shows the expected fluctuations in (3.20). Their strategy was to build a classically stochastic model for the fluctuations in Λ, working with a single universe rather than a superposition. We will now review this model whose properties will be our focus for the remainder of this chapter.

From the causal set point of view, and due to the number-volume correspondence, (3.20) is equivalent to

$$\delta\Lambda = \frac{8\pi\alpha}{\ell_{cs}^2}\frac{1}{\sqrt{N}}\,, \tag{3.22}$$

where we have defined

$$\alpha \equiv \frac{1}{2}\left(\frac{\ell_p}{\ell_{cs}}\right)^2\,. \tag{3.23}$$

We next make the ansatz that (3.22) is true even element-wise, in the sense that each causal set element contributes a random variable with standard deviation $8\pi\alpha/\ell_{cs}^2$ to Λ. The simplest such scenario would be to think of the single element contribution to Λ as a Bernoulli random variable:

$$\Lambda|_{1\text{ element}} = \Lambda|_x = \begin{cases} \frac{8\pi\alpha}{\ell_{cs}^2}\,, & \text{probability} = \frac{1}{2} \\ -\frac{8\pi\alpha}{\ell_{cs}^2}\,, & \text{probability} = \frac{1}{2} \end{cases}\,. \tag{3.24}$$

(In fact any other distribution with mean 0 and standard deviation $8\pi\alpha/\ell_{cs}^2$ would work just as well for our purposes). The cosmological constant part of the action is obtained by summing over all the above stochastic variables in the N-element (sub)causal set corresponding to the region with volume V:

$$S_\Lambda = \frac{\ell_{cs}^4}{8\pi\ell_p^2}\sum_{x\in C_V}\Lambda|_x\,. \tag{3.25}$$

We have replaced the integral $\int d^4x\sqrt{-g}$ with the discrete sum $\ell_{cs}^4\sum_x$. Also, a minus sign has been absorbed into the definition of the cosmological constant. Whatever the distribution of $\Lambda|_x$ is, if its standard deviation is $8\pi\alpha/\ell_{cs}^2$, S_Λ will follow a normal distribution $\mathcal{N}(0,\sqrt{N}/2)$ according to the central limit theorem (recall that $N, V \gg 1$).

As already mentioned, by summing over the elements x corresponding to a given region of spacetime in (3.25), we obtain an effective Λ for the region. So far we have not discussed how to choose this region. We will follow [64] and choose it at any time of interest to be the volume V of the past light cone of a point x at that time. By making this choice, not only is V causally defined, it is also conveniently parametrised by the point as $V(x)$. In the same way we can also express the effective Λ's volume-dependence as a function $\Lambda(x)$ (not to be confused with the element-wise $\Lambda|_x$).

We interpret $\Lambda(x)$ that appears in (3.17) then as the effective cosmological constant that is ascribed to the past light cone of the point x such that

$$S_\Lambda|_{\text{past of } x} = \frac{1}{8\pi G}\Lambda(x)V(x) = \frac{\ell_{cs}^2}{16\pi\alpha}\Lambda(x)N(x)\,. \tag{3.26}$$

$V(x)$ is the volume of the past light cone at x, and $N(x)$ is the corresponding number of causal set

elements. Therefore, the distribution of $\Lambda(x)$ turns out to be

$$\mathcal{N}(0, \frac{8\pi\alpha}{\sqrt{V}}) \,. \tag{3.27}$$

Thus we have provided a stochastic dynamics that satisfies (3.20), as desired.

Now we can obtain a recursive formula for $\Lambda(x_n)$. Imagine we have found $S_\Lambda|_{\text{past of } x_{n-1}}$ and $\Lambda(x_{n-1})$ for some point x_{n-1} in the past of x_n, close enough to x_n to capture the variations in $\Lambda(x)$. The additional contribution to $S_\Lambda|_{\text{past of } x_n}$ compared to $S_\Lambda|_{\text{past of } x_{n-1}}$ is from all random variables $\Lambda|_z$ such that z is in the past of x_n but not in the past of x_{n-1}. The number of such points is proportional to the volume they occupy, $\Delta V_n = V(x_n) - V(x_{n-1})$, and the sum of their contributions to S_Λ is, again from the central limit theorem, a random variable $\sim \sqrt{\Delta V_n}/(2\ell_{cs}^2)\mathcal{N}(0,1)$. Therefore if we generate a standard normal random variable ξ_n, we can write

$$\Lambda(x_n) = \frac{8\pi\ell_p^2 S_\Lambda|_{\text{past of } x_{n-1}} + 8\pi\alpha\sqrt{\Delta V_n}\,\xi_n}{V(x_n)} \tag{3.28}$$

$$= \frac{\Lambda(x_{n-1})V(x_{n-1}) + 8\pi\alpha\sqrt{\Delta V_n}\,\xi_n}{V(x_n)} \,. \tag{3.29}$$

Summing this recursive relation back to some initial point x_0 in the far past of x_n, we find

$$\boxed{\Lambda(x_n) = \frac{8\pi G S_0 + 8\pi\alpha\sum_{i=1}^{n}\sqrt{\Delta V_i}\,\xi_i}{V(x_n)}} \,, \tag{3.30}$$

where $8\pi G S_0$ is $\Lambda(x_0)V(x_0)$, or possibly an initial condition coming from a pre-geometric regime (see Section 3.4.3).

In the above derivation we have explicitly kept track of ℓ_p and ℓ_{cs}, such that we can interpret the parameter α that sets the magnitude of the fluctuations as the square of the ratio of the Planck length to the discreteness length, as in (3.23).[3] We should only take this interpretation with a grain of salt, however, as we started with the assumption that each causal set element contributes a random value to the total Λ, while in fact the true causal set mechanism that gives rise to Everpresent Λ is not fully understood. Nonetheless, our account of α as the ratio of scales is consistent with the predictions from the path integral that led to (3.18). For example, we see that if we take the continuum limit and send $\ell_{cs} \to 0$, then $\alpha \to \infty$. Since in the continuum we have definite volumes with no uncertainty, this limit corresponds to $\delta V \to 0$, and therefore from (3.19) we expect infinite fluctuations in the cosmological constant. Thinking of α in this way is also useful in that it guides the kinds of values we would expect it to take, since we believe the discreteness scale to be close to the Planck scale.

3.3 Methodology

Before we study some characteristic features emerging from simulations of Everpresent Λ based on Model 1, we must first say a few more words regarding our methodology and implementation of the

[3] In [64–66] the scale of the fluctuations was more loosely defined, i.e. not in terms of any explicit expression connected to other scales in the problem.

model. This section will cover these details.

3.3.1 An isotropic and homogeneous universe

In the previous section we ended by providing a stochastic model that describes the evolution of Λ over discrete time steps. This evolution was expressed as a function of a random parameter and the volume V of a point's past light cone at those instances. Next we need to know the metric with which to calculate V, and how this metric evolves as a result of Λ.

We will follow [64] by assuming we have a spatially flat, isotropic and homogeneous cosmology, represented by the FRW metric (in conformal time coordinates)

$$ds^2 = a(\eta)^2(-d\eta^2 + dr^2 + r^2 d\Omega^2)\,. \tag{3.31}$$

In standard GR the evolution of the scale factor a is described by the Friedmann equations:

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\left(\rho_m + \rho_r\right) + \frac{\Lambda}{3}\,, \tag{3.32}$$

$$\frac{2\ddot{a}}{a} + \left(\frac{\dot{a}}{a}\right)^2 = -8\pi G(P_r + P_\Lambda)\,, \tag{3.33}$$

where the derivatives are with respect to the cosmic time t, related to the conformal time η in (3.31) by $dt = a\, d\eta$.

Any implementation of Everpresent Λ represents a departure from classical GR, because the Einstein equation does not admit a varying cosmological constant (see Appendix C.2 for a discussion on this). Our strategy (similar to what was done in [64]) will be to use only (3.32) to generate our Everpresent Λ histories.

We can obtain V by integrating the volume element coming from (3.31), $dV = a^4 r^2 \sin\theta\; d\eta\, dr\, d\theta\, d\phi$, over the past light cone of a given point x at conformal time η, to the future of some initial time η_i. The 3d integration of the volume element at a fixed time slice η' gives $d\eta'\; \frac{4\pi}{3} a^4 r(\eta')^3$, where $r(\eta') = \eta - \eta'$ is the radius of the past light cone on that hypersurface. Therefore the η'-integral remains to be calculated:[4]

$$V(\eta) = \frac{4\pi}{3}\int_{\eta_i}^{\eta} d\eta' a(\eta')^4(\eta - \eta')^3 \tag{3.34}$$

$$= \frac{4\pi}{3}\left[\eta^3\int_{\eta_i}^{\eta} d\eta' a(\eta')^4 - 3\eta^2\int_{\eta_i}^{\eta} d\eta' a(\eta')^4\eta' + 3\eta\int_{\eta_i}^{\eta} d\eta' a(\eta')^4\eta'^2 - \int_{\eta_i}^{\eta} d\eta' a(\eta')^4\eta'^3\right]. \tag{3.35}$$

3.3.2 Numerical recipe

We now describe our numerical procedure for evolving Λ and V. Although we have thus far described quantities as evolving with time, we will find it useful to instead use the scale factor a as our "time"

[4] The expression in (3.34) can be seen to be equivalent to the integral $V(t) = \frac{4\pi}{3}\int_{t_i}^{t} dt'\; a(t')^3\left(\int_{t'}^{t} dt''\; \frac{1}{a(t'')}\right)^3$ in [64] through the coordinate transformation $d\eta = \frac{dt}{a}$.

parameter,[5] and calculate the evolution of the conformal time η. We divide a into n equal parts so that at step $i \geq i_{initial}$ of our simulation, the value of a is i/n. In this picture $a = 0$ represents the time of the big bang, and $a = 1$ the present day. The convention to use equal size steps in a is a choice that gives reasonable results in our analysis. However, one may experiment in varying the step size, e.g. step sizes that start out small and gradually become larger.

Before we can start a simulation, we must choose our initial parameters. These parameters are S_0, V_0, the number of time steps n, and the time step from which to start the evolution $i_{initial}$. Since we do not yet know the details of the quantum gravitational physics that preceded the Everpresent Λ era, we must make some assumptions in these choices. See Section 3.4.3 for further discussion on this. We also choose the matter and radiation densities ρ_m and ρ_r so that they are consistent with their present day values in the expression for $d\eta/da$ below. We also have to pick a seed number for our random number generator. Any seed number contains in itself a unique sequence of Gaussian random numbers ξ_i.

At each time step i in our simulation, each of our variables update as follows:

1. Increment $a_{i-1} \to a_i$.

2. Using Λ_{i-1} from the previous time step, we calculate $d\eta/da$ as a function of (arbitrary) a as

$$\frac{d\eta}{da} = \frac{1}{a\dot{a}} = \sqrt{\frac{1}{a^4 H_i^2}} \tag{3.36}$$

$$= \sqrt{\frac{3}{8\pi G a^4 (\rho_m + \rho_r) + a^4 \Lambda_{i-1}}} \,. \tag{3.37}$$

3. Given a_i, we calculate V_i by changing the integration variable in (3.35) to a and computing the sum

$$V_i = \sum_{j=0}^{3} c_j \eta(a_i)^j V_{j,i} \,, \tag{3.38}$$

where

$$V_{j,i} = V_{j,i-1} + \int_{a_{i-1}}^{a_i} da\, \eta^{3-j} \frac{d\eta}{da} a^4 \,, \tag{3.39}$$

i.e., we add the change in the four integrals of (3.35) in the interval $[a_{i-1}, a_i]$ to their previous values, where c_j are the appropriate prefactors. The time function $\eta(a)$ is itself defined as the integral

$$\eta(a) = \eta(a_{i-1}) + \int_{a_{i-1}}^{a} da \frac{d\eta}{da} \,, \tag{3.40}$$

where $\eta(a_{i-1})$ is saved from the previous time step.

4. Next, we generate a random number ξ_i (which is determined by our choice of the seed number) and update the action from the numerator of (3.30) to obtain

$$S_i = S_{i-1} + \alpha \sqrt{V_i - V_{i-1}}\, \xi_i / G \,, \tag{3.41}$$

[5] This is because the Friedmann equation (3.32) is already in terms of a, and in practice the simulations also tend to run more quickly when working with a.

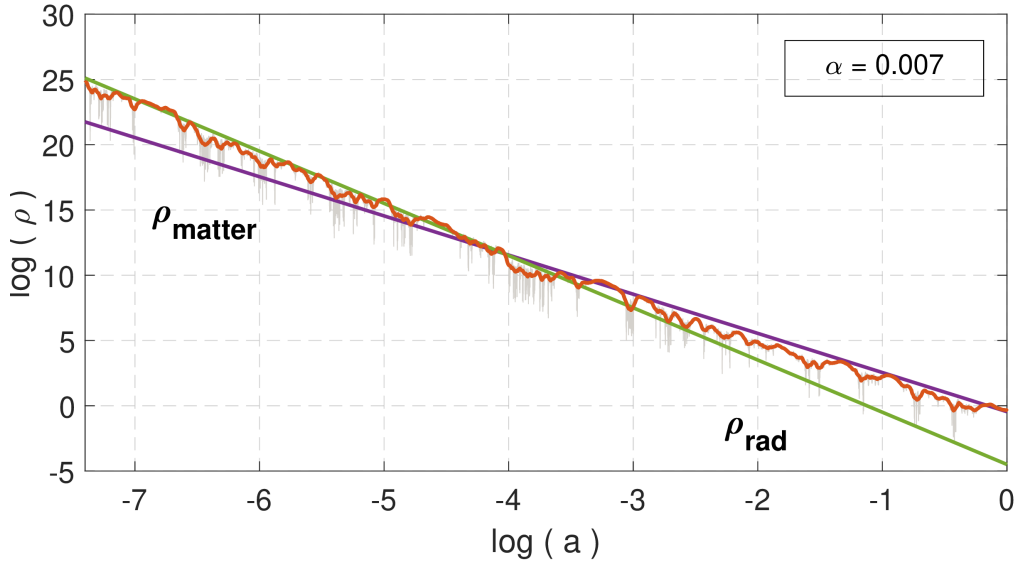


Figure 3.1: The evolution of matter (purple) and radiation (green) density in the universe is shown. The thin grey curve is the absolute value of the energy density of (Everpresent Λ) dark energy for a particular realisation of our simulation. A smoothed-out version of the dark energy density is shown in red. The plot shows that the dark energy density is always of the order of the ambient density, be it radiation (early on) or matter (later). All the densities are normalised by the critical density, such that the present-time total energy density is 1. Here the dimensionless model parameter α is set to 0.007.

and update Λ as

$$\Lambda_i = 8\pi G S_i / V_i \,. \tag{3.42}$$

The integrals in (3.39) and (3.40) are numerically evaluated using Romberg's method, to speed up the convergence.

Figure 3.1 shows an example profile for a fluctuating Λ produced in this way. As can be seen in this figure, the magnitude of ρ_Λ tracks the ambient energy density when following the above implementation, as expected from $|\Lambda| \sim 1/\sqrt{V} \sim H^2$.

It is also an implicit assumption that a must always have a positive time derivative in order to perform the volume integrals, even though plausibly we could end up with a collapsing universe. However, given that our universe is expanding, this is not an important detail.

Having all the necessary equations at hand, we can perform simulations of the background cosmology. The statements we can make are no longer deterministic. We should instead look at the final distribution of various physical quantities after performing many realisations of the stochastic process.

3.3.3 $H^2 < 0$ crashed runs

It is important to note that since there are no additional constraints on the random numbers ξ_i, it is possible for the fluctuations of Λ in (3.30) to cause the total energy density ρ_{tot}, and therefore also H^2, to become negative. When this happens $d\eta/da$ becomes imaginary in (3.37), and the spacetime volume a complex number, which does not make sense classically.

If one takes Model 1 seriously, then it should be regarded as a legitimate possibility that H^2 can become

negative. It would not be correct to view this occurrence as a sign of unphysical input parameters. This is because, statistically speaking an evolution will almost surely reach a point where the random fluctuations in Λ will cause H^2 to become negative, and it is only a matter of how long the evolution according to this model lasts before this happens.

A negative H^2 and ρ_{tot} is a sign that our effective picture is no longer valid, and is due to the use of the Friedmann equation in (3.36) and (3.37). The Friedmann equation must be replaced by some other (as yet unknown) dynamics at this point, and one can speculate what the evolution of the universe will look like when this occurs. It may be that a different effective but still classical evolution equation should be used under such circumstances. Alternatively, a quantum dynamics may be needed.

As the dynamics in the $H^2 < 0$ regime is unknown at present, we will disregard universes that reach $H^2 < 0$ before the present, and terminate their evolutions at this point. When this happens, we refer to it as an $H^2 < 0$ *crash*, as the number in the square root of (3.37) becomes negative, and would cause a computer code requiring real variables to crash. We again emphasize the *physical* origin of these crashes, as they are not due to a programming error.

There is a higher likelihood for an $H^2 < 0$ crash to happen at earlier times, for larger values of α, and when the simulations are performed over longer time scales. The fact that the $H^2 < 0$ crashing frequency becomes greater the earlier the start time of the simulations is, seems to suggest that there might be a link between these crashes and the breakdown of the model's semi-classical approximations. These $H^2 < 0$ crashes are an aspect of Model 1 that needs to be understood better and we devote Section 3.4.2 below to further discuss some aspects of these simulations. See also Section 3.4.3.

3.3.4 ρ_m as a scaling parameter at late times

Model 1 shows a very curious scaling invariance as an artifact of treating the seed number as a parameter. For a fixed seed number s, although $(\Omega_m^0 h^2, \alpha)$ are independent parameters of the model, (Ω_m, α) are perfectly correlated in the matter-dominated era. This is because if we keep the seed and α fixed, and then make the change

$$\rho_m^0 \to c_0 \rho_m^0 \,, \tag{3.43}$$

the quantities Λ, H^2, and $1/\sqrt{V}$ also get re-scaled by c_0 (assuming that $S_0 = 0$), as can be seen from (3.36) - (3.42). In particular, this means that the ratio ρ_m/Λ and hence Ω_m and Ω_Λ remain unchanged for fixed α and random number series $\{\xi_i\}$ in (3.30). In addition, Ω_Λ remains unchanged after a re-scaling of ρ_m in the radiation domination era as well, but this time simply because ρ_m is sub-dominant in the background. Therefore, for a fixed seed, Ω_m is only a function of α and not of ρ_m; (Ω_m, α) are perfectly correlated.

Using the above fact, there is a simple way to reconstruct the expansion history in the matter-dominated era for $(s, \Omega_m^0 h^2, \alpha)$ if we know the expansion history for $(s, 0.14, \alpha)$, where 0.14 is just a fiducial choice for the matter density. The latter gives us the knowledge of Ω_Λ at different redshifts. Together with $\rho_m(a)$, this is enough to find H and hence the expansion history, without having to perform numerically expensive volume integrals. In addition, if the realisation crashes due to encountering a negative H^2 for some $(s_1, 0.14, \alpha_1)$, then the same would happen for every $(s_1, \Omega_m^0 h^2, \alpha_1)$, because of the scaling property of $\Omega_m^0 h^2$, and also for every $(s_1, 0.14, \alpha_2)$ with α_2 slightly bigger than α_1. Intuitively, this can

be understood by remembering from (3.30) that α controls the magnitude of the fluctuations in Λ, although not in a linear way since V in the denominator also depends on α. Nevertheless, an $H^2 < 0$ crash means that ρ_Λ has become equal to $-\rho_m$ at some time, and now a slight increase in α makes ρ_Λ more negative. The $H^2 < 0$ crash will therefore occur in this case as well. This was confirmed by our numerical simulations. Hence, α_1 is an upper bound on α to avoid $H^2 < 0$ crashes, in the fixed seed s_1.[6]

3.4 Results and further investigations

We summarise below some results highlighting both qualitative properties and quantitative statistics of Model 1 simulations of Everpresent Λ.

3.4.1 Statistical distributions of simulation outcomes

We carry out a set of simulations to explore the statistical properties of Model 1. In Model 1, the free background parameters are $(\Omega_m^0 h^2, \alpha)$. In this subsection, we fix $\Omega_m^0 h^2 = 0.14$ and report the statistics from three values for α: 0.005, 0.01 and 0.012. For each one, we run 10,000 simulations with different seeds and subsequently study the runs that did not encounter an $H^2 < 0$ crash. For $\alpha = 0.012$, only 535/10000 runs successfully completed. We do not consider larger values of α as we would have lacked a sufficient number of completed runs. The results are shown in Figure 3.2. The first column shows the distribution of H_0. For $\alpha = 0.005$, around 99.9% of the simulations completed the runs. However, none of them come close to the observed value of H_0 of $\sim 72\,\mathrm{km/s/Mpc}$. The peak of the distribution is about 7.9 standard deviations away from this value. For $\alpha = 0.01$ and 0.012, the peaks are about 3.44 and 2.74 standard deviations away, respectively, and we see that a few realisations give reasonable values of H_0.

It is interesting to note that the mean of H_0 in all three cases is close to $40\,\mathrm{km/s/Mpc}$. This is because, at any time, Λ is a Gaussian random number with a mean of zero. Therefore, we expect the Hubble parameter to peak at the value where the universe has no dark energy. If we had a universe filled with only matter and radiation, the present-time Hubble parameter would be around 37.7km/s/Mpc. As we are excluding the runs that encounter an $H^2 < 0$ crash, the peak of the distribution shifts to higher values. Also as α increases, the number of rejected samples increases, shifting the peak of the Hubble parameter further to the right.

The second column shows $\Omega_\Lambda^0 = \rho_\Lambda^0/(\rho_\Lambda^0 + \rho_m^0)$. The value of Ω_Λ must at all times be less than 1 and have mean zero, and both of these features can be seen in the plots. With increasing α, the standard deviation of the distribution increases.

The final column shows the distribution of $H_0^4 V$, where V is the present volume of the past light cone. We expect V to be proportional to H_0^{-4}. Here also, the standard deviation of the distribution increases with increasing α. Again, the peak of the distribution roughly corresponds to a zero dark energy universe, for which $H_0^4 V = 0.033$. This value is computed for a pure matter-dominated universe. It

[6] This α_1 is not a global bound, however. As explained in Section 3.4.3, for some seeds there is another regime of large values of α, beyond their corresponding α_1, for which the simulations do not crash early on. This regime, however, is not of our interest at present, as it describes expansion histories that are strongly dominated by dark energy.

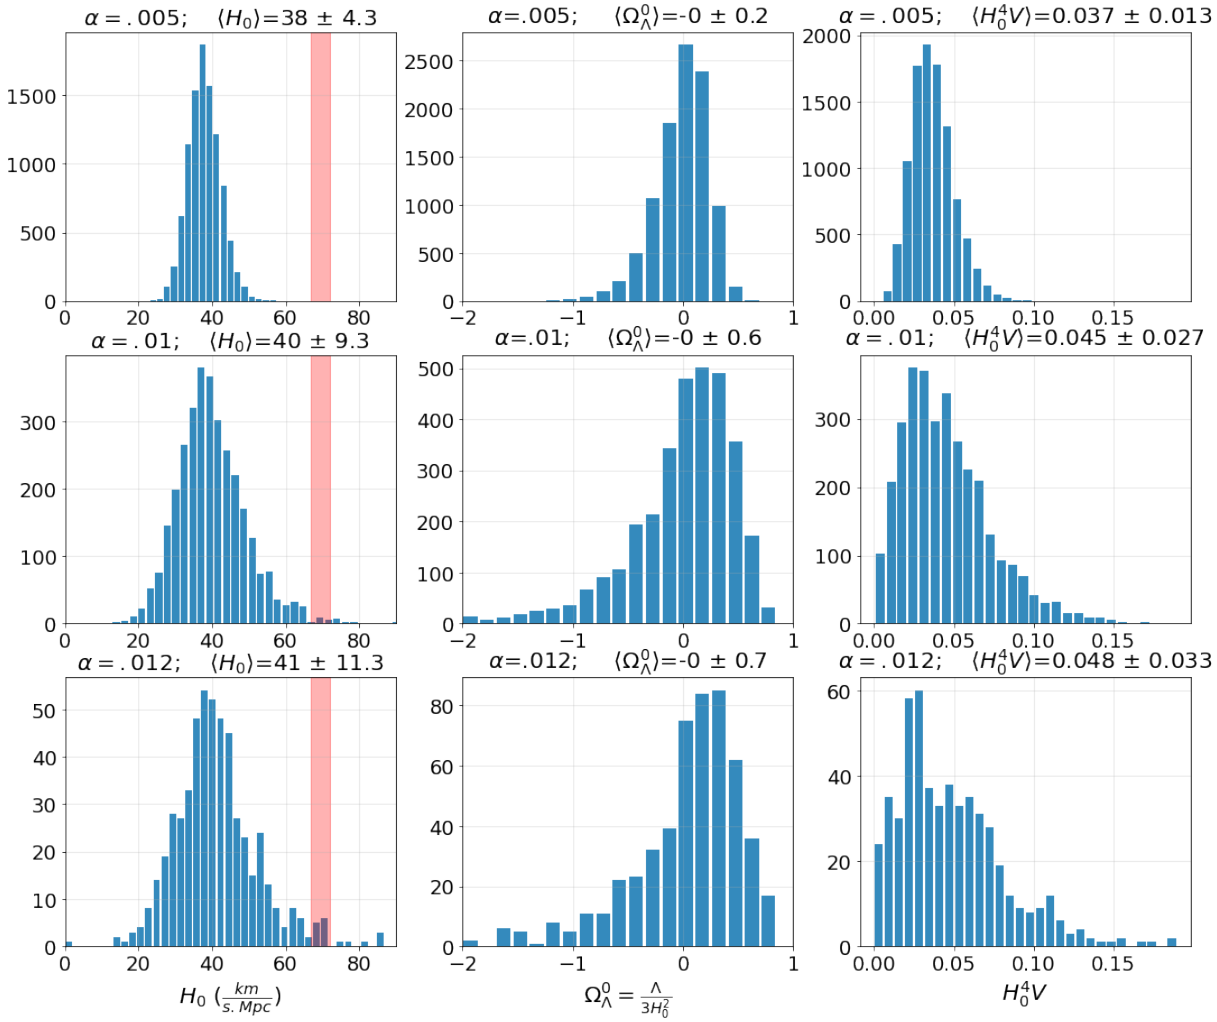


Figure 3.2: The plot shows the distribution of H_0, Ω^0_Λ, and H_0^4V (where V is the present volume) for Model 1 with $\Omega^0_m h^2 = 0.14$. The three rows show the plots for $\alpha = 0.005$, 0.01, and 0.012, respectively. The interval $H_0 \in [67, 72]$km/s/Mpc is shaded in red in the first column to show the presently accepted range of the Hubble parameter. As α increases, more runs terminate because of reaching negative values of H^2. The number of completed runs for the three α's are 9990, 3228, and 535, respectively, each out of 10000.

also connects to the rough values of α required for a dark energy-dominated universe today. To see this, note that we require $\Lambda/3 \simeq 0.7H_0^2$. From the distribution (3.27), we get

$$\alpha \sim \frac{3 \times 0.7}{8\pi}\sqrt{VH_0^4}. \tag{3.44}$$

Substituting $H_0^4V = 0.033$ (which, strictly speaking, only holds in matter domination), gives $\alpha \sim 0.015$. This is the value for which a present-day dark energy fraction of about 70% would be probable, although in practice, many runs would encounter an $H^2 < 0$ crash before reaching that point. This justifies why we are primarily interested in values of $\alpha \sim 0.01$.

3.4.2 Frequency of completed runs

As a stochastic model, Model 1 produces a unique expansion history for every seed (which in turn gives a series of random numbers) and α. As mentioned in Section 3.3.3, in many cases we obtain a

negative total energy density before reaching the present. This leads to an imaginary Hubble parameter, which in turn would result in a complex spacetime volume if the steps described in Section 3.3.2 are followed. We therefore discard these runs as the correct dynamics beyond this point is not known. There are a number of factors which affect the likelihood to encounter an $H^2 < 0$ crash. One is the value of α, which controls the magnitude of fluctuations of the dark energy density. Increasing α raises the magnitude of fluctuations, increasing the probability of a particular run to crash. The choice of time step during the expansion history calculation is another critical factor. The dark energy density remains constant between every two consecutive time steps and it updates at the end of every time step. Even though ideally we would like the time step to be of the order of the discreteness scale, it is not computationally feasible to implement it that way. Instead we choose a more modest time step in terms of the scale factor (a), time (t), conformal time (η), or some other suitable quantity. Decreasing the value of the time step increases the number of them and with it the number of times the dark energy density fluctuates. Furthermore, increasing the number of fluctuations increases the chances for positive and negative contributions to cancel one another, decreasing the possibility of an $H^2 < 0$ crash.

We run a set of simulations to confirm the observations of the previous paragraph and further analyse the statistics of $H^2 < 0$ crashes. We use a constant time step in the conformal time domain, $\Delta\eta$, and for three choices of α values calculate the fraction of completed runs (i.e. those that do not encounter an $H^2 < 0$ crash before the present) for different values of $\Delta\eta$. The results are shown in Figure 3.3. Each data point was produced using 10,000 simulations. When $\alpha = 0.02$ (not shown in the figure), less than 10 out of 10,000 for any given $\Delta\eta$ do not end up crashing, showing that such values for α are not practical in Model 1 (if we take $S_0 = 0$).

The results in Figure 3.3 are for $\alpha = 0.005$, 0.01 and 0.015. As anticipated, we see that as α increases, the fraction of completed runs drops considerably. The trends with $\Delta\eta$ are more visible for the curves corresponding to the two smaller values of α, as a large fraction of runs crashed for $\alpha = 0.015$. Once again as anticipated from the discussion in the first paragraph of this subsection, as we decrease the size of the time step, the fraction reaching the present day increases (in fact it starts to plateau for the smallest values of $\Delta\eta$).

We also explored a few other types of time steps and observed similar behaviour for them. For example, we considered constant time steps in η, t, a, $\log a$. etc. In all cases, as we decrease the size of the time steps, the crashing fraction decreases and finally becomes almost constant.

The $H^2 < 0$ crashing frequency is also sensitive to the starting time of the simulations. The volume of the past light cone decreases as we go back in time, and with that the dark energy density can become very large as it is of the order of $1/\sqrt{V}$. For the results we report on, we start the simulations from $z \approx 10^4$, and except in the next subsection, we calculate the volume up to this time using only the matter and the radiation component of the universe. We have seen that if we start increasing the starting redshift, the crashing fraction also increases. In fact, for the analysis shown in Figure 3.1, where we take a logarithmic time step in a and a starting redshift $z \sim 10^8$, the crashing fraction is almost 100% for $\alpha = 0.01$.

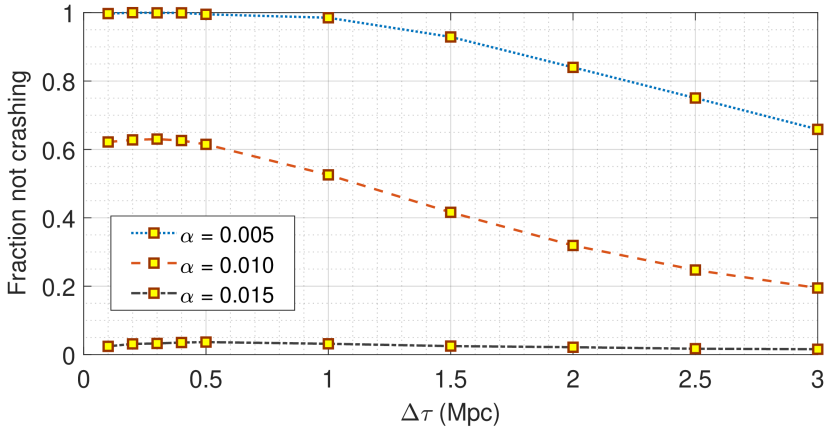


Figure 3.3: The plot shows how the time step affects the crashing fraction for $\alpha = 0.005$, 0.001, and 0.015. As α increases, the fraction of completed runs drops considerably. For higher α the crashing fraction almost tends to 100%. As we decrease the time step, $\Delta\eta$, the crashing fraction decreases, eventually saturating to a fixed value.

3.4.3 Initial conditions

As mentioned at the beginning of Section 3.3.2, we need to choose initial values for the volume, V_0, and the action, S_0, for Model 1. In this subsection, we discuss some of the trends that result from different choices for these initial conditions.

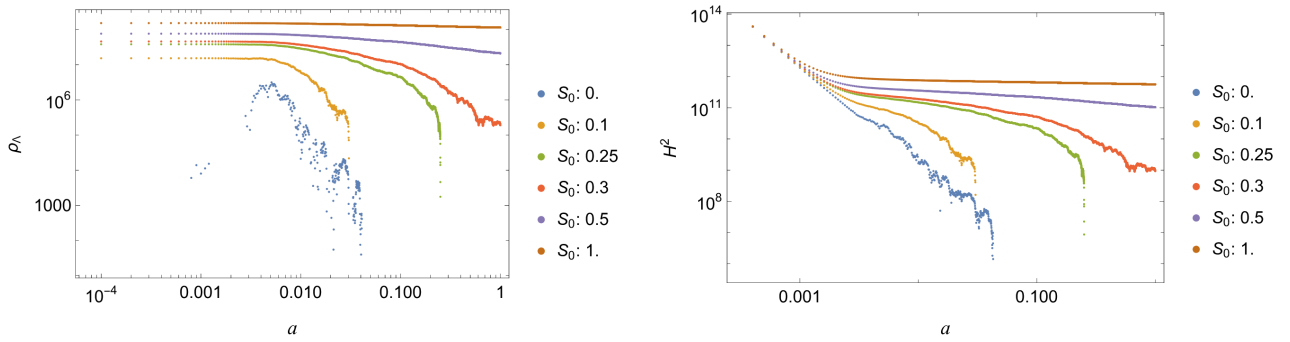
There are several approaches to setting these initial values. Everpresent Λ does not tell us about the evolution of the very early universe, before a number-volume correspondence becomes meaningful. Nevertheless, one approach is to (in a more or less ad hoc manner) attribute a V_0 and S_0 to the early non-geometric causal set, when it has presumably reached the size it will begin to grow into an FRW manifoldlike causal set. From a purely pragmatic point of view, setting $S_0 > 0$ adds a positive bias to Λ that makes it less likely to become too negative, and the simulation less likely to crash. Therefore, $S_0 > 0$ values are more interesting, and we will focus on this case below. Note that the initial volume V_0 only enters the numerical process described in Section 3.3.2 in the denominator of (3.42). In other words, we have:

$$\Lambda(x_n) = 8\pi \frac{GS_0 + \alpha \sum_{i=1}^{n} \sqrt{\Delta V_i}\, \xi_i}{V(x_n) + V_0} \,. \tag{3.45}$$

V_0 does not take part in the fluctuating numerator, which represents the effect of the random action of Model 1. The contribution of the pre-geometric phase in the numerator is encapsulated in $8\pi G S_0$.

Guided by the results summarised in Figure 3.3, we will set α to a value for which most simulations encounter an $H^2 < 0$ crash, in order to investigate whether a larger S_0 can prevent some of these crashes. In Figures 3.4 and 3.5 we show the results. In the simulations, we set $\alpha = 0.015$, $V_0 = 1$ (in units of Mpc^4/c), $i_{initial} = 1$, $n = 10000$, and increase S_0 from 0 to 1 (in units of $c^3\mathrm{Mpc}^2/G$). We see that as S_0 increases, the (initially crashing) simulation eventually stops crashing, although the behaviour is drastically different in the new regime.

The presence of S_0 and V_0 gives rise to a different expansion history with different epochs compared to the tracking behaviour of Everpresent Λ that we found in Figure 3.1.



(a) (b)

Figure 3.4: Left: A log-log plot of ρ_Λ (normalised by present critical density of our Universe) versus a when increasing S_0, with $\alpha = 0.015$. The expected behaviour of Everpresent Λ is that it follows the same scaling as ρ_m at late times, i.e. $\sim a^{-3}$. For $S_0 = 0$, ρ_Λ does indeed follow that scaling, however the scaling becomes closer to zero, and therefore Λ becomes more and more dominant the larger S_0 becomes. Right: A log-log plot of H^2 (in $(\mathrm{km/s/Mpc})^2$) versus a when increasing S_0, with $\alpha = 0.015$. Similar to the case of ρ_Λ, we expect a decay of $\sim a^{-3}$ from normal Everpresent Λ cosmology, which we roughly see for the S_0 curve before it crashes. Furthermore, if the evolution were to predict a value of H_0 similar to that of our universe, at $a = 1$ we would expect $H^2 \sim 10^3$. However, the completed runs have far higher values, consistent with a large dark energy provided by the initial S_0.

Early radiation or matter domination

Due to a non-zero V_0 in the denominator (which dominates over V), the contribution from Λ is suppressed at early times. So H^2 is dominated by radiation or matter density as it can be seen in Figure 3.4b at small a. For $S_0 = 0$, Λ behaves like $\Lambda \sim \pm 8\pi\alpha\sqrt{V}/V_0$ which fluctuates with an increasing trend. For non-zero S_0, we have $\Lambda \simeq 8\pi G S_0/V_0$, which gives a constant but sub-dominant dark energy. This can be seen in Figure 3.4a for $a \lesssim 3 \times 10^{-3}$.

Dark energy domination followed by slow-roll

As H^2 keeps decreasing, it will finally become of the order of the almost-constant Λ, and then dark energy takes over. Meanwhile, with the expansion of the universe, the volume of the past light cone grows and eventually becomes comparable to V_0. At this point, Λ is given by [7]

$$\Lambda \simeq \frac{8\pi G S_0}{V + V_0}\,. \tag{3.46}$$

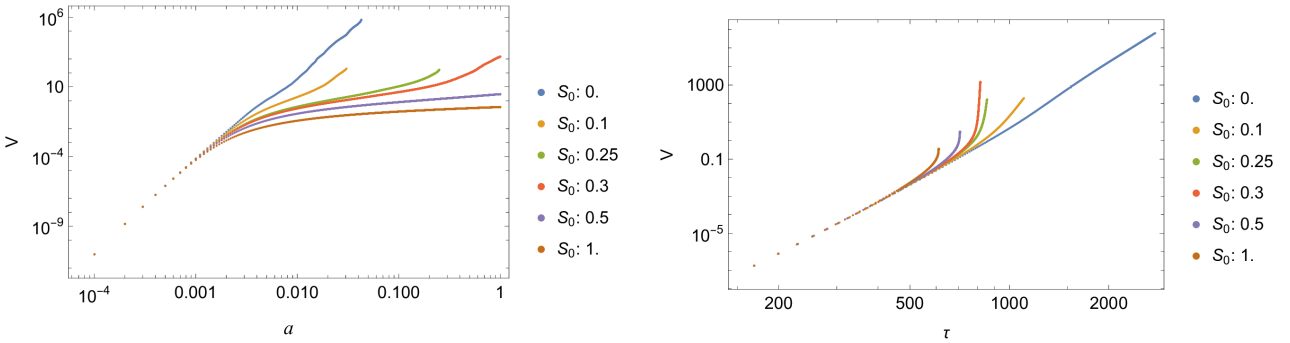
Together with the expression for the volume and the Friedmann equation

$$V(\eta) = \frac{4\pi}{3}\int_{\eta_i}^{\eta} d\eta' a(\eta')^4(\eta - \eta')^3\,, \tag{3.47}$$

$$H^2 = \frac{\Lambda}{3}, \tag{3.48}$$

these give an integro-differential equation for the scale factor in terms of η. This is difficult to solve analytically. The numerical solution in Figure 3.4 indicates a slow-roll-like behaviour. This is backed up by Figure 3.5 which shows an inflation-like growth of scale factor with respect to V and V with respect to η.

[7] For $S_0 = 0$, when V catches up with V_0, we simply see the onset of Everpresent Λ tracking behaviour, as in the blue curve in Figure 3.4 for $a \gtrsim 10^{-2}$.



(a) (b)

Figure 3.5: Left: A log-log plot of V (in Mpc^4) versus a when increasing S_0, with $\alpha = 0.015$. Right: A log-log plot of V (in Mpc^4) versus η (in Mpc) when increasing S_0, with $\alpha = 0.015$. We see that the slow growth of V in units of a in Figure 3.5a is counteracted by the slow growth of η with a, so that V has an explosive growth in physical time, as one would expect when there is a highly dominant dark energy.

Everpresent Λ tracking

The inflationary growth does not last forever. With the growth of the volume V, the fluctuating term in Equation 3.45 grows like $\alpha\sqrt{V}$. This eventually takes over the GS_0 term, and the fluctuations in Λ and H^2 start to appear as in Figure 3.4 for larger values of a. Eventually, negative fluctuations will decrease H^2 by many orders of magnitude until it reaches the energy density of radiation (if this was in early universe) or matter. From that point, we would get the tracking behaviour of the Everpresent Λ.

The results of the simulations above were for a particular value of α, and a seed for which the runs reach an $H^2 < 0$ crash when $S_0 = 0$. If we instead consider the behaviour for a smaller α where the simulation with $S_0 = 0$ does not crash, the effect of increasing S_0 leads to similar results to above, except that for intermediate values of S_0 the runs will actually encounter an $H^2 < 0$ crash. We therefore have three "regimes" when setting S_0: small S_0, where the effect on the final results are modest; intermediate S_0, where fluctuations become large enough to crash almost all runs[8]; and large S_0, where Λ dominates and we get scalings that are unlike our universe.

It is also interesting to note that S_0 and α induce similar behaviour in Model 1. The same observation as above holds for α, namely that for small α almost all runs succeed, and have the correct scalings. For intermediate α (larger than 0.015 when $S_0 = 0$, as seen in Figure 3.3) almost all runs crash. For large enough α, fewer simulations crash early on, since ρ_Λ dominates from the onset.

Finally, note that we performed our simulations in the late universe with large values for S_0 and V_0, but in fact, the regime of interest for the above behaviours is the early universe. This investigation suggests that a nonlocal model of Λ as given by Equation (3.46) is capable of giving a slow-roll-like inflationary behaviour. When supplemented by stochastic terms as for Everpresent Λ, the inflationary growth finally ends and dark energy will track the ambient radiation density. A closer look at this transition as well as the naturalness of the parameters S_0 and V_0 needed for generating a viable inflation model is left for future work.

[8]We can actually see this effect in the figures above, since the example run with $S_0 = 0.1$ crashes earlier than the run with $S_0 = 0$.

3.4.4 Understanding the auto-correlation time

An important property of Model 1 is that the auto-correlation time of Λ is of the order of the Hubble time. Intuitively, this means that it takes roughly one Hubble time for Λ to become independent of its past values and change sign. By one Hubble time, we mean the time it takes for the scale factor to roughly double.
To see why this is true, we will calculate the auto-correlation $\langle \Omega_\Lambda(a_1)\Omega_\Lambda(a_2)\rangle$, where $\Omega_\Lambda = \rho_\Lambda/\rho_{tot} = \Lambda/3H^2$. Combining (3.25) and (3.26) gives

$$\Lambda(a_1) = \frac{4\pi\ell_p^2}{V(a_1)} \sum_{e_i \in \mathcal{V}(a_1)} b(e_i)\,. \tag{3.49}$$

The summation is performed on all the causal set elements e_i that are in the past light cone of a point at a_1, and b is a Bernoulli random number:

$$b(e_i) = \begin{cases} 1\,, & \text{probability} = \frac{1}{2} \\ -1\,, & \text{probability} = \frac{1}{2} \end{cases}\,. \tag{3.50}$$

We assume the random variables are independent, such that $\langle b(e_i)b(e_j)\rangle = \delta_{ij}$; therefore, if $a_1 < a_2$ we have

$$\langle V_1\, V_2\, \Lambda(a_1)\Lambda(a_2)\rangle = 16\pi^2\ell_p^4 \sum_{e_i \in \mathcal{V}(a_1)} 1 = 16\pi^2 \frac{\ell_p^4}{\ell_{cs}^4}\langle V_1\rangle = (8\pi\alpha)^2\langle V_1\rangle\,, \tag{3.51}$$

where $V_i \equiv V(a_i)$ and the nonzero contribution to the sum is the common past of the two points which is $\mathcal{V}(a_2)\cap\mathcal{V}(a_1) = \mathcal{V}(a_1)$.

For obtaining the correlation $\langle \Omega_\Lambda(a_1)\Omega_\Lambda(a_2)\rangle$, we can replace $\Lambda(a) = 3\Omega_\Lambda H^2$ in (3.51) and assume that we can take V and H in and out of the expectation values. This assumption is certainly not correct, but for the purpose of a rough understanding of the correlation time, it suffices. A few algebraic manipulations lead us to

$$\langle \Omega_\Lambda(a_1)\,\Omega_\Lambda(a_2)\rangle = \frac{(8\pi\alpha/3)^2}{\langle H_1^2\rangle\langle H_2^2\rangle\langle V_2\rangle}\,. \tag{3.52}$$

Next we write $\langle V(a)\rangle = \beta/\langle H(a)^2\rangle^2$, where we expect $\beta \sim \mathcal{O}(10^{-2})$ based on the third column of Figure 3.2. So we find

$$\langle \Omega_\Lambda(a_1)\,\Omega_\Lambda(a_2)\rangle = \frac{(8\pi\alpha)^2}{9\beta}\frac{\langle H_2^2\rangle}{\langle H_1^2\rangle}\,. \tag{3.53}$$

Since $\langle\Lambda\rangle = 0$, and $\langle H^2\rangle \propto \sum_i\langle\rho_i\rangle$ from the Friedmann equation (3.32), we find that in the matter dominated era

$$\langle \Omega_\Lambda(a_1)\,\Omega_\Lambda(a_2)\rangle = \frac{(8\pi\alpha)^2}{9\beta}\left(\frac{a_1}{a_2}\right)^3\,. \tag{3.54}$$

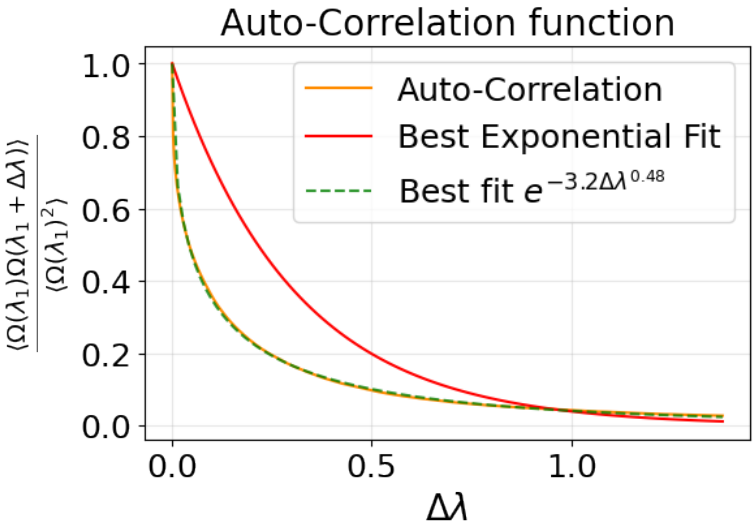


Figure 3.6: The normalised auto-correlation function $\frac{\langle\Omega_\Lambda(\lambda_1)\,\Omega_\Lambda(\lambda_1+\Delta\lambda)\rangle}{\langle\Omega_\Lambda(\lambda_1)^2\rangle}$ versus $\Delta\lambda$. To make this plot, 3000 realisations of Model 1 were simulated with the parameters $\alpha = 0.01$ and $\Omega_m^0 h^2 = 0.14$. For each simulation, $\langle\Omega_\Lambda(\lambda_1)\,\Omega_\Lambda(\lambda_1+\Delta\lambda)\rangle$ was calculated by averaging over λ_1. The resulting function of $\Delta\lambda$ was then averaged over the 3000 realisations, and suitably normalised so that it equals 1 at $\Delta\lambda = 0$. The final auto-correlation function is the orange curve, and the red curve is the best fit of an exponential $e^{-c|\lambda_1-\lambda_2|}$ to it. The dashed curve is the best fit of a function of the form $e^{-c_1|\lambda_1-\lambda_2|^{c_2}}$.

Changing the time variable from the scale factor a_1 to $\lambda_1 \equiv \ln a_1$ we have shown that

$$\langle\Omega_\Lambda(\lambda_1)\,\Omega_\Lambda(\lambda_2)\rangle = \frac{(8\pi\alpha)^2}{9\beta} e^{-3|\lambda_1-\lambda_2|}\,, \tag{3.55}$$

which suggests that the auto-correlation time is of the order of one Hubble time.

We have made some naive assumptions in deriving (3.55), hence we do not expect the auto-correlation function to be exactly an exponential function. In Figure 3.6 we have plotted the auto-correlation of Ω_Λ (for a given α) versus $\Delta\lambda$ along with a best-fit exponential function. Indeed, one can see that the auto-correlation function is not an exponential. The best fit to the orange curve turns out to be approximately $e^{-3.2(\Delta\lambda)^{0.48}}$.

Nevertheless, the dependence on $\Delta\lambda$ observed in the simulation data confirms that when the scale factor is multiplied by an $\mathcal{O}(1)$ number, the correlation of Λ with its previous values gets highly suppressed. Also note that the auto-correlation curve of Model 1 is dropping very fast close to $\Delta\lambda \sim 0$. This means that in very small time scales, Ω_Λ fluctuates a lot with a small magnitude. This is why the plot of Ω_Λ (Figure 3.7, in blue) looks jagged at small intervals.

Model 2

An alternative phenomenological model of Everpresent Λ, known as Model 2, was put forth in [66]. Model 2 is designed to incorporate the fluctuating nature of dark energy in Everpresent Λ as well as the Hubble time auto-correlation emergent in Model 1. This is done by making the assumption that $-1 \leq \Omega_\Lambda(\lambda) \leq 1$, which allows for the definition of a new variable $\hat{\Omega}_\Lambda(\lambda) = \tanh^{-1}(\Omega_\Lambda(\lambda))$. In fact, $\Omega_\Lambda(\lambda) \leq 1$ is always true since the other components of the universe have positive energy density, but $\Omega_\Lambda(\lambda) \geq -1$ is an assumption that is made to allow for the convenient definition of an $\hat{\Omega}_\Lambda$ that

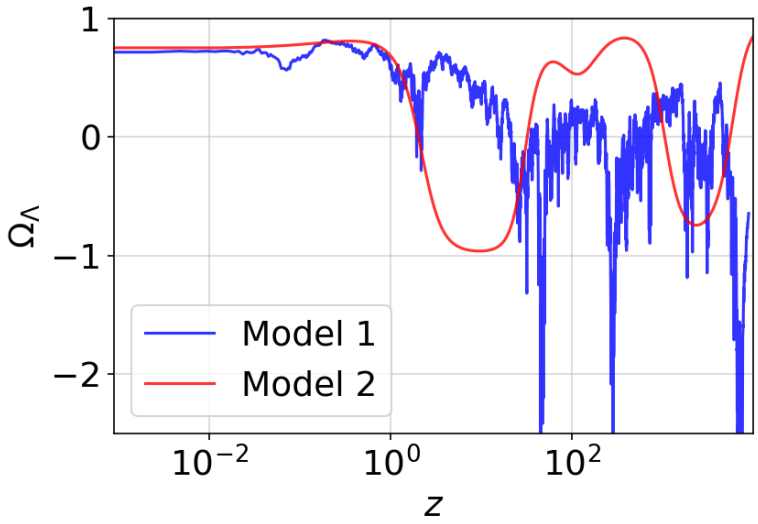


Figure 3.7: Plot of Ω_Λ versus redshift for one realisation of Model 1 (in blue) and one realisation of Model 2 (in red). The seeds for this figure were chosen such that for the given set of parameters of the models ($\alpha = 0.01$, $\Omega_m^0 h^2 = 0.14$ for Model 1 and $\tilde{\alpha} = 1.0$, $\mu = 1.108$ for Model 2), $\Omega_\Lambda^0 \sim 0.7$; hence they are not necessarily among the typical realisations of the models.

can be generated by a Gaussian process. Given that $\hat{\Omega}_\Lambda(\lambda)$ is unbounded, it can be generated as a *Gaussian process*. Over a discrete domain, this is done by treating the set of $\hat{\Omega}_\Lambda$ at different λ_i's as variables over which a multivariate Gaussian distribution is defined. The random vector drawn from this distribution then describes the variation of $\hat{\Omega}_\Lambda$ over λ. The precise way in which $\hat{\Omega}_\Lambda$ auto-correlates is dictated by the covariance matrix $\langle \hat{\Omega}_\Lambda(\lambda_1)\hat{\Omega}_\Lambda(\lambda_2)\rangle$ in the Gaussian distribution. Hence, one can instantly generate a history of $\Omega_\Lambda(\lambda)$ using the covariance matrix of $\hat{\Omega}_\Lambda(\lambda)$. As this model does not involve volume calculations, it is computationally much faster than Model 1. In [66] the ansatz was made that the covariance matrix is a squared-exponential

$$\langle \hat{\Omega}_\Lambda(\lambda_1)\hat{\Omega}_\Lambda(\lambda_2)\rangle = \tilde{\alpha}^2 e^{-\frac{\mu}{2}(\lambda_1-\lambda_2)^2}\,, \tag{3.56}$$

where $\tilde{\alpha}$ and μ are free parameters of the model.

For $\mu \sim \mathcal{O}(1)$, the auto-correlation time of the fluctuations in dark energy are of the order of one Hubble time.[9] However, one essential difference between the realisations of Model 1 and those of Model 2 is that the histories of Ω_Λ in Model 1 are more jagged at smaller time scales (see Figure 3.7). The reason is that the square root exponential auto-correlation function of Model 1 drops faster close to $\Delta\lambda = 0$.

Taking these considerations into account, more models can be built based on the Everpresent Λ idea that are computationally as fast as Model 2 but better resemble the realisations of Model 1. For this purpose, (3.56) should be replaced with a square root exponential function, or even with the numerical orange curve from Figure 3.6.

[9]Even though the covariance matrix in Model 2 is for $\hat{\Omega}_\Lambda$, for values of $\tilde{\alpha} \approx 1$, the auto-correlation of $\Omega_\Lambda(\lambda)$ closely follows that of $\hat{\Omega}_\Lambda(\lambda)$.

3.5 Discussion

We have laid out and reviewed the foundations of Everpresent Λ cosmology. We also investigated key properties of simulations based on the original phenomenological implementation of it which is in terms of a random action per causal set element.

We found that in Model 1 the value of H_0 for the noncrashing simulations is approximately normally distributed about a mean that is close to the value it would have if Λ were zero. While for small values of the fluctuation parameter α the observed value of H_0 is highly unlikely, we found that increasing this parameter increases the width of the distribution of H_0, and today's value becomes more likely.

We further found that when the initial parameters V_0 and S_0, which could originate from a deep quantum gravity regime in the early universe, are non-zero, the scale factor undergoes rapid growth, and the dark energy density no longer tracks the radiation or matter density. An interesting avenue for future work is to investigate whether, by tuning the parameters S_0 and α, one could have an alternative inflationary model for the early universe with a graceful exit to radiation domination. Some signs of a slow-roll type of behaviour followed by an exit to a decreasing Hubble parameter can already be seen in Figure 3.4b.

Finally, we found that in Model 1 the sign of Λ is likely to change roughly every Hubble time. An observational evidence of sign change in dark energy would therefore be a strong indication of Everpresent Λ in our universe.

While we have aimed to provide a complete exposition of the topics relevant to Everpresent Λ, there are a number of things we did not delve into. Let us briefly mention three of these: (*i*) One of the assumptions of the model is that the mean about which Λ fluctuates is 0; motivating this assumption with a Lorentzian argument remains an open question. Ideas from the Lorentzian path integral [263] and analogue gravity [264] can potentially be fruitful in this direction. (*ii*) There are several classical aspects of the model that likely must be replaced by their quantum counterparts. Of course this is easier said than done, in the absence of the complete quantum gravity theory. However, there are hints at how this can begin to be explored. For example, one can replace the classical stochastic random walk of Model 1 with a quantum random walk inspired by [265, 266]. See also [267, 268]. As well, insights from quantum unimodular gravity [269, 270] can be consulted for indications of a modified dynamics. (*iii*) Mathematically, it may be worth doing an analysis of the Friedmann equations, treating them as stochastic differential or difference equations, in which Λ is the random component. Such an analysis may offer insights on the statistical evolution of the scale factor as well as how to handle negative energy densities and derivatives of Λ.

It would also be interesting to explore connections to [230, 271] where a stochastic cosmological constant from fluctuations of vacuum energy was studied.

In the next chapter we will report on our results from using Models 1 and 2 to confront supernova and cosmic microwave background data. It is perhaps a bit premature to test these simple models against the full range of data, but we can nevertheless learn from the analysis.

As mentioned in the introduction, many of our conventional theoretical models have fallen short of

shedding light on cosmological open questions. The cosmological constant puzzle has been a longstanding example of such an open question, and more recently, there have emerged new puzzles such as the Hubble tension. Moreover, in an era where *understanding* cosmological data (which will, in turn, relieve tensions) takes priority over merely fitting it, we believe that models such as Everpresent Λ should be further developed and explored.

Chapter 4

Stochastic dark energy: cosmological tests of Everpresent Λ

In the previous chapter, we considered the theoretical motivations and statistical properties of Everpresent Λ and its phenomenological implementation. Here, we will focus on the observational signatures and evidence for and against these models. As we show in the subsequent sections, the model successfully fits SN Ia for a fraction of its realisations, while at least in its present form, it struggles to explain the CMB data. The reason is partly due to the ISW effect. Its contribution in Everpresent Λ is very different from its contribution in ΛCDM as it is non-vanishing in the matter dominated epoch. The SW effect will also turn out to be modified non-trivially in Everpresent Λ.

The primary aims are to determine to what extent the existing implementations of Everpresent Λ are able to explain cosmological data, and whether they can be modified to improve the agreement. While we use the ΛCDM fit to data as a reference point that we compare the Everpresent Λ fits to, it is not necessarily our goal to quantitatively improve on it at this stage. As the Everpresent Λ models are still works in progress, we wish to find the qualitative features of those realisations that fit the data well or relatively better.

In Section 4.1, we compare the predictions of Model 1 of Everpresent Λ with supernova observations. In Section 4.2, we compare the model against CMB data. In Section 4.3, we provide a theoretical analysis of the constraint on Λ during recombination. We then numerically study arbitrary fluctuations in dark energy, allowing for variations as large as the ambient density, and constrain it before and after recombination separately, showing that only small fluctuations in dark energy are allowed. Finally, we discuss our findings in Section 4.4.

4.1 Fitting Model 1 to supernova data

The cosmological distance modulus is an important observable in astronomy. It measures the difference between the apparent magnitude and the absolute magnitude of an object and is an indicator of distance

in astronomy. The distance modulus is defined by

$$\mu = m_B - M_B = 5 \log_{10} d_L + 25 \,, \tag{4.1}$$

where d_L is the luminosity distance of an object (in Mpc) and can be calculated as

$$d_L = (1+z) \int_0^z \frac{dz'}{H(z')} = \frac{1}{a} \int_a^1 \frac{da'}{a'^2 H(a')} \,. \tag{4.2}$$

The most crucial data set for testing the background expansion history of the Universe comes from type Ia Supernova (SN Ia). All SN Ia have approximately the same luminosity; hence, they are widely used as standard candles for cosmological distance measurements [22, 33, 272, 273], and they can provide a strong constraint on the Hubble parameter. We use the Pantheon+SH0ES sample of the apparent magnitude and redshift of 1550 SN Ia events with redshift $z \geq 0.01$ together with 42 calibrator events [22]. As the expansion history for Everpresent Λ differs from the standard ΛCDM model, it affects the distance modulus μ.

Since the fluctuations in Λ are of the order of the inverse square root of the spacetime volume, which in an FLRW universe is of the order of the Hubble parameter squared, the fluctuations will always be of the order of the dominant energy density. A priori, one might expect that since the fluctuations of Λ for small redshifts track the matter density, Everpresent Λ is more like CDM rather than ΛCDM, and should not improve on ΛCDM in describing the supernova data. Although this is statistically correct, we shall surprisingly see that there exist realisations of Model 1 that fit the observations better than ΛCDM. We will then investigate the characteristics of these realisations.

4.1.1 MCMC analysis

Everpresent Λ is a stochastic model of the Universe. Given a pair of values for the parameters $(\Omega_m^0 h^2, \alpha)$, and a series of Gaussian random numbers, Model 1 of Everpresent Λ provides a unique expansion history of the Universe (note that the dimensionless Hubble parameter h represents $h = H_0/(100\text{km/s/Mpc})$ and is not a function of redshift). The series of random numbers is characterised by a seed s. That is, given an integer number s as a seed parameter, we uniquely generate a series $\xi_1, \xi_2, ...$ of samples from the standard normal distribution, and use them in discrete time steps to evolve Λ according to (3.30). Therefore, given $(\Omega_m^0 h^2, \alpha, s)$, we can get a unique expansion history. We also add M_B, the absolute magnitude of SN Ia, as a nuisance parameter in the MCMC. The constraint on M_B arises from measurements of the distance modulus using Cepheids located in the same host galaxies as the 42 calibrator Type Ia supernovae, as conducted by the SH0ES team.

In Appendix C.3, we outline possible approaches for performing MCMC with a stochastic model and explain our method for handling the seed number. Specifically, we choose $N = 90,000$ different seeds, and for each of them, we run an MCMC process independently and determine how often we obtain a likelihood as good as ΛCDM for the supernova data.

We use flat priors on $\Omega_m^0 h^2$, α, and M_B. We allow the parameters to vary in a large range, namely, $\Omega_m^0 h^2 \in [0.01, 1]$, $\alpha \in [0.001, 0.025]$, and $M_B \in [-30.0, -10.0]$. The range for α was chosen based on

the behaviour observed in the previous chapter, so that the chains would not hit the boundaries of the ranges. Values of α around 0.012 produced the most runs with H_0 comparable to the observed range (when dividing the history of the Universe into 10,000 timesteps). The majority of runs encounter an $H^2 < 0$ crash when $\alpha > 0.015$, and a value of H_0 that is too small when $\alpha < 0.005$. Therefore the best fit α should be obtainable within the given range, without hitting the boundaries. The desired value of $\Omega_m^0 h^2$ rounded to 0.14 is well within the given range and we shall see that the distribution obtained in this parameter does not overlap with the boundaries of its range either.

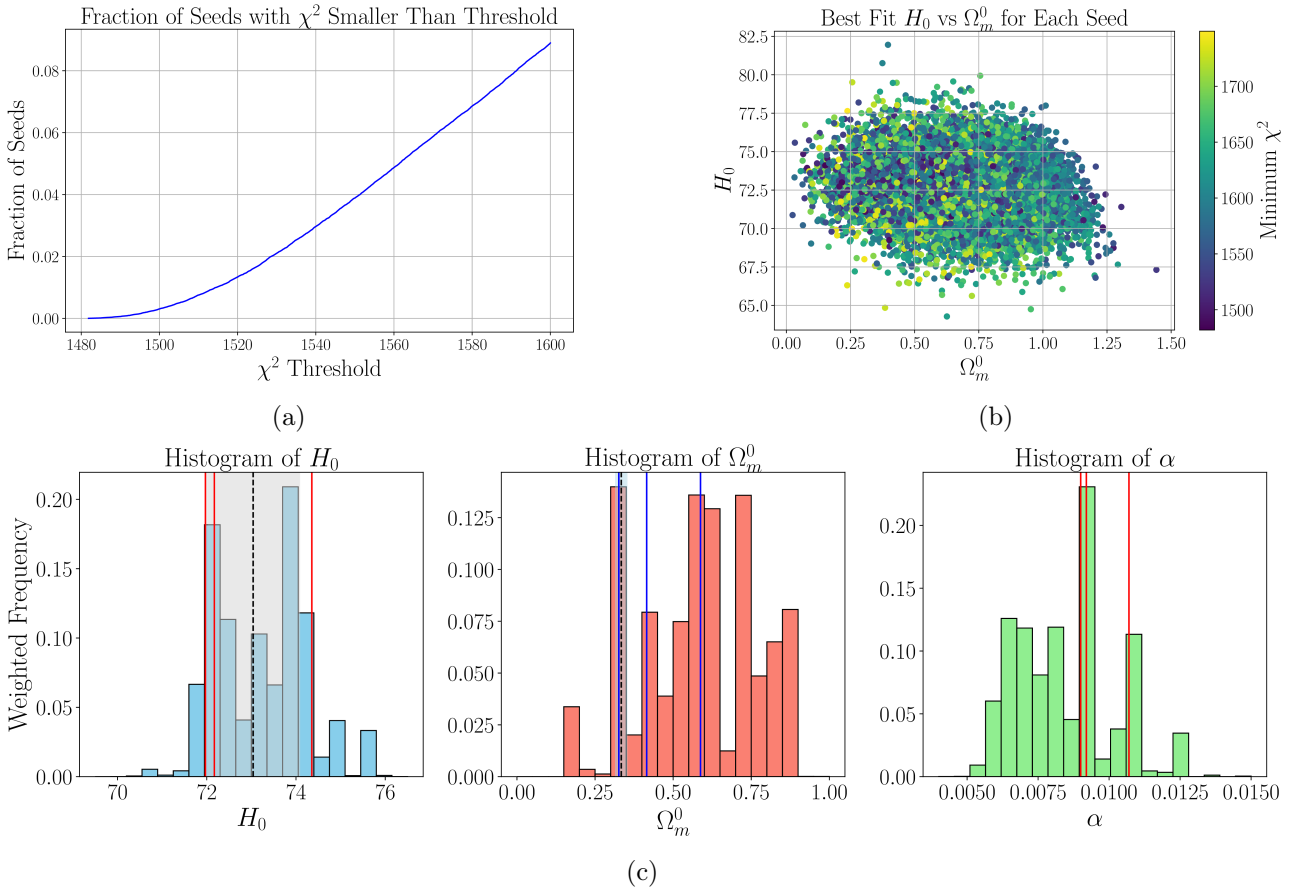


Figure 4.1: (a): The fraction of seeds that have a best-fit χ^2 smaller than some threshold on the x axis. For example, a fraction of 0.003 of seeds have a best-fit $\chi^2 < 1500$. (b): The scatter plot of best-fit (Ω_m^0,H_0) for 16767 seeds with a best-fit χ^2 smaller than 1750. The plot has a mild Pearson correlation coefficient of -0.13, showing that the seeds that can reduce H_0 and have a reasonable $\Omega_m^0 \sim 0.3$ are rare. Also the figure shows clearly that the likelihood surface is not smooth if we change the seed. (c): The histograms represent the posterior averages of the parameters H_0, Ω_m^0, and α from different seeds. For each seed, averages are taken after the burn-in period, and each seed is weighted by its posterior $P(s_i|D)$. The three vertical coloured lines show the average parameters of the three best seeds. The dashed vertical lines and the shaded regions represent the mean and 1σ confidence interval as reported by the SH0ES model-independent measurement for H_0 [22] and the flat ΛCDM fit to Pantheon+ for Ω_m^0 [167].

We try to find seeds that provide fits to the data close to or better than the standard model does. ΛCDM gives $\chi^2 \approx 1485.3$ for the Pantheon+SH0ES data. We therefore consider only seeds with a χ^2 close to that, and find that out of the 90,000 seeds, 55 give a χ^2 in the range [1481.0, 1490.0]. 16 of these seeds give a better likelihood than ΛCDM. This is summarized in Figure 4.1a; for each value x on the horizontal axis, the plot shows the fraction of the seeds that have a minimum χ^2 smaller than x.

Figures 4.1b and 4.1c show the summary statistics of the best-fit and average parameters of different seeds. Figure 4.1b shows the scatter of (Ω_m^0, H_0), where the best-fit point of each seed is coloured

according to its χ^2. The same parameters can give a very good or bad fit depending on the choice of the seed; hence, the likelihood surface is extremely discontinuous. This further justifies running each MCMC chain for a fixed seed number. However, the variance of this scatter plot should not be interpreted as the variance of the parameters themselves. To properly bound the parameters, it is essential to weight each seed's average parameters with the corresponding likelihood. This approach is known as Bayesian Model Averaging [274–276]. It can be employed for summarizing the posterior behaviour of any stochastic model. For example, the Bayesian average and variance of H_0 for Model 1 are given by

$$\langle H_0 \rangle = \sum_{i=1}^{N} H_{0i} \mathbb{P}(s_i|D), \quad \mathrm{var}(H_0) = \sum_{i=1}^{N} \Big(\mathrm{var}(H_0|s_i, D) + H_{0i}^2 \Big) \mathbb{P}(s_i|D) - \langle H_0 \rangle^2 \,, \tag{4.3}$$

where the sum is over all the considered seeds, $H_{0i} = \mathbb{E}(H_0|s_i, D)$ is the posterior average of H_0 for seed s_i, $\mathrm{var}(H_0|s_i, D)$ is the variance of H_0 for seed s_i, D represents the data set, and

$$\mathbb{P}(s_i|D) = \frac{\mathbb{P}(D|s_i)\mathbb{P}(s_i)}{\sum_j \mathbb{P}(D|s_j)\mathbb{P}(s_j)} \tag{4.4}$$

is the posterior probability of the seed s_i. $\mathbb{P}(D|s_i)$ is the marginalized likelihood of the data for a given seed, and is estimated by the harmonic mean of the likelihoods in a given MCMC chain [275], see Appendix C.3. Note that $\mathbb{P}(s_i)$ is the probability of a particular seed, which we assume to be uniform across all considered seeds. Given that the value of Ω_m^0 has been robustly measured using various independent methods such as cluster abundance and weak lensing experiments (e.g., [277–279]), there is little flexibility for modifying Ω_m^0. However, since Model 1 allows for a broad range of Ω_m^0 values, $\mathbb{P}(s_i)$ could theoretically be adjusted to prefer seeds that produce an Ω_m^0 consistent with these established measurements; nevertheless, we do not pursue this adjustment further. The result of the Bayesian Model Averaging is

$$\begin{aligned} H_0 &= 73.27 \pm 1.27 \,, \\ \Omega_m^0 &= 0.58 \pm 0.19 \,, \\ \alpha &= 0.0085 \pm 0.0017 \,. \end{aligned} \tag{4.5}$$

The corresponding weighted 1-d distributions are plotted in Figure. 4.1c. Note that, due to the relatively small number of seeds with a good χ^2, the distributions appear less smooth. For smoother plots, it is necessary to utilise a pool of at least a million seeds.

Our H_0 is consistent with the model-independent result by SH0ES [22]:

$$H_0^{\mathrm{SH0ES}} = 73.04 \pm 1.04 \,. \tag{4.6}$$

The model independence is achieved by Taylor expanding H around its present value and constraining the expansion coefficients. Strictly speaking, this method does not accommodate stochastic models where H cannot be assumed to be smooth, including Everpresent Λ. Nevertheless, we find a similar constraint to that of SH0ES (although with a larger error bar) owing to the fact that the fluctuations in H are relatively small at low redshift, see the right panel in Figure 4.5.

Unlike deterministic models, Ω_m^0 is pretty much unconstrained for Everpresent Λ. To better understand

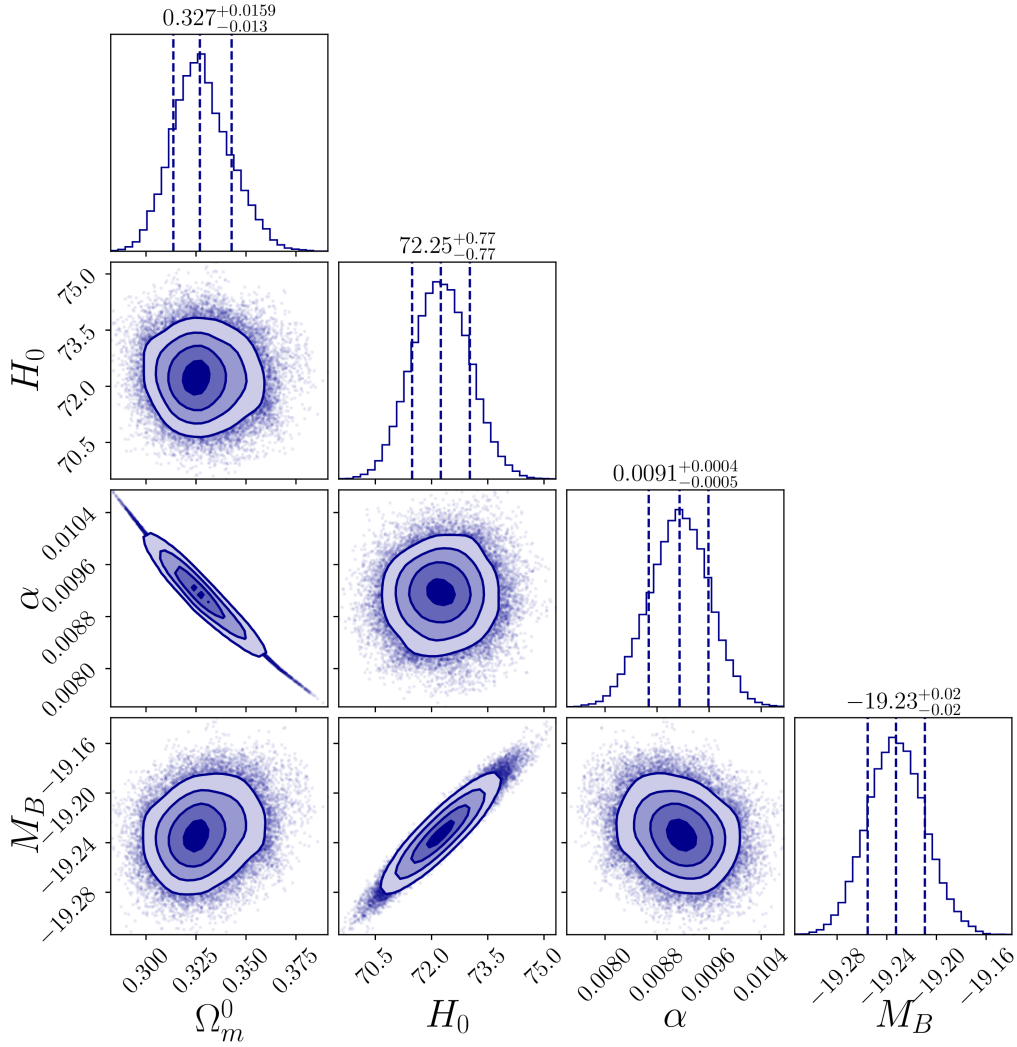


Figure 4.2: The corner plot for the best seed, with 0.5σ, 1σ, 1.5σ, and 2σ contour lines. The mean and 1σ error of each parameter is above its 1-d distribution. The MCMC parameters are $(\Omega_m^0 h^2, \alpha, M_B)$. H_0 and Ω_m^0 are derived parameters. Note that the perfect correlation between α and Ω_m^0 is an artefact of having fixed the seed number (see Section 3.3.4).

why Ω_m^0 does not get constrained by the Pantheon+ data, let us look at the expression for χ^2, ignoring the likelihood for calibrator events and also the covariance between different events for the sake of argument:

$$\chi^2 = \sum_i \left(\frac{1}{\Delta m_i^2}\right) \left[5\log_{10}\left((1+z_i)\int_0^{z_i} \frac{dz'}{H\,(z')}\right) + 25 + M_B - m_i\right]^2 \tag{4.7}$$

$$= \sum_i \left(\frac{1}{\Delta m_i^2}\right) \left[5\log_{10}\left((1+z_i)\int_0^{z_i} \frac{dz'}{\sqrt{\Omega_m^0(1+z)^3 + (1-\Omega_m^0)g(z)}}\right) - 5\log_{10}(H_0) + 25 + M_B - m_i\right]^2 .$$

Here, the summation is over all the supernova data points, and m_i is the apparent magnitude of the i-th supernova. $g(z)$ is some function by which the dark energy density is varying in redshift for the particular seed. $g(z)$ is completely different for different seeds. The supernova data roughly follows a straight line in the μ vs $\log(z)$ plot at low redshift. This is because $g(z) \to 1$ as $z \to 0$, and hence the integrand in the second line does not depend on Ω_m^0. Hence, the y-intercept of the line μ vs $\log(z)$ for small redshifts ($z \ll 1$) solely depends on H_0. Therefore, the overall intercept of the plot fixes H_0. Ω_m^0,

on the other hand, controls the deviation from the straight line at high redshifts ($z > 1$) through the integral. The best fit Ω_m^0 should then be found from $\frac{\partial \chi^2}{\partial \Omega_m^0} = 0$. As each seed gives a completely different $g(z)$, the value of Ω_m^0 also changes significantly from one seed to another, as seen in our results.

In Figure 4.2, we show the results from the best seed that has a best-fit $\chi^2 = 1481.9$. Its best-fit parameters are

$$\begin{aligned} H_0 &= 72.18\,, \\ \Omega_m^0 &= 0.33\,, \\ \alpha &= 0.0092\,. \end{aligned} \tag{4.8}$$

The figure shows the 2-dimensional contour plots and 1-d likelihood distributions for Ω_m^0, H_0, α, and M_B. H_0 is a derived parameter for the analysis. The plots show that Ω_m^0 and H_0 are almost uncorrelated. As α is a function of Ω_m^0, they are inversely related (see Section 3.3.4).

4.1.2 Diagnosing good seeds by a single parameter

Our analysis shows that it is possible to get expansion histories from Model 1 that fit the supernova data better than the standard ΛCDM. However, the probability of getting such seeds is small (16 out of 90000, or about 0.017%). The small number of these runs raises the question of the practical ability of Model 1 to generate data that can be compared with observations. For a deterministic and non-random model, it is straightforward to assess its success with the data as it comes with a fixed χ^2 value. However, more care is needed to judge the results from a stochastic model such as Everpresent Λ.

From a Bayesian perspective, the predominance of seeds that give a poor fit to SN Ia creates a large volume in the parameter space with a very low likelihood. Assuming that all seeds have the same prior probability, this would suppress the evidence for the model. This could lead to a small Bayes factor when compared to ΛCDM, hence undermining the model from a Bayesian model comparison perspective. One way to rescue the model would be to find theoretical reasons to reject some of the seeds a priori. In its current form, however, Everpresent Λ does not prohibit any seeds apart from those that encounter an $H^2 < 0$ crash before reaching the present time. In this section, we aim to distinguish between the seeds with the lowest χ^2 ("good" seeds) and those with a relatively high χ^2 ("bad" seeds) a posteriori. An important question in this regard is: what are the physical features of those realisations that give the best fits to the supernova data, and why are such seeds rare? We leave the question of whether there are compelling theoretical reasons for preferring the seeds with such physical features *a priori* to future research, along with a rigorous model comparison.

In Figure 4.3, 4 of the seeds with the lowest χ^2 (good seeds) and 4 with a relatively high χ^2 (bad seeds) are shown for comparison. The solid fluctuating lines are the energy density in Λ (normalised using a fixed constant $h_0 = 70.2$km/s/Mpc, not equal to their Hubble constant), and the dashed curves are the corresponding matter density. Visual inspection provides a distinction between these two extremes: In the range $z < 1.5$, a good seed generally has a smaller Λ than a bad seed. Moreover, bad seeds are more closely tracking the matter density, while for good seeds Λ fluctuates as if it is independent from matter density in the range $z < 1.5$. A more quantitative distinction can be made using the so-called $Om(z)$ diagnostic.

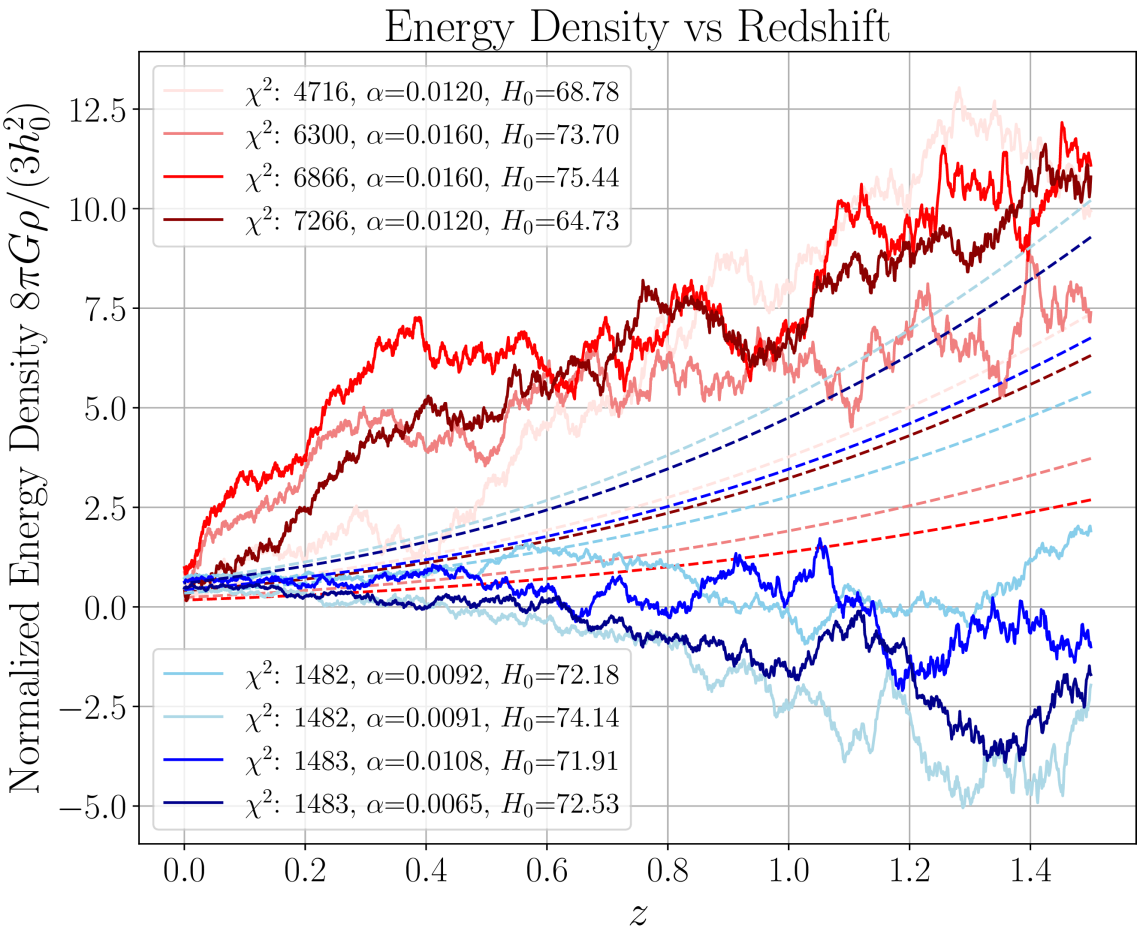


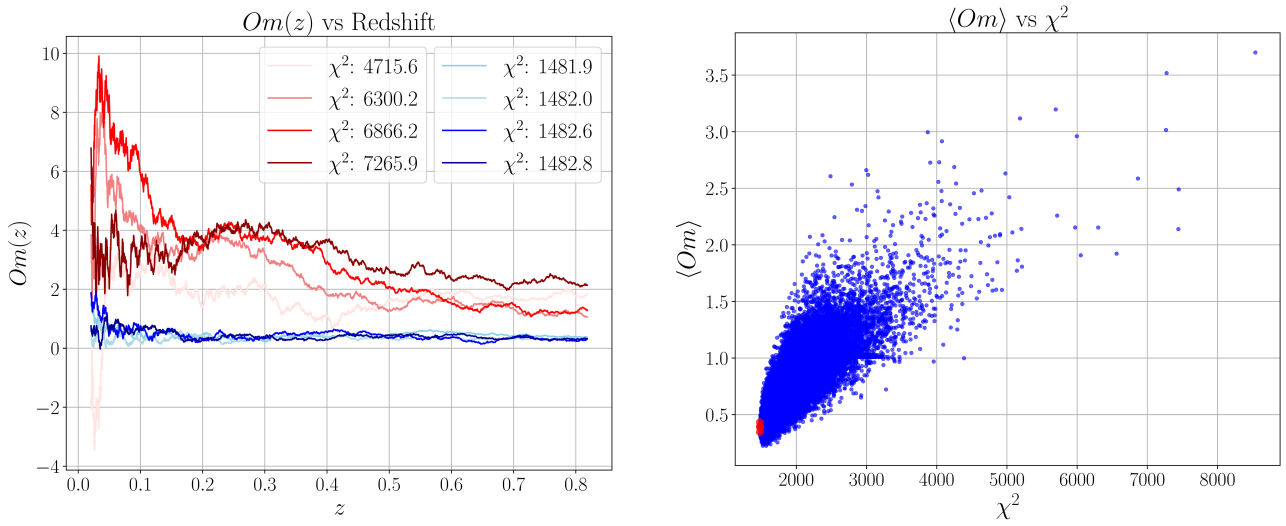
Figure 4.3: $\Lambda/(3h_0^2)$ versus redshift is plotted for the 4 realisations of Model 1 with smallest χ^2 for Pantheon+SH0ES data set, with blue solid curves. The dashed curves are the corresponding $8\pi G\rho_m/(3h_0^2)$. Note that $h_0 = 70.2$km/s/Mpc is a fixed constant we use for normalisation. For comparison, the same quantities are plotted in red for 4 realisations that give very high χ^2 for the supernova data. In the 4 "good" realisations, Λ tends to cross the corresponding matter density fewer times, and stays smaller than the matter density for a larger redshift interval.

The $Om(z)$ diagnostic was first introduced in [280] as a discriminator between different dark energy models. It is defined as

$$Om(z) = \frac{\left(\frac{H(z)}{H_0}\right)^2 - 1}{(1+z)^3 - 1}\,. \tag{4.9}$$

For ΛCDM, $Om(z) = \Omega_m^0$, and for pure CDM, $Om(z) = 1$. It has the property that at small redshifts $z < 1$, it magnifies any deviations from the standard dark energy equation of state, $w = -1$. For example, for a quintessence dark energy (i.e. $-1 < w < 0$), as a function of increasing redshift z, Om is decreasing and positively curved, while for a phantom dark energy (i.e. $w < -1$) it is increasing and negatively curved. In both cases, Om converges to Ω_m^0 for $z \gg 1$.

In the left plot of Figure 4.4, we have plotted the Om function versus z for 4 of the best and worst seeds. The Om diagnostic makes their distinction even more apparent. Bad seeds have a bump around $z \sim 0.3$ and then decrease with redshift. This is like a combination of quintessence and matter-like behaviour. For good seeds, however, Om fluctuates around some horizontal line and has much smaller values. The plot on the right shows the scatter plot of the average of Om in the interval $z \in [0.1, 0.8]$



(a) (b)

Figure 4.4: Left plot: $Om(z)$ versus redshift for the same 8 realisations as in Figure 4.3. The distinction between good and bad seeds is clear here. After large oscillations for $z < 0.1$, the seeds with a good fit to the supernova data settle and fluctuate around some horizontal line, while the seeds with a bad fit have a bump-like feature around $z \sim 0.3$. Right plot: For all 90000 seeds, we compute $\langle Om \rangle = \frac{1}{0.7}\int_{.1}^{.8} Om(z)dz$. A scatter plot of $\langle Om \rangle$ versus the χ^2 is shown. There is a strong correlation between these two quantities. The Pearson correlation coefficient of the plot is 0.85. The 4 best seeds are marked with red dots.

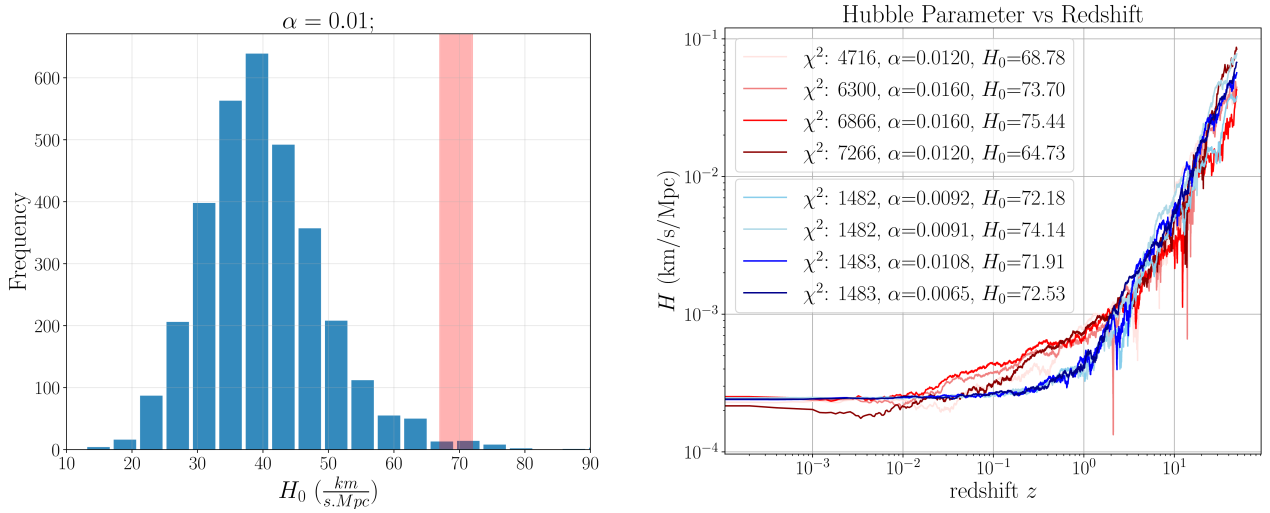
versus the best fit χ^2 for all 90,000 seeds. There is a strong correlation between these two quantities. The good seeds tend to have a smaller $\langle Om \rangle$. In fact, we find that for the 70 seeds with $\chi^2 < 1060$, $\langle Om \rangle = 0.27 \pm 0.03$. Therefore, not only is the Om behaviour of a good seed reminiscent of constant Λ, it also fluctuates around the Om value of ΛCDM and not around Ω_m^0 of the same run (recall that the value Ω_m^0 of each seed, according to Figure 4.1, is not concentrated around 0.3).

Regarding the practicality of the model, this result suggests that to obtain good seeds with a higher probability, one should focus on a sub-space of seeds that give smaller values of $\langle Om \rangle$ in the low-redshift range $z \lesssim 0.8$. It further indicates that to find good fits to data for a stochastic model such as Everpresent Λ, an understanding of the physical constraints that have to be satisfied is required in order to strategically limit one's search to seed numbers that are more likely to satisfy those constraints. For the supernova data, the single number $\langle Om \rangle$ has turned out be a good constraint parameter.

Next, given that only a small fraction of the seeds produce better fits to the data than ΛCDM, we would like to further understand if their behaviour is markedly different from the typical behaviour of Everpresent Λ.

4.1.3 Understanding the results with mock data

To confirm the exceptionality of the good seeds, we perform a mock data analysis. Our procedure for generating mock supernova data is as follows: for a realisation of Model 1, we find the value of the distance modulus from (4.1) at 1580 redshifts greater than 0.01, corresponding to the non-calibrator SN Ia events in the Pantheon+ dataset. Using their full covariance matrix, we add Gaussian random noise to the corresponding distance moduli. In this way we obtain supernova data that could occur in a universe described by a particular Everpresent Λ realisation.



(a) (b)

Figure 4.5: Left plot: The distribution of H_0 for Model 1 with $\Omega_m^0 h^2 = 0.14$, $\alpha = 0.01$ (taken from the previous chapter). The average and standard deviation is given by 40.3 ± 9.3. Notice the difference between this distribution and the posterior distribution in Figure 4.1. The range $[67, 72]$ for H_0 is shaded in red. Right plot: The Hubble parameter H as a function of redshift for four of the best and four of the worst seeds. Good seeds show much smaller fluctuations at $z \lesssim 1.0$.

As an interesting first example, consider the following case: Using an arbitrary seed and fixing the parameters $(\alpha, H_0, \Omega_m^0) = (0.01, 70.4, 0.28)$, we generated mock data. After running the MCMC for this mock data set, the best fit for Model 1 yielded the parameters $(\alpha, H_0, \Omega_m^0) = (0.0047, 71.0, 0.81)$. This again confirms the observation from the previous subsection: The supernova data does not constrain Ω_m^0 for Model 1. The model can fit the data, whether real or mock, with almost any value of Ω_m^0. In other words, the model does not make a prediction for Ω_m^0, even for a data set generated by itself.

We are then interested in the following fraction for this mock data set, D:

$$f_D := \frac{\#\{seed = s | \text{Model 1 with } seed = s \text{ gives a better fit to } D \text{ than } \Lambda\text{CDM}\}}{\text{Total number of seeds}}. \tag{4.10}$$

We previously saw that if we take D to be the Pantheon+SH0ES data set, then $f_D = 16/90000$. The question then is what the fraction f_D is if we take D to be a mock data set generated by Model 1. Furthermore, we are interested to know if the mock data generated by an arbitrary seed has different characteristics from the mock data generated by a good seed.

First, let us investigate the question for an arbitrary seed. We try to be realistic with the choice of the seed s for generating the mock data. In the left plot of Figure 4.5, taken from the previous chapter, we show the distribution of the present day Hubble parameter for Everpresent Λ realisations with parameters $\Omega_m^0 h^2 = 0.14$ and $\alpha = 0.01$. Most of the expansion histories give a value of the Hubble parameter close to 40 km/s/Mpc. Only a few realisations give Hubble parameters close to 72 km/s/Mpc, i.e. close to the present cosmology. To better mimic the real data, we pick a set of 10 seed numbers from the tail of this distribution, as the value of $H_0 \sim 72$km/s/Mpc lies in the tail. The expansion history for each of these seeds give values of H_0 in the range $\sim [69, 72]$ km/s/Mpc. There is no further constraint on the choice of seed. We then do a similar analysis as before. We take 10,000 arbitrary

seeds and then try to find the best fit χ^2 for each of these seeds with each of the mock data sets. For the 10 mock data sets, we find that f_D lies in the interval $[0.026, 0.205]$, with an average of 0.076. Thus it is around three orders of magnitude larger than the f_D for the real data.

However, if we instead use one of the good seeds to generate the mock data, f_D turns out to be $1/10000$, which is compatible with the real data. This analysis further confirms that the seed numbers that are in the lowest corner of the right plot in Figure 4.4 are distinguished, and they more closely resemble the late-time expansion history of our universe. This result indicates that Everpresent Λ is easily capable of producing a momentary value of H_0 that fits present observations, but that it is more difficult to produce a history that has the right character to fit all data.

4.2 Fitting Model 1 to CMB data

Supernova data give a direct measurement of the Hubble parameter. However, they can constrain dark energy only at low redshifts. To constrain the expansion history at larger redshifts, we can utilise cosmological perturbation data from the CMB. In this section, we discuss our analysis of Model 1 (and briefly Model 2) of Everpresent Λ with CMB data.

4.2.1 Methodology

To investigate Model 1, we use a similar technique as used in our supernova data analysis in the previous section. Given a seed number s, the remaining free parameters are the standard 6 parameters $\Omega_m^0 h^2$, $\Omega_b^0 h^2$, h, τ_{reio}, n_s, and A_s. Notice that, according to our discussion in Sections 3.3.4 and 4.1.1, for a given seed, α and Ω_m^0 are related. So α is not a free parameter, since we can infer it using $\Omega_m^0 h^2$ and h. We then feed the dark energy density as a function of scale factor, together with the standard 6 free parameters, to a modified version of the `CAMB` Boltzmann code [281]. In the current form of Everpresent Λ, dark energy is not a fluid, so we turn off the dark energy perturbations. See Appendix C.2 for a discussion on the fluid interpretation.

From Figure 4.1b, we know that if we change s at every MCMC step, the likelihood surface would not be continuous. So, similar to our SN Ia analysis, we select 100,000 seeds, and for each one, we run a separate MCMC chain using `SCoPE` [282] along with the Planck likelihood code (high-ℓ, low-ℓ, lensing likelihood) [39, 158].

Running the MCMC for Model 1 is time-consuming, especially if we recalculate the expansion history for every choice of parameters. Note that the nature of the expansion history for a given seed does not depend on $\Omega_m^0 h^2$. For a given seed, quantities like Λ, H^2, and $1/\sqrt{V}$ scale linearly with $\Omega_m^0 h^2$ as discussed in Section 3.3.4. So we can make the process more efficient with the following simplifications: First, we calculate the expansion history for a fixed $\Omega_m^0 h^2$, say 0.14. The histories with different values of $\Omega_m^0 h^2$ can be determined by the scaling property mentioned above. Secondly, $H^2(a)$ is a slowly varying function of α. Therefore, we calculate the expansion histories for different α values at a fixed interval and store them. These values are then interpolated for calculating the expansion history for any other α and then scaled according to $\Omega_m^0 h^2$. Note that we start the Everpresent Λ simulations from about $z = 10^5$. Therefore, there is a significant redshift range which is radiation dominated. For any Ω_m^0 we calculate the matter-radiation equality as $a_{eq} = \Omega_r^0/\Omega_m^0$ and scale the dark energy separately

during the matter and radiation dominated era. However, the scaling relation holds, approximately, for a completely matter dominated or completely radiation dominated era. It does not hold during matter and radiation equality. Nevertheless, this region is small in conformal time and has negligible impact on C_ℓ, as we verified numerically.

Therefore, for any given seed, we make the fiducial choice $\Omega_m^0 h^2 = 0.14$. We then generate expansion histories for 20 different α values starting from $\alpha = 0.001$ to 0.02 in steps of $\Delta\alpha = 0.001$. The smallest α for which we get an $H^2 < 0$ crash is denoted by $\alpha^{\max}$ for that seed. As α and Ω_m^0 are inversely related, it corresponds to some ${\Omega_m^0}^{\min}$. For any value of α (Ω_m^0) greater (less) than $\alpha^{\max}$ (${\Omega_m^0}^{\min}$), the expansion history for that particular seed results in an $H^2 < 0$ crash (see Section 3.3.4). To ensure realistic outputs, we only keep seeds that (i) yield $H_0 > 60$ km/s/Mpc for at least one choice of α, and (ii) have ${\Omega_m^0}^{\min} < 0.45$. These also guarantee a sufficiently high $\alpha^{\max}$, improving the chances of matching H_0. This reduces our seed set from 100,000 to 3106. Finally, for ΛCDM, the best fit χ^2 for the CMB data is about 2339.0 for the given likelihood. Therefore, we also discard seeds that yield $\chi^2 > 3000.00$ after the burn-in period.

4.2.2 Results

With the reduced set of seeds, our MCMC shows that there are 91 seeds that give a χ^2 less than 3000.0. The best fit χ^2 is 2495.56, about 157 units larger than that of ΛCDM. In the first row of Figure 4.6, we have summarized the statistics of our analysis. The first and the second columns show the distribution of the best fit values of $\Omega_m^0 h^2$ and h for the 91 seeds. In all of these cases, the best fit H_0 is between 40 and 51 km/sec/Mpc. The value of $\Omega_m^0 h^2$ in the MCMC chains lingers around 0.16 to 0.2 in most of the cases (82 out of 91 chains), occasionally going below 0.16 such that their best fit values are in many cases near 0.14. Nevertheless, this indicates that the MCMC analysis favours a small dark energy fraction, typically $\Omega_\Lambda^0 \leq 0.2$. The last column shows the best fit χ^2 of the final 91 seeds. Except for one seed, all others have a χ^2 greater than 2500.

Figure 4.7 compares the best-fit power spectrum from Model 1 with that of ΛCDM. Model 1 shows notable deviations at low multipoles and near acoustic peaks, with similar behaviour seen in two other representative seeds.

Motivated by the claim in [66] that the Everpresent Λ models need sufficiently small dark energy at the last scattering surface to explain CMB, we repeat our analysis with an artificial suppression of dark energy at the last scattering surface using a factor $1 - \exp\left(-\frac{1}{2}\left(\frac{z-1100}{\Delta z}\right)^2\right)$ for $\Delta z = 100$ and 200. For $\Delta z = 100$ and 200, we get 248 and 539 seeds that yield a χ^2 less than 3000, respectively, with a ~ 136 unit improvement in best-fit χ^2. The best fit parameters and χ^2 for both the analyses are given in Table 4.1. The histograms for $\Omega_m^0 h^2$ and h, and χ^2 for the best fit for each seed are shown in the second and the third row of Figure 4.6. Increasing Δz in the suppression factor has the effect of giving more $\Omega_m^0 h^2$ values closer to 0.14, as well as stretching the tail of h to higher values. These results partially support [66], however, the MCMC process still tries to remove the dark energy component and yields too small values for H_0. In its current formulation, Model 1 does not explain the CMB data at a level comparable to ΛCDM.

We also carried out tests using Model 2 of Everpresent Λ, with 3000 seeds. All the seeds yield $\tilde{\alpha} < 0.1$

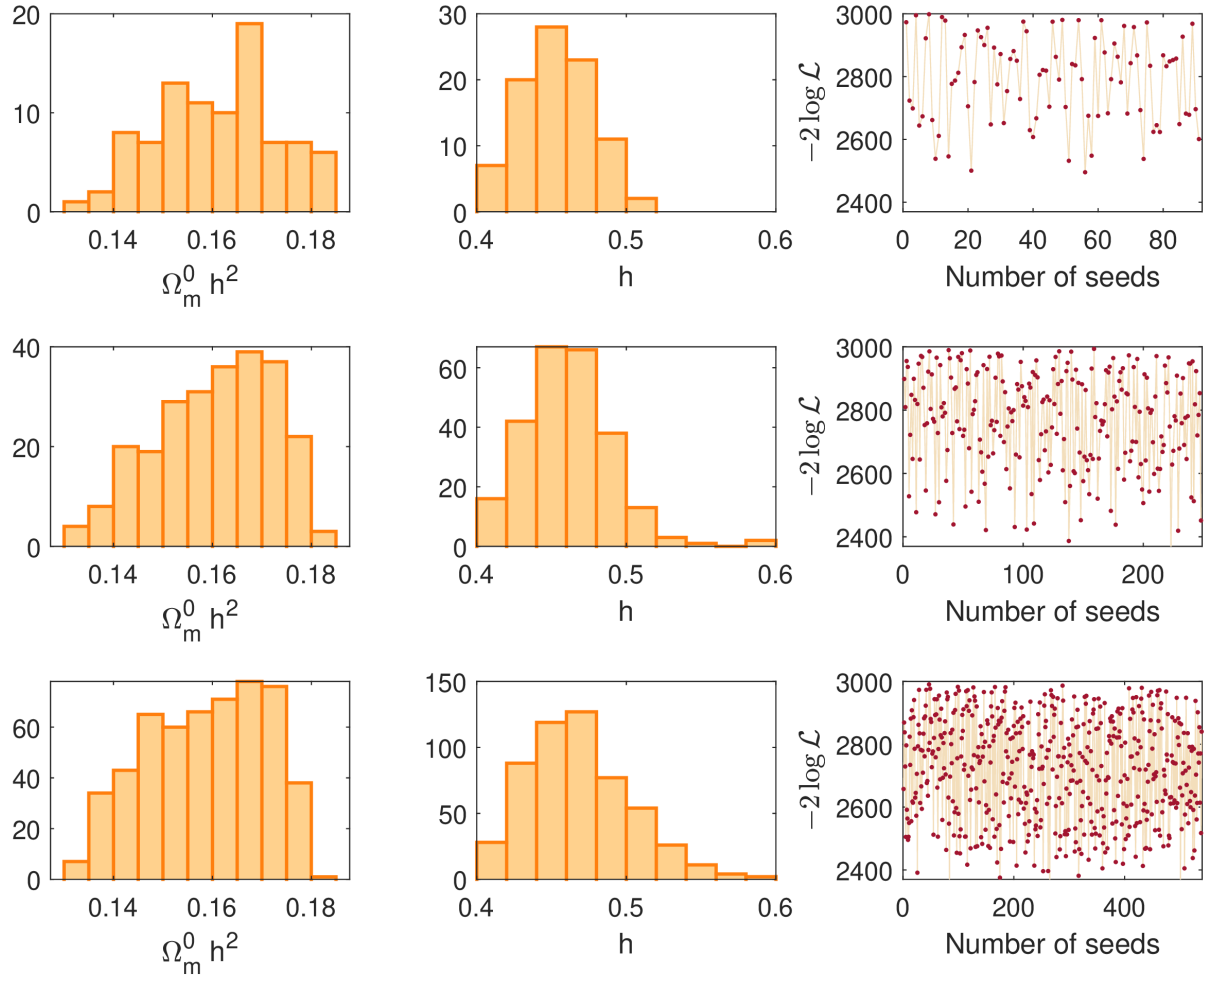


Figure 4.6: The results of our CMB analysis with Planck likelihood from 100,000 seeds. The first row is the output from Model 1; the second and the third row are from Model 1 with last scattering surface suppression with $\Delta z = 100,\ 200$, respectively. The first two columns are the histograms of the best fit $\Omega_m^0 h^2$ and h for the seeds that give $\chi^2 < 3000.0$. The histograms show that $\Omega_m^0 h^2$ is often higher than the standard ΛCDM value of 0.14, and most of the best fit h are within the range $0.4 - 0.5$, and it rarely goes to 0.6.

(or alternately $\Omega_m^0 > 0.9$) and μ fails to converge. The Hubble parameter stays close to 40, and the model generally gives poor fits.

In [66] Model 2 was implemented using a modified equation of state parameter in `CAMB`. Their equation of state parameter $w(z)$ was restricted to a narrow negative range [284]. This is a major point of departure from our treatment, in which we allow the dark energy density to fluctuate both in sign and in magnitude within the full range that the model permits. Therefore, our results should not be compared to those of this previous work.

4.3 Constraining dark energy using CMB

A key property of Everpresent Λ is that the dark energy density is of the order the background energy density throughout the expansion history of the Universe. Models 1 and 2 capture this behaviour in a practical manner. However, both of these models are phenomenological, as we require a more developed

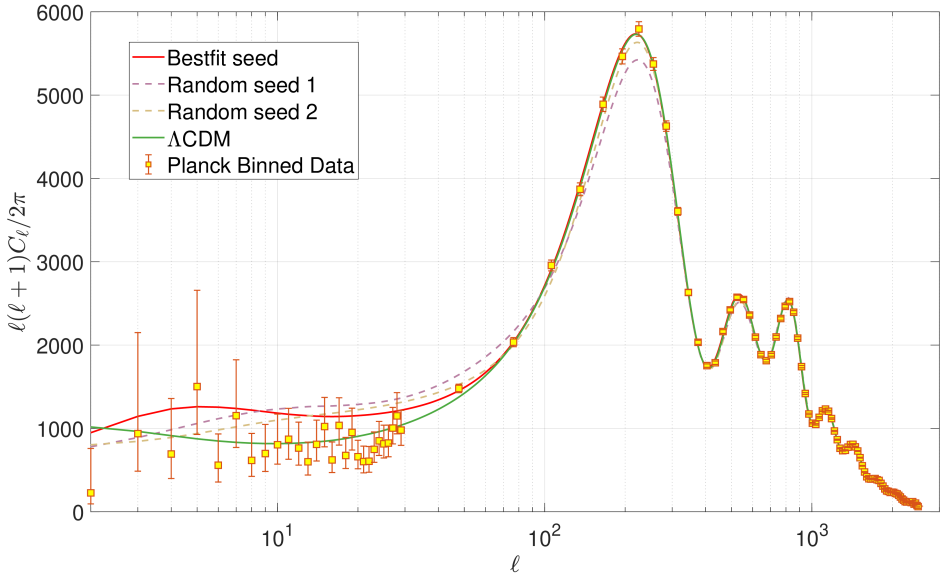


Figure 4.7: We show the CMB power spectrum data for different models. The green C_ℓ curve is the ΛCDM power spectrum. We show the best fit power spectrum for Model 1 out of 100,000 seeds in red. Two other dashed lines show the best fit power spectrum for two other arbitrarily chosen seeds from the 91 seeds for which $\chi^2 < 3000.0$. The binned data from the Planck satellite is also displayed in the plot. For values of ℓ ranging from 2 to 29, the data are shown individually. After that, the data is presented in bins based on the value of ℓ. The Planck data is taken from [283].

	Model 1	Model 1 with last scattering surface suppression ($\Delta z = 100$)	Model 1 with last scattering surface suppression ($\Delta z = 200$)
$\#seed_{-2\log\mathcal{L}<3000.0}$	91	248	539
$-2\log\mathcal{L}$	2495.6	2360.5	2358.5
	Best-fit parameters		
$\Omega_m^0 h^2$	0.1408	0.1409	0.1430
$\Omega_b^0 h^2$	0.0223	0.0222	0.0221
H_0	43.27	43.20	42.68
$\tau_{\rm reio}$	0.0866	0.1319	0.1285
n_s	0.9622	0.9777	0.9632
$\log(10^{10}A_s)$	3.0892	3.1918	3.1875
α (derived)	0.0040457	0.0036563	0.0025567

Table 4.1: Tabulated results of best-fit values from 100,000 seeds with Planck likelihood. We show the results for Model 1 and Model 1 with last scattering surface suppression. For comparison, $-2\log\mathcal{L}$ for ΛCDM is ~ 2339.0.

theory of quantum gravity to inform the correct dynamics for Everpresent Λ. Furthermore, as discussed in the previous section, the present models do not fit the CMB data. Therefore, in this section, we try to understand how the variation in dark energy in different eras affect the CMB power spectrum. These studies can give an indication of the constraints on a fluctuating dark energy density. It can also inform how future models may need to be modified.

According to the line-of-sight integration method [285], the l-th moment in a multipole expansion of the temperature fluctuation Θ of CMB photons is given by

$$\Theta_\ell(\eta_0, k) = \int_0^{\eta_0} d\eta \; j_\ell \left(k(\eta_0 - \eta)\right) S(\eta, k)\,, \tag{4.11}$$

where j_ℓ is the spherical Bessel function, k is the wavenumber, and S is the source function defined by

$$S(\eta, k) = g(\eta)(\psi + \Theta_0) + \frac{1}{k^2}(g\theta_b)' + e^{-\tau_{\rm opt}}(\psi' + \phi')\,. \tag{4.12}$$

Here, ψ and ϕ are the Newtonian gravitational potentials in the perturbed FRW metric (1.41), θ_b is the velocity divergence of baryons, $\tau_{\rm opt}$ is the optical depth at η, and $g = -\tau'_{\rm opt}\exp(-\tau_{\rm opt})$ is the visibility function. The first term in the source function is the SW term, and it captures the temperature fluctuations and gravitational redshift during recombination. The second term accounts for the Doppler effect during recombination. While g is non-zero only during recombination and reionization, the third term– the ISW effect–accounts for the gravitational redshift of the photons throughout their free-streaming. The effect of polarization fluctuations is subdominant, and we have ignored it.

Therefore, the expansion history of the Universe can be divided into three phases; before, during, and after recombination. In the following three subsections, we first discuss the qualitative changes in the CMB power spectrum when a dark energy component is added in each of these phases. We then numerically constrain the allowed variation in the expansion history before and after the last scattering surface. While other analyses have examined the permitted variations in the expansion history at certain epochs [286–288], here we provide a model-independent analysis of the allowed fluctuations, positive or negative, in dark energy across the entire expansion history.

4.3.1 During recombination

Here, our main aim is to understand why our fits of Model 1 improved when we suppressed dark energy around the last scattering surface. Suppose we add a non-negligible dark energy component during the recombination era. Specifically, we keep $\Omega_\Lambda(z)$ constant around recombination between redshifts $z = 900$ and 1300. For other redshifts, we keep Λ constant such that $\Omega^0_\Lambda = 0.7$. We first examine how it affects the position of the peak and the width of g.

Decoupling of photons happens when the free electron density drops due to recombination. From the Saha equation, the free electron density n_e depends only on the photons' temperature, which is inversely proportional to the scale factor. Changing the dark energy fraction does not change the dependence of n_e on the scale factor. The position of the visibility peak, therefore, remains the same in a. In conformal time, however, it shifts to lower (higher) τ with a positive (negative) dark energy density.

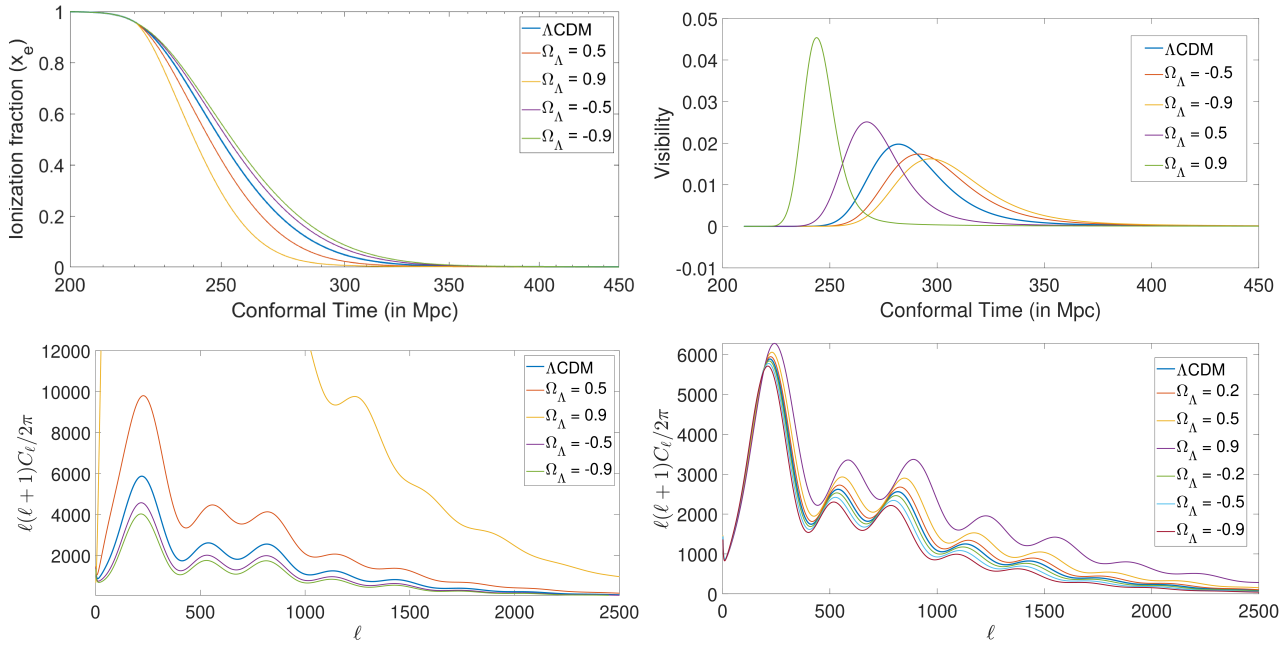


Figure 4.8: Top-left: The ionization fraction (x_e) as a function of conformal time for different dark energy fractions. Top-right: The visibility function (g) as a function of conformal time, Bottom-left: C_ℓ for different dark energy fractions between $z = 900$ and 1300. Bottom-right: C_ℓ for different dark energy fractions between $z = 900$ and 1300 provided the dark energy only changes the ionization fraction.

The width of the visibility function in conformal time is approximately given by [289]:

$$\Delta\eta_g = \frac{z_*^2}{1.44 \times 10^4 H_*} \,. \tag{4.13}$$

The starred quantities are evaluated at the peak of the visibility function. Therefore, by adding a positive (negative) dark energy component during recombination, H_* increases (decreases). As a result, the width of g decreases (increases), with the area under the curve remaining fixed. These effects are shown in Figure 4.8. In the top-left plot of Figure 4.8, we show how the ionization fraction as a function of conformal time changes. The top-right plot shows how the visibility function changes as a function of the conformal time, which is in line with the above expectations.

To see how such changes to the visibility function affect the CMB power spectrum, consider the following semi-analytic approximation to C_ℓ that comes from SW and Doppler terms, as these are the dominant contributions for $\ell \gtrsim 200$ [289]:

$$\ell(\ell+1)C_\ell \simeq \frac{B}{\pi}\left[\mathcal{A}_0 + \sqrt{\frac{\pi}{\ell\theta_s}}\Big(\mathcal{A}_1 \cos(\ell\theta_s + \pi/4) + \mathcal{A}_2 \cos(2\ell\theta_s + \pi/4)\Big) e^{-\ell^2/\ell_D^2}\right] , \tag{4.14}$$

where B is a constant from the primordial power spectrum, $\mathcal{A}_0$ is a non-oscillating function of ℓ, and $\mathcal{A}_1$ and $\mathcal{A}_2$ are slowly varying functions of ℓ. θ_s is given by

$$\theta_s = \frac{r_s}{\eta_0} \,, \tag{4.15}$$

and is approximately the angular size of the sound horizon at decoupling with r_s being the sound horizon. The oscillatory part is a direct consequence of the acoustic oscillations in the photon-baryon

plasma that gave rise to a WKB solution to the temperature fluctuations $\Theta_0 \sim \exp(ikr_s)$. Finally, ℓ_D represents the damping scale, and is given by

$$\frac{1}{\ell_D^2} = 2\frac{\lambda_S^2}{\eta_0^2} + 2\frac{\lambda_L^2}{\eta_0^2}\,. \tag{4.16}$$

The first term gives the Silk damping which is responsible for damping inhomogeneities on scales smaller than the comoving diffusion scale of photons, λ_S. The second term gives the Landau damping which is due to the smoothing of fluctuations resulting from photons not all being last scattered instantaneously; specifically giving a comoving scale $\lambda_L \propto \Delta\eta_g$. Both of these effects are responsible for the suppression of the power spectrum at $\ell > 1000$.

Now, for a positive dark energy contribution during recombination, the resulting smaller conformal time of decoupling leads to a smaller sound horizon and θ_s, which from the above solution (4.14) amplifies oscillations, shifts peaks to smaller scales (higher ℓ), and increases peak separation. At the same time, a smaller diffusion length of photons and a smaller width of g both reduce the exponential suppression at large ℓ. The opposite happens for a negative dark energy. These trends are confirmed by the bottom two panels in Figure 4.8.

Therefore, the position of the peaks in the TT power spectrum, their amplitude, and their high-ℓ damping strongly constrains the fraction of dark energy during recombination. To have a chance at explaining the CMB data, a fluctuating dark energy has to be suppressed around the last scattering surface, as it was imposed on Model 1.

4.3.2 After recombination

Here we aim to understand how much temporal fluctuation in dark energy is permitted after recombination. One of the main contributions in this phase comes from the ISW effect. The ISW term contains the time derivative of the gravitational potentials. For ΛCDM, ϕ and ψ are constant during pure matter domination. So ISW consists of early ISW from just before recombination when radiation energy density is not yet negligible, and late ISW at $z \lesssim 1$ when dark energy becomes important. The latter only contributes to low multipoles. However, for Everpresent Λ, there never exists a pure matter domination era, as dark energy contributes with the same order of magnitude but with fluctuations that change sign. Therefore, we expect a significantly different ISW contribution in this case. We can already see the resulting enhanced low-ℓ power for Model 1 realisations in Figure 4.7. The question is whether there exist realisations of a stochastic model like Everpresent Λ such that the integrated effect of the Λ fluctuations cancel each other in the ISW term. We evaluate this using a model-independent approach.

For this, we treat the values of a dimensionless function f(z) at five selected redshifts 1, 10, 50, 100 and 500, as free parameters. This function parametrises the dark energy density through

$$\rho_\Lambda(z) = \rho_c\Omega_\Lambda^0 + \rho_c \mathrm{f}(z)(1+z)^3\,, \tag{4.17}$$

which is a constant plus a fluctuating term that tracks the matter density. Note that ρ_c is the critical density today. A spline is then constructed to interpolate smoothly between these values, with boundary

conditions of f and its derivative being zero at redshifts 0 and 1100. Since we are only considering redshifts after the last scattering surface, the radiation energy density can be ignored compared to matter.

For an MCMC analysis, a total of 11 parameters are considered, which include the standard six ΛCDM parameters $\Omega_m^0 h^2$, $\Omega_b^0 h^2$, h, A_s, n_s, and $\tau_{\rm opt}$; as well as the five f(z_i) values at the redshifts mentioned above. We let $\Omega_\Lambda^0 = 1 - \Omega_m^0$. The prior on f($z_i$) allows the dark energy density to be of the order of the ambient matter density.

The full posterior plots are shown in Appendix C.4. The best fit and the average values of the parameters are listed in Table 4.2. The fitted values f(z_i) are small, indicating that the magnitude of dark energy fluctuations must remain below 5% of the matter density. The rejection rate is high, reflecting the complex likelihood landscape with multiple minima, which also implies that the mean function f(z) plotted in solid lines in Figure 4.9 and Figure 4.10 does not necessarily give a good fit.

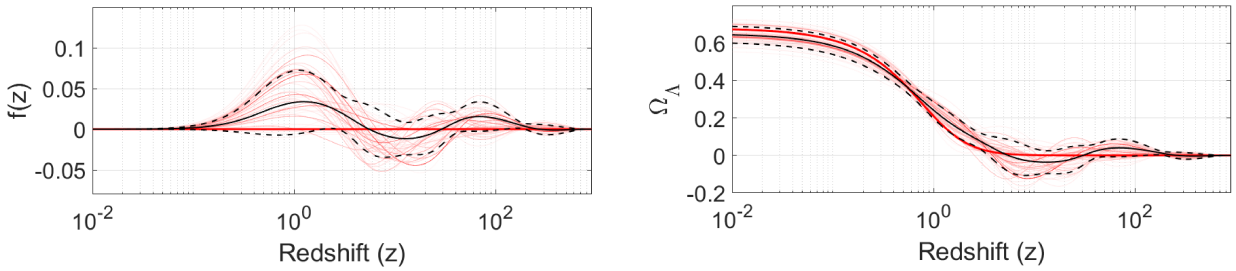


Figure 4.9: A few random expansion histories from the 200 best fit χ^2 are shown. The opacity of each curve is proportional to its likelihood. The thick red line shows the standard ΛCDM values for reference. The mean and the standard deviation are shown in solid black and dashed black curves. Left: plot of f(z). We keep the f(z) and its derivatives at redshifts $z = 0$ and $z = 1100$ to be 0. A bump at low redshift and a dip at higher redshift give a better likelihood than standard ΛCDM. Right: Ω_Λ is shown for f(z) from the left plot.

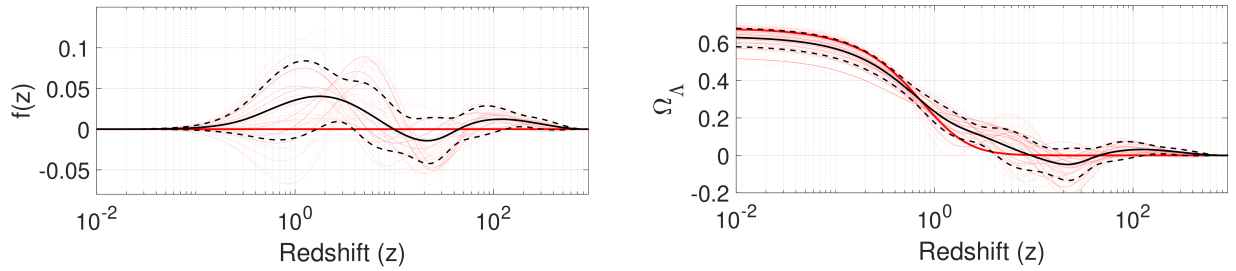


Figure 4.10: A few arbitrary expansion histories from the 200 best fit χ^2 for a different set of z points (i.e. the points where f(z) is a free parameter) are shown. Left: plot of f(z). Even though the z points are different, the qualitative behaviour of these features is similar to that in Figure 4.9. Right: Ω_Λ is shown for f(z) from the left plot.

Figure 4.9 shows representative plots of f(z) (left) and Ω_Λ (right) from the top 200 fits. A consistent pattern emerges: a low-redshift bump and a high-redshift dip (around $z = 5 - 100$) in f(z) often improve the fit relative to ΛCDM. This feature was anticipated in [290] and is echoed here numerically. The corresponding plots of Ω_Λ indicate that most well-fitting models closely follow the standard evolution. These results are consistent with the results in [291, 292].

To test robustness, we repeated the analysis using a different set of redshift control points (0.5, 5, 25,

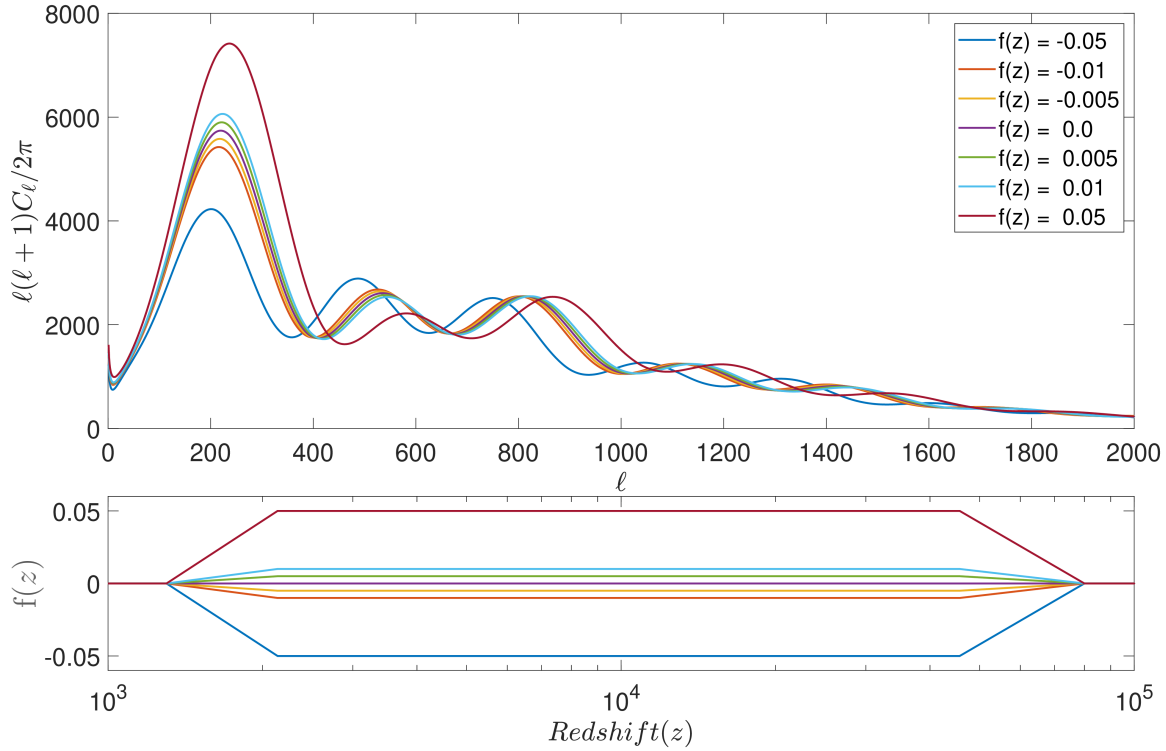


Figure 4.11: C_ℓ's for different fixed f(z) before recombination are shown. Positive f(z) moves the power spectrum towards higher ℓ and gives a higher power to the odd peaks, while negative f(z) pushes the spectrum towards lower ℓ and gives a higher power to the even peaks.

50, 250). As shown in Figure 4.10, the qualitative behaviour of f(z) remains unchanged, with a similar low$-z$ bump and high$-z$ dip, though the dip's exact location depends slightly on the chosen redshift points.

4.3.3 Before recombination

In this subsection we investigate how a fluctuating dark energy prior to the last scattering surface affects the CMB power spectrum. Before recombination, the major contribution to the CMB power spectrum comes from the SW and Doppler terms in (4.12). Similar to our discussion in Section 4.3.1, adding positive dark energy before the last scattering surface decreases r_s and subsequently lowers the value of θ_s. This then moves the CMB power spectrum towards higher multipoles. Conversely, a negative dark energy causes the power spectrum to shift towards lower multipoles. This is shown in Figure 4.11 where we add a positive or negative dark energy contribution before the last scattering surface. We keep f(z) nearly fixed before recombination$-$holding it constant, but allowing it to go continuously to 0 near $z = 10^3$ and 10^5.

Now, similar to Section 4.3.2, we study the amount of variation allowed in the dark energy. Before the last scattering surface, we write the dark energy density fraction as

$$\rho_\Lambda(z) = \rho_c \left[\Omega_{\Lambda_0} + \left((1+z)^3 + \frac{\Omega_r^0}{\Omega_m^0}(1+z)^4\right) \mathrm{f}(z)\right] . \tag{4.18}$$

We have added a term so that it models a fluctuating dark energy that tracks radiation energy density.

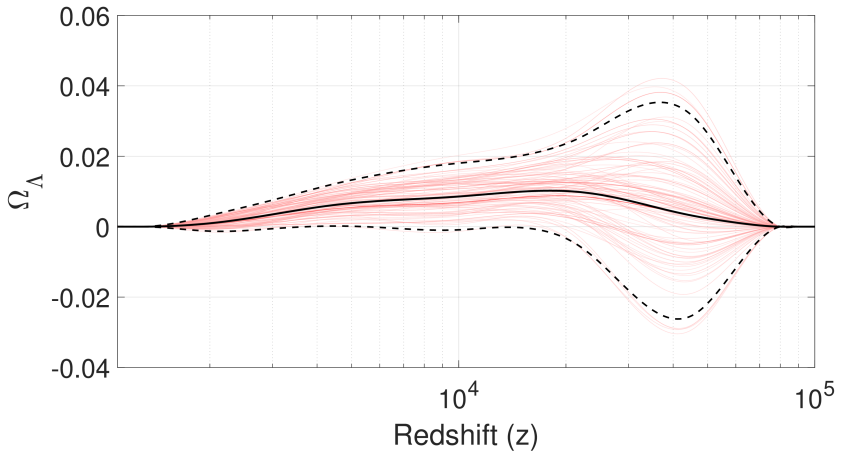


Figure 4.12: A few arbitrary realisations of Ω_Λ from the 200 best-fit χ^2 are shown. The solid black and the dashed black curves show the mean and the standard deviation of the curves. The opacity of a curve is proportional to its likelihood. Most of the best fit values have a positive Ω_Λ at high redshift.

As in the previous section, we select $\mathrm{f}(z_j)$ at five logarithmically spaced redshifts between $z = 1100$ and $z = 10^5$: $z_j = 2140$, 4570, 10000, 21380, 45710 as free parameters, and fit a spline $\mathrm{f}(z)$ through them. We fix $\mathrm{f}(z)$ and its derivatives to be 0 at $z = 1100$ and 10^5. Together with the six ΛCDM parameters, we have 11 free parameters for the MCMC analysis.

Table 4.2 and Figure C.3 in Appendix C.4 summarize the results. All average and best-fit $\mathrm{f}(z_i)$ values are positive, suggesting that small positive early dark energy is mildly favoured. The best fit χ^2 is about 9.02 lower than the best fit ΛCDM model. The contour plots show that the $\mathrm{f}(z_j)$ parameters are positively correlated, especially at high redshifts, indicating that $\mathrm{f}(z)$ cannot fluctuate arbitrarily. As shown in Figure 4.12, most of the best-fit expansion histories yield a small but positive $\Omega_\Lambda(z)$ at all redshifts, in contrast with Everpresent Λ where larger and more variable early-time Ω_Λ is typical.

While an extensive model comparison between our sign-agnostic dark energy fluctuations and the early dark energy literature [293, 294] is beyond the scope of this thesis, we note, as an example, that our results are consistent with those of [287], where the authors used two phenomenological scalar field models to implement early dark energy. They found that an extra 5% energy density fraction in the form of dark energy at $z \sim 5000$ which later dilutes faster than radiation gives a better fit to CMB than ΛCDM, and furthermore, alleviates the Hubble tension. We have found that with the phenomenological model (4.18), we can have a dark energy fraction of $(0.6 \pm 0.6)\%$ at $z \sim 5000$, and $(1.0 \pm 1.5)\%$ at $z \sim 20000$ which is consistent with the "$n = 2$" model in [287] within 2σ uncertainty. At the same time, our result is consistent with zero dark energy within 1σ. The reason for a less significant dark energy fraction in our case could be attributed to neglecting dark energy perturbations.

Introducing early dark energy in the form of (4.18) also increases the value of the Hubble parameter, reducing the Hubble tension with the SH0ES measurement [22] from 5σ to 3.5σ. This can be understood by noting that an additional dark energy contribution prior to recombination decreases the sound horizon r_s, and so, to maintain the same θ_s as measured by Planck2018, η_0 must decrease, or equivalently, H_0 must increase.

Modified DE after last scattering surface			Modified DE before last scattering surface		
	Average	Best Fit		Average	Best Fit
$\Omega_m^0 h^2$	0.1415 ± 0.0015	0.1415	$\Omega_m^0 h^2$	0.1426 ± 0.0021	0.1430
$\Omega_b^0 h^2$	0.0224 ± 0.00016	0.0224	$\Omega_b^0 h^2$	0.0227 ± 0.00026	0.0226
h	0.6370 ± 0.0384	0.6275	h	0.6910 ± 0.0121	0.6866
τ_{reio}	0.1238 ± 0.0295	0.1231	τ_{reio}	0.1303 ± 0.0287	0.1245
n_s	0.9650 ± 0.0045	0.9655	n_s	0.9758 ± 0.01	0.9713
$\log\left(10^{10} A_s\right)$	3.17 ± 0.0571	3.1340	$\log\left(10^{10} A_s\right)$	3.1910 ± 0.0550	3.1811
$\mathrm{f}_{z=1}$	0.0248 ± 0.028	0.0467	$\mathrm{f}_{z=2140}$	0.00037 ± 0.00078	0.00029
$\mathrm{f}_{z=10}$	0.0018 ± 0.0289	-0.0249	$\mathrm{f}_{z=4570}$	0.0019 ± 0.0018	0.0016
$\mathrm{f}_{z=50}$	0.0034 ± 0.0201	-0.0015	$\mathrm{f}_{z=10000}$	0.0026 ± 0.0028	0.0021
$\mathrm{f}_{z=100}$	0.0133 ± 0.0160	0.0224	$\mathrm{f}_{z=21380}$	0.0030 ± 0.0045	0.0026
$\mathrm{f}_{z=500}$	0.00075 ± 0.0057	-0.0028	$\mathrm{f}_{z=45710}$	0.0010 ± 0.0084	0.0009
$-2\Delta \log \mathcal{L}$		5.68	$-2\Delta \log \mathcal{L}$		9.02

Table 4.2: Average, standard deviation, and the best fit values of the parameters from the 11 dimensional parameter estimation. Results for modified expansion histories after and before recombination are shown on the left and right sides of the table, respectively.

4.4 Summary and outlook

In this chapter, we tested a particular implementation of Everpresent Λ, Model 1, against SN Ia and CMB data. Our analysis of the supernova data indicates that a small fraction of seeds ($\sim 0.14\%$) can produce expansion histories with a similar χ^2 value to the ΛCDM model. An even smaller fraction ($\sim 0.015\%$) can do better, assuming all other parameters (α and $\Omega_m^0 h^2$) are chosen appropriately. To better understand this, we generated mock supernova datasets using Model 1 realisations. We carried out a similar analysis on these mock data sets, i.e. fitting Model 1 to the mock data. Our results indicate that almost 10% of the seeds can provide better fits to the mock data set than ΛCDM, compared to only 0.015% for the real data. This suggests that the good realisation are statistically distinguished. Our analysis also revealed that at low redshifts, the atypical good realisations of Model 1 have smaller dark energy variation and magnitude than in the case of the bad realisations.

Given that SN Ia data only probes the late-time behaviour of a cosmological model, we also considered CMB data. Neither of the Everpresent Λ models (Model 1 and Model 2) were able to improve on ΛCDM. For the CMB data, the χ^2 values for the best fit Model 1 realisations in our investigations were more than 150 greater than ΛCDM. Considering only those realisations with χ^2 below a threshold of $\chi^2 < 3000$, fits to the CMB data led to unrealistically small best-fit values for H_0, thus favouring a matter dominated universe. This indicates that Model 1 as is has difficulty accounting for the present energy density of the Universe. Tests of Model 2 yielded similar behaviour. However, as suggested in [66], we found that suppressing the dark energy around the last scattering surface leads to significantly

improved fits to the data (with χ^2 about 20.0 greater than ΛCDM), and increases H_0 (up to ~ 60). This, together with the supernova results, suggests that for a stochastic model to yield a realistic expansion history, the fluctuations must have the right properties at the relevant times.

To understand how much variation in expansion history is allowed by CMB data, we constrained the fluctuations before and after recombination separately. Our results suggest that while some variation in dark energy density is possible, for the generic kinds of variation we considered, it is significantly smaller than typical fluctuations in Everpresent Λ. Interestingly, these smaller variations were able to yield better χ^2 values than ΛCDM when fitting them to the CMB spectra and even alleviate the Hubble tension for the case of pre-recombination dark energy. This is consistent with earlier results in the literature on early dark energy models.

Therefore, incorporating some variation in the cosmological constant can be a valuable avenue for model development. While these findings are suggestive, it remains an open question whether larger fluctuations are possible if one considers a larger number of representative redshifts. Namely, it may be in principle possible to have larger fluctuations if they have the right character to lead to overall cancellations in some of the integrated effects. As mentioned in the introduction, observational data continue to hint that Λ is time dependent [227, 228].

In conclusion, the basic implementations of Everpresent Λ struggle to match both the SN Ia and CMB data. Nevertheless, promising results from specific Model 1 realisations and generic dark energy fluctuations suggest that further developments of stochastic dark energy models is warranted. There are several avenues in which this can be pursued. One approach is to modify existing models by imposing further constraints on the amplitude of Λ fluctuations. For example, a Gaussian enhancing factor near $z = 0$ could be used to impose a larger H_0. From a more fundamental perspective, Everpresent Λ would be best described by a stochastic correction to the Einstein equations. Although this problem has been approached from different perspectives [190, 194, 218], a complete and satisfactory formulation is missing. This is not merely a problem of defining a stochastic process on a fixed background, but formulating a full stochastic theory of the metric and geometry. Such a formulation would resolve the conservation issues of Λ in Everpresent Λ and restore consistency between the Friedmann equations. Moreover, spatial perturbations in dark energy−which we ignored in this chapter−and their dynamics would be consistently derived. Such a theory could also link the stochastic geodesic equations in Chapters 1 and 2 to a stochastic dark energy component. Ultimately, since Everpresent Λ relies on quantum uncertainty, its final formulation will demand a mature understanding of quantum gravity.

Epilogue

Stochasticity, spacetime uncertainty, nonlocality, and Lorentz invariance are the core physical principles underpinning the quantum gravity phenomenology explored in this thesis. Lorentz invariance imposes strong constraints on the form of effective stochastic terms. Stochasticity and nonlocality together enable Planck-scale effects to accumulate and manifest on cosmological scales. These features allowed us to test and confront our phenomenological models against a range of cosmological observations.

The above principles are deeply interconnected in causal set theory. Spacetime uncertainty is encoded kinematically through discreteness. The total number of discrete elements corresponds to spacetime 4-volume, leading naturally to nonlocality. Furthermore, there are compelling reasons, including Lorentz invariance, to believe that the embedding of a causal set into the continuum follows a Poisson process. The intrinsic randomness, together with the requirement of quasi-local laws of motion rather than global, give rise to stochastic behaviour. Thus, causal set theory naturally motivates phenomenological models such as Everpresent Λ and covariant Brownian motion. This is, however, by no means, a one-to-one correspondence. If any other theory of quantum gravity embodies spacetime uncertainty, nonlocality, and Lorentz invariance, then the stochastic phenomenology developed here could be attributed to the theory.

Regarding our understanding of the dark sector and cosmological tensions, the coming years are especially promising. Next-generation surveys such as those by the Rubin Observatory [295] and Euclid [296, 297] will deliver an unprecedented combination of sky coverage and galaxy sample size. These data will place powerful constraints on the growth of cosmic structure, the clustering properties of dark matter, and the evolution of dark energy. In particular, they are expected to play a decisive role in settling the status of the S_8 and H_0 tensions [298–301]. The final data release from DESI will be crucial in determining the fate of the present preference for dynamical dark energy, seen in the current BAO data [227, 228].

While these forthcoming observational results may well be explainable within known physics, they could also open the door to new paradigms. Quantum gravity is one of the unresolved frontiers in theoretical physics, and the possibility of a connection between its resolution and the big open problems on the nature of dark sector is a compelling case for continued exploration.

Bibliography

[1] R. M. Williams. "Discrete quantum gravity". In: *J. Phys. Conf. Ser.* 33 (2006). Ed. by Mariano Cadoni, Marco Cavaglia, and Jeanette E. Nelson, pp. 38–48. DOI: 10.1088/1742-6596/33/1/004.

[2] Kirill V Krasnov. "Counting surface states in the loop quantum gravity". In: *Phys. Rev. D* 55.CGPG-96-9-6 (1997), p. 3505. arXiv: gr-qc/9603025.

[3] Ludovico Machet and Jinzhao Wang. "On the horizon entropy of a causal set". In: *Classical and Quantum Gravity* 38.8 (2021), p. 085004. arXiv: 2012.06212.

[4] Rafael D Sorkin and Yasaman K Yazdi. "Entanglement entropy in causal set theory". In: *Classical and Quantum Gravity* 35.7 (2018), p. 074004. arXiv: 2212.13586.

[5] Renate Loll. "Quantum gravity from causal dynamical triangulations: a review". In: *Classical and Quantum Gravity* 37.1 (2019), p. 013002. arXiv: 1905.08669.

[6] S. P. Loomis and S. Carlip. "Suppression of non-manifold-like sets in the causal set path integral". In: *Class. Quant. Grav.* 35.2 (2018), p. 024002. DOI: 10.1088/1361-6382/aa980b. arXiv: 1709.00064 [gr-qc].

[7] Edmund Dable-Heath et al. "Algebraic classical and quantum field theory on causal sets". In: *Physical Review D* 101.6 (2020), p. 065013. arXiv: 1908.01973.

[8] Emma Albertini et al. "In-in correlators and scattering amplitudes on a causal set". In: *Physical Review D* 109.10 (2024), p. 106014. arXiv: 2402.08555.

[9] Andrea Addazi et al. "Quantum gravity phenomenology at the dawn of the multi-messenger era—A review". In: *Progress in Particle and Nuclear Physics* 125 (2022), p. 103948. arXiv: 2111.05659.

[10] Andrew G. Cohen, David B. Kaplan, and Ann E. Nelson. "Effective field theory, black holes, and the cosmological constant". In: *Phys. Rev. Lett.* 82 (1999), pp. 4971–4974. DOI: 10.1103/PhysRevLett.82.4971. arXiv: hep-th/9803132.

[11] Laurent Freidel et al. "The Vacuum Energy Density and Gravitational Entropy". In: (Dec. 2022). arXiv: 2212.00901 [hep-th].

[12] Joseph Bramante and Elizabeth Gould. "Material matter effects in gravitational UV/IR mixing". In: *Phys. Rev. D* 101.8 (2020), p. 084022. DOI: 10.1103/PhysRevD.101.084022. arXiv: 1910.07905 [hep-th].

[13] Dmitry Krotov and Alexander M. Polyakov. "Infrared Sensitivity of Unstable Vacua". In: *Nucl. Phys. B* 849 (2011), pp. 410–432. DOI: 10.1016/j.nuclphysb.2011.03.025. arXiv: 1012.2107 [hep-th].

[14] Niayesh Afshordi and João Magueijo. "Lower bound on the cosmological constant from the classicality of the early Universe". In: *Phys. Rev. D* 106.12 (2022), p. 123518. DOI: 10.1103/PhysRevD.106.123518. arXiv: 2209.07914 [hep-th].

[15] Alejandro Perez, Daniel Sudarsky, and James D. Bjorken. "A microscopic model for an emergent cosmological constant". In: *Int. J. Mod. Phys. D* 27.14 (2018), p. 1846002. DOI: 10.1142/S0218271818460021. arXiv: 1804.07162 [gr-qc].

[16] Denis Bernard and Andre LeClair. "Scrutinizing the Cosmological Constant Problem and a possible resolution". In: *Phys. Rev. D* 87.6 (2013), p. 063010. DOI: 10.1103/PhysRevD.87.063010. arXiv: 1211.4848 [hep-th].

[17] Bahram Mashhoon. "Nonlocal Gravity". In: *14th Brazilian School of Cosmology and Gravitation*. Jan. 2011. arXiv: 1101.3752 [gr-qc].

[18] Bahram Mashhoon. *Nonlocal Gravity*. International Series of Monographs on Physics. Oxford University Press, May 2017. ISBN: 978-0-19-880380-5.

[19] Rafael D. Sorkin. "Does locality fail at intermediate length-scales". In: (Mar. 2007), pp. 26–43. arXiv: gr-qc/0703099.

[20] S. Perlmutter et al. "Measurements of Ω and Λ from 42 high redshift supernovae". In: *Astrophys. J.* 517 (1999), pp. 565–586. DOI: 10.1086/307221. arXiv: astro-ph/9812133.

[21] Adam G. Riess et al. "Observational evidence from supernovae for an accelerating universe and a cosmological constant". In: *Astron. J.* 116 (1998), pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201.

[22] Adam G. Riess et al. "A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s^{-1} Mpc^{-1} Uncertainty from the Hubble Space Telescope and the SH0ES Team". In: *Astrophys. J. Lett.* 934.1 (2022), p. L7. DOI: 10.3847/2041-8213/ac5c5b. arXiv: 2112.04510 [astro-ph.CO].

[23] Eleonora Di Valentino et al. "In the realm of the Hubble tension—a review of solutions". In: *Classical and Quantum Gravity* 38.15 (2021), p. 153001. arXiv: 2103.01183.

[24] Ali Rida Khalife et al. "Review of Hubble tension solutions with new SH0ES and SPT-3G data". In: *Journal of Cosmology and Astroparticle Physics* 2024.04 (2024), p. 059. arXiv: 2312.09814.

[25] Alexandra Amon et al. "Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration". In: *Physical Review D* 105.2 (2022), p. 023514. arXiv: 2105.13543.

[26] Lucas Frozza Secco et al. "Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to modeling uncertainty". In: *Physical Review D* 105.2 (2022), p. 023515. arXiv: 2105.13544.

[27] Shun-Sheng Li et al. "KiDS-1000: Cosmology with improved cosmic shear measurements". In: *Astronomy & Astrophysics* 679 (2023), A133. arXiv: 2306.11124.

[28] Luis A Anchordoqui et al. "Dissecting the H0 and S8 tensions with Planck+ BAO+ supernova type Ia in multi-parameter cosmologies". In: *Journal of High Energy Astrophysics* 32 (2021), pp. 28–64. arXiv: 2107.13932.

[29] Sunny Vagnozzi. "Seven hints that early-time new physics alone is not sufficient to solve the Hubble tension". In: *Universe* 9.9 (2023), p. 393. arXiv: 2308.16628.

[30] Elcio Abdalla et al. "Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies". In: *Journal of High Energy Astrophysics* 34 (2022), pp. 49–211. arXiv: 2203.06142.

[31] Eleonora Di Valentino et al. "In the realm of the Hubble tension - a review of solutions". In: *Classical and Quantum Gravity* 38.15 (July 2021), p. 153001. DOI: 10.1088/1361-6382/ac086d. URL: https://doi.org/10.1088%2F1361-6382%2Fac086d.

[32] Adam G. Riess et al. "New Parallaxes of Galactic Cepheids from Spatially Scanning the Hubble Space Telescope: Implications for the Hubble Constant". In: *Astrophys. J.* 855.2 (2018), p. 136. DOI: 10.3847/1538-4357/aaadb7. arXiv: 1801.01120 [astro-ph.SR].

[33] Adam G. Riess et al. "Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond ΛCDM". In: *The Astrophysical Journal* 876.1 (May 2019), p. 85. DOI: 10.3847/1538-4357/ab1422. URL: https://doi.org/10.3847%2F1538-4357%2Fab1422.

[34] Adam G. Riess. "The Expansion of the Universe is Faster than Expected". In: *Nature Rev. Phys.* 2.1 (2019), pp. 10–12. DOI: 10.1038/s42254-019-0137-0. arXiv: 2001.03624 [astro-ph.CO].

[35] Kenneth C. Wong et al. "H0LiCOW – XIII. A 2.4 per cent measurement of H0 from lensed quasars: 5.3σ tension between early- and late-Universe probes". In: *Mon. Not. Roy. Astron. Soc.* 498.1 (2020), pp. 1420–1439. DOI: 10.1093/mnras/stz3094. arXiv: 1907.04869 [astro-ph.CO].

[36] Eleonora Di Valentino. "A combined analysis of the H_0 late time direct measurements and the impact on the Dark Energy sector". In: *Mon. Not. Roy. Astron. Soc.* 502.2 (2021), pp. 2065–2073. DOI: 10.1093/mnras/stab187. arXiv: 2011.00246 [astro-ph.CO].

[37] Adam G. Riess et al. "Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with ΛCDM". In: *Astrophys. J. Lett.* 908.1 (2021), p. L6. DOI: 10.3847/2041-8213/abdbaf. arXiv: 2012.08534 [astro-ph.CO].

[38] G. E. Addison et al. "Elucidating ΛCDM: Impact of Baryon Acoustic Oscillation Measurements on the Hubble Constant Discrepancy". In: *Astrophys. J.* 853.2 (2018), p. 119. DOI: 10.3847/1538-4357/aaa1ed. arXiv: 1707.06547 [astro-ph.CO].

[39] N. Aghanim et al. "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641 (2020). [Erratum: Astron.Astrophys. 652, C4 (2021)], A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[40] Mikhail M Ivanov, Marko Simonović, and Matias Zaldarriaga. "Cosmological parameters from the BOSS galaxy power spectrum". In: *Journal of Cosmology and Astroparticle Physics* 2020.05 (2020), p. 042. arXiv: 1909.05277.

[41] Mikhail M Ivanov. "Cosmological constraints from the power spectrum of eBOSS emission line galaxies". In: *Physical Review D* 104.10 (2021), p. 103514. arXiv: 2106.12580.

[42] Oliver HE Philcox and Mikhail M Ivanov. "BOSS DR12 full-shape cosmology: Λ CDM constraints from the large-scale galaxy power spectrum and bispectrum monopole". In: *Physical Review D* 105.4 (2022), p. 043517. arXiv: 2112.04515.

[43] GW Pratt et al. "The galaxy cluster mass scale and its impact on cosmological constraints from the cluster population". In: *Space Science Reviews* 215 (2019), pp. 1–82. arXiv: 1902.10837.

[44] David Benisty. "Quantifying the S8 tension with the Redshift Space Distortion data set". In: *Physics of the Dark Universe* 31 (2021), p. 100766. arXiv: 2005.03751.

[45] Rafael C Nunes and Sunny Vagnozzi. "Arbitrating the S 8 discrepancy with growth rate measurements from redshift-space distortions". In: *Monthly Notices of the Royal Astronomical Society* 505.4 (2021), pp. 5427–5437. arXiv: 2106.01208.

[46] Ted Jacobson, Stefano Liberati, and David Mattingly. "Lorentz violation at high energy: Concepts, phenomena and astrophysical constraints". In: *Annals Phys.* 321 (2006), pp. 150–196. DOI: 10.1016/j.aop.2005.06.004. arXiv: astro-ph/0505267.

[47] David Mattingly. "Modern tests of Lorentz invariance". In: *Living Rev. Rel.* 8 (2005), p. 5. DOI: 10.12942/lrr-2005-5. arXiv: gr-qc/0502097.

[48] G. Amelino-Camelia et al. "The Search for quantum gravity signals". In: *AIP Conf. Proc.* 758.1 (2005). Ed. by Alfredo Macias, Claus Lämmerzahl, and Dario Nuñez, pp. 30–80. DOI: 10.1063/1.1900507. arXiv: gr-qc/0501053.

[49] John Ellis, Nick E. Mavromatos, and Dimitri V. Nanopoulos. "Comments on Graviton Propagation in Light of GW150914". In: *Mod. Phys. Lett. A* 31.26 (2016), p. 1675001. DOI: 10.1142/S0217732316750018. arXiv: 1602.04764 [gr-qc].

[50] V. Vasileiou et al. "Constraints on Lorentz Invariance Violation from Fermi-Large Area Telescope Observations of Gamma-Ray Bursts". In: *Phys. Rev. D* 87.12 (2013), p. 122001. DOI: 10.1103/PhysRevD.87.122001. arXiv: 1305.3463 [astro-ph.HE].

[51] John Ellis et al. "Robust Constraint on Lorentz Violation Using Fermi-LAT Gamma-Ray Burst Data". In: *Phys. Rev. D* 99.8 (2019), p. 083009. DOI: 10.1103/PhysRevD.99.083009. arXiv: 1807.00189 [astro-ph.HE].

[52] Zhen Cao et al. "Stringent Tests of Lorentz Invariance Violation from LHAASO Observations of GRB 221009A". In: *Phys. Rev. Lett.* 133.7 (2024), p. 071501. DOI: 10.1103/PhysRevLett.133.071501. arXiv: 2402.06009 [astro-ph.HE].

[53] Tsvi Piran and Dmitry D. Ofengeim. "Lorentz invariance violation limits from GRB 221009A". In: *Phys. Rev. D* 109.8 (2024), p. L081501. DOI: 10.1103/PhysRevD.109.L081501. arXiv: 2308.03031 [astro-ph.HE].

[54] J. Albert et al. "Probing Quantum Gravity using Photons from a flare of the active galactic nucleus Markarian 501 Observed by the MAGIC telescope". In: *Phys. Lett. B* 668 (2008), pp. 253–257. DOI: 10.1016/j.physletb.2008.08.053. arXiv: 0708.2889 [astro-ph].

[55] A. Abramowski et al. "Search for Lorentz Invariance breaking with a likelihood fit of the PKS 2155-304 Flare Data Taken on MJD 53944". In: *Astropart. Phys.* 34 (2011), pp. 738–747. DOI: 10.1016/j.astropartphys.2011.01.007. arXiv: 1101.3650 [astro-ph.HE].

[56] Don Colladay and V. Alan Kostelecky. "Lorentz violating extension of the standard model". In: *Phys. Rev. D* 58 (1998), p. 116002. DOI: 10.1103/PhysRevD.58.116002. arXiv: hep-ph/9809521.

[57] V. Alan Kostelecky. "Gravity, Lorentz violation, and the standard model". In: *Phys. Rev. D* 69 (2004), p. 105009. DOI: 10.1103/PhysRevD.69.105009. arXiv: hep-th/0312310.

[58] Fay Dowker, Joe Henson, and Rafael D Sorkin. "Quantum gravity phenomenology, Lorentz invariance and discreteness". In: *Modern Physics Letters A* 19.24 (2004), pp. 1829–1840. arXiv: gr-qc/0311055.

[59] Rafael D. Sorkin. "Stochastic evolution on a manifold of states". In: *Annals Phys.* 168 (1986), pp. 119–147. DOI: 10.1016/0003-4916(86)90113-2.

[60] Lydia Philpott, Fay Dowker, and Rafael D Sorkin. "Energy-momentum diffusion from spacetime discreteness". In: *Physical Review D* 79.12 (2009), p. 124047. arXiv: 0810.5591.

[61] Rafael D. Sorkin. "A Modified Sum-Over-Histories for Gravity reported in the article by D. Brill and L. Smolin: "Workshop on quantum gravity and new directions"". In: *Highlights in gravitation and cosmology: Proceedings of the International Conference on Gravitation and Cosmology, Goa, India, 14–19 December 1987*. Ed. by B. R. Iyer et al. 1988, pp. 184–186.

[62] Rafael D. Sorkin. "Is the cosmological 'constant' a nonlocal quantum residue of discreteness of the causal set type?" In: *AIP Conf. Proc.* 957.1 (2007). Ed. by Arttu Rajantie et al., pp. 142–153. DOI: 10.1063/1.2823750. arXiv: 0710.1675 [gr-qc].

[63] Rafael D. Sorkin. "On the Role of Time in the Sum Over Histories Framework for Gravity". In: *Int. J. Theor. Phys.* 33 (1994), pp. 523–534. DOI: 10.1007/BF00670514.

[64] Maqbool Ahmed et al. "Everpresent Λ". In: *Physical Review D* 69.10 (May 2004). DOI: 10.1103/physrevd.69.103523. arXiv: astro-ph/0209274. URL: https://doi.org/10.1103%5C%2Fphysrevd.69.103523.

[65] Maqbool Ahmed and Rafael D. Sorkin. "Everpresent Λ II: Structural stability". In: *Physical Review D* 87.6 (Mar. 2013). DOI: 10.1103/physrevd.87.063515. URL: https://doi.org/10.1103%5C%2Fphysrevd.87.063515.

[66] Nosiphiwo Zwane, Niayesh Afshordi, and Rafael D Sorkin. "Cosmological tests of Everpresent Λ". In: *Classical and Quantum Gravity* 35.19 (Sept. 2018), p. 194002. DOI: 10.1088/1361-6382/aadc36. arXiv: 1703.06265[gr-qc]. URL: https://doi.org/10.1088%5C%2F1361-6382%5C%2Faadc36.

[67] Richard M Dudley. "Lorentz-invariant Markov processes in relativistic phase space". In: *Arkiv för Matematik* 6.3 (1966), pp. 241–268.

[68] Jacques Franchi and Yves Le Jan. "Relativistic diffusions and Schwarzschild geometry". In: *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* 60.2 (2007), pp. 187–251. arXiv: math/0410485.

[69] Nemanja Kaloper and David Mattingly. "Low energy bounds on Poincaré violation in causal set theory". In: *Physical Review D—Particles, Fields, Gravitation, and Cosmology* 74.10 (2006), p. 106001. arXiv: astro-ph/0607485.

[70] Carlo R Contaldi, Fay Dowker, and Lydia Philpott. "Polarization diffusion from spacetime uncertainty". In: *Classical and Quantum Gravity* 27.17 (July 2010), p. 172001. ISSN: 1361-6382. DOI: 10.1088/0264-9381/27/17/172001. arXiv: 1001.4545. URL: http://dx.doi.org/10.1088/0264-9381/27/17/172001.

[71] Fabrice Debbasch, Kirone Mallick, and Jean-Pierre Rivet. "Relativistic ornstein–uhlenbeck process". In: *Journal of statistical physics* 88 (1997), pp. 945–966.

[72] Jörn Dunkel and Peter Hänggi. "Theory of relativistic Brownian motion: The (1+ 3)-dimensional case". In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 72.3 (2005), p. 036106.

[73] Jose Antonio Alcantara Felix and Simone Calogero. "On a relativistic Fokker-Planck equation in kinetic theory". In: *Kin. Rel. Mod.* 4 (2011), pp. 401–426. arXiv: 1011.5429 [math-ph].

[74] Z. Haba. "Relativistic diffusion". In: *Phys. Rev. E* 79 (2009), p. 021128. DOI: 10.1103/PhysRevE.79.021128. arXiv: 0809.1340 [hep-th].

[75] Z. Haba. "Relativistic diffusive transport". In: (Nov. 2009). arXiv: 0911.3126 [hep-th].

[76] Z. Haba. "Relativistic diffusion with friction on a pseudoriemannian manifold". In: *Class. Quant. Grav.* 27 (2010), p. 095021. DOI: 10.1088/0264-9381/27/9/095021. arXiv: 0909.2880 [gr-qc].

[77] F. Debbasch and C. Chevalier. "Relativistic Stochastic Processes". In: *AIP Conf. Proc.* 913 (2007), pp. 42–48. DOI: 10.1063/1.2746722.

[78] Jörn Dunkel and Peter Hänggi. "Relativistic Brownian motion". In: *Phys. Rept.* 471 (2009), pp. 1–73. DOI: 10.1016/j.physrep.2008.12.001. arXiv: 0812.1996 [cond-mat.stat-mech].

[79] Joachim Herrmann. "Diffusion in the general theory of relativity". In: *Phys. Rev. D* 82 (2010), p. 024026. DOI: 10.1103/PhysRevD.82.024026. arXiv: 1003.3753 [gr-qc].

[80] Yifan Cai, Tao Wang, and Liu Zhao. "Relativistic stochastic mechanics I: Langevin equation from observer's perspective". In: *Journal of Statistical Physics* 190.12 (2023), p. 193. arXiv: 2306.01982.

[81] Yifan Cai, Tao Wang, and Liu Zhao. "Relativistic stochastic mechanics II: Reduced Fokker-Planck equation in curved spacetime". In: *Journal of Statistical Physics* 190.11 (2023), p. 181. arXiv: 2307.07805.

[82] Jürgen Angst and Camille Tardif. "On the Poisson boundary of the relativistic Brownian motion". In: (2020). arXiv: 1812.11250.

[83] JL Garcia-Palacios. "Introduction to the theory of stochastic processes and Brownian motion problems". In: (2007). arXiv: cond-mat/0701242.

[84] Robert Graham. "Path integral formulation of general diffusion processes". In: *Zeitschrift für Physik B Condensed Matter* 26.3 (1977), pp. 281–290.

[85] Rubén O. Acuña-Cárdenas, Carlos Gabarrete, and Olivier Sarbach. "An introduction to the relativistic kinetic theory on curved spacetimes". In: *General Relativity and Gravitation* 54.3 (Feb. 2022). ISSN: 1572-9532. DOI: 10.1007/s10714-022-02908-5. arXiv: 2106.09235. URL: http://dx.doi.org/10.1007/s10714-022-02908-5.

[86] Subrahmanyan Chandrasekhar. "Stochastic problems in physics and astronomy". In: *Rev. Mod. Phys.* 15 (1943). Ed. by K. C. Wali, pp. 1–89. DOI: 10.1103/RevModPhys.15.1.

[87] Folkert Kuipers. *Stochastic mechanics: The unification of quantum mechanics with Brownian motion.* Springer Nature, 2023. arXiv: 2301.05467.

[88] Folkert Kuipers. "Quantum mechanics from stochastic processes". In: *The European Physical Journal Plus* 138.6 (2023), pp. 1–12. arXiv: 2304.07524.

[89] Luca Bombelli et al. "Space-Time as a Causal Set". In: *Phys. Rev. Lett.* 59 (1987), pp. 521–524. DOI: 10.1103/PhysRevLett.59.521.

[90] G. 't Hooft. "Quantum Gravity: A Fundamental Problem and Some Radical Ideas". In: *Recent Developments in Gravitation: Cargèse 1978*. Ed. by Maurice Lévy and S. Deser. Boston, MA: Springer US, 1979, pp. 323–345. ISBN: 978-1-4613-2955-8. DOI: 10.1007/978-1-4613-2955-8_8. URL: https://doi.org/10.1007/978-1-4613-2955-8_8.

[91] Jan Myrheim. *Statistical geometry.* Tech. rep. Geneva: CERN, 1978. URL: https://cds.cern.ch/record/293594.

[92] Rafael D. Sorkin. "Causal sets: Discrete gravity". In: *School on Quantum Gravity*. Sept. 2003, pp. 305–327. DOI: 10.1007/0-387-24992-3_7. arXiv: gr-qc/0309009.

[93] Fay Dowker. "Causal sets and the deep structure of spacetime". In: *100 Years Of Relativity: space-time structure: Einstein and beyond*. Ed. by Abhay Ashtekar. 2005, pp. 445–464. DOI: 10.1142/9789812700988_0016. arXiv: gr-qc/0508109.

[94] Joe Henson. "The Causal set approach to quantum gravity". In: (Jan. 2006), pp. 393–413. arXiv: gr-qc/0601121.

[95] Sumati Surya. "The causal set approach to quantum gravity". In: *Living Rev. Rel.* 22.1 (2019), p. 5. DOI: 10.1007/s41114-019-0023-1. arXiv: 1903.11544 [gr-qc].

[96] Luca Bombelli, Joe Henson, and Rafael D. Sorkin. "Discreteness without symmetry breaking: A Theorem". In: *Mod. Phys. Lett. A* 24 (2009), pp. 2579–2587. DOI: 10.1142/S0217732309031958. arXiv: gr-qc/0605006.

[97] T. Regge. "General Relativity Without Coordinates". In: *Nuovo Cim.* 19 (1961), pp. 558–571. DOI: 10.1007/BF02733251.

[98] Charles W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. San Francisco: W. H. Freeman, 1973. ISBN: 978-0-7167-0344-0, 978-0-691-17779-3.

[99] Jan Ambjorn and R. Loll. "Nonperturbative Lorentzian quantum gravity, causality and topology change". In: *Nucl. Phys. B* 536 (1998), pp. 407–434. DOI: 10.1016/S0550-3213(98)00692-0. arXiv: hep-th/9805108.

[100] R. Loll. "Quantum Gravity from Causal Dynamical Triangulations: A Review". In: *Class. Quant. Grav.* 37.1 (2020), p. 013002. DOI: 10.1088/1361-6382/ab57c7. arXiv: 1905.08669 [hep-th].

[101] Erik Christopher Zeeman. "Causality implies the Lorentz group". In: *Journal of Mathematical Physics* 5.4 (1964), p. 490.

[102] E. H. Kronheimer and Roger Penrose. "On the structure of causal spaces". In: *Proc. Cambridge Phil. Soc.* 63 (1967), pp. 481–501. DOI: 10.1017/S030500410004144X.

[103] S. W. Hawking, A. R. King, and P. J. Mccarthy. "A New Topology for Curved Space-Time Which Incorporates the Causal, Differential, and Conformal Structures". In: *J. Math. Phys.* 17 (1976), pp. 174–181. DOI: 10.1063/1.522874.

[104] David B Malament. "Class of continuous timelike curves determines the topology of spacetime.[Causal, topological, differential, and conformal structure]". In: *J. Math. Phys.(NY);(United States)* 18.7 (1977).

[105] Alessio Belenchia, Dionigi M. T. Benincasa, and Fay Dowker. "The continuum limit of a 4-dimensional causal set scalar d'Alembertian". In: *Class. Quant. Grav.* 33.24 (2016), p. 245018. DOI: 10.1088/0264-9381/33/24/245018. arXiv: 1510.04656 [gr-qc].

[106] Alessio Belenchia, Dionigi M. T. Benincasa, and Stefano Liberati. "Nonlocal Scalar Quantum Field Theory from Causal Sets". In: *JHEP* 03 (2015), p. 036. DOI: 10.1007/JHEP03(2015)036. arXiv: 1411.6513 [gr-qc].

[107] Steven Paul Johnston. "Quantum Fields on Causal Sets". Other thesis. Oct. 2010. arXiv: 1010.5514 [hep-th].

[108] Niayesh Afshordi et al. "A Ground State for the Causal Diamond in 2 Dimensions". In: *JHEP* 10 (2012), p. 088. DOI: 10.1007/JHEP10(2012)088. arXiv: 1207.7101 [hep-th].

[109] Rafael D. Sorkin. "Scalar Field Theory on a Causal Set in Histories Form". In: *J. Phys. Conf. Ser.* 306 (2011). Ed. by Lajos Diosi et al., p. 012017. DOI: 10.1088/1742-6596/306/1/012017. arXiv: 1107.0698 [gr-qc].

[110] Ian Jubb. "Interacting Quantum Scalar Field Theory on a Causal Set". In: 2024. DOI: 10.1007/978-981-19-3079-9_76-1. arXiv: 2306.12484 [hep-th].

[111] J. Myrheim. "Statistical Geometry". In: (Aug. 1978).

[112] Graham Brightwell and Ruth Gregory. "The Structure of random discrete space-time". In: *Phys. Rev. Lett.* 66 (1991), pp. 260–263. DOI: 10.1103/PhysRevLett.66.260.

[113] Raluca Ilie, Gregory B. Thompson, and David D. Reid. "A Numerical study of the correspondence between paths in a causal set and geodesics in the continuum". In: *Class. Quant. Grav.* 23 (2006), pp. 3275–3286. DOI: 10.1088/0264-9381/23/10/002. arXiv: gr-qc/0512073.

[114] Lydia Philpott. "Causal Set Phenomenology". In: *PhD Thesis* (2010). arXiv: 1009.1593.

[115] Jorn Dunkel, Peter Talkner, and Peter Hanggi. "Relativistic diffusion processes and random walk models". In: *Phys. Rev. D* 75 (2007), p. 043001. DOI: 10.1103/PhysRevD.75.043001. arXiv: cond-mat/0608023.

[116] Gökçe Başar et al. "Stochastic relativistic advection diffusion equation from the Metropolis algorithm". In: *Phys. Rev. C* 110.4 (2024), p. 044903. DOI: 10.1103/PhysRevC.110.044903. arXiv: 2403.04185 [nucl-th].

[117] Jay Bhambure, Rajeev Singh, and Derek Teaney. "Stochastic relativistic viscous hydrodynamics from the Metropolis algorithm". In: *Phys. Rev. C* 111.6 (2025), p. 064909. DOI: 10.1103/15y6-6gmz. arXiv: 2412.10306 [nucl-th].

[118] Chung-Pei Ma and Edmund Bertschinger. "Cosmological Perturbation Theory in the Synchronous and Conformal Newtonian Gauges". In: *The Astrophysical Journal* 455 (Dec. 1995), p. 7. ISSN: 1538-4357. DOI: 10.1086/176550. arXiv: astro-ph/9506072v1. URL: http://dx.doi.org/10.1086/176550.

[119] Diego Blas, Julien Lesgourgues, and Thomas Tram. "The cosmic linear anisotropy solving system (CLASS). Part II: Approximation schemes". In: *Journal of Cosmology and Astroparticle Physics* 2011.07 (2011), p. 034. arXiv: 1104.2933.

[120] Julien Lesgourgues and Thomas Tram. "The Cosmic Linear Anisotropy Solving System (CLASS) IV: efficient implementation of non-cold relics". In: *Journal of Cosmology and Astroparticle Physics* 2011.09 (2011), p. 032. arXiv: 1104.2935.

[121] Caio Bastos de Senna Nascimento. "Generalized Boltzmann hierarchy for massive neutrinos in cosmology". In: *Physical Review D* 104.8 (2021), p. 083535. arXiv: 2104.00703v2.

[122] Alexander Grigor'yan and Masakazu Noguchi. "The heat kernel on hyperbolic space". In: *Bulletin of the London Mathematical Society* 30.6 (1998), pp. 643–650.

[123] Scott Dodelson and Lawrence M Widrow. "Sterile neutrinos as dark matter". In: *Physical Review Letters* 72.1 (1994), p. 17. arXiv: 1807.07938.

[124] D Boyanovsky. "Clustering properties of a sterile neutrino dark matter candidate". In: *Physical Review D* 78.10 (2008), p. 103505. arXiv: 0807.0646.

[125] Daniel Boyanovsky and Jun Wu. "Small scale aspects of warm dark matter: power spectra and acoustic oscillations". In: *Physical Review D* 83.4 (2011), p. 043524. arXiv: 1008.0992.

[126] Christian M Mueller. "Cosmological bounds on the equation of state of dark matter". In: *Physical Review D—Particles, Fields, Gravitation, and Cosmology* 71.4 (2005), p. 047302. arXiv: astro-ph/0410621.

[127] Suresh Kumar and Lixin Xu. "Observational constraints on variable equation of state parameters of dark matter and dark energy after Planck". In: *Physics Letters B* 737 (2014), pp. 244–247. arXiv: 1207.5582.

[128] Mang Feng. "Complete solution of the Schrödinger equation for the time-dependent linear potential". In: *Physical Review A* 64.3 (2001), p. 034101.

[129] Amir Kadivar. *Techniques for Inhomogeneous Linear PDEs*. https://amirkdv.ca/files/inhom-pde.pdf. 2018.

[130] Hermano Velten and Thiago RP Caramês. "To conserve, or not to conserve: A review of nonconservative theories of gravity". In: *Universe* 7.2 (2021), p. 38. arXiv: 2102.03457.

[131] Thibaut Josset, Alejandro Perez, and Daniel Sudarsky. "Dark Energy from Violation of Energy Conservation". In: *Physical Review Letters* 118.2 (Jan. 2017). ISSN: 1079-7114. DOI: 10.1103/physrevlett.118.021102. arXiv: 1604.04183. URL: http://dx.doi.org/10.1103/PhysRevLett.118.021102.

[132] Alejandro Perez and Daniel Sudarsky. "Dark energy from quantum gravity discreteness". In: *Physical review letters* 122.22 (2019), p. 221302. arXiv: 1711.05183.

[133] Cristóbal Corral, Norman Cruz, and Esteban González. "Diffusion in unimodular gravity: Analytical solutions, late-time acceleration, and cosmological constraints". In: *Physical Review D* 102.2 (2020), p. 023508. arXiv: 2005.06052.

[134] Alejandro Perez, Daniel Sudarsky, and Edward Wilson-Ewing. "Resolving the H 0 tension with diffusion". In: *General Relativity and Gravitation* 53.1 (2021), p. 7. arXiv: 2001.07536.

[135] Francisco X Linares Cedeno and Ulises Nucamendi. "Revisiting cosmological diffusion models in Unimodular Gravity and the H0 tension". In: *Physics of the Dark Universe* 32 (2021), p. 100807. arXiv: 2009.10268.

[136] Júlio C Fabris et al. "Nonconservative unimodular gravity: a viable cosmological scenario?" In: *The European Physical Journal C* 82.6 (2022), p. 522. arXiv: 2112.06644.

[137] Marco de Cesare and Edward Wilson-Ewing. "Interacting dark sector from the trace-free Einstein equations: Cosmological perturbations with no instability". In: *Physical Review D* 106.2 (July 2022). ISSN: 2470-0029. DOI: 10.1103/physrevd.106.023527. arXiv: 2112.12701. URL: http://dx.doi.org/10.1103/PhysRevD.106.023527.

[138] Lautaro Amadei and Alejandro Perez. "Planckian discreteness as seeds for cosmic structure". In: *Physical Review D* 106.6 (2022), p. 063528. arXiv: 2104.08881.

[139] Gabriel Leon. "Inflation and the cosmological (not-so) constant in unimodular gravity". In: *Classical and Quantum Gravity* 39.7 (2022), p. 075008. arXiv: 2202.04029.

[140] Jussi Väliviita, Elisabetta Majerotto, and Roy Maartens. "Large-scale instability in interacting dark energy and dark matter fluids". In: *Journal of Cosmology and Astroparticle Physics* 2008.07 (2008), p. 020. arXiv: 10804.0232.

[141] M. B. Gavela et al. "Dark coupling". In: *JCAP* 07 (2009). [Erratum: JCAP 05, E01 (2010)], p. 034. DOI: 10.1088/1475-7516/2009/07/034. arXiv: 0901.1611 [astro-ph.CO].

[142] Kevin J Ludwick. "The viability of phantom dark energy: A review". In: *Modern Physics Letters A* 32.28 (2017), p. 1730025. arXiv: 1708.06981.

[143] Mark S Madsen. "Scalar fields in curved spacetimes". In: *Classical and quantum gravity* 5.4 (1988), p. 627. arXiv: 1401.2026.

[144] L. O. Pimentel. "Energy Momentum Tensor in the General Scalar - Tensor Theory". In: *Class. Quant. Grav.* 6 (1989), pp. L263–L265. DOI: 10.1088/0264-9381/6/12/005.

[145] Valerio Faraoni and Jeremy Côté. "Imperfect fluid description of modified gravities". In: *Physical Review D* 98.8 (2018), p. 084019. arXiv: 1808.02427.

[146] Ignacy Sawicki et al. "Consistent perturbations in an imperfect fluid". In: *Journal of Cosmology and Astroparticle Physics* 2013.01 (2013), p. 004. arXiv: 1208.4855.

[147] Winfried Zimdahl, Hermano ES Velten, and William C Algoner. "Matter growth in imperfect fluid cosmology". In: *Universe* 5.3 (2019), p. 68. arXiv: 1903.03383.

[148] Vivian Poulin, Pasquale D Serpico, and Julien Lesgourgues. "A fresh look at linear cosmological constraints on a decaying dark matter component". In: *Journal of Cosmology and Astroparticle Physics* 2016.08 (2016), p. 036. arXiv: 1606.02073.

[149] Lucca Matteo. "Dark energy–dark matter interactions as a solution to the S8 tension Phys". In: *Dark Univ* 34 (2021), p. 100899. arXiv: 2105.09249.

[150] Shankar Agarwal and Hume A Feldman. "The effect of massive neutrinos on the matter power spectrum". In: *Monthly Notices of the Royal Astronomical Society* 410.3 (2011), pp. 1647–1654. arXiv: 1006.0689.

[151] Jose Beltrán Jiménez et al. "Probing elastic interactions in the dark sector and the role of S 8". In: *Physical Review D* 104.10 (2021), p. 103503. arXiv: 2106.11222.

[152] Adam He et al. "S8 tension in the context of dark matter–baryon scattering". In: *The Astrophysical Journal Letters* 954.1 (2023), p. L8. arXiv: 2301.08260.

[153] Guillermo F Abellán et al. "Implications of the S 8 tension for decaying dark matter with warm decay products". In: *Physical Review D* 105.6 (2022), p. 063525. arXiv: 2008.09615.

[154] Mei-Yu Wang and Andrew R Zentner. "Effects of unstable dark matter on large-scale structure and constraints from future surveys". In: *Physical Review D—Particles, Fields, Gravitation, and Cosmology* 85.4 (2012), p. 043514. arXiv: 1201.2426.

[155] Shohei Aoyama et al. "Evolution of perturbations and cosmological constraints in decaying dark matter models with arbitrary decay mass products". In: *Journal of Cosmology and Astroparticle Physics* 2014.07 (2014), p. 021. arXiv: 1402.2972.

[156] Thejs Brinckmann and Julien Lesgourgues. "MontePython 3: boosted MCMC sampler and other features". In: (2018). arXiv: 1804.07261.

[157] Benjamin Audren et al. "Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code". In: *JCAP* 1302 (2013), p. 001. DOI: 10.1088/1475-7516/2013/02/001. arXiv: 1210.7183.

[158] Nabila Aghanim et al. "Planck 2018 results-V. CMB power spectra and likelihoods". In: *Astronomy & Astrophysics* 641 (2020), A5. arXiv: 1907.12875.

[159] Cullan Howlett et al. "The clustering of the SDSS main galaxy sample–II. Mock galaxy catalogues and a measurement of the growth of structure from redshift space distortions at z= 0.15". In: *Monthly Notices of the Royal Astronomical Society* 449.1 (2015), pp. 848–866. arXiv: 1409.3238.

[160] Shadab Alam et al. "The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample". In: *Monthly Notices of the Royal Astronomical Society* 470.3 (2017), pp. 2617–2652. arXiv: 1607.03155.

[161] Julian E Bautista et al. "The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red galaxy sample from the anisotropic correlation function between redshifts 0.6 and 1". In: *Monthly Notices of the Royal Astronomical Society* 500.1 (2021), pp. 736–762. arXiv: 2007.08993.

[162] Héctor Gil-Marín et al. "The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red galaxy sample from the anisotropic power spectrum between redshifts 0.6 and 1.0". In: *Monthly Notices of the Royal Astronomical Society* 498.2 (2020), pp. 2492–2531. arXiv: 2007.08994.

[163] Jiamin Hou et al. "The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: BAO and RSD measurements from anisotropic clustering analysis of the quasar sample in configuration space between redshift 0.8 and 2.2". In: *Monthly Notices of the Royal Astronomical Society* 500.1 (2021), pp. 1201–1221. arXiv: 2007.08998.

[164] Richard Neveux et al. "The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: BAO and RSD measurements from the anisotropic power spectrum of the quasar sample between redshift 0.8 and 2.2". In: *Monthly Notices of the Royal Astronomical Society* 499.1 (2020), pp. 210–229. arXiv: 2007.08999.

[165] Shadab Alam et al. "Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory". In: *Physical Review D* 103.8 (2021), p. 083533. arXiv: 2007.08991.

[166] Helion Du Mas Des Bourboux et al. "The completed SDSS-IV extended baryon oscillation spectroscopic survey: baryon acoustic oscillations with Lyα forests". In: *The Astrophysical Journal* 901.2 (2020), p. 153. arXiv: 2007.08995.

[167] Dillon Brout et al. "The Pantheon+ analysis: cosmological constraints". In: *The Astrophysical Journal* 938.2 (2022), p. 110. arXiv: 2202.04077.

[168] Tilman Tröster et al. "KiDS-1000 Cosmology: constraints beyond flat ΛCDM". In: *Astronomy & Astrophysics* 649 (2021), A88. arXiv: 2010.16416.

[169] TMC Abbott et al. "Dark Energy Survey Year 3 results: Constraints on extensions to Λ CDM with weak lensing and galaxy clustering". In: *Physical Review D* 107.8 (2023), p. 083504. arXiv: 2207.05766.

[170] Benjamin Stölzner et al. "KiDS-Legacy: Consistency of cosmic shear measurements and joint cosmological constraints with external probes". In: (Mar. 2025). arXiv: 2503.19442 [astro-ph.CO].

[171] Fay Dowker. "Introduction to causal sets and their phenomenology". In: *Gen. Rel. Grav.* 45.9 (2013), pp. 1651–1667. DOI: 10.1007/s10714-013-1569-y.

[172] Stacy Y. Kim, Annika H. G. Peter, and Jonathan R. Hargis. "Missing Satellites Problem: Completeness Corrections to the Number of Satellite Galaxies in the Milky Way are Consistent with Cold Dark Matter Predictions". In: *Phys. Rev. Lett.* 121.21 (2018), p. 211302. DOI: 10.1103/PhysRevLett.121.211302. arXiv: 1711.06267 [astro-ph.CO].

[173] EO Nadler et al. "Constraints on dark matter properties from observations of Milky Way satellite galaxies". In: *Physical review letters* 126.9 (2021), p. 091101. arXiv: 2008.00022.

[174] Bruno Villasenor et al. "New constraints on warm dark matter from the Lyman-α forest power spectrum". In: *Physical Review D* 108.2 (2023), p. 023502. arXiv: 2209.14220.

[175] Deanna C Hooper et al. "One likelihood to bind them all: Lyman-α constraints on non-standard dark matter". In: *Journal of Cosmology and Astroparticle Physics* 2022.10 (2022), p. 032. arXiv: 2206.08188.

[176] WJG De Blok. "The Core-Cusp Problem". In: *Advances in Astronomy* 2010.1 (2010), p. 789293. arXiv: 0910.3538.

[177] James S Bullock and Michael Boylan-Kolchin. "Small-scale challenges to the Λ CDM paradigm". In: *Annual Review of Astronomy and Astrophysics* 55.1 (2017), pp. 343–387. arXiv: 1707.04256.

[178] Andrew Pontzen and Fabio Governato. "Cold dark matter heats up". In: *Nature* 506.7487 (2014), pp. 171–178. arXiv: 1402.1764.

[179] Paolo Salucci. "The distribution of dark matter in galaxies". In: *The Astronomy and Astrophysics Review* 27.1 (2019), p. 2. arXiv: 1811.08843.

[180] Jesús Zavala and Carlos S Frenk. "Dark matter haloes and subhaloes". In: *Galaxies* 7.4 (2019), p. 81. arXiv: 1907.11775.

[181] Hongwan Liu et al. "Lyman-α constraints on cosmic heating from dark matter annihilation and decay". In: *Physical Review D* 104.4 (2021), p. 043514. arXiv: 2008.01084.

[182] Digvijay Wadekar and Zihui Wang. "Strong constraints on decay and annihilation of dark matter from heating of gas-rich dwarf galaxies". In: *Physical Review D* 106.7 (2022), p. 075007. arXiv: 2111.08025.

[183] Gilad Gour and L Sriramkumar. "Will small particles exhibit Brownian motion in the quantum vacuum?" In: *Foundations of Physics* 29.12 (1999), pp. 1917–1949. arXiv: quant-ph/9808032.

[184] Philip R Johnson and Bei-Lok Hu. "Stochastic theory of relativistic particles moving in a quantum field: Scalar Abraham-Lorentz-Dirac-Langevin equation, radiation reaction, and vacuum fluctuations". In: *Physical Review D* 65.6 (2002), p. 065015. arXiv: quant-ph/0101001.

[185] Carlos HG Béssa, Valdir B Bezerra, and LH Ford. "Brownian motion in Robertson–Walker spacetimes from electromagnetic vacuum fluctuations". In: *Journal of mathematical physics* 50.6 (2009). arXiv: 0804.1360.

[186] DJ Bedingham. "Single particle energy diffusion from relativistic spontaneous localization". In: *Physical Review D—Particles, Fields, Gravitation, and Cosmology* 88.4 (2013), p. 045032. arXiv: 1306.5625.

[187] Roderich Tumulka. "A relativistic version of the Ghirardi–Rimini–Weber model". In: *Journal of Statistical Physics* 125 (2006), pp. 821–840. arXiv: quant-ph/0406094.

[188] Daniel J Bedingham. "Relativistic state reduction dynamics". In: *Foundations of Physics* 41 (2011), pp. 686–704. arXiv: 1003.2774.

[189] Jonathan Oppenheim. "A Postquantum Theory of Classical Gravity?" In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041040. DOI: 10.1103/PhysRevX.13.041040. arXiv: 1811.03116. URL: https://link.aps.org/doi/10.1103/PhysRevX.13.041040.

[190] Jonathan Oppenheim, Emanuele Panella, and Andrew Pontzen. "Emergence of phantom cold dark matter from spacetime diffusion". In: (2024). arXiv: 2407.13820 [gr-qc]. URL: https://arxiv.org/abs/2407.13820.

[191] Cedomir Crnkovic and Edward Witten. "Covariant description of canonical formalism in geometrical theories." In: *Three hundred years of gravitation* (1987), pp. 676–684.

[192] Alessio Belenchia, Dionigi MT Benincasa, and Stefano Liberati. "Nonlocal scalar quantum field theory from causal sets". In: *Journal of High Energy Physics* 2015.3 (2015), pp. 1–26. arXiv: 1411.6513.

[193] Arad Nasiri. "Synge's world function applied to causal diamonds and causal sets". In: (2023). arXiv: 2304.00088.

[194] Bei Lok Hu and Enric Verdaguer. "Stochastic gravity: Theory and applications". In: *Living Reviews in Relativity* 11 (2008), pp. 1–112. arXiv: 0802.0658.

[195] Archil Kobakhidze, Cyril Lagger, and Adrian Manning. "Constraining noncommutative spacetime from GW150914". In: *Phys. Rev. D* 94.6 (2016), p. 064033. DOI: 10.1103/PhysRevD.94.064033. arXiv: 1607.03776 [gr-qc].

[196] Andrea Maselli et al. "From micro to macro and back: probing near-horizon quantum structures with gravitational waves". In: *Class. Quant. Grav.* 36.16 (2019), p. 167001. DOI: 10.1088/1361-6382/ab30ff. arXiv: 1811.03689 [gr-qc].

[197] Steven B. Giddings, Seth Koren, and Gabriel Treviño. "Exploring strong-field deviations from general relativity via gravitational waves". In: *Phys. Rev. D* 100.4 (2019), p. 044005. DOI: 10.1103/PhysRevD.100.044005. arXiv: 1904.04258 [gr-qc].

[198] Ivan Agullo et al. "Potential Gravitational Wave Signatures of Quantum Gravity". In: *Phys. Rev. Lett.* 126.4 (2021), p. 041302. DOI: 10.1103/PhysRevLett.126.041302. arXiv: 2007.13761 [gr-qc].

[199] Angel Garcia-Chung et al. "Propagation of quantum gravity-modified gravitational waves on a classical FLRW spacetime". In: *Phys. Rev. D* 103.8 (2021), p. 084053. DOI: 10.1103/PhysRevD.103.084053. arXiv: 2012.09366 [gr-qc].

[200] Gianluca Calcagni. "Lorentz violations in multifractal spacetimes". In: *Eur. Phys. J. C* 77.5 (2017), p. 291. DOI: 10.1140/epjc/s10052-017-4841-6. arXiv: 1603.03046 [gr-qc].

[201] Sai Wang. "Exploring the CPT violation and birefringence of gravitational waves with ground- and space-based gravitational-wave interferometers". In: *Eur. Phys. J. C* 80.4 (2020), p. 342. DOI: 10.1140/epjc/s10052-020-7812-2. arXiv: 1712.06072 [gr-qc].

[202] Nicolas Yunes, Kent Yagi, and Frans Pretorius. "Theoretical Physics Implications of the Binary Black-Hole Mergers GW150914 and GW151226". In: *Phys. Rev. D* 94.8 (2016), p. 084002. DOI: 10.1103/PhysRevD.94.084002. arXiv: 1603.08955 [gr-qc].

[203] Gianluca Calcagni et al. "Gravitational-wave luminosity distance in quantum gravity". In: *Phys. Lett. B* 798 (2019), p. 135000. DOI: 10.1016/j.physletb.2019.135000. arXiv: 1904.00384 [gr-qc].

[204] Gianluca Calcagni et al. "Quantum gravity and gravitational-wave astronomy". In: *JCAP* 10 (2019), p. 012. DOI: 10.1088/1475-7516/2019/10/012. arXiv: 1907.02489 [gr-qc].

[205] Enis Belgacem et al. "Testing modified gravity at cosmological distances with LISA standard sirens". In: *JCAP* 07 (2019), p. 024. DOI: 10.1088/1475-7516/2019/07/024. arXiv: 1906.01593 [astro-ph.CO].

[206] Chiara Caprini et al. "Science with the space-based interferometer eLISA. II: Gravitational waves from cosmological phase transitions". In: *JCAP* 04 (2016), p. 001. DOI: 10.1088/1475-7516/2016/04/001. eprint: 1512.06239.

[207] Nicola Bartolo et al. "Science with the space-based interferometer LISA. IV: Probing inflation with gravitational waves". In: *JCAP* 12 (2016), p. 026. DOI: 10.1088/1475-7516/2016/12/026. eprint: 1609.04791.

[208] Nicola Bartolo et al. "Science with the space-based interferometer LISA. V: Probing dark matter and primordial black holes". In: *JCAP* 11 (2018), p. 055. DOI: 10.1088/1475-7516/2018/11/055. eprint: 1712.01219.

[209] Oleg A. Fonarev. "Wigner function and quantum kinetic theory in curved space-time and external fields". In: *J. Math. Phys.* 35 (1994), pp. 2105–2129. DOI: 10.1063/1.530542. arXiv: gr-qc/9309005.

[210] Hugo Garcia-Compean and Francisco J. Turrubiates. "Ground-state Wigner functional of linearized gravitational field". In: *Int. J. Mod. Phys. A* 26 (2011), pp. 5241–5259. DOI: 10.1142/S0217751X11054899. arXiv: 1109.1036 [hep-th].

[211] Sybren Ruurds Groot et al. *Relativistic kinetic theory: principles and applications*. North Holland, 1980.

[212] Nicola Bartolo et al. "Probing anisotropies of the Stochastic Gravitational Wave Background with LISA". In: *JCAP* 11 (2022), p. 009. DOI: 10.1088/1475-7516/2022/11/009. arXiv: 2201.08782 [astro-ph.CO].

[213] Chiara Caprini et al. "Gravitational waves from first-order phase transitions in LISA: reconstruction pipeline and physics interpretation". In: *JCAP* 10 (2024), p. 020. DOI: 10.1088/1475-7516/2024/10/020. arXiv: 2403.03723 [astro-ph.CO].

[214] Marek Lewicki and Ville Vaskonen. "Gravitational waves from bubble collisions and fluid motion in strongly supercooled phase transitions". In: *Eur. Phys. J. C* 83.2 (2023), p. 109. DOI: 10.1140/epjc/s10052-023-11241-3. arXiv: 2208.11697 [astro-ph.CO].

[215] Anson Hook, Gustavo Marques-Tavares, and Davide Racco. "Causal gravitational waves as a probe of free streaming particles and the expansion of the Universe". In: *JHEP* 02 (2021), p. 117. DOI: 10.1007/JHEP02(2021)117. arXiv: 2010.03568 [hep-ph].

[216] N. Bartolo et al. "Testing primordial black holes as dark matter with LISA". In: *Phys. Rev. D* 99.10 (2019), p. 103521. DOI: 10.1103/PhysRevD.99.103521. arXiv: 1810.12224 [astro-ph.CO].

[217] Fay Dowker. "Introduction to causal sets and their phenomenology". In: *Gen. Rel. Grav.* 45.9 (2013), pp. 1651–1667. DOI: 10.1007/s10714-013-1569-y.

[218] Santiago Agui Salcedo, Thomas Colas, and Enrico Pajer. "An Open Effective Field Theory for light in a medium". In: *JHEP* 03 (2025), p. 138. DOI: 10.1007/JHEP03(2025)138. arXiv: 2412.12299 [hep-th].

[219] Z. Haba. "Relativistic diffusion of elementary particles with spin". In: *J. Phys. A* 42 (2009), p. 445401. DOI: 10.1088/1751-8113/42/44/445401. arXiv: 0907.0364 [hep-th].

[220] Andreu Font-Ribera et al. "Quasar-Lyman α Forest Cross-Correlation from BOSS DR11 : Baryon Acoustic Oscillations". In: *JCAP* 05 (2014), p. 027. DOI: 10.1088/1475-7516/2014/05/027. arXiv: 1311.1767 [astro-ph.CO].

[221] Timothée Delubac et al. "Baryon acoustic oscillations in the Lyα forest of BOSS DR11 quasars". In: *Astron. Astrophys.* 574 (2015), A59. DOI: 10.1051/0004-6361/201423969. arXiv: 1404.1801 [astro-ph.CO].

[222] Yuting Wang et al. "Evolution of dark energy reconstructed from the latest observations". In: *Astrophys. J. Lett.* 869 (2018), p. L8. DOI: 10.3847/2041-8213/aaf238. arXiv: 1807.03772 [astro-ph.CO].

[223] Éric Aubourg et al. "Cosmological implications of baryon acoustic oscillation measurements". In: *Phys. Rev. D* 92.12 (2015), p. 123516. DOI: 10.1103/PhysRevD.92.123516. arXiv: 1411.1074 [astro-ph.CO].

[224] Eoin O Colgáin et al. *Putting Flat ΛCDM In The (Redshift) Bin*. 2023. arXiv: 2206.11447 [astro-ph.CO].

[225] Mohammad Malekjani et al. *Negative Dark Energy Density from High Redshift Pantheon+ Supernovae*. 2023. arXiv: 2301.12725 [astro-ph.CO].

[226] Eoin O Colgáin et al. *MCMC Marginalisation Bias and ΛCDM tensions*. 2023. arXiv: 2307.16349 [astro-ph.CO].

[227] A. G. Adame et al. "DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations". In: *JCAP* 02 (2025), p. 021. DOI: 10.1088/1475-7516/2025/02/021. arXiv: 2404.03002 [astro-ph.CO].

[228] M. Abdul Karim et al. "DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints". In: (Mar. 2025). arXiv: 2503.14738 [astro-ph.CO].

[229] Marco de Cesare, Fedele Lizzi, and Mairi Sakellariadou. "Effective cosmological constant induced by stochastic fluctuations of Newton's constant". In: *Phys. Lett. B* 760 (2016), pp. 498–501. DOI: 10.1016/j.physletb.2016.07.015. arXiv: 1603.04170 [gr-qc].

[230] Samuel S. Cree et al. "Can the fluctuations of the quantum vacuum solve the cosmological constant problem?" In: *Phys. Rev. D* 98.6 (2018), p. 063506. DOI: 10.1103/PhysRevD.98.063506. arXiv: 1805.12293 [gr-qc].

[231] Rafael D Sorkin. "Spacetime and causal sets." In: *Relativity and Gravitation* (1991), p. 150.

[232] John D. Barrow. "A Strong Constraint on Ever-Present Lambda". In: *Phys. Rev. D* 75 (2007), p. 067301. DOI: 10.1103/PhysRevD.75.067301. arXiv: gr-qc/0612128.

[233] Joseph A. Zuntz. "The cosmic microwave background in a causal set universe". In: *Phys. Rev. D* 77 (2008), p. 043002. DOI: 10.1103/PhysRevD.77.043002. arXiv: 0711.2904 [astro-ph].

[234] Rafael Sorkin. *Supplementary considerations on Everpresent Lambda*. en. PIRSA:19110069 see, https://pirsa.org/19110069. Nov. 2019. DOI: 10.48660/19110069. URL: https://pirsa.org/19110069.

[235] Joe Henson et al. "Onset of the Asymptotic Regime for (Uniformly Random) Finite Orders". In: *Exper. Math.* 26.3 (2016), pp. 253–266. DOI: 10.1080/10586458.2016.1158134.

[236] Luca Bombelli, Joe Henson, and Rafael D. Sorkin. "Discreteness without symmetry breaking: A Theorem". In: *Mod. Phys. Lett. A* 24 (2009), pp. 2579–2587. DOI: 10.1142/S0217732309031958. arXiv: gr-qc/0605006.

[237] Fay Dowker and Rafael D. Sorkin. "Symmetry-breaking and zero-one laws". In: *Class. Quant. Grav.* 37.15 (2020), p. 155007. DOI: 10.1088/1361-6382/ab81cd. arXiv: 1909.06070 [gr-qc].

[238] Mehdi Saravani and Siavash Aslanbeigi. "On the Causal Set-Continuum Correspondence". In: *Class. Quant. Grav.* 31.20 (2014), p. 205013. DOI: 10.1088/0264-9381/31/20/205013. arXiv: 1403.6429 [hep-th].

[239] D. P. Rideout and R. D. Sorkin. "A Classical sequential growth dynamics for causal sets". In: *Phys. Rev. D* 61 (2000), p. 024002. DOI: 10.1103/PhysRevD.61.024002. arXiv: gr-qc/9904062.

[240] Stav Zalel. "Discrete Random Spacetimes: Covariance and Quantization in Growth Dynamics for Causal Sets". PhD thesis. Imperial Coll., London, 2021. DOI: 10.25560/93625.

[241] Dionigi M. T. Benincasa and Fay Dowker. "The Scalar Curvature of a Causal Set". In: *Phys. Rev. Lett.* 104 (2010), p. 181301. DOI: 10.1103/PhysRevLett.104.181301. arXiv: 1001.2725 [gr-qc].

[242] Fay Dowker and Lisa Glaser. "Causal set d'Alembertians for various dimensions". In: *Class. Quant. Grav.* 30 (2013), p. 195016. DOI: 10.1088/0264-9381/30/19/195016. arXiv: 1305.2588 [gr-qc].

[243] M. Henneaux and C. Teitelboim. "The Cosmological Constant and General Covariance". In: *Phys. Lett. B* 222 (1989), pp. 195–199. DOI: 10.1016/0370-2693(89)91251-3.

[244] W. Buchmuller and N. Dragon. "Gauge Fixing and the Cosmological Constant". In: *Phys. Lett. B* 223 (1989), pp. 313–317. DOI: 10.1016/0370-2693(89)91608-0.

[245] Steven Weinberg. "The Cosmological Constant Problem". In: *Rev. Mod. Phys.* 61 (1989). Ed. by Jong-Ping Hsu and D. Fine, pp. 1–23. DOI: 10.1103/RevModPhys.61.1.

[246] Rafael D. Sorkin. "Forks in the road, on the way to quantum gravity". In: *Int. J. Theor. Phys.* 36 (1997), pp. 2759–2781. DOI: 10.1007/BF02435709. arXiv: gr-qc/9706002.

[247] J. David Brown and James W. York. "Jacobi's action and the recovery of time in general relativity". In: *Phys. Rev. D* 40 (10 Nov. 1989), pp. 3312–3318. DOI: 10.1103/PhysRevD.40.3312. URL: https://link.aps.org/doi/10.1103/PhysRevD.40.3312.

[248] W. G. Unruh. "A Unimodular Theory of Canonical Quantum Gravity". In: *Phys. Rev. D* 40 (1989), p. 1048. DOI: 10.1103/PhysRevD.40.1048.

[249] William G. Unruh and Robert M. Wald. "Time and the Interpretation of Canonical Quantum Gravity". In: *Phys. Rev. D* 40 (1989), p. 2598. DOI: 10.1103/PhysRevD.40.2598.

[250] Enrique Alvarez and Eduardo Velasco-Aja. "A Primer on Unimodular Gravity". In: Jan. 2023. arXiv: 2301.07641 [gr-qc].

[251] Luca Bombelli, W. E. Couch, and R. J. Torrence. "Time as spacetime four-volume and the Ashtekar variables". In: *Phys. Rev. D* 44 (8 Oct. 1991), pp. 2589–2592. DOI: 10.1103/PhysRevD.44.2589. URL: https://link.aps.org/doi/10.1103/PhysRevD.44.2589.

[252] Alan Daughton, Jorma Louko, and Rafael D Sorkin. "Instantons and unitarity in quantum cosmology with fixed four volume". In: *Phys. Rev. D* 58 (1998), p. 084008. DOI: 10.1103/PhysRevD.58.084008. arXiv: gr-qc/9805101.

[253] Mario Herrero-Valea and Raquel Santos-Garcia. "Non-minimal Tinges of Unimodular Gravity". In: *JHEP* 09 (2020), p. 041. DOI: 10.1007/JHEP09(2020)041. arXiv: 2006.06698 [hep-th].

[254] George F R Ellis. "The Trace-Free Einstein Equations and inflation". In: *Gen. Rel. Grav.* 46 (2014), p. 1619. DOI: 10.1007/s10714-013-1619-5. arXiv: 1306.3021 [gr-qc].

[255] Steven Carlip. "Quantum gravity: A Progress report". In: *Rept. Prog. Phys.* 64 (2001), p. 885. DOI: 10.1088/0034-4885/64/8/301. arXiv: gr-qc/0108040.

[256] Rafael Sorkin. *Introduction to Everpresent Lambda*. en. PIRSA:19110065 see, https://pirsa.org/19110065. Nov. 2019. DOI: 10.48660/19110065. URL: https://pirsa.org/19110065.

[257] Taichiro Kugo, Ryuichi Nakayama, and Nobuyoshi Ohta. "Covariant BRST quantization of unimodular gravity: Formulation with antisymmetric tensor ghosts". In: *Phys. Rev. D* 105.8 (2022), p. 086006. DOI: 10.1103/PhysRevD.105.086006. arXiv: 2202.03626 [hep-th].

[258] S. W. Hawking. "The Cosmological Constant Is Probably Zero". In: *Phys. Lett. B* 134 (1984), p. 403. DOI: 10.1016/0370-2693(84)91370-4.

[259] Sidney Coleman. "Why there is nothing rather than something: A theory of the cosmological constant". In: *Nuclear Physics B* 310.3 (Dec. 1988), pp. 643–668. DOI: 10.1016/0550-3213(88)90097-1.

[260] Y. Jack Ng and H. van Dam. "Possible solution to the cosmological constant problem". In: *Phys. Rev. Lett.* 65 (1990), pp. 1972–1974. DOI: 10.1103/PhysRevLett.65.1972.

[261] Steffen Gielen and João Magueijo. "Quantum resolution of the cosmological singularity without new physics". In: *EPL* 141.5 (2023), p. 59001. DOI: 10.1209/0295-5075/acbc48. arXiv: 2204.01771 [hep-th].

[262] Steffen Gielen, Daniele Oriti, and Lorenzo Sindoni. "Homogeneous cosmologies as group field theory condensates". In: *JHEP* 06 (2014), p. 013. DOI: 10.1007/JHEP06(2014)013. arXiv: 1311.1238 [gr-qc].

[263] Job Feldbrugge, Jean-Luc Lehners, and Neil Turok. "Lorentzian Quantum Cosmology". In: *Phys. Rev. D* 95.10 (2017), p. 103508. DOI: 10.1103/PhysRevD.95.103508. arXiv: 1703.02076 [hep-th].

[264] Joseph Samuel and Supurna Sinha. "Surface tension and the cosmological constant". In: *Phys. Rev. Lett.* 97 (2006), p. 161302. DOI: 10.1103/PhysRevLett.97.161302. arXiv: cond-mat/0603804.

[265] Xavier Martin, Denjoe O'Connor, and R. D. Sorkin. "The Random walk in generalized quantum theory". In: *Phys. Rev. D* 71 (2005), p. 024029. DOI: 10.1103/PhysRevD.71.024029. arXiv: gr-qc/0403085.

[266] Salvador Elias Venegas-Andraca. "Quantum walks: a comprehensive review". In: *Quant. Inf. Proc.* 11 (2012), pp. 1015–1106. DOI: 10.1007/s11128-012-0432-5. arXiv: 1201.4780 [quant-ph].

[267] Will Cunningham. *Classical and Quantum Growth Models for Discrete Spacetime*. https://willcunningham.net/trunk/lanl2019.pdf. 2019.

[268] Will Cunningham. *Quantum Dynamics of Total Orders*. https://willcunningham.net/trunk/odense2020-2.pdf. 2020.

[269] Antonio Padilla and Ippocratis D. Saltas. "A note on classical and quantum unimodular gravity". In: *Eur. Phys. J. C* 75.11 (2015), p. 561. DOI: 10.1140/epjc/s10052-015-3767-0. arXiv: 1409.3573 [gr-qc].

[270] Astrid Eichhorn. "On unimodular quantum gravity". In: *Class. Quant. Grav.* 30 (2013), p. 115016. DOI: 10.1088/0264-9381/30/11/115016. arXiv: 1301.0879 [gr-qc].

[271] Qingdi Wang, Zhen Zhu, and William G. Unruh. "How the huge energy of quantum vacuum gravitates to drive the slow accelerating expansion of the Universe". In: *Phys. Rev. D* 95.10 (2017), p. 103504. DOI: 10.1103/PhysRevD.95.103504. arXiv: 1703.00543 [gr-qc].

[272] D. M. Scolnic et al. "The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample". In: *Astrophys. J.* 859.2 (2018), p. 101. DOI: 10.3847/1538-4357/aab9bb. arXiv: 1710.00845 [astro-ph.CO].

[273] Dan Scolnic et al. "The Pantheon+ analysis: the full data set and light-curve release". In: *The Astrophysical Journal* 938.2 (2022), p. 113.

[274] Jennifer A Hoeting et al. "Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and EI George, and a rejoinder by the authors". In: *Statistical science* 14.4 (1999), pp. 382–417.

[275] S. Paradiso et al. "A convenient approach to characterizing model uncertainty with application to early dark energy solutions of the Hubble tension". In: *Mon. Not. Roy. Astron. Soc.* 528.2 (2024), pp. 1531–1540. DOI: 10.1093/mnras/stae101. arXiv: 2310.06747 [astro-ph.CO].

[276] S. Paradiso, G. McGee, and W. J. Percival. "Evaluating extensions to LCDM: an application of Bayesian model averaging and selection". In: *JCAP* 10 (2024), p. 021. DOI: 10.1088/1475-7516/2024/10/021. arXiv: 2403.02120 [astro-ph.CO].

[277] Neta A. Bahcall and Renyue Cen. "Galaxy Clusters and Cold Dark Matter: A Low-Density Unbiased Universe?" In: *The Astrophysical Journal Letters* 398 (Oct. 1992), p. L81. DOI: 10.1086/186582.

[278] M Costanzi, E Rozo, and et.al. "Methods for cluster cosmology and application to the SDSS in preparation for DES Year 1 release". In: *Monthly Notices of the Royal Astronomical Society* 488.4 (July 2019), pp. 4779–4800. ISSN: 0035-8711. DOI: 10.1093/mnras/stz1949. eprint: https://academic.oup.com/mnras/article-pdf/488/4/4779/29155032/stz1949.pdf. URL: https://doi.org/10.1093/mnras/stz1949.

[279] Mohamed H. Abdullah, Anatoly Klypin, and Gillian Wilson. "Cosmological Constraints on Ω_m and σ_8 from Cluster Abundances Using the GalWCat19 Optical-spectroscopic SDSS Catalog". In: *The Astrophysical Journal* 901.2, 90 (Oct. 2020), p. 90. DOI: 10.3847/1538-4357/aba619. arXiv: 2002.11907 [astro-ph.CO].

[280] Varun Sahni, Arman Shafieloo, and Alexei A. Starobinsky. "Two new diagnostics of dark energy". In: *Phys. Rev. D* 78 (2008), p. 103502. DOI: 10.1103/PhysRevD.78.103502. arXiv: 0807.3548 [astro-ph].

[281] Antony Lewis and Anthony Challinor. "Camb: Code for anisotropies in the microwave background". In: *Astrophysics source code library* (2011), ascl–1102.

[282] Santanu Das and Tarun Souradeep. "SCoPE: An efficient method of Cosmological Parameter Estimation". In: *JCAP* 07 (2014), p. 018. DOI: 10.1088/1475-7516/2014/07/018. arXiv: 1403.1271 [astro-ph.CO].

[283] Planck Collaboration. *Planck Legacy Archive*. https://pla.esac.esa.int/#cosmology. 2018.

[284] Nosiphiwo Zwane. private communication.

[285] Uros Seljak and Matias Zaldarriaga. "A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies". In: *The Astrophysical Journal* 469 (Oct. 1996), p. 437. DOI: 10.1086/177793. URL: https://doi.org/10.1086%2F177793.

[286] Jose Luis Bernal, Licia Verde, and Adam G. Riess. "The trouble with H_0". In: *JCAP* 10 (2016), p. 019. DOI: 10.1088/1475-7516/2016/10/019. arXiv: 1607.05617 [astro-ph.CO].

[287] Vivian Poulin et al. "Early Dark Energy Can Resolve The Hubble Tension". In: *Phys. Rev. Lett.* 122.22 (2019), p. 221301. DOI: 10.1103/PhysRevLett.122.221301. arXiv: 1811.04083 [astro-ph.CO].

[288] Edmund J. Copeland et al. "Scaling solutions as Early Dark Energy resolutions to the Hubble tension". In: *JCAP* 05 (2024), p. 078. DOI: 10.1088/1475-7516/2024/05/078. arXiv: 2309.15295 [astro-ph.CO].

[289] V. Mukhanov. *Physical Foundations of Cosmology*. Oxford: Cambridge University Press, 2005. ISBN: 978-0-521-56398-7. DOI: 10.1017/CBO9780511790553.

[290] Santanu Das and Tarun Souradeep. "Suppressing CMB low multipoles with ISW effect". In: *Journal of Cosmology and Astroparticle Physics* 2014.02 (Feb. 2014), pp. 002–002. DOI: 10.1088/1475-7516/2014/02/002. arXiv: 1312.0025[astro-ph.CO]. URL: https://doi.org/10.1088%2F1475-7516%2F2014%2F02%2F002.

[291] Yuting Wang et al. "Evolution of Dark Energy Reconstructed from the Latest Observations". In: *The Astrophysical Journal Letters* 869.1, L8 (Dec. 2018), p. L8. DOI: 10.3847/2041-8213/aaf238. arXiv: 1807.03772 [astro-ph.CO].

[292] Gong-Bo Zhao et al. "Dynamical dark energy in light of the latest observations". In: *Nature Astronomy* 1 (Aug. 2017), pp. 627–632. DOI: 10.1038/s41550-017-0216-z. arXiv: 1701.08165 [astro-ph.CO].

[293] Marc Kamionkowski and Adam G. Riess. "The Hubble Tension and Early Dark Energy". In: *Ann. Rev. Nucl. Part. Sci.* 73 (2023), pp. 153–180. DOI: 10.1146/annurev-nucl-111422-024107. arXiv: 2211.04492 [astro-ph.CO].

[294] Vivian Poulin, Tristan L. Smith, and Tanvi Karwal. "The Ups and Downs of Early Dark Energy solutions to the Hubble tension: A review of models, hints and constraints circa 2023". In: *Phys. Dark Univ.* 42 (2023), p. 101348. DOI: 10.1016/j.dark.2023.101348. arXiv: 2302.09032 [astro-ph.CO].

[295] Alexandra Abate et al. "Large Synoptic Survey Telescope: Dark Energy Science Collaboration". In: (Nov. 2012). arXiv: 1211.0310 [astro-ph.CO].

[296] Y. Mellier et al. "Euclid. I. Overview of the Euclid mission". In: *Astron. Astrophys.* 697 (2025), A1. DOI: 10.1051/0004-6361/202450810. arXiv: 2405.13491 [astro-ph.CO].

[297] H Aussel et al. "Euclid Quick Data Release (Q1)–Data release overview". In: *arXiv preprint arXiv:2503.15302* (2025).

[298] V. Ajani et al. "Euclid Preparation. XXVIII. Forecasts for ten different higher-order weak lensing statistics". In: *Astron. Astrophys.* 675 (2023), A120. DOI: 10.1051/0004-6361/202346017. arXiv: 2301.12890 [astro-ph.CO].

[299] Ottavia Truttero et al. "Baryon-free S_8 tension with Stage IV cosmic shear surveys". In: (Oct. 2024). DOI: 10.33232/001c.129965. arXiv: 2410.18191 [astro-ph.CO].

[300] Carlos A. P. Bengaly, Chris Clarkson, and Roy Maartens. "The Hubble constant tension with next-generation galaxy surveys". In: *JCAP* 05 (2020), p. 053. DOI: 10.1088/1475-7516/2020/05/053. arXiv: 1908.04619 [astro-ph.CO].

[301] B. Jain et al. "The Whole is Greater than the Sum of the Parts: Optimizing the Joint Science Return from LSST, Euclid and WFIRST". In: (Jan. 2015). arXiv: 1501.07897 [astro-ph.IM].

[302] Jun John Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Quantum physics, quantum information and quantum computation. Cambridge University Press, Oct. 2020. ISBN: 978-0-8053-8291-4, 978-1-108-52742-2, 978-1-108-58728-0. DOI: 10.1017/9781108587280.

[303] Scott Dodelson and Fabian Schmidt. *Modern Cosmology*. Elsevier Science, 2020. DOI: 10.1016/C2017-0-01943-2.

Appendix A

Technical supplement to Chapter 1

A.1 General covariance of Fokker-Planck equation

In this appendix, we give a brief summary of Sorkin's arguments in [59] for deriving the Fokker-Planck equation from first physical principles. An alternative approach is the standard textbook derivation in stochastic processes, but the following method offers useful insights into the transformation properties of the drift and diffusion coefficients.

Given any region Ω of the phase space,

$$\int_\Omega dZ \; f_\tau(Z,\tau) \tag{A.1}$$

is equal to the number of particles that are in Ω when their proper time is τ [1]. Since f_τ is not integrated with any volume measure, it acts as a density. In other words, under an arbitrary coordinate transformation $Z^A \to Z'^M(Z)$ of the phase space, f_τ changes as

$$f'_\tau(Z') = \Big|\frac{\partial Z}{\partial Z'}\Big| f_\tau(Z)\,, \tag{A.2}$$

where $|\frac{\partial Z'}{\partial Z}|$ is the determinant of the Jacobian matrix of the transformation, and primed quantities are in the new coordinate system.

The basic ansatz for the evolution equation of f_τ is

$$\frac{\partial f_\tau}{\partial \tau} = K^{AB}\partial_A\partial_B f_\tau + F^A\partial_A f_\tau - C f_\tau\,, \tag{A.3}$$

where $K^{AB}(X)$ is a yet-unknown symmetric matrix in the phase space. $F^A(X)$ and $C(X)$ are unknown functions as well.

This is the most general second-order equation that can be written for f_τ. One might question the generality of this ansatz by considering the inclusion of higher-order derivatives of f_τ. The first-order

[1] In Sorkin's notation in [59], f_τ is a probability density, so its integral over the whole phase space is 1. This differs from our notation by a normalisation factor equal to the total number of particles.

derivative in τ is justified by the need for an evolution equation for f_τ. In Appendix A in [59], Sorkin proves that higher phase space derivatives on the right-hand side of this equation will lead to a negative f_τ for some particular choices of the initial distribution at $\tau = 0$. Hence such terms cannot be present unless one keeps an infinite series of these terms. The same positivity requirement from f_τ at all (Z, τ) forces K^{AB} to be a positive semi-definite matrix.

Now the assumption of conservation of particle number can be written as

$$\frac{d}{d\tau}\int_{\mathcal{S}} dZ\ f_\tau(Z,\tau) = 0 \quad \Rightarrow \quad \int_{\mathcal{S}} dZ\ \frac{\partial f_\tau}{\partial \tau}(Z,\tau) = 0\,. \tag{A.4}$$

Using (A.3) and integrating by parts leads to

$$\int_{\mathcal{S}} dZ\ (\partial_A\partial_B K^{AB} - \partial_A F^A - C) f_\tau = 0\,. \tag{A.5}$$

Since this holds for any distribution f_τ, one concludes that

$$C = \partial_A\partial_B K^{AB} - \partial_A F^A\,. \tag{A.6}$$

Defining $v^A = 2\partial_B K^{AB} - F^A$ transforms the ansatz into

$$\frac{\partial f_\tau}{\partial \tau} = -\partial_A J^A\,, \tag{A.7}$$

where

$$J^A = v^A f_\tau - \partial_B(K^{AB} f_\tau)\,, \tag{A.8}$$

which leads to (1.1).

Given that f_τ is a density, the transformation properties of the other quantities can be derived. Equation (A.7) says that J^A must be a vector density; i.e. it transforms like

$$J'^M(Z') = \left|\frac{\partial Z}{\partial Z'}\right| \frac{\partial Z'^M}{\partial Z^A} J^A(Z)\,. \tag{A.9}$$

This transformation rule for J is valid for *any* distribution f_τ. Expanding (A.9) using the definition (A.8) together with the transformation rule (A.2), one finds

$$K'^{MN} = \frac{\partial Z'^M}{\partial Z^A}\frac{\partial Z'^N}{\partial Z^B} K^{AB}\,, \tag{A.10}$$

while v mixes with K when transformed;

$$v'^M = \frac{\partial Z'^M}{\partial Z^A} v^A + \frac{\partial^2 Z'^M}{\partial Z^A \partial Z^B} K^{AB}\,. \tag{A.11}$$

A.1.1 Detailed derivation of covariant diffusion equation

Here, we outline the steps in deriving the covariant diffusion equation from the Fokker-Planck equation, which were omitted in section 1.1.1.

Having found the most general form of u^A and K^{AB} (equations (1.17) and (1.14)) on the phase space Γ_m^+, we substitute them back into the Fokker-Planck equation to find

$$\frac{\partial f_\tau}{\partial \tau} = -\partial_A\Big(\frac{1}{m}L^A f_\tau\Big) + \kappa\frac{\partial}{\partial q^i}\bigg(\sqrt{|\hat{g}|}\mathsf{h}^{ij}\frac{\partial}{\partial q^j}\Big(\frac{f_\tau}{\sqrt{|\hat{g}|}}\Big)\bigg)\,. \tag{A.12}$$

Now, we have to integrate out the proper time since it is not observable. Using the definition

$$\sqrt{|\hat{g}|}f(x,p) = \int_{-\infty}^{\infty} d\tau\ f_\tau\,, \tag{A.13}$$

one finds

$$\frac{1}{\sqrt{|\hat{g}|}}\partial_A\Big(\frac{1}{m}L^A\sqrt{|\hat{g}|}f\Big) = \kappa\frac{1}{\sqrt{|\hat{g}|}}\frac{\partial}{\partial q^i}\Big(\sqrt{|\hat{g}|}\mathsf{h}^{ij}\frac{\partial}{\partial q^j}f\Big)\,. \tag{A.14}$$

To proceed further, we need to be explicit about the choice of the metric $\hat{g}$ on the phase space Γ_m^+. The natural choice is the Sasaki metric on the cotangent bundle [85], for which the volume measure, when restricted to Γ_m^+, is $d^4x\,d^3q\,\sqrt{-g}\,\sqrt{\mathsf{h}}$. As a result, the metric determinant separates into spacetime and momentum contributions, i.e.

$$\sqrt{|\hat{g}|} = \sqrt{-g}\ \sqrt{\mathsf{h}}\,. \tag{A.15}$$

Let $\hat{\nabla}$ be the covariant derivative corresponding to the Sasaki metric. The left-hand side of (A.14) is now equal to

$$\frac{1}{m}\hat{\nabla}_A(fL^A) = \frac{1}{m}f\hat{\nabla}_A L^A + \frac{1}{m}L[f]\,. \tag{A.16}$$

Liouville's theorem states that $\hat{\nabla}_A L^A = div(L) = 0$ [85]. Therefore, (A.14) can be written as

$$L[f] = m\kappa\frac{1}{\sqrt{\mathsf{h}}}\frac{\partial}{\partial q^i}\Big(\sqrt{\mathsf{h}}\mathsf{h}^{ij}\frac{\partial}{\partial q^j}f\Big)\,. \tag{A.17}$$

h_{ij} being the hyperbolic metric on $\mathbb{H}_3$, the right-hand side of this equation gives us the Laplacian operator on $\mathbb{H}_3$. We have found the desired covariant diffusion equation.

A.1.2 Covariance of Liouville operator

On the cotangent bundle of the spacetime, the Liouville operator is defined as

$$L = p^\mu\frac{\partial}{\partial x^\mu} + \Gamma^\sigma_{\mu\nu}p_\sigma p^\mu\frac{\partial}{\partial p_\nu}\,. \tag{A.18}$$

The spacetime components give the definition of momentum, and its momentum components yield the geodesic motion. In other words, for a particle of mass m with no diffusion, the Liouville operator gives the total derivative of its phase space motion, $dZ/d\tau = L/m$. Here we want to show that L is indeed a true vector under the set of admissible coordinate transformations (1.9). First, we change the coordinates of the cotangent bundle from (x^μ, p_μ) to $(x^\mu, q_a = e_a{}^\mu p_\mu)$, where $a = 0, ..., 3$. The

coordinate vector fields are written in the new coordinates as

$$\begin{aligned}\frac{\partial}{\partial x^\mu} &\to \frac{\partial}{\partial x^\mu} + p_\nu e_a{}^\nu{}_{,\mu}\frac{\partial}{\partial q_a}\,,\\ \frac{\partial}{\partial p_\mu} &\to e_a{}^\mu\frac{\partial}{\partial q_a}\,.\end{aligned} \tag{A.19}$$

Therefore,

$$L = p^\mu\frac{\partial}{\partial x^\mu} + p.\nabla_p e_a\frac{\partial}{\partial q_a}\,. \tag{A.20}$$

Now, we demonstrate that this is indeed a true vector. In components form, L is expressed as

$$L = \Big(L^\mu = p^\mu,\ L_a = p.\nabla_p e_a\Big)\,. \tag{A.21}$$

In a new coordinate system (x'^α, q'_r), this would be

$$L' = \Big(L'^\alpha = p'^\alpha,\ L'_r = p.\nabla_p e'_r\Big)\,, \tag{A.22}$$

and we have to show that

$$L'^\alpha = \frac{\partial x'^\alpha}{\partial x^\mu}L^\mu + \frac{\partial x'^\alpha}{\partial q_a}L_a\,, \tag{A.23}$$

$$L'_r = \frac{\partial q'_r}{\partial x^\mu}L^\mu + \frac{\partial q'_r}{\partial q_a}L_a\,. \tag{A.24}$$

The Jacobian for the coordinate transformation is

$$\begin{pmatrix} \frac{\partial x'^\alpha}{\partial x^\mu} & \frac{\partial x'^\alpha}{\partial q_a} = 0 \\ \frac{\partial q'_r}{\partial x^\mu} = \partial_\mu\big(e'_r(x).e^a(x)\big)q_a & \frac{\partial q'_r}{\partial q_a} = e'_r(x).e^a(x)\end{pmatrix}\,. \tag{A.25}$$

Since $\partial x'/\partial q = 0$, (A.23) reduces to

$$p'^\alpha = \frac{\partial x'^\alpha}{\partial x^\mu}p^\mu\,, \tag{A.26}$$

which is trivially valid due to p^μ being a covariant vector on the spacetime. For equation (A.24), the right-hand side equals

$$\begin{aligned}p^\mu\partial_\mu(e'_r.e^a)q_a + (e'_r.e^a)p.\nabla_p e_a &= p.\big(p^\mu\partial_\mu(e'_r.e^a)e_a + (e'_r.e^a)\nabla_p e_a\big)\\ &= p.\nabla_p\big((e'_r.e^a)e_a\big) = p.\nabla_p e'_r\,,\end{aligned} \tag{A.27}$$

and therefore, (A.24) is also true. We have shown that L is a true vector field under coordinate transformations of the cotangent bundle of the form $(x^\mu, q_a) \to (x'^\alpha, q'_r)$, where $a, r = 0, ..., 3$. Now note that since L/m is the Hamiltonian vector field corresponding to the Hamiltonian $H = \frac{1}{2m}p_\mu p_\nu g^{\mu\nu}$, it is tangent to the surfaces of constant H. In particular, L is also tangent to the phase space Γ_m^+ and can be thought of as a vector field on the phase space, covariant under admissible coordinate transformations (1.9). Since L is tangent to Γ_m^+, restricting it to this submanifold is done just by dropping $\partial/\partial q_0$. In other words, when restricted to (y^μ, q_i), $i = 1, 2, 3$, L is written as

$$L = p^\mu\frac{\partial}{\partial y^\mu} + p.\nabla_p e_i\frac{\partial}{\partial q_i}\,. \tag{A.28}$$

A.2 Unravelling the covariant diffusion equation

A.2.1 Equivalence between Fokker-Planck equation and Itô SDEs

The evolution of a stochastic system can be described equivalently in two ways. On the one hand, if one is interested purely in the ensemble properties of the system, the natural description is in terms of the evolution of the probability distribution on the space of the dynamical degrees of freedom Z^A. In the case of a Markovian, almost surely continuous evolution, the evolution equation for f_τ is of the form of a Fokker-Planck equation:

$$\partial_\tau f_\tau = -\partial_A\Big(v^A f_\tau\Big) + \partial_A\partial_B\Big(K^{AB} f_\tau\Big)\,, \tag{A.29}$$

Following Sorkin's argument in [59], it is possible to gain insight into the meaning of K^{AB} and v^A by using a particular substitution for f_τ. Take the special case where the particles are prepared at a particular point Z_0^A in the phase space at their zero proper time:

$$f_\tau(Z,\tau=0) = N\delta(Z-Z_0)\,. \tag{A.30}$$

Here N is the total number of particles. Then he computes the average of quantities Z^A and Z^AZ^B, and differentiates with respect to τ at $\tau=0$:

$$\frac{d\langle Z^A\rangle}{d\tau}|_{\tau=0} = \frac{d}{d\tau}\int_{\mathcal{S}} dZ\ f_\tau(\tau,X)Z^A = \int_{\mathcal{S}} dZ\ \frac{\partial f_\tau}{\partial\tau}Z^A = -\int_{\mathcal{S}} dZ\ \partial_B(J^B)Z^A\,. \tag{A.31}$$

Integration by parts and use of (A.30) reveals that

$$\frac{d\langle Z^A\rangle}{d\tau}\Big|_{\tau=0} = v^A(Z_0)\,. \tag{A.32}$$

In a similar fashion,

$$\frac{d}{d\tau}\Big\langle (Z^A-Z_0^A)(Z^B-Z_0^B)\Big\rangle\Big|_{\tau=0} = 2K^{AB}(Z_0)\,. \tag{A.33}$$

Since Z_0 was an arbitrary point, the last two equations are valid everywhere in the phase space independent of the choice of the distribution function. These are essentially the same equations as (1.2) written in a slightly different form:

$$\begin{aligned} v^A &= \lim_{\delta\tau\to0}\frac{1}{\delta\tau}\Big\langle \delta Z^A\Big\rangle\,,\\ K^{AB} &= \lim_{\delta\tau\to0}\frac{1}{2\delta\tau}\Big\langle \delta Z^A\delta Z^B\Big\rangle\,. \end{aligned} \tag{A.34}$$

If one is interested in a sample realisation of the stochastic dynamics, the Fokker-Planck equation is limiting. It is then more convenient to *unravel* the dynamics in terms of individual stochastic trajectories $Z^A(\tau)$. Indeed, there exists an equivalent description of the stochastic process in terms of Itô SDEs. Here one should determine the differential increments dZ^A along a trajectory. Separating the mean, this can be described by the following SDE:

$$dZ^A = v^A d\tau + dW^A\,. \tag{A.35}$$

Furthermore, to satisfy (A.34), the stochastic process $W^A(\tau)$ should satisfy the following statistics

$$\Big\langle dW^A \Big\rangle = 0 \,, \qquad \Big\langle dW^A dW^B \Big\rangle = 2K^{AB} d\tau \,, \tag{A.36}$$

with dW^A being defined as the infinitesimal limit of the increment in W^A. We assume sufficient coarse graining so that the dynamics is Markovian, with each increment consisting of numerous independent random kicks. So motivated by the central limit theorem, we take the increments to follow a Gaussian distribution with the above statistics. It is useful to think of dW as order $\sqrt{d\tau}$, with Itô's lemma making this precise. In practice, this means that the increments satisfy Itô's rules:

$$dW^A dW^B = 2K^{AB} d\tau \,, \quad dW^A d\tau = 0 \,. \tag{A.37}$$

A common, though imprecise, way these equations appear in the physics literature is in the Langevin form:

$$\partial_\tau Z^A = v^A(Z,\tau) + \xi^A(\tau) \,, \tag{A.38}$$

where ξ^A, formally the distributional time derivative of the stochastic process $W^A(\tau)$, is a collection of Gaussian white noise fields with the following stochastic properties:

$$\Big\langle \xi^A(\tau) \Big\rangle = 0 \,, \qquad \Big\langle \xi^A(\tau)\xi^B(\tau') \Big\rangle = 2K^{AB}\delta(\tau-\tau') \,. \tag{A.39}$$

Therefore, moving from the Fokker-Planck description to the SDE description is straightforward. After determining the true vector u^A and the true tensor K^{AB}, we identify the drift pseudo-vector v^A from (1.5). Then, we can easily write down the corresponding SDE (A.35), together with the correlations of the noise term (A.36). In the next appendix, we apply this to the covariant diffusion equation.

A.2.2 From covariant diffusion equation to stochastic trajectories

Up to zeroth order in a curvature expansion, we have already identified the matrix K^{AB}; see (1.14). The spacetime components of v^A have also been found to be proportional to the momentum. For the momentum components of v, we write

$$\begin{aligned} v^i &= u^i + \frac{\kappa}{\sqrt{\mathsf{h}}} \frac{\partial}{\partial q^j} \Big(\sqrt{\mathsf{h}} \mathsf{h}^{ij} \Big) \\ &= p.\nabla_p e^i + \kappa E \frac{\partial}{\partial q^j} \Big(\frac{\delta^{ij}}{E} + \frac{q^i q^j}{m^2 E} \Big) \\ &= p.\nabla_p e^i + \frac{3\kappa}{m^2} q^i \,, \end{aligned} \tag{A.40}$$

where in the second line, we have used the representation of the Liouville operator in the (x^μ, q_i) coordinates (A.28). We can already see the similarity between $3\kappa q^i/m^2$ and the bias 4-vector b^μ in

(1.33). With v^i and K^{ij} determined, we can now write the following SDE

$$dx^\mu = \frac{1}{m}p^\mu d\tau\,, \tag{A.41}$$

$$dq^i = p.\nabla_p e^i d\tau + \frac{3\kappa}{m^2}q^i d\tau + dW^i\,, \tag{A.42}$$

where the noise term satisfies

$$\Big\langle dW^i \Big\rangle = 0\,, \qquad \Big\langle dW^i dW^j \Big\rangle = 2\kappa\Big(\delta^{ij} + \frac{1}{m^2}q^i q^j\Big)d\tau\,. \tag{A.43}$$

To make this look manifestly covariant, we next find the differential increment of energy $E = \sqrt{q^2 + m^2}$ using Itô's Lemma:

$$dE = \frac{\partial E}{\partial q^i}dq^i + \frac{1}{2}\frac{\partial^2 E}{\partial q^i \partial q^j}dq^i dq^j \tag{A.44}$$

$$= \frac{1}{E}q^i p.\nabla_p e_i d\tau + \frac{3\kappa q^2}{Em^2}d\tau + \frac{1}{E}q_i dW^i + \frac{3\kappa}{E}d\tau \tag{A.45}$$

$$= p.\nabla_p e^0 d\tau + \frac{3\kappa}{m^2}E d\tau + dW^0\,, \tag{A.46}$$

where in the last line we have defined

$$dW^0 \coloneqq \frac{1}{E}q_i dW^i\,, \tag{A.47}$$

and used the fact that $q^a p.\nabla_p e_a = 0$. This holds whenever $q_a = p.e_a$ and the e_a form a complete set of vectors.

Therefore, we can now extend (A.42) to include dq^0 as well:

$$dq^a = p.\nabla_p e^a d\tau + \frac{3\kappa}{m^2}q^a d\tau + dW^a\,, \tag{A.48}$$

such that the noise has the following statistics

$$\Big\langle dW^a \Big\rangle = 0\,, \qquad \Big\langle dW^a dW^b \Big\rangle = 2\kappa\Big(\eta^{ab} + \frac{1}{m^2}q^a q^b\Big)d\tau\,. \tag{A.49}$$

Note that by definition (A.47) satisfies $q_a dW^a = 0$. Finally, transforming from q^a coordinates to p^μ is straightforward by contracting with the tetrad field. If one defines

$$dW^\mu \coloneqq e_a{}^\mu dW^a\,, \tag{A.50}$$

one can easily check that the above SDE gives the stochastic geodesic equation (1.34). We have shown that the covariant diffusion equation at the level of the distribution function is equivalent to the Langevin equation of covariant Brownian motion.

A.2.3 Relation between spacetime covariance and phase space covariance

Our derivation of the covariant diffusion equation (1.19) and the corresponding Langevin equation (1.34) followed from different lines of reasoning. For the former, we required the Fokker-Planck equation to

respect general covariance under admissible coordinate transformations of the phase space Γ_m^+. For the Langevin equation, however, we required manifest general covariance of the correction to the geodesic equation, together with conservation of mass. In the previous appendix, we showed that the final results are two equivalent descriptions of the same phenomenon. It is illuminating to also find interconnections between the two sets of first principles we used to derive the equations. This is the purpose of this appendix.

First of all, notice that in section A.2.1, we showed that the matrix tensor K^{AB} is the covariance matrix of the differentials of the motion. On the other hand, in section 1.1.2, we had $D^{\mu\nu}$ as the covariance matrix of the noise terms in the geodesic equation. To connect the two, we only have to remember that our phase space coordinate system uses a projection on some given tetrad field and only keeps the spatial components of the physical momentum as independent variables. Therefore,

$$\boxed{K^{ij} = e^i{}_\mu e^j{}_\nu D^{\mu\nu}}\,. \tag{A.51}$$

Similarly, in the Fokker-Planck equation, we had a pseudo-vector v^A that we found to be the drift of the differentials of motion. On the other hand, we had the drift 4-vector b^μ in our discussion of the SDEs. Again, the two are related by a projection on the tetrad field:

$$\boxed{v^i = e^i{}_\mu b^\mu}\,, \tag{A.52}$$

and this can be used to find the true vector u using (1.5). Now we prove the following two statements:

Statement 1. The matrix

$$K = \begin{pmatrix} 0 & 0 \\ 0 & K^{ij} \end{pmatrix} \tag{A.53}$$

is a tensor on the phase space Γ_m^+ if and only if $D^{\mu\nu}$, which satisfies $K^{ij} = e^i{}_\mu e^j{}_\nu D^{\mu\nu}$, is a symmetric spacetime tensor orthogonal to p:

$$p_\mu D^{\mu\nu} = 0\,, \tag{A.54}$$

whenever $p^2 = -m^2$.

The proof is straightforward. Remember that the set of admissible coordinate transformations of the phase space are

$$\left(x^\mu, q_i\right) \to \left(x'^\alpha(x),\ q'_I(x,q) = e'_I(x).e^a(x)\ q_a\right), \tag{A.55}$$

which have the following Jacobian in the momentum coordinates

$$\frac{\partial q'^I}{\partial q^i} = e'^I.e_i + \frac{q_i}{q^0} e'^I.e_0\,. \tag{A.56}$$

The tensor condition on K reduces to

$$e'^I{}_\alpha e'^J{}_\beta D^{\alpha\beta} = \frac{\partial q'^I}{\partial q^i}\frac{\partial q'^J}{\partial q^j} e^i{}_\mu e^j{}_\nu D^{\mu\nu}\,. \tag{A.57}$$

On the other hand, with some algebra and using the completeness of the vector fields e_a, one can show that

$$\frac{\partial q'^I}{\partial q^i} e^i{}_\mu = e'^I{}_\mu + \frac{1}{q^0} e'^I . e_0 \; p_\mu \,. \tag{A.58}$$

Therefore, the right-hand side of (A.57) equals the left-hand side, plus three additional terms involving the contraction of $D^{\mu\nu}$ with either p_μ, p_ν, or both. Since this equation has to be valid for any choice of the tetrad fields, one concludes that $p.D = 0$. Note that this only needs to be true on the mass shell.

The same argument works for vectors as well. Given a vector field on the spacetime, its projection on the tetrad gives a vector field on the phase space if and only if it is orthogonal to p. This will be used in the proof of the next statement.

Statement 2. Assume that $p_\nu D^{\mu\nu} = 0$ on the mass shell, and that $K^{ij} = e^i{}_\mu e^j{}_\nu D^{\mu\nu}$. Then u^A, given by

$$u^A = \left(u^\mu = 0, \;\; u^i = v^i - \frac{1}{\sqrt{|\hat{g}|}} \frac{\partial}{\partial q^j} \left(K^{ij} \sqrt{|\hat{g}|} \right) \right) , \tag{A.59}$$

is a vector field on the phase space Γ_m^+ if and only if b^μ, which satisfies $v^i = e^i{}_\mu b^\mu$, is a vector field on the spacetime that satisfies

$$p_\mu b^\mu + D^\mu{}_\mu = 0 \,, \tag{A.60}$$

whenever $p^2 = -m^2$.

For the proof, first let us find u^i in terms of b and D using its definition (1.5):

$$u^i = e^i{}_\mu b^\mu - \frac{1}{\sqrt{\mathsf{h}}} \frac{\partial}{\partial q^j} \left(\sqrt{\mathsf{h}} D^{\mu\nu} \right) e^i{}_\mu e^j{}_\nu \,. \tag{A.61}$$

Next we switch from the $\partial/\partial q_i$ basis to $\partial/\partial p_\mu$ using

$$\frac{\partial}{\partial q^i} = \left(e_i{}_\mu + \frac{q_i}{E} e_0{}_\mu \right) \frac{\partial}{\partial p_\mu} \,. \tag{A.62}$$

Using the fact that both $p_\nu D^{\mu\nu}$ and its tangential derivative along the mass shell vanish on the mass shell, one can write u as the projection of a spacetime vector on the tetrad:

$$u^i = e^i{}_\mu B^\mu \,, \tag{A.63}$$

$$B^\mu = b^\mu - \left(g^{\nu\sigma} + \frac{1}{m^2} p^\nu p^\sigma \right) \frac{\partial}{\partial p^\sigma} D^\mu{}_\nu \,. \tag{A.64}$$

Similar to our proof for the previous statement, since the spacetime components of u are assumed to be zero, the vector condition on u reduces to B being orthogonal to p. In order to show this, first note that the following vector field is tangent to the mass shell:

$$\left(g^{\nu\sigma} + \frac{1}{m^2} p^\nu p^\sigma \right) \frac{\partial}{\partial p^\sigma} \,, \tag{A.65}$$

which implies that when acted on $p_\mu D^\mu{}_\nu$ gives zero. Therefore, on the mass shell,

$$p_\mu B^\mu = p_\mu b^\mu + D^\mu{}_\mu \,. \tag{A.66}$$

So u is a vector field if and only if the right-hand side is zero on the mass shell. The statement follows.

We highlight that the covariance of the Liouville operator cannot be inferred from this statement, as its spacetime components are non-zero. It requires an independent proof, which we provided in Appendix (A.1.2). Furthermore, using (A.63), it can be easily shown that the u vector corresponding to (1.33) is identically zero.

A.3 Generalised Boltzmann hierarchy equations for stochastic dark matter

The generalised Boltzmann hierarchy provides a systematic way of doing non-relativistic approximation, without the need to assume that the stochastic dark matter is a fluid. It provides a hierarchy of equations both at the background and perturbations level, in which the fluid quantities constitute only the lowest moments of the distribution function. In this appendix, we will see how the moments are defined and what their evolution equations are for the case of stochastic dark matter that is governed by the covariant diffusion equation.

At the background level, the pressure moments are defined as follows [121]:

$$\bar{P}_n = \frac{4\pi}{3}\int_0^\infty dq\; q^2 E\Big(\frac{q}{E}\Big)^{2n} f_0 = w_n \bar{\rho}_{\rm dm} \,. \tag{A.67}$$

Therefore, the zeroth moment gives the energy density $\bar{P}_0 = \bar{\rho}_{\rm dm}/3$, and $w_0 = 1/3$. The first moment is the pressure itself $\bar{P}_1 = \bar{P}_{\rm dm}$, and $w_1 = w_{\rm dm}$ is the equation of state of dark matter. These two quantities are coupled to higher-order moments via their evolution equation. In practice, as we are only interested in the regime $\Gamma \ll H_0$ of the diffusion parameter, even the pressure has a negligible impact on the background dynamics. Nevertheless, writing down the coupled hierarchy of equations sets the stage for the analogous calculations for perturbations.

Multiplying the background covariant diffusion equation (1.47) by q^{2n+2}/E^{2n-1} and integrating dq gives

$$\bar{P}_n' + (2n+3)\mathcal{H}\bar{P}_n - (2n-1)\mathcal{H}\bar{P}_{n+1} = \frac{a\Gamma}{9}\Big(2n(2n+1)m\bar{N}_{n-1} - (8n^2-3)m\bar{N}_n + 2n(2n-1)m\bar{N}_{n+1}\Big)\,, \tag{A.68}$$

where N_n's are the higher moment analogues of the number density, defined via

$$\bar{N}_n = 4\pi\int_0^\infty dq\; q^2\Big(\frac{q}{E}\Big)^{2n} f_0 \,. \tag{A.69}$$

$\bar{N}_0 = \bar{n}_{\rm dm}$ is the background number density. These number density moments arise from the diffusion term. Now in order to have a closed system of equations, one might try to find the evolution equations

for $\bar{N}_n$, but those will also contain integrals that do not simplify to either $\bar{P}_n$ or $\bar{N}_n$. Instead, similar to (1.46), we expand $\bar{N}_n$ in terms of higher $\bar{P}_n$ moments. To do so, first use $E^2 = m^2 + q^2$ to write

$$m = -E \sum_{j=0}^{\infty} \frac{(2j-3)!!}{j!} \left(\frac{q^2}{2E^2}\right)^j . \tag{A.70}$$

We plug this into (A.69) to find

$$m\bar{N}_n = -3 \sum_{j=0}^{\infty} \frac{(2j-3)!!}{2^j j!} \bar{P}_{n+j} . \tag{A.71}$$

The equations (A.68) and (A.71) together give an infinite hierarchy of equations for the background. In principle, by including high enough moments in this hierarchy, one could account for relativistic contributions to an arbitrary precision. In practice, we set all the pressure moments beyond some n_{max} to zero.

In the non-relativistic approximation, though, one could directly use the approximate solution (1.49) to find all higher-order pressure moments. The result turns out to be

$$w_n \simeq \frac{(2n+1)!!}{3} w_1^n , \tag{A.72}$$

where w_1 is given by (1.51). We use CDM as the initial condition for the stochastic dark matter; therefore, at any non-zero scale factor, we use (1.51) and (A.72) for setting the initial conditions of the hierarchy.

For perturbations, one could similarly define the higher-order density perturbations as [121]:

$$\delta P_n = \frac{2\pi}{3} \int_0^{\infty} dq \int_{-1}^{1} d\mu \; q^2 E \Big(\frac{q}{E}\Big)^{2n} \delta f , \tag{A.73}$$

where the total distribution function has been expanded as $f = f_0 + \delta f$, and $\mu = \hat{k}.\hat{q}$, k being the wavevector. In addition, at the perturbations level, the angular moments of the distribution function become important as well [121]:

$$(\bar{\rho}_{\rm dm} + \bar{P}_{\rm dm}) f_{n,\ell} = \frac{2\pi \ell! i^\ell}{(2\ell-1)!!} \int_0^{\infty} dq \int_{-1}^{1} d\mu \; P_\ell(\mu) q^2 E \Big(\frac{q}{E}\Big)^{2n+\ell} \delta f . \tag{A.74}$$

Therefore, we have a hierarchy characterised by two indices (n, ℓ), the generalised Boltzmann hierarchy. Define the n-th order density contrast and velocity divergence as $\delta_n = \delta P_n / \bar{\rho}_{\rm dm}$ and $k f_{n,1} = \theta_n$, respectively. The lowest order moments correspond to the fluid quantities as follows

$$\delta_0 = \frac{1}{3}\delta_{\rm dm}, \quad \delta_1 = \frac{\delta P}{\bar{\rho}_{\rm dm}} = c_s^2 \delta_{\rm dm} , \tag{A.75}$$

$$k f_{0,1} = \theta, \quad k f_{0,2} = \sigma . \tag{A.76}$$

To find the coupled equations, we first need to write down the covariant diffusion equation (1.19) for

the perturbed distribution δf in FRW:

$$\delta f' - q\mathcal{H}\frac{\partial \delta f}{\partial q} + i\frac{q}{E}k\mu\,\delta f + \frac{\partial f_0}{\partial q}(q\phi' - iEk\mu\psi) \\ = \frac{a\Gamma m}{3}\frac{1}{q^2}\Big[\psi\frac{\partial}{\partial q}\Big(Eq^2\frac{\partial}{\partial q}f_0\Big) + \frac{\partial}{\partial q}\Big(Eq^2\frac{\partial}{\partial q}\delta f\Big) + \frac{m^2}{E}\frac{1}{\sin\vartheta}\frac{\partial}{\partial\vartheta}\Big(sin\vartheta\frac{\partial}{\partial\vartheta}\delta f\Big)\Big]\,, \tag{A.77}$$

where $\mu = cos\vartheta$. By integrating this equation with suitable pre-factors, we get the evolution equations for δ_n and $f_{n,\ell}$. Similar to what happened in the background, here the right-hand side of the above equation gives integrals that are not readily in the form of either δ_n or $f_{n,\ell}$. To have a closed system, we need to insert expansions similar to (A.70) into such integrals. After some algebra, we find

$$\delta_n' + \frac{1}{3}(1+w_1)\theta_n + (2n-3w_1)\mathcal{H}\delta_n - (2n-1)\mathcal{H}\delta_{n+1} - \Big((2n+3)w_n - (2n-1)w_{n+1}\Big)\phi' \\ = \frac{a\Gamma}{3}\left(\sum_{j=-1}^{\infty} c_j^{(n)} w_{n+j}\psi - 3\sum_{j=-1}^{\infty} c_j^{(0)} w_j\delta_n + \sum_{j=-1}^{\infty} c_j^{(n)}\delta_{n+j}\right)\,, \tag{A.78}$$

and

$$f_{n,\ell}' + \Big(2n+\ell+\frac{-5w_1+w_2}{1+w_1}\Big)\mathcal{H}f_{n,\ell} - (2n+\ell-1)\mathcal{H}f_{n+1,\ell} - \frac{\ell^2}{4\ell^2-1}kf_{n+1,\ell-1} + kf_{n,\ell+1} \\ - \delta_{\ell,1}\Big(\frac{(2n+3)w_n - (2n-1)w_{n+1}}{1+w_1}\Big)k\psi = a\Gamma\left(-\sum_{j=0}^{\infty}(c_j^{(0)} + \frac{1}{3}c_{j-1}^{(1)})\frac{w_j}{1+w_1}f_{n,\ell} + \sum_{j=-1}^{\infty} e_j^{(n,\ell)} f_{n+j,\ell}\right)\,. \tag{A.79}$$

Note that $\delta_{\ell,1}$ is a Kronecker delta. The coefficients are defined by

$$e_j^{(n,\ell)} = \frac{1}{3}c_j^{(n+\ell/2)} - \ell(\ell+1)\frac{(2j-3)!!}{(j+1)!2^{j+1}}\,, \tag{A.80}$$

together with

$$c_j^{(n)} = -2n(2n+1)\frac{(2j-1)!!}{2^{j+1}(j+1)!} + (8n^2-3)\frac{(2j-3)!!}{2^j j!} - 2n(2n-1)\frac{(2j-5)!!}{2^{j-1}(j-1)!}\,, \tag{A.81}$$

with the exception that for $j=-1$, only the first and for $j=0$ only the first two terms on the right should be taken into account, and that $(-5)!! = 1/3$ and $(-3)!! = -1$.

Whenever we truncate this hierarchy of equations at (n_{max}, ℓ_{max}), we should use an ansatz for higher-order moments. We follow the truncation scheme of [121]. As for the initial conditions, we use the same initial conditions as CDM, setting higher moments to zero at the start of numerical integration. Strictly speaking, one should find the superhorizon solution to the distribution function and use it it set the initial conditions as we did for background quantities; See Appendix A.3.1 for the case of δ_1. However, it turns out that even by setting the higher moments to zero, they quickly catch up with their attractor solution, see Figure 1.3b.

Finally, one might be worried that the diffusion source term on the right-hand side of (A.78) and (A.79) that are now appearing as a series of higher-order moments would instantly couple lower-order moments

to all of the moments, causing fast propagation of numerical errors. However, as long as the stochastic dark matter particles remain non-relativistic, these higher moments are suppressed, and even restricting oneself to the first two terms of each of these infinite sums results in the same numerical solutions.

A.3.1 Superhorizon solutions for stochastic dark matter

As a simple demonstration of the perturbation equations, let us solve for the perturbed pressure density for small Γ in the superhorizon limit $k \ll \mathcal{H}$. δ_1 is already of first order in a Γ/H_0 expansion. If we are interested in $c_s^2 = \delta_1/\delta$, we only need δ up to zeroth order in Γ, which is well-known:

$$\delta' - 3\phi' = 0 \quad \rightarrow \quad \delta = 3\phi + const\,. \tag{A.82}$$

Using adiabatic initial conditions of CDM for stochastic dark matter and ignoring the early-time contribution of neutrinos, the above equation gives the frozen superhorizon solution

$$\delta = -\frac{3}{2}\phi\,. \tag{A.83}$$

The evolution equation for δ_1, up to first order in Γ/H_0, is

$$\delta_1' + 2\mathcal{H}\delta_1 = \frac{2}{9}a\Gamma\delta\,. \tag{A.84}$$

Changing to the $\mathcal{T}$ time variable (1.48), this can be solved:

$$\delta_1 = \frac{2}{9}\frac{\Gamma\mathcal{T}}{a^2}\delta\,. \tag{A.85}$$

Therefore, comparing to (1.51), the superhorizon c_s^2 is

$$c_s^2 = \frac{2}{9}\frac{\Gamma\mathcal{T}}{a^2} = \frac{1}{3}w\,. \tag{A.86}$$

We can see this behaviour in Figure 1.3b briefly before the mode crosses the horizon. Finally, note that for stochastic dark matter, the effective sound speed c_s^2 does not have its usual interpretation of the speed of propagation of acoustic waves through dark matter; it is only a measure of typical velocities of dark matter particles at a given time as a result of diffusion.

A.4 Dark energy equations when coupled to stochastic dark matter

As discussed in section 1.2.3, at the background level we have $\bar{P}_x = -\bar{\rho}_x$. All other imperfect fluid quantities are 1st order in perturbations. Now, the balance equation (1.65) gives

$$\bar{\rho}_x' = -a\Gamma\bar{\rho}_{\rm dm}\,, \tag{A.87}$$

where a term proportional to Γw is omitted because w itself is linear in Γ, and terms proportional to Γ^2 can be ignored in the background since $\Gamma \ll H_0$. Assuming $\bar{P}_{\rm dm} \ll \bar{\rho}_{\rm dm}$, one can check that the

following is the solution to the above equation in the matter-dark energy era:

$$\bar{\rho}_x = \bar{\rho}_x^0 + \frac{2}{3}\Gamma(H - H_0)\,. \tag{A.88}$$

This explicitly confirms that when $\Gamma \ll H_0$, the deviation from a constant dark energy density $\bar{\rho}_x^0$ is negligible. Note that in this solution, we are implicitly assuming that $8\pi G/3$ is absorbed in the energy density.

For the perturbations, it is crucial to note that the relation $c_x^2 = \delta P_x/\delta\rho_x$ holds only in the rest frame of the dark energy. In general, one should account for the velocity of the dark energy fluid [140]

$$\delta P_x = c_x^2\delta\rho_x + (c_x^2 - c_a^2)V_x\bar{\rho}_x'\,, \tag{A.89}$$

where $c_a^2 = w_x = -1$ defines the adiabatic sound speed of dark energy, and V_x appears in the velocity vector of dark energy by $U^\mu = \frac{1}{a}(1-\psi, \partial_i V_x)$. Using our assumption $c_x^2 = 1$ and the tight coupling between dark matter and dark energy,

$$\delta P_x = \delta\rho_x + 2a\Gamma\bar{\rho}_{\rm dm}\frac{\theta}{k^2}\,. \tag{A.90}$$

Furthermore, we write the heat flux vector as $\mathcal{Q}_x^\mu = \frac{1}{a}(0, \bar{\rho}_x\beta^i)$, and we define $\beta = \partial_i\beta^i$. Note that this is a perturbation quantity. Now, the zeroth component of the balance equation (1.65) gives

$$\delta_x' + 6\mathcal{H}\delta_x + \beta = a\Gamma\frac{\bar{\rho}_{\rm dm}}{\bar{\rho}_x}\Big(\delta_x - \delta_{\rm dm} - \psi - 6\mathcal{H}\frac{\theta}{k^2}\Big)\,. \tag{A.91}$$

The spatial components of the balance equation determine the dynamics of the heat flux:

$$\beta' + 4\mathcal{H}\beta - k^2\delta_x = a\Gamma\frac{\bar{\rho}_{\rm dm}}{\bar{\rho}_x}\Big(\beta + \theta\Big)\,. \tag{A.92}$$

The heat flux of dark energy also enters into Einstein's equations. For example, the $0i$ components of the Einstein equations give

$$k^2(\phi' + \mathcal{H}\psi) = \frac{3}{2}\mathcal{H}^2\Big(\frac{\bar{\rho}_x}{\bar{\rho}_{tot}}\beta + \sum_i\frac{\bar{\rho}_i}{\bar{\rho_{tot}}}(1+w_i)\theta_i\Big)\,. \tag{A.93}$$

The summation is over all species. The relativistic Poisson equation also gets modified accordingly [147].

A.5 Derivation of semi-analytic solutions to stochastic dark matter perturbations in FRW

Starting from the covariant diffusion equation (1.19), the modified Boltzmann equation for δf is

$$\begin{aligned}\frac{\partial}{\partial\eta}\delta f - q\mathcal{H}\frac{\partial}{\partial q}\delta f + ik\frac{q}{E}\mu\,\delta f + \frac{\partial f_0}{\partial q}\Big(q\phi' - iEk\mu\psi\Big)\\ = \frac{a\Gamma m^3}{3}\Big[\psi\frac{1}{m^2q^2}\frac{\partial}{\partial q}\big(Eq^2\partial_q f_0\big) + \frac{\partial}{\partial q^j}\Big(\big(\frac{\delta^{ij}}{E} + \frac{1}{Em^2}q^iq^j\big)\frac{\partial}{\partial q^i}\delta f\Big)\Big]\,.\end{aligned} \tag{A.94}$$

This is the same equation as (A.77), but this time the hyperbolic Laplacian has been expanded in Cartesian momentum coordinates instead of spherical. The most important approximation is to assume that the stochastic dark matter particles remain non-relativistic up to the present time. We apply this by writing $E \simeq m$. By ignoring the second term in the inverse metric on the momentum space (which is suppressed by $(q/E)^2$), the second term on the right-hand side simply becomes proportional to the flat Laplacian in the momentum space. Also, our experience with numerical solutions of the generalised Boltzmann hierarchy shows that the first term on the right is negligible compared to the second. Moreover, we focus on the small-scale solutions where $k \gg \mathcal{H}$. Therefore, the term $q\phi'$ is suppressed by both q/E and $\mathcal{H}/k$ compared to $-iEk\mu\psi$, and we can safely ignore it. Now, using the comoving momentum $\tilde{q} = aq$ and comoving energy $\epsilon = \sqrt{\tilde{q}^2 + m^2a^2}$, the Hubble friction term gets absorbed into the time derivative, and we find

$$\frac{1}{a^3}\frac{\partial}{\partial\eta}\delta f + \frac{1}{a^4 m}i\vec{k}.\vec{\tilde{q}}\,\delta f - \frac{1}{a^2}\frac{\partial f_0}{\partial\tilde{q}}imk\mu\psi = \frac{\Gamma m^2}{3}\delta^{ij}\frac{\partial^2}{\partial\tilde{q}^i\partial\tilde{q}^j}\delta f\,. \tag{A.95}$$

The second term on the left-hand side can be absorbed through a redefinition of δf

$$F \coloneqq e^{-ik\mu y(\tilde{q},\eta)}\delta f\,. \tag{A.96}$$

Here y is the free streaming length: the comoving length travelled by a particle with comoving momentum $\tilde{q}$ from η to some reference conformal time η_f. It is only defined through its derivative:

$$\frac{dy}{d\eta}(\tilde{q},\eta) = -\frac{\tilde{q}}{m}\frac{1}{a(\eta)}\,. \tag{A.97}$$

More on the choice for η_f will be said later. If we also change the time variable to $d\mathcal{T} = a^3 d\eta$, the equation for F becomes

$$\frac{\partial F}{\partial\mathcal{T}} - e^{-ik\mu y}\frac{\partial f_0}{\partial\tilde{q}}ik\mu\frac{m}{a^2}\phi = \frac{\Gamma m^2}{3}\Big(\delta^{ij}\frac{\partial^2}{\partial\tilde{q}^i\partial\tilde{q}^j}F + 2i\frac{y}{\tilde{q}}k^i\frac{\partial}{\partial\tilde{q}^i}F - \big(\frac{y}{\tilde{q}}\big)^2k^2F\Big)\,. \tag{A.98}$$

Now, to eliminate the term $\partial_{\tilde{q}^i}F$ on the right, we introduce yet another momentum variable

$$\vec{Q} \coloneqq \vec{\tilde{q}} - \frac{2}{3}\Gamma m s(\mathcal{T})i\vec{k}\,, \tag{A.99}$$

such that s is a time function that is defined by

$$\frac{ds}{d\mathcal{T}} = -\mathcal{Y}(\mathcal{T})\,, \quad \mathcal{Y}(\mathcal{T}) \coloneqq \frac{m}{\tilde{q}}y\,. \tag{A.100}$$

Writing F as a function of Q rather than $\tilde{q}$ simplifies its equation to

$$\frac{\partial F}{\partial\mathcal{T}} - e^{-ik\mu y}\frac{\partial f_0}{\partial\tilde{q}}ik\mu\frac{m}{a^2}\phi = \frac{\Gamma m^2}{3}\Big(\delta^{ij}\frac{\partial^2}{\partial Q^i\partial Q^j}F - \frac{k^2}{m^2}\mathcal{Y}^2F\Big)\,. \tag{A.101}$$

The last term on the right can be absorbed into the time derivative with yet another redefinition

$$\hat{F}(Q,\mathcal{T}) \coloneqq e^{-\frac{1}{3}\Gamma k^2\mathcal{G}(\mathcal{T})}F(Q,\mathcal{T})\,, \tag{A.102}$$

where $\mathcal{G}$ is also a time function only defined by its derivative

$$\frac{d\mathcal{G}}{d\mathcal{T}}(\mathcal{T}) = -\mathcal{Y}(\mathcal{T})^2\,. \tag{A.103}$$

We also define $\mathcal{G}(\mathcal{T},\mathcal{T}') \equiv -\mathcal{G}(\mathcal{T}) + \mathcal{G}(\mathcal{T}')$. For $\hat{F}$ we have

$$\frac{\partial \hat{F}}{\partial \mathcal{T}} + \mathcal{J}(\mathcal{T},Q) = \frac{\Gamma m^2}{3}\delta^{ij}\frac{\partial^2}{\partial Q^i \partial Q^j}\hat{F}\,, \tag{A.104}$$

This is a simple heat equation with an additional non-homogeneous term:

$$\mathcal{J} = -e^{-ik\mu y}\frac{\partial f_0}{\partial \tilde{q}} ik\mu \frac{m}{a^2}\phi\,. \tag{A.105}$$

It can be solved using Duhamel's principle [129] by first solving for the homogeneous equation with an arbitrary initial condition:

$$\begin{aligned} \frac{\partial F_{\rm hom}}{\partial \mathcal{T}} &= \frac{\Gamma m^2}{3}\delta^{ij}\frac{\partial^2}{\partial Q^i \partial Q^j}F_{\rm hom}\,, \\ F_{\rm hom}\big|_{\mathcal{T}_1} &= \zeta(Q)\,. \end{aligned} \tag{A.106}$$

Let us call this solution $F_{\rm hom}\left(\mathcal{T},Q;\zeta,\mathcal{T}_1\right)$. Using the flat heat kernel, this can be solved as

$$F_{\rm hom}\left(\mathcal{T},Q;\zeta,\mathcal{T}_1\right) = \int d^3Q'\ \frac{e^{-\frac{|Q-Q'|^2}{\frac{4}{3}\Gamma m^2(\mathcal{T}-\mathcal{T}_1)}}}{\left(\frac{4}{3}\pi\Gamma m^2(\mathcal{T}-\mathcal{T}_1)\right)^{3/2}}\zeta(Q')\,. \tag{A.107}$$

Duhamel's principle says that the solution to (A.104) is the following:

$$\hat{F}(\mathcal{T},Q) = F_{\rm hom}\left(\mathcal{T},Q;\hat{F}\big|_{\mathcal{T}_i},\mathcal{T}_i\right) + \int_{\mathcal{T}_i}^{\mathcal{T}} d\mathcal{T}'\ F_{\rm hom}\left(\mathcal{T},Q;\mathcal{J}\big|_{\mathcal{T}'},\mathcal{T}'\right). \tag{A.108}$$

Using the heat kernel solution to $F_{\rm hom}$, one can recover the solution (1.56) to δf by tracing back all the change of variables:

$$\begin{aligned} \delta f(\mathcal{T},k,Q) = &\int d^3Q' \left(e^{-\frac{1}{3}\Gamma k^2 \mathcal{G}(\mathcal{T},\mathcal{T}_i)} \frac{e^{-\frac{|Q-Q'|^2}{\frac{4}{3}\Gamma m^2(\mathcal{T}-\mathcal{T}_i)}}}{\left(\frac{4}{3}\pi\Gamma m^2(\mathcal{T}-\mathcal{T}_i)\right)^{3/2}} e^{-i\vec{k}.\vec{y}(Q',\mathcal{T}_i)+i\vec{k}.\vec{y}(Q,\mathcal{T})}\,\delta f(\mathcal{T}_i,k,Q') \right) \\ &+ \int_{\mathcal{T}_i}^{\mathcal{T}} d\mathcal{T}' \int d^3Q' \left(e^{-\frac{1}{3}\Gamma k^2 \mathcal{G}(\mathcal{T},\mathcal{T}')} \frac{e^{-\frac{|Q-Q'|^2}{\frac{4}{3}\Gamma m^2(\mathcal{T}-\mathcal{T}')}}}{\left(\frac{4}{3}\pi\Gamma m^2(\mathcal{T}-\mathcal{T}')\right)^{3/2}} e^{-i\vec{k}.\vec{y}(Q',\mathcal{T}')+i\vec{k}.\vec{y}(Q,\mathcal{T})} \frac{m}{a(\mathcal{T}')^2} i\vec{k}.\vec{q}' \frac{\partial f_0}{\partial \tilde{q}}\phi_k(\mathcal{T}') \right). \end{aligned} \tag{A.109}$$

This might seem like an intractable integral solution. However, the Q' integrals are Gaussian and can be performed exactly; remember that in the non-relativistic approximation, f_0 too is Gaussian (1.49).

Note that in the limit of negligible diffusion, $\Gamma \to 0$, the heat kernel becomes a delta function, identifying $Q' = Q$. Hence, one would recover the integral solution to the CDM distribution function (see [125]). For the rest of this appendix, we are interested in the opposite limit: The strong diffusion regime. In

this regime, $\Gamma k^2 \mathcal{G}(\mathcal{T}) \gg 1$. As a result, the first term in (A.109) would be exponentially suppressed. This is just the mathematical formulation that when the diffusion is strong, the initial conditions become unimportant and the dynamics of the stochastic dark matter particles becomes irreversible. Besides, the integrand of the remaining time integral in (A.109) becomes highly peaked around $\mathcal{T}' \sim \mathcal{T}$, so that one could use a combination of saddle point approximation and numerical techniques to tackle the integral.

Before performing the integrals, we determine the behaviour of the time functions $\mathcal{T}$, $\mathcal{Y}$, s, and $\mathcal{G}$ during matter domination. Integrating $d\mathcal{T} = a^3 d\eta$, one finds

$$\mathcal{T} = \frac{1}{7} a^3 \eta \,. \tag{A.110}$$

For simplicity, we set $\eta_f \to \infty$ to get

$$y(\tilde{q}, \eta) = \frac{\tilde{q}}{m} \frac{\eta}{a} \,. \tag{A.111}$$

Therefore, $\mathcal{Y}(\eta) = \eta/a$. Finally, integrating (A.100) and (A.103), one finds

$$s(\eta) = -\frac{1}{6} a^2 \eta^2, \quad \mathcal{G}(\eta) = -\frac{1}{5} a \eta^3 \,. \tag{A.112}$$

Now everything is ready to perform the integrations. First, we look at the energy density

$$\delta\rho(\eta, k) = \frac{m}{a^3} \int d^3\tilde{q} \; \delta f(\eta, k, \tilde{q}) \,. \tag{A.113}$$

After performing the Gaussian Q' integrals implicit in δf, the $\tilde{q}$ integral also turns out to be Gaussian and can be done exactly. Some algebra reveals that

$$\delta(\eta, k) = -k^2 \int_{\eta_i}^{\eta} d\eta' \; \eta' \big(1 - \frac{\eta'}{\eta}\big) \phi(\eta', k) e^{-\frac{1}{315} a\Gamma k^2 \eta^3 \left(1 - 6(\frac{\eta'}{\eta})^5 + 5(\frac{\eta'}{\eta})^6\right)} \,. \tag{A.114}$$

This is where the strong diffusion regime helps us to approximate the integral. It is defined by

$$\frac{1}{315} a\Gamma k^2 \eta^3 \gg 1 \,. \tag{A.115}$$

In this regime, almost all of the contribution to the integral comes from the region where $\eta' \sim \eta$. Besides, one can numerically check that $\phi \sim 1/\eta$ is a good approximation for the gravitational potential in the strong diffusion regime. Note that this is unlike the usual case where $\phi \sim const$ approximates the behaviour of the gravitational potential during matter domination. Using this approximate solution, (A.114) can be approximated by

$$\delta(\eta, k) = -(k\eta)^2 \phi(\eta, k) \int_0^1 dt \; (1 - t) e^{-\frac{1}{315} a\Gamma k^2 \eta^3 \left(1 - 6t^5 + 5t^6\right)} \,. \tag{A.116}$$

Now one can either do a saddle point approximation $1 - 6t^5 + 5t^6 \simeq 15(1-t)^2$ around $t = 1$ and perform a Gaussian integral or instead perform the integral in (A.116) numerically for different values

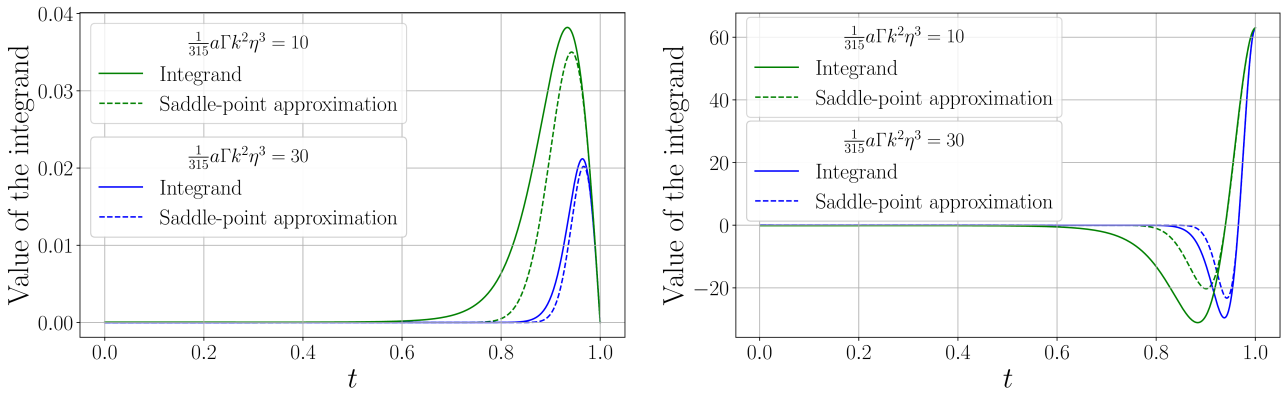


Figure A.1: Comparison of the integrands in (A.116) and (A.121) with their saddle-point approximations for two different values of $a\Gamma k^2\eta^3$. Note that $t = 1$ is not necessarily the saddle point of the whole integrand but just the exponential factor. So what we mean by the saddle-point approximation of the integral is in fact just using the 2nd order Taylor expansion of the exponent $1 - 6t^5 + 5t^6 \simeq 15(1-t)^2$.

of $a\Gamma k^2\eta^3$ and find the power law dependence for large values of $a\Gamma k^2\eta^3$. The result of the latter is

$$\delta(\eta, k) \approx \frac{-11.5}{a\Gamma\eta}\phi(\eta, k)\,. \tag{A.117}$$

Notice that this is more accurate than (1.62) which uses the saddle-point approximation. The difference is visible in Figure A.1a. By increasing $a\Gamma k^2\eta^3$, one can see that the integrand gets increasingly concentrated at its endpoint, and furthermore, the saddle-point approximation becomes better.

The same numerical approach can be applied to the perturbed pressure

$$\delta P = \frac{1}{3m}\frac{1}{a^5}\int d^3\tilde{q}\ \tilde{q}^2\,\delta f \tag{A.118}$$

to find

$$\frac{\delta P}{\bar{\rho}} \simeq -\Big(1.06 + \mathcal{O}\big(\frac{1}{\sqrt{a\Gamma k^2\eta^3}}\big)\Big)\phi\,. \tag{A.119}$$

Perhaps the most illuminating case is that of the perturbed velocity. Starting from

$$(\bar{\rho} + \bar{P})\theta = \frac{1}{a^4}\int d^3\tilde{q}\ i(\vec{k}.\tilde{q})\,\delta f\,, \tag{A.120}$$

and performing the Gaussian Q' and $\tilde{q}$ integrals, one finds

$$\big(1 + w(\eta)\big)\frac{\theta(\eta, k)}{k} = 0.016\,k\eta\,\phi(\eta, k)\int_0^1 dt\ t\Big(63 - a\Gamma k^2\eta^3\big(\frac{1}{t} - 1 - t^5 + t^6\big)\Big)e^{-\frac{1}{315}a\Gamma k^2\eta^3\left(1-6t^5+5t^6\right)}\,. \tag{A.121}$$

The interesting thing about this integrand is that due to the polynomial prefactor, it changes sign once near $t = 1$. The result of the integral lies in the fine difference between the integral of the positive and negative parts of the integrand. Therefore, if one adopts the saddle point approximation, this fine difference gets lost, and one would even get the sign of the integral wrong. This can be clearly seen in Figure A.1b. The true integrand has a larger tail in its negative region, while the saddle-point approximation cannot capture this behaviour. As a result, in this case, one has no choice but to use

the numerical approach. The result is

$$(1+w)\frac{\theta}{k} \approx -48.6\frac{1}{a\Gamma k\eta^2}\phi\,. \tag{A.122}$$

A.6 Full posteriors for stochastic dark matter

We present the 68% and 95% confidence levels of marginalised posteriors for the 7 main parameters used in our MCMC analysis, alongside Ω_m and S_8 for illustrative purposes. The plots include both stochastic dark matter and ΛCDM. Figure A.2 is based on Planck2018 data only. Figure A.3 compares the Planck2018 posteriors to the Low-z Baseline. Figure A.4 contrasts the Baseline with Baseline+DES, thereby highlighting the effect of the S_8 prior from DES.

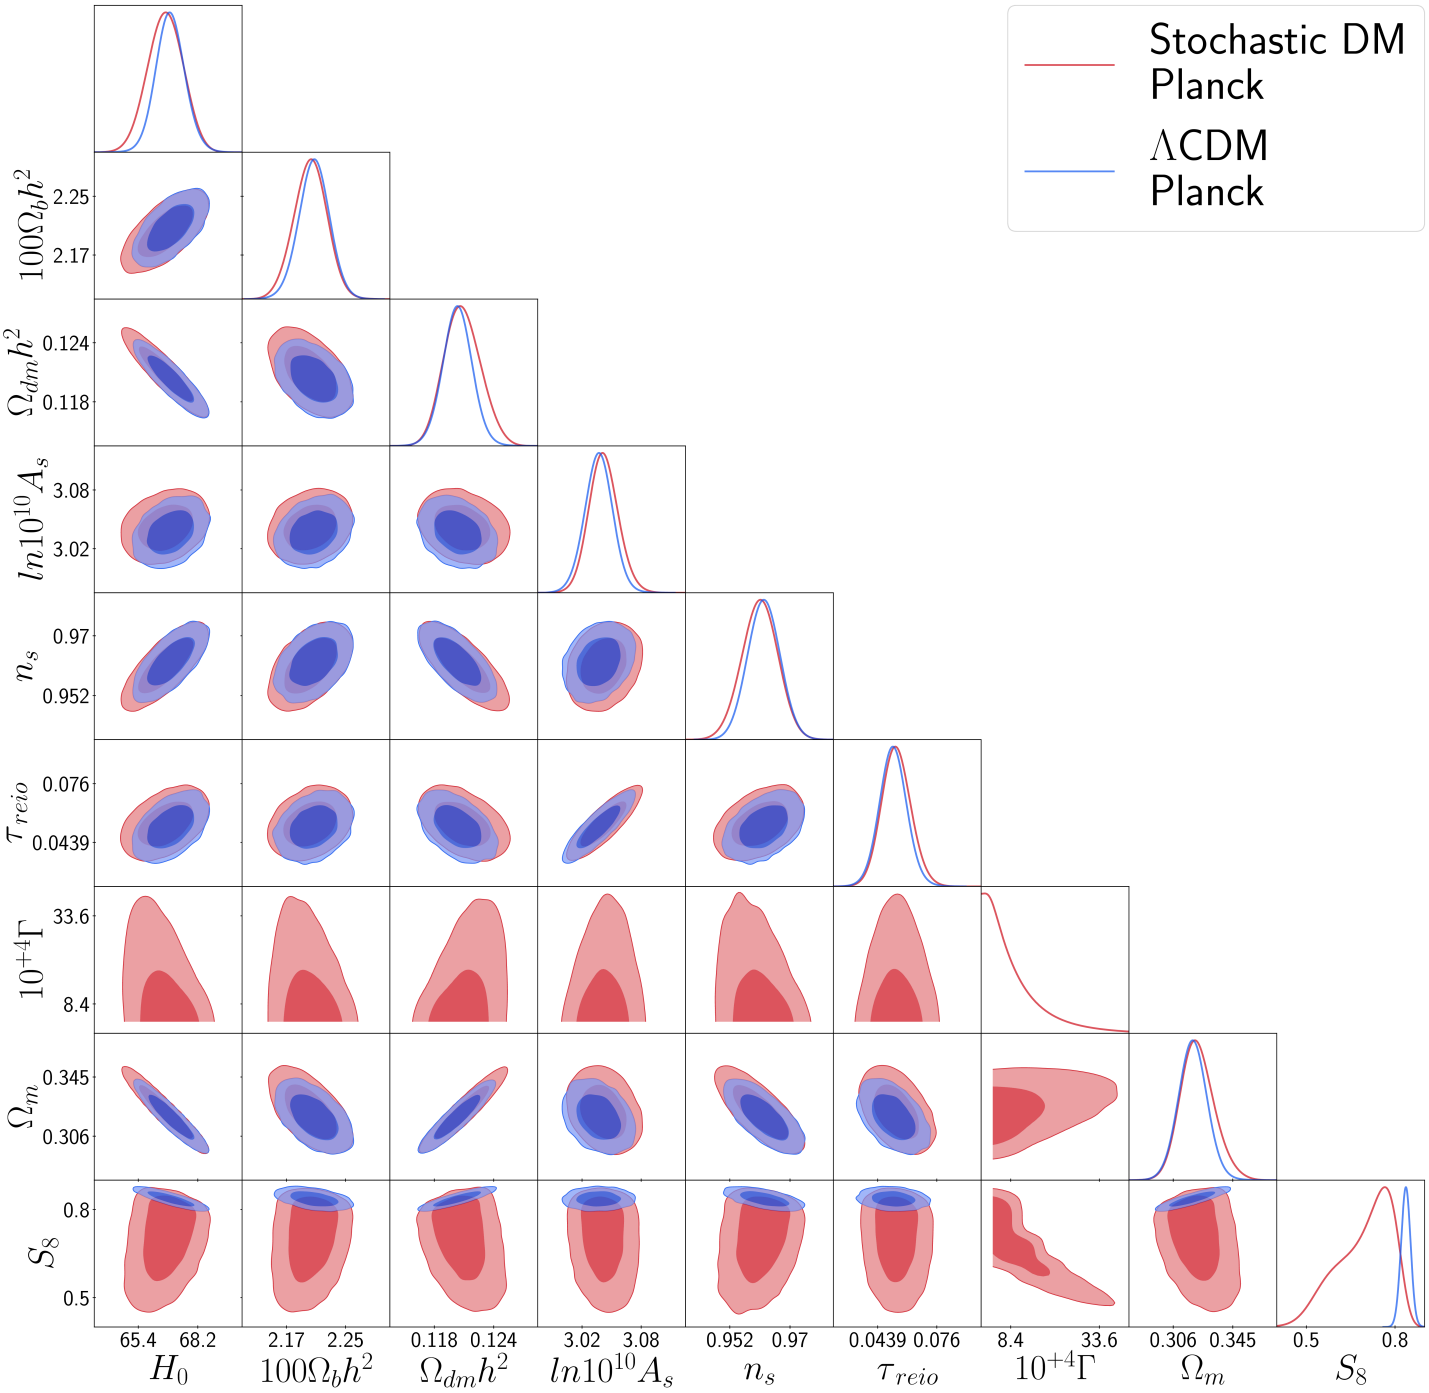


Figure A.2: 68% and 95% confidence-level posterior distributions of all cosmological parameters for stochastic dark matter and ΛCDM from Planck. Γ specifically target structures at small scales, giving a very wide S_8 posterior while leaving other cosmological parameters with roughly the same constraints.

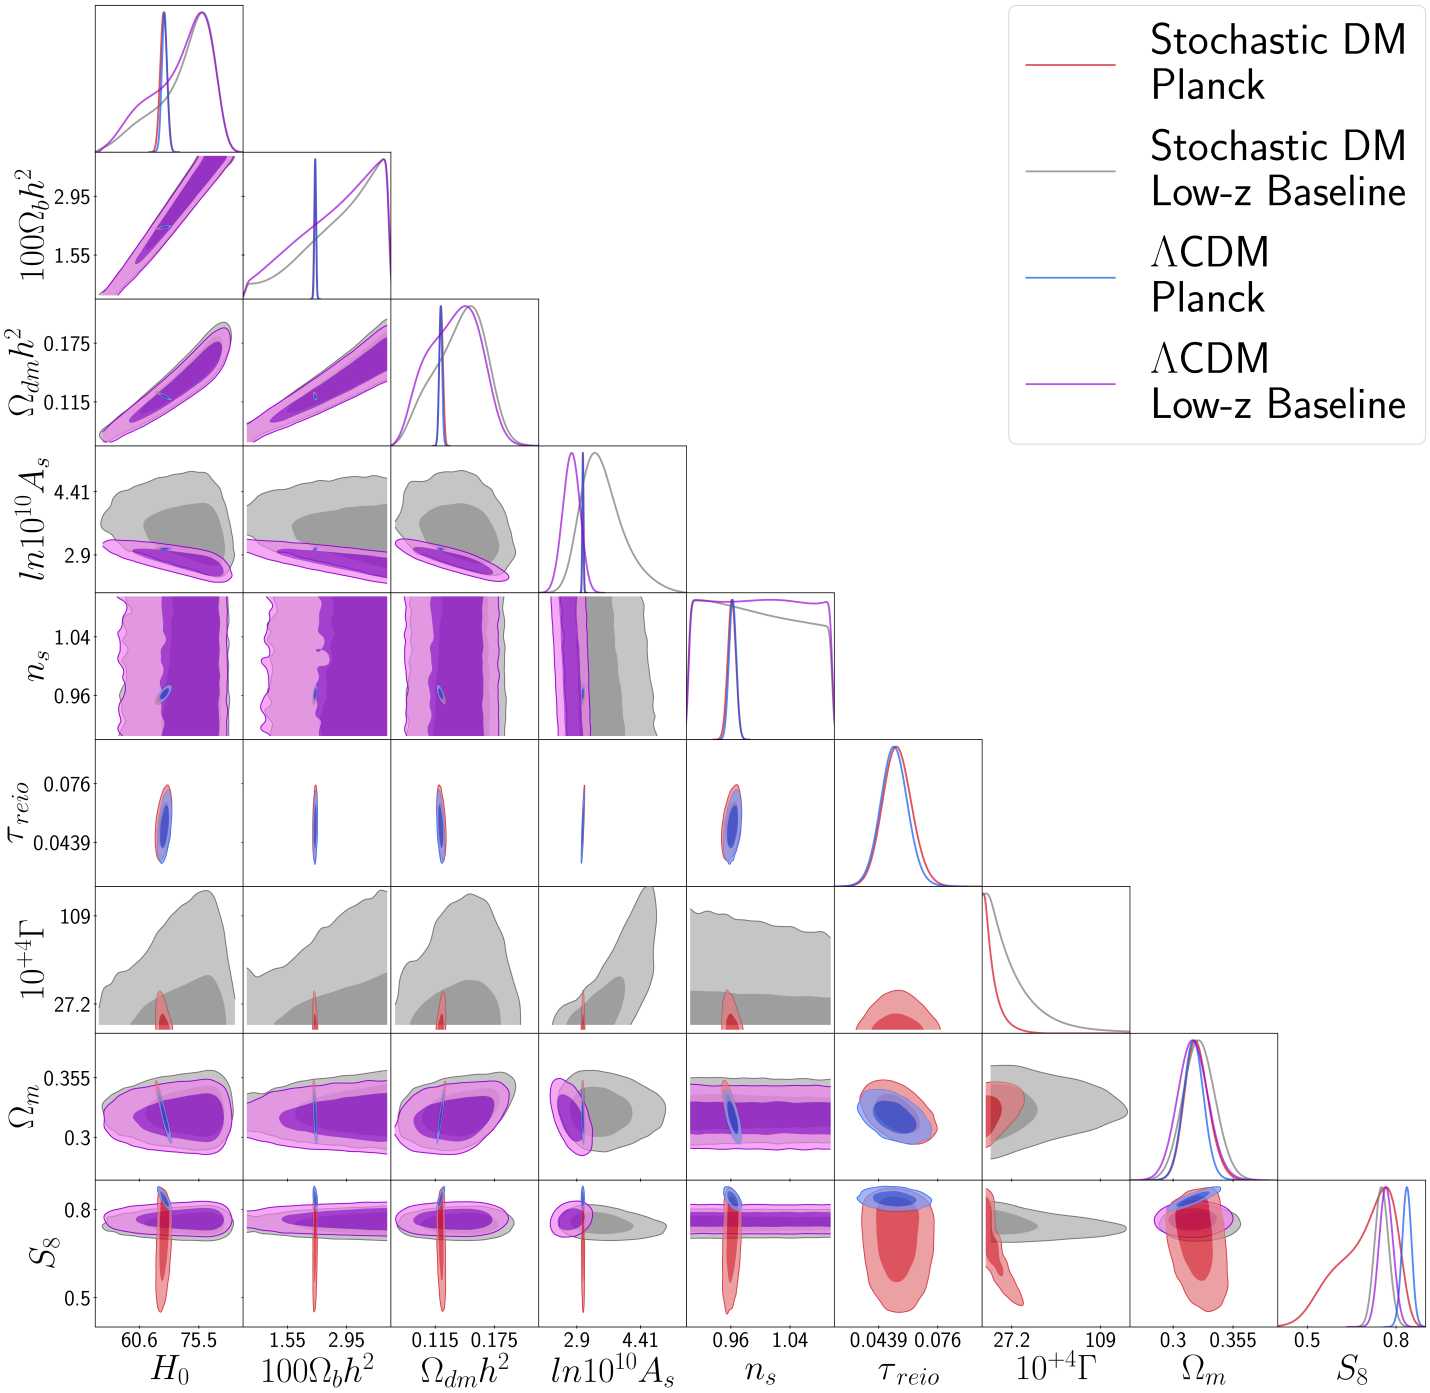


Figure A.3: 68% and 95% confidence-level posterior distributions of all cosmological parameters for stochastic dark matter and ΛCDM, from Planck and low-z Baseline. The constraints on the amplitude of primordial fluctuations, A_s, are very loose from the low-z Baseline. As a result, for stochastic dark matter, Γ and A_s are positively correlated: larger Γ strongly dampens the perturbations in dark matter, which gets compensated by a larger primordial amplitude A_s to yield the same S_8. So S_8 has no correlation with A_s or Γ. In contrast, Planck data puts a strong constraint on A_s and therefore yields a stronger constraint on Γ as well. Since these values of Γ cannot get compensated by A_s, they still significantly decrease S_8, leading to a very broad posterior for S_8.

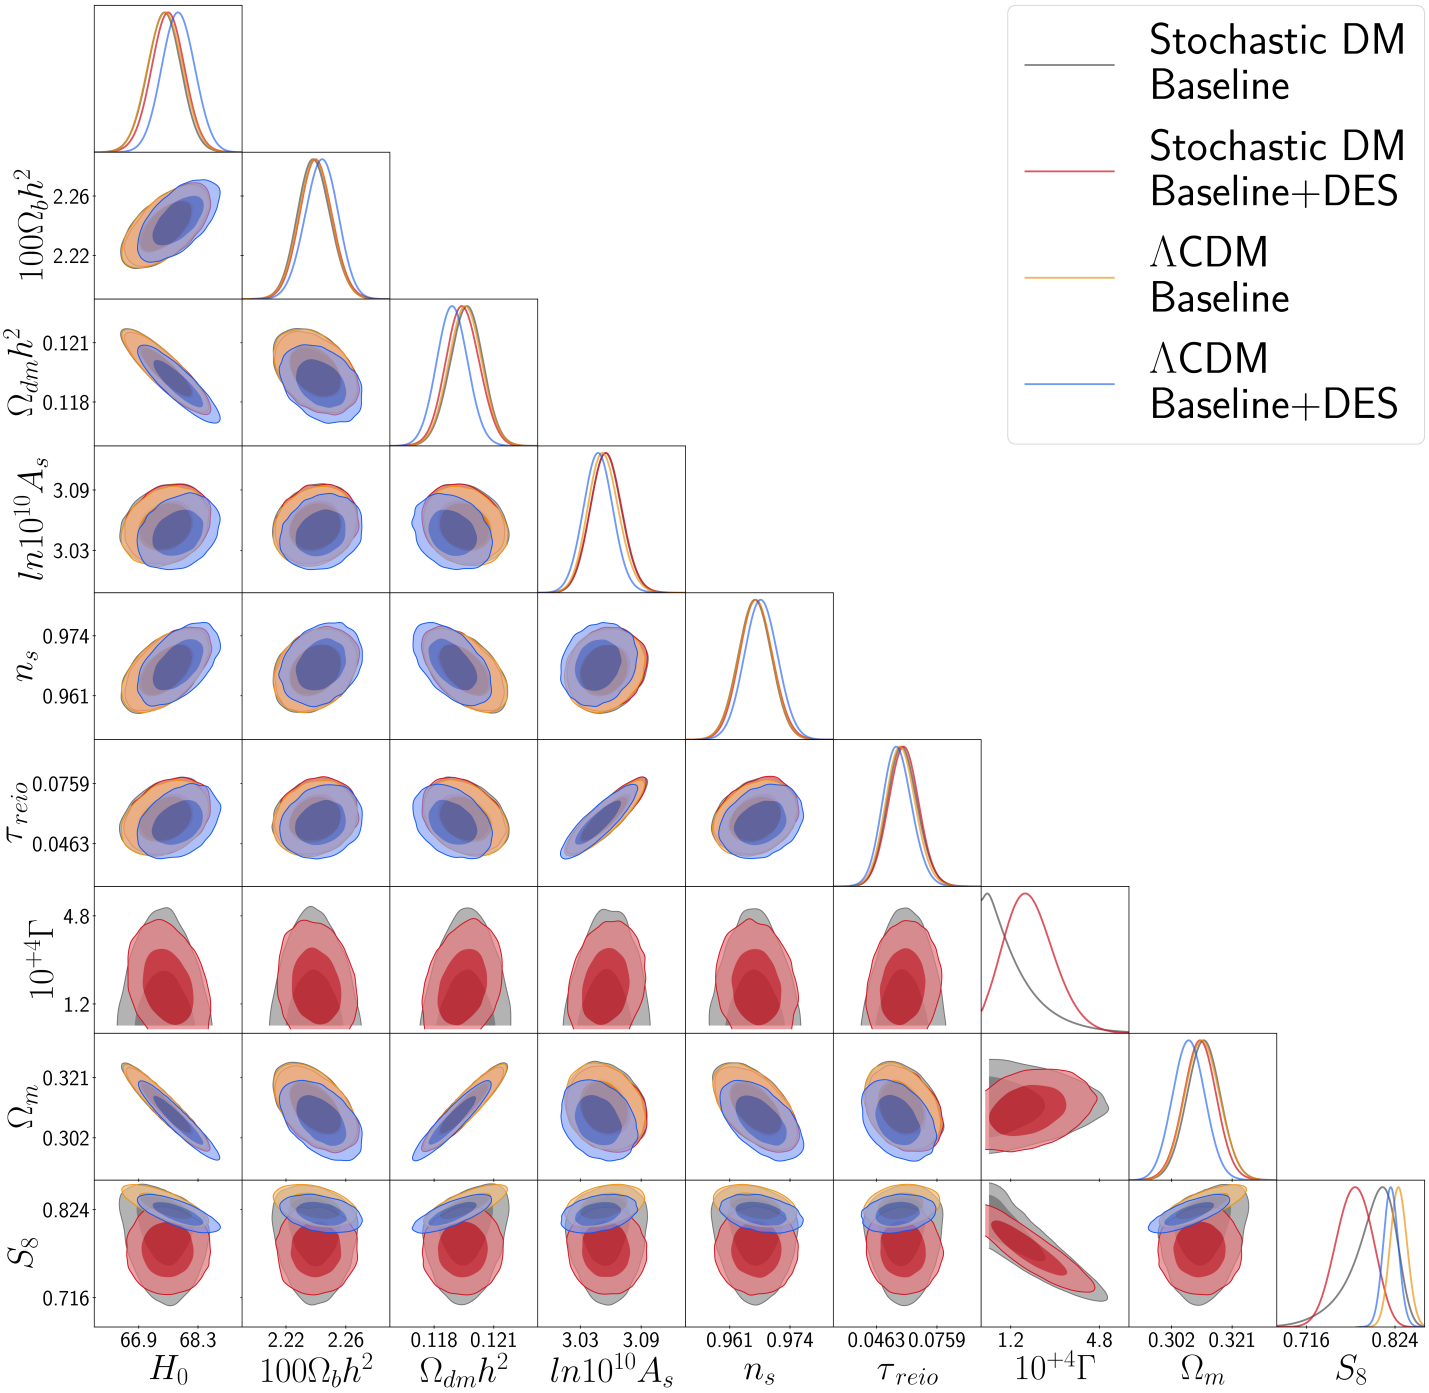


Figure A.4: 68% and 95% confidence-level posterior distributions of all cosmological parameters for stochastic dark matter and ΛCDM, from Baseline and Baseline+DES. In the case of CDM, decreasing S_8 to satisfy the DES prior is obtained by very slight shifts in Ω_m and A_s to smaller values; CDM provides no other mechanism for less structure formation. The stochastic dark matter naturally does this by a non-zero Γ.

Appendix B

Technical supplement to Chapter 2

B.1 Analytic solution to background massless diffusion equation in FRW

Starting from equation (2.24), we define the comoving energy as $\epsilon = aE$. It combines the two terms on the left into one single derivative. Defining a new time parameter $d\mathsf{T} = a^2 d\eta$ further simplifies the left-hand side so that we get

$$\frac{\partial f}{\partial \mathsf{T}} = \left(3\kappa_1 - \kappa_2\right)\frac{1}{\epsilon^2}\frac{\partial}{\partial \epsilon}\left(\epsilon^2 f\right) + \kappa_1 \frac{1}{\epsilon^2}\frac{\partial}{\partial \epsilon}\left(\epsilon^3 \frac{\partial}{\partial \epsilon} f\right). \tag{B.1}$$

Now we define a new function

$$g(\epsilon, \mathsf{T}) = \epsilon^{3-\frac{\kappa_2}{2\kappa_1}} f(\epsilon, \mathsf{T}). \tag{B.2}$$

This removes the first-order derivatives in ϵ:

$$\frac{\partial g}{\partial \mathsf{T}} = \kappa_1 \epsilon \partial_\epsilon^2 g + \kappa_3 g/\epsilon, \tag{B.3}$$

where

$$\kappa_3 = \frac{\kappa_2}{4\kappa_1}(2\kappa_1 - \kappa_2). \tag{B.4}$$

One can see that the above PDE for g is scale-invariant with respect to the following transformations

$$g \to \xi^c g, \quad \mathsf{T} \to \xi^a \mathsf{T}, \quad \epsilon \to \xi^a \epsilon. \tag{B.5}$$

A similarity solution is one that is invariant with respect to this re-scaling (note that the PDE being symmetric does not imply that all of its solutions must be symmetric as well). Therefore, the class of symmetric solutions is as follows

$$g(\epsilon, \mathsf{T}) = \mathsf{T}^d \mathcal{G}(\epsilon/\mathsf{T}), \tag{B.6}$$

where $d = c/a$ is an arbitrary real number. $\mathcal{G}$ satisfies an ODE with variable $y = \epsilon/\mathsf{T}$:

$$\kappa_1 \mathcal{G}'' + \mathcal{G}' + (\frac{\kappa_3}{y^2} - \frac{d}{y})\mathcal{G}(y) = 0. \tag{B.7}$$

Once again, we can remove the first derivative by a redefinition

$$G(y) = e^{y/(2\kappa_1)}\mathcal{G}\,, \tag{B.8}$$

to get

$$\kappa_1 G'' + (\frac{\kappa_3}{y^2} - \frac{d}{y} - \frac{1}{4\kappa_1})G = 0\,. \tag{B.9}$$

This equation can be mapped to a radial Schrodinger equation with Coulomb potential if we make the following identification

$$\ell(\ell+1) = -\frac{\kappa_3}{\kappa_1}\,, \tag{B.10}$$

$$2me^2/\hbar^2 = -d/\kappa_1\,, \tag{B.11}$$

$$2m\mathcal{E}/\hbar^2 = -\frac{1}{4\kappa_1^2}\,. \tag{B.12}$$

The last equation informs that the energy level of the equivalent hydrogen atom has to be negative, $\mathcal{E} < 0$, so we should indeed use the bound states. The first equation implies that ℓ needs not to be an integer in this matching. The bound-state solution of the equivalent problem is [302]

$$G(y) = \Big(\frac{y}{2\kappa_1}\Big)^{\ell+1} e^{-y/(2\kappa_1)} F(\ell+1+d; 2(\ell+1); y/\kappa_1)\,, \tag{B.13}$$

where F is the confluent hypergeometric function. We have to select an orthonormal subset of these functions by suitably choosing d. One way is to enforce

$$d = -(\ell+1) - n\,, \tag{B.14}$$

where n is a natural number. Then F will reduce to a generalised Laguerre polynomial of degree n.

$$F\Big(-n; 1+(2\ell+1); \frac{y}{\kappa_1}\Big) = (-1)^n n! L_n^{(2\ell+1)}\Big(\frac{y}{\kappa_1}\Big)\,. \tag{B.15}$$

The orthogonality relation of these polynomials is known. Therefore, we have found a complete basis of solutions to our diffusion equation, which we can use to expand any solution (given that it has similar limiting behaviour as $\epsilon \to 0$):

$$f(\mathsf{T},\epsilon) = e^{-\frac{\epsilon}{\kappa_1 \mathsf{T}}} \frac{\epsilon^{\frac{\kappa_2}{2\kappa_1}-3}}{\mathsf{T}^{\ell+1}} \Big(\frac{\epsilon}{2\kappa_1\mathsf{T}}\Big)^{\ell+1} \sum_{n=0}^{\infty} (-1)^n n! c_n \frac{1}{\mathsf{T}^n} L_n^{(2\ell+1)}\Big(\frac{\epsilon}{\kappa_1\mathsf{T}}\Big)\,. \tag{B.16}$$

The coefficients c_n depend on the initial condition at $\mathsf{T} = \mathsf{T}_i$ and can be found using the orthogonality of Laguerre polynomials. The result is

$$c_n = \Big(\frac{2}{\kappa_1}\Big)^{\ell+1} (-1)^n \frac{\mathsf{T}_i^n}{\Gamma(n+2\ell+2)} \int_0^\infty d\epsilon\; f(\mathsf{T}_i,\epsilon)\epsilon^{\ell+3-\frac{\kappa_2}{2\kappa_1}} L_n^{(2\ell+1)}\Big(\frac{\epsilon}{\kappa_1\mathsf{T}_i}\Big)\,. \tag{B.17}$$

This series expansion of the solution is of little use for initial conditions that are not in the form of a simple exponential function times a rational function of ϵ. In particular, it poorly approximates Gaussian functions in ϵ, requiring many terms and highly accurate integration to determine the coefficients c_n.

Therefore, the above solution is of little practical use. However, it can be used to construct heat kernels for the diffusion equation. These heat kernels, in turn, can easily give us the general solution as an integral over energy rather than a series expansion.

We define the heat kernel of the massless diffusion equation to be $K(\mathsf{T}, \epsilon; \epsilon_i)$ such that it satisfies

$$\frac{\partial}{\partial \mathsf{T}} K = (3\kappa_1 - \kappa_2)\frac{1}{\epsilon^2}\frac{\partial}{\partial \epsilon}(\epsilon^2 K) + \kappa_1 \frac{1}{\epsilon^2}\frac{\partial}{\partial \epsilon}\Big(\epsilon^3 \frac{\partial}{\partial \epsilon} K\Big)\,, \tag{B.18}$$

with the initial condition

$$K(\mathsf{T}_i, \epsilon; \epsilon_i) = \delta(\epsilon - \epsilon_i)\,. \tag{B.19}$$

Substituting this delta function instead of $f(\mathsf{T}_i, \epsilon)$ in (B.17) gives the expansion coefficients $c_n[K]$ of K

$$c_n[K] = \Big(\frac{2}{\kappa_1}\Big)^{\ell+1} \frac{(-1)^n}{\Gamma(n+2\ell+2)} \mathsf{T}_i^n \epsilon_i^{\ell+3-\frac{\kappa_2}{2\kappa_1}} L_n^{(2\ell+1)}\Big(\frac{\epsilon_i}{\kappa_1 \mathsf{T}_i}\Big)\,. \tag{B.20}$$

Plugging this expression for the coefficients back into the expansion (B.16) for K gives

$$K(\mathsf{T}, \epsilon; \epsilon_i) = e^{-\frac{\epsilon}{\kappa_1 \mathsf{T}}} \frac{1}{\epsilon_i}\Big(\frac{\epsilon_i}{\epsilon}\Big)^{3-\frac{\kappa_2}{2\kappa_1}} \Big(\frac{\epsilon \epsilon_i}{\kappa_1^2 \mathsf{T}^2}\Big)^{\ell+1} \sum_{n=0}^{\infty} \frac{n!}{\Gamma(n+2\ell+2)} \Big(\frac{\mathsf{T}_i}{\mathsf{T}}\Big)^n L_n^{(2\ell+1)}\Big(\frac{\epsilon}{\kappa_1 \mathsf{T}}\Big) L_n^{(2\ell+1)}\Big(\frac{\epsilon_i}{\kappa_1 \mathsf{T}_i}\Big)\,. \tag{B.21}$$

This cumbersome expression would not be useful unless we have an analytic expression for the series. Surprisingly enough, this exists and is given by the Hardy-Hille formula:

$$\sum_{n=0}^{\infty} \frac{n!}{\Gamma(n+1+\alpha)} t^n L_n^{(\alpha)}(x) L_n^{(\alpha)}(y) = \frac{1}{(xyt)^{\alpha/2}(1-t)} e^{-\frac{(x+y)t}{1-t}} I_\alpha\Big(\frac{2\sqrt{xyt}}{1-t}\Big)\,, \tag{B.22}$$

where I_α is the modified Bessel function. Therefore, we now have a closed-form expression for the heat kernel K. However, note that since equation (B.10) has two solutions

$$2\ell_\pm + 1 = \pm(1 - \frac{\kappa_2}{\kappa_1})\,, \tag{B.23}$$

we have actually found two corresponding heat kernels $K_\pm$ given in equation (2.25). Now we demand the total particle number to be finite and conserved. In terms of the heat kernels, we demand $\int_0^\infty d\epsilon\ \epsilon^2 K_\pm$ to be equal to ϵ_i^2 at all times. To see whether either of $K_\pm$ satisfies this condition, we should perform a complicated weighted integral over the modified Bessel function. A simpler way is to use the PDE governing $K_\pm$. Multiplying equation (B.18) by ϵ^2 and integrating gives total derivative integrals on the right. So we only get contributions from $\epsilon = 0, \infty$. One can see that at $\epsilon \to \infty$, both $K_\pm$ go to zero. Hence

$$\partial_{\mathsf{T}} \int_0^\infty d\epsilon\ \epsilon^2 K_\pm = -(3\kappa_1 - \kappa_2)\epsilon^2 K_\pm\Big|_{\epsilon\to 0} - \kappa_1 \epsilon^3 \partial_\epsilon K_\pm\Big|_{\epsilon\to 0}\,. \tag{B.24}$$

Therefore, we would have to know the limiting behaviour of $K_\pm$ as $\epsilon \to 0$. Using the limit of the

modified Bessel function, one can show that for $\epsilon \ll \kappa_1 \Delta\mathsf{T}$,

$$K_+ \sim \begin{cases} \dfrac{1}{\Gamma(2-\kappa_2/\kappa_1)} \left(\dfrac{\epsilon_i}{\kappa_1 \Delta\mathsf{T}}\right)^{2-\kappa_2/\kappa_1} e^{-\frac{\epsilon+\epsilon_i}{\kappa_1\Delta\mathsf{T}}} \epsilon_i/\epsilon^2\,, & \text{if } \kappa_2/\kappa_1 \notin \{1,2,3,...\} \\ \dfrac{1}{\Gamma(\kappa_2/\kappa_1)} \left(\dfrac{1}{\kappa_1 \Delta\mathsf{T}}\right)^{\kappa_2/\kappa_1} e^{-\frac{\epsilon+\epsilon_i}{\kappa_1\Delta\mathsf{T}}} \epsilon_i^2 \epsilon^{\kappa_2/\kappa_1-3}\,, & \text{if } \kappa_2/\kappa_1 \in \{1,2,3,...\} \end{cases},$$

$$K_- \sim \begin{cases} \dfrac{1}{\Gamma(2-\kappa_2/\kappa_1)} \left(\dfrac{\epsilon_i}{\kappa_1 \Delta\mathsf{T}}\right)^{2-\kappa_2/\kappa_1} e^{-\frac{\epsilon+\epsilon_i}{\kappa_1\Delta\mathsf{T}}} \epsilon_i/\epsilon^2\,, & \text{if } \kappa_2/\kappa_1 \in \{0,-1,-2,...\} \\ \dfrac{1}{\Gamma(\kappa_2/\kappa_1)} \left(\dfrac{1}{\kappa_1 \Delta\mathsf{T}}\right)^{\kappa_2/\kappa_1} e^{-\frac{\epsilon+\epsilon_i}{\kappa_1\Delta\mathsf{T}}} \epsilon_i^2 \epsilon^{\kappa_2/\kappa_1-3}\,, & \text{if } \kappa_2/\kappa_1 \notin \{0,-1,-2,...\} \end{cases}.$$

Using these limit forms in (B.24), we find simple ODEs for the particle number as a function of time, which can easily be solved and give the results in the main text (2.28) for the total particle number.

B.2 Fisher analysis

The Fisher analysis involves performing the Taylor expansion of the exponent in the likelihood from equation (2.34) as a function of the theoretical parameters θ around its minimum[1]. The parameter set that minimises the chi-squared is given by the fiducial values assumed in the analysis, and is denoted by $\bar{\theta}$. The gradient of the chi-squared at $\bar{\theta}$ vanishes, and therefore the first non-vanishing term of the expansion comes from its Hessian, which is called the Fisher matrix, evaluated at the fiducial values

$$\begin{aligned} \chi^2(\theta) &\simeq (\theta-\bar{\theta})^\alpha \cdot \mathcal{F}_{\alpha\beta}(\bar{\theta}) \cdot (\theta-\bar{\theta})^\beta\,, \\ \mathcal{F}_{\alpha\beta} &= \frac{1}{2}\frac{\partial^2\chi^2(\theta)}{\partial\theta^\alpha \partial\theta^\beta}_{|\theta=\bar{\theta}} = T_{\rm obs}\int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \frac{\partial\Omega_{\rm GW}(\nu;\theta)}{\partial\theta^\alpha}_{|\theta=\bar{\theta}} \frac{\partial\Omega_{\rm GW}(\nu;\theta)}{\partial\theta^\beta}_{|\theta=\bar{\theta}} \frac{4\pi}{\Omega_n^2(\nu)}\,, \end{aligned} \tag{B.25}$$

where the last equality holds since we are expanding the chi-squared around its minimum. In our model, we identify the set of parameters of the unperturbed model as θ_0, and the two newly introduced diffusion and drift parameters are $\kappa_{1,2}$. The explicit evaluation of the Fisher matrix simplifies in the limit where $\kappa_{1,2} \ll 1$:

$$\begin{aligned} \mathcal{F}_{\theta_{0\,a},\theta_{0\,b}} &= T_{\rm obs}\int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \frac{\partial\Omega_0(\nu;\theta_0)}{\partial\theta_{0\,a}}_{|\theta_0=\bar{\theta}_0} \frac{\partial\Omega_0(\nu;\theta_0)}{\partial\theta_{0\,b}}_{|\theta_0=\bar{\theta}_0} \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\theta_{0\,a},\kappa_i} &= T_{\rm obs}\int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \frac{\partial\Omega_0(\nu;\theta_0)}{\partial\theta_{0\,a}}_{|\theta_0=\bar{\theta}_0} \Omega_i(\nu;\theta_0)_{|\theta_0=\bar{\theta}_0} \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\kappa_i,\kappa_j} &= T_{\rm obs}\int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \Omega_i(\nu;\theta_0)_{|\theta_0=\bar{\theta}_0} \Omega_j(\nu;\theta_0)_{|\theta_0=\bar{\theta}_0} \frac{4\pi}{\Omega_n^2(\nu)}\,. \end{aligned} \tag{B.26}$$

[1] For a detailed derivation of the Fisher analysis for cosmological observables, we refer to [303]

The forecast likelihood function of the theoretical parameters can be approximated by a multivariate normal distribution with $\bar{\theta}$ as its mean and the inverse of the Fisher matrix as its covariance matrix:

$$\begin{aligned} \mathcal{L}(\theta) &\simeq \mathcal{N}(\bar{\theta}, \Sigma^F)\,, \\ \Sigma^F &= \mathcal{F}^{-1}\,. \end{aligned} \tag{B.27}$$

For the case of the GW background considered in 2.4.1, the Fisher matrix is

$$\begin{aligned} \mathcal{F} &= \begin{pmatrix} \mathcal{F}_{\Omega_*,\Omega_*} & \mathcal{F}_{\Omega_*,\nu_b} & \mathcal{F}_{\Omega_*,\kappa_1} & \mathcal{F}_{\Omega_*,\kappa_2} \\ \mathcal{F}_{\nu_b,\Omega_*} & \mathcal{F}_{\nu_b,\nu_b} & \mathcal{F}_{\nu_b,\kappa_1} & \mathcal{F}_{\nu_b,\kappa_2} \\ \mathcal{F}_{\kappa_1,\Omega_*} & \mathcal{F}_{\kappa_1,\nu_b} & \mathcal{F}_{\kappa_1,\kappa_1} & \mathcal{F}_{\kappa_1,\kappa_2} \\ \mathcal{F}_{\kappa_2,\Omega_*} & \mathcal{F}_{\kappa_2,\nu_b} & \mathcal{F}_{\kappa_2,\kappa_1} & \mathcal{F}_{\kappa_2,\kappa_2} \end{pmatrix}, \\ \mathcal{F}_{\Omega_*,\Omega_*} &= T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \psi_0^2(\nu;\bar{\nu}_b)\, \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\Omega_*,\nu_b} &= \bar{\Omega}_* T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \psi_0(\nu;\bar{\nu}_b) \left(\frac{\partial \psi_0(\nu;\nu_b)}{\partial \nu_b}_{|_{\nu_b=\bar{\nu}_b}} \right) \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\nu_b,\nu_b} &= \bar{\Omega}_*^2 T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu \left(\frac{\partial \psi_0(\nu;\nu_b)}{\partial \nu_b}_{|_{\nu_b=\bar{\nu}_b}} \right)^2 \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\kappa_i,\nu_b} &= \bar{\Omega}_*^2 T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \psi_i(\nu;\bar{\nu}_b) \left(\frac{\partial \psi_0(\nu;\nu_b)}{\partial \nu_b}_{|_{\nu_b=\bar{\nu}_b}} \right) \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\kappa_i,\kappa_j} &= \bar{\Omega}_*^2 T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \psi_i(\nu;\bar{\nu}_b)\, \psi_j(\nu;\bar{\nu}_b)\, \frac{4\pi}{\Omega_n^2(\nu)}\,, \\ \mathcal{F}_{\Omega_*,\kappa_i} &= \bar{\Omega}_* T_{\rm obs} \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu\, \psi_0(\nu;\bar{\nu}_b)\, \psi_i(\nu;\bar{\nu}_b)\, \frac{4\pi}{\Omega_n^2(\nu)}\,, \end{aligned} \tag{B.28}$$

with $\bar{\theta} = \{\bar{\Omega}_*, \bar{\nu}_b, \bar{\kappa}_1, \bar{\kappa}_2\}$ being the fiducial parameters. The forecast error on κ_1 or κ_2 after marginalising the posterior distribution over the other parameters is the square root of the (κ_1,κ_1) or (κ_2,κ_2) element of the inverse Fisher matrix, i.e.

$$\begin{aligned} \sigma_{\kappa_1}^{\rm marginal} &= \sqrt{\Sigma^F_{\kappa_1\kappa_1}}\,, \\ \sigma_{\kappa_2}^{\rm marginal} &= \sqrt{\Sigma^F_{\kappa_2\kappa_2}}\,. \end{aligned} \tag{B.29}$$

Appendix C

Technical supplement to Everpresent Λ

C.1 On the equations of motion of global unimodular gravity

C.1.1 An alternative derivation of the equations of motion in unimodular gravity

In order to find the stationary point of the action in (3.4), let us perform the variation $g_{\mu\nu} \to g_{\mu\nu} + \delta g_{\mu\nu}$. The volume constraint in the path integral gives

$$\int d^4x\sqrt{-g} = V \quad \Rightarrow \quad \int d^4x\sqrt{-g}\ \delta g_{\mu\nu}g^{\mu\nu} = 0\,. \tag{C.1}$$

Now the stationary-action principle requires $\delta S = 0$. On the other hand, we know

$$\delta S \propto \int d^4x\sqrt{-g}\ \delta g_{\mu\nu}E^{\mu\nu}\,, \tag{C.2}$$

where $E_{\mu\nu} = G_{\mu\nu} - 8\pi G T_{\mu\nu} + \tilde{\Lambda}g_{\mu\nu}$. Note that we can no longer infer the full Einstein equations $E_{\mu\nu} = 0$ from (C.2) because $\delta g_{\mu\nu}$ should also satisfy (C.1). The condition (C.1) is satisfied if we rewrite the metric $g_{\mu\nu}$ in terms of an unconstrained metric $\hat{g}_{\mu\nu}$ as follows[1]

$$g_{\mu\nu}(x) = \sqrt{\frac{V}{\int d^4y\sqrt{-\hat{g}(y)}}}\ \hat{g}_{\mu\nu}(x)\,. \tag{C.3}$$

Performing a variation on (C.3) gives

$$\delta g_{\mu\nu}(x) = \sqrt{\frac{V}{\int d^4y\sqrt{-\hat{g}(y)}}}\left(\delta\hat{g}_{\mu\nu}(x) - \frac{1}{4}\frac{\int d^4y\sqrt{-\hat{g}(y)}\ \delta\hat{g}_{\alpha\beta}(y)g^{\alpha\beta}(y)}{\int d^4y\sqrt{-\hat{g}(y)}}g_{\mu\nu}(x)\right)\,. \tag{C.4}$$

Then for finding the classical equations of motion, we substitute (C.4) into (C.2) and then set the

[1] In the variational approach one normally obtains an equation of motion by setting the coefficient of the varied function to zero. Here, since we have an additional constraint, it turns out to be helpful to define our function g in terms of an auxiliary function $\hat{g}$ such that g satisfies the constraint, while the coefficient of $\hat{g}$ in the varied action gives the equation of motion. Notice that since $\hat{g}$ is defined by a rescaling of g, we will have relations such as $\hat{g}^{\alpha\beta}(y)\hat{g}_{\mu\nu}(x) = g^{\alpha\beta}(y)g_{\mu\nu}(x)$.

coefficient of $\delta\hat{g}_{\mu\nu}(x)$ in δS equal to zero. A bit of algebra reveals

$$E_{\mu\nu}(x) - \frac{1}{4}g_{\mu\nu}(x)\langle E\rangle = 0\,, \tag{C.5}$$

where $\langle E\rangle$ is the spacetime average of the trace of $E_{\mu\nu}$:

$$\langle E\rangle = \frac{\int d^4y\sqrt{-g}\; E(y)}{\int d^4y\sqrt{-g}}\,. \tag{C.6}$$

For an alternative derivation of (C.5), see Appendix C.1.2. Taking the trace of (C.5) reveals that $E = \langle E\rangle$, i.e., E must be a constant:

$$-R - 8\pi GT + 4\tilde{\Lambda} = \text{const}\,. \tag{C.7}$$

In particular, (C.5) can now be rewritten as

$$E_{\mu\nu}(x) - \frac{1}{4}g_{\mu\nu}(x)E = 0 \quad \Rightarrow \quad G_{\mu\nu} - 8\pi GT_{\mu\nu} + \frac{1}{4}g_{\mu\nu}(R + 8\pi GT) = 0\,, \tag{C.8}$$

and as argued in Section 3.1.2, $(R + 8\pi GT)$ must be constant, and can be set to 4Λ. We therefore obtain

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi GT_{\mu\nu}\,, \tag{C.9}$$

as we did in (3.12).

C.1.2 An Alternative derivation of (C.5)

In the usual variational principle, say for the Einstein-Hilbert action, one can argue that since $\delta g_{\mu\nu}$ is arbitrary, it can be chosen to be $\epsilon\delta^{\mu_0}_{\mu}\delta^{\nu_0}_{\nu}\delta(x - x_0)$, where ϵ is some small number. Then from (C.2), the stationary condition $\delta S \propto \int d^4x\sqrt{-g}\;\delta g_{\mu\nu}E^{\mu\nu} = 0$ gives $E^{\mu_0\nu_0}(x_0) = 0$, which is the Einstein equation for any point x_0 and for any component $\mu_0\nu_0$. In unimodular gravity, however, $\delta g_{\mu\nu}$ is not totally arbitrary.

One might naively think that we can satisfy (C.1), i.e.,

$$\int d^4x\sqrt{-g}\;\delta g_{\mu\nu}g^{\mu\nu} = 0\,, \tag{C.10}$$

by making a change in $\delta g_{\mu\nu}$ at infinity to compensate for a delta function at the point x_0, and then conclude that the same equation of motion as GR can be derived from global unimodular gravity. However, if the variation at infinity is done with care, it will lead to the same conclusion as equation (C.5) (i.e., $E_{\mu\nu}(x) - \frac{1}{4}g_{\mu\nu}(x)\langle E\rangle = 0$). Let us expand on this point. To satisfy (C.1), we can choose a point x_1 far away from x_0 and set

$$\delta g_{\mu\nu} = \epsilon\left(\delta^{\mu_0}_{\mu}\delta^{\nu_0}_{\nu}\delta(x - x_0) - \delta^{\mu_1}_{\mu}\delta^{\nu_1}_{\nu}\delta(x - x_1)\frac{g^{\mu_0\nu_0}(x_0)}{g^{\mu_1\nu_1}(x_1)}\right)\,. \tag{C.11}$$

It is easy to check that this variation satisfies (C.1). Now δS should be zero for such a variation. This

gives

$$E^{\mu_0\nu_0}(x_0) - E^{\mu_1\nu_1}(x_1)\frac{g^{\mu_0\nu_0}(x_0)}{g^{\mu_1\nu_1}(x_1)} = 0\,. \tag{C.12}$$

This may seem to be a nonlocal equation, since it is relating tensors that can be spacelike separated. But in fact it is just telling us that $\frac{E^{\mu_0\nu_0}(x_0)}{g^{\mu_0\nu_0}(x_0)}$ is a constant which does not depend on x_0 or $\mu_0\nu_0$. This is exactly the statement of (C.5). It is equivalent to the statement $E_{\mu\nu} - \frac{1}{4}g_{\mu\nu}E = 0$, which is the traceless Einstein equation. The argument that starts from (C.3) is essentially a more formal derivation of this fact.

C.2 On the fluid interpretation of Everpresent Λ

In our implementation of Model 1 to create background histories for Everpresent Λ, we only used the first Friedmann equation (3.32), which contains the energy density ρ_Λ but not the pressure P_Λ. However, for doing perturbation theory, it is necessary to work with the full Einstein equations, which contain both ρ_Λ and P_Λ. As mentioned in Section 3.3.1, if we insist on Everpresent Λ being indeed a "cosmological constant" in the sense that $P_\Lambda = -\rho_\Lambda$, then the fluctuating Λ makes the Friedmann equations (3.32) and (3.33), and the continuity equation $\dot{\rho}_\Lambda + 3H(\rho_\Lambda + P_\Lambda) = 0$ inconsistent with each other. As first suggested in [65], one way out of this dilemma is to relax the condition $P_\Lambda = -\rho_\Lambda$, and allow the equation of state parameter $w_\Lambda \equiv P_\Lambda/\rho_\Lambda$ to be realisation-dependent. This is the fluid interpretation of Everpresent Λ.

There are two options for modifying the equation of state in order to satisfy all three equations:

- Scheme 1: Set $\rho_\Lambda = \frac{1}{8\pi G}\Lambda$. The continuity equation then implies that

$$P_\Lambda = -\frac{1}{8\pi G}\left(\frac{\dot{\Lambda}}{3H} + \Lambda\right)\,. \tag{C.13}$$

- Scheme 2: Set $P_\Lambda = -\frac{1}{8\pi G}\Lambda$. By the continuity equation, ρ_Λ should then be the solution to the first order differential equation

$$\frac{\dot{\rho}_\Lambda}{3H} + \rho_\Lambda = \frac{\Lambda}{8\pi G}\,. \tag{C.14}$$

In both (C.13) and (C.14), the derivative is with respect to the cosmic time t.

An advantage of the fluid interpretation is that there is no need to alter the Einstein equations, and hence it may seem that the perturbation theory for the Everpresent Λ cosmology would be more straightforward. However, we shall now see that both of the above schemes are problematic.

In Scheme 1, whenever Λ crosses zero, we have $\rho_\Lambda = 0$. However, from (C.13), $\dot{\Lambda}$ being non-zero implies that $P_\Lambda \neq 0$. Therefore, the equation of state parameter w_Λ diverges at these zero-crossing instants, as can be seen in Figure C.1. This is not necessarily an issue from the point of view of the background evolutions, but problems arise at the level of perturbations. For example, the equation governing the density contrast of the dark energy fluid, δ_Λ, is

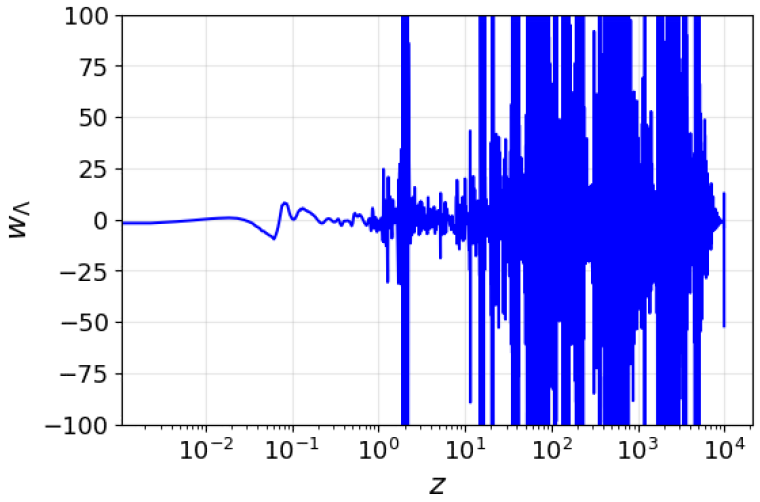


Figure C.1: We show w_Λ (in Scheme 1) versus redshift for one realisation of Everpresent Λ with parameters $\alpha = 0.01$, $\Omega_m h^2 = 0.14$. Each divergence of w_Λ corresponds to a zero-crossing of Λ. (Note that the value obtained for w_Λ is dependent on the method used to differentiate the stochastic time series of Λ).

$$\delta'_\Lambda + (1 + w_\Lambda)(\theta_\Lambda - 3\phi') + 3\mathcal{H}\left(\frac{\delta P_\Lambda}{\delta \rho_\Lambda} - w_\Lambda\right)\delta_\Lambda = 0\,. \tag{C.15}$$

The above equation is written in the Newtonian gauge. Now, according to (C.15), a divergence in w_Λ is detrimental to the perturbation variable δ_Λ. Whenever Λ crosses zero, the perturbation theory of Scheme 1 breaks down.

Scheme 2 is problematic in a different way. As shown in [65], solving the first Friedmann equation (3.32) with ρ_Λ from (C.14) is essentially equivalent to solving the second Friedmann equation (3.33). The issue with this is numerical instability. Small initial numerical errors grow uncontrollably and render the numerical integration unreliable. In addition, the problem of divergences in w_Λ persists in this scheme at zero-crossings of ρ_Λ instead of Λ.

In this thesis, we do not use the fluid interpretation. We try to remain agnostic by only employing the first Friedmann equation for the background, and not adding any perturbation equations for Everpresent Λ. A completely consistent treatment (that is potentially derived from modifications to the Einstein equations to accommodate a stochastic Λ) remains to be investigated in the future.

C.3 MCMC analysis with a seed number

Everpresent Λ is a stochastic model of the Universe. To calculate the expansion history from an Everpresent Λ model, in addition to the continuous parameters $(\alpha, \Omega_m^0 h^2)$, we need the discrete parameter s, the seed, to generate the series of random numbers. Unlike the other continuous parameters, s can assume any integer value in $[0, \infty)$, and we get an entirely different expansion history and a different CMB power spectrum for each value of s.

In this appendix we summarise three methods for doing an MCMC analysis to determine the posterior distribution of model parameters. The main difference between these three methods is in how the seed number that produces the unique histories in Models 1 and 2 is handled. Method 1 attempts to find

the distributions of model parameters by considering a large but finite selection of seeds. Method 2, incorporated in [66], treats the seed itself as one of the model parameters that the MCMC iterates over. Finally, Method 3 is how the MCMC is implemented in Chapter 4, where the distribution is determined for each seed separately, before combining the results for the model as a whole.

Method 1

For a stochastic model, the seed s needs to be marginalised to get the posterior of model parameters θ. We can write the likelihood function as

$$\mathcal{L}(\theta) = \mathbb{P}(D|\theta) = \sum_s \mathbb{P}(D|\theta, s)\,\mathbb{P}(s) = \frac{1}{N}\sum_s \mathbb{P}(D|\theta, s) = \frac{1}{N}\sum_s \mathcal{L}(\theta, s)\ . \tag{C.16}$$

Here D is the observational data and N is the number of seeds in the sample space. We assume that all the seeds are equally likely, and thus $\mathbb{P}(s) = \frac{1}{N}$. Hence the likelihood $\mathcal{L}(\theta)$ is an average over $\mathcal{L}(\theta, s) \propto \exp(-\chi^2(\theta, s)/2)$.

For sampling this likelihood using a Metropolis-Hastings MCMC, we can set the transition probability as

$$T(\theta_{i+1}|\theta_i) = \min\left[1,\ \frac{g(\theta_i|\theta_{i+1})\sum_{s'}\exp\left(-\frac{1}{2}\chi^2(\theta_{i+1}, s')\right)}{g(\theta_{i+1}|\theta_i)\sum_s \exp\left(-\frac{1}{2}\chi^2(\theta_i, s)\right)}\right], \tag{C.17}$$

where g is the proposal distribution. If $g(\theta_{i+1}|\theta_i) = g(\theta_i|\theta_{i+1})$, we can run the MCMC analysis using the following recipe. At each step of the MCMC, for the proposed θ_{i+1}, take all N seeds, and compute the χ^2 for each seed. Then sum these likelihoods to find the transition probability according to (C.17). If we run the MCMC chain long enough while imposing the relevant priors, we can sample the posterior distribution, $\mathbb{P}(\theta|D)$.

However, the problem with this method is that it is extremely slow as we are calculating the likelihood for all of the seeds at each step.

Method 2

A second option, proposed in [66], is to use the seed s as an independent parameter and keep changing the seed value at each step. In such a scenario, the transition probability is given by

$$T(\theta_{i+1}, s_{i+1}|\theta_i, s_i) = \min\left[1,\ \frac{g(\theta_i|\theta_{i+1})g(s_i|s_{i+1})\exp\left(-\frac{1}{2}\chi^2(\theta_{i+1}, s_{i+1})\right)}{g(\theta_{i+1}|\theta_i)g(s_{i+1}|s_i)\exp\left(-\frac{1}{2}\chi^2(\theta_i, s_i)\right)}\right]. \tag{C.18}$$

To sample the posterior, we need to choose the seeds from a finite unbiased sample of seeds. Assuming all the seeds are equally likely, $g(s_{i+1}|s_i) = g(s_i|s_{i+1})$. Assuming $g(\theta_{i+1}|\theta_i) = g(\theta_i|\theta_{i+1})$, the transition

probability simplifies to

$$T\left(\theta_{i+1}, s_{i+1}|\theta_i, s_i\right) = \min\left[1,\ \frac{\exp\left(-\frac{1}{2}\chi^2\left(\theta_{i+1}, s_{i+1}\right)\right)}{\exp\left(-\frac{1}{2}\chi^2\left(\theta_i, s_i\right)\right)}\right]. \tag{C.19}$$

However, the problem is that χ^2 is not a slowly varying function of s. In most of the cases, for each new proposal s_{i+1}, the transition probability is either 1 or negligibly small. Thus the Markov chain does not effectively explore the parameter space of θ. While, in principle, it may be possible to sample the full posterior distribution by running the chains for an infinite amount of time, in practice this is computationally infeasible.

Method 3

A more efficient option is to run the MCMC for each of the seeds independently, i.e. to keep s constant and only update θ. The transition probability in that case is given by

$$T\left(\theta_{i+1}, s|\theta_i, s\right) = \min\left[1,\ \frac{\exp\left(-\frac{1}{2}\chi^2\left(\theta_{i+1}, s\right)\right)}{\exp\left(-\frac{1}{2}\chi^2\left(\theta_i, s\right)\right)}\right]. \tag{C.20}$$

In this way, we can get the posterior distribution of θ for that particular seed, $\mathbb{P}(\theta|D, s)$. We then calculate the posterior over all seeds by

$$\mathbb{P}(\theta|D) = \sum_s \mathbb{P}(\theta|s, D)\mathbb{P}(s|D)\,. \tag{C.21}$$

Therefore, we take the average of the posteriors, each weighted by the marginalised posterior probability of that seed, $\mathbb{P}(s|D)$, to get $\mathbb{P}(\theta|D)$. To find the marginalised posterior of a seed, note that according to Bayes' rule:

$$\mathbb{P}(s|D) = \frac{\mathbb{P}(D|s)\mathbb{P}(s)}{\sum_{s'}\mathbb{P}(D|s')\mathbb{P}(s')}\,. \tag{C.22}$$

We take $\mathbb{P}(s)$ to be uniform on the given sample of seeds, so now we need to find the *marginalised likelihood* of a seed s, $\mathbb{P}(D|s)$. Once again, Bayes' rule gives

$$\mathbb{P}(D|s) = \frac{\mathbb{P}(D|\theta, s)\mathbb{P}(\theta|s)}{\mathbb{P}(\theta|D, s)} \quad \Rightarrow \quad \sum_\theta \frac{\mathbb{P}(\theta|D, s)}{\mathbb{P}(D|\theta, s)} = \frac{1}{\mathbb{P}(D|s)}\,. \tag{C.23}$$

Now, if we have an MCMC sample of points θ_i, $1 \leq i \leq N_s$, they are distributed according to the posterior $\mathbb{P}(\theta|D, s)$. Given the likelihood of each one, $\mathbb{P}(D|\theta_i, s)$, we can therefore estimate $\mathbb{P}(D|s)$ using

$$\frac{1}{\mathbb{P}(D|s)} = \frac{1}{N_s}\sum_{i=1}^{N_s}\frac{1}{\mathbb{P}(D|\theta_i, s)}\,. \tag{C.24}$$

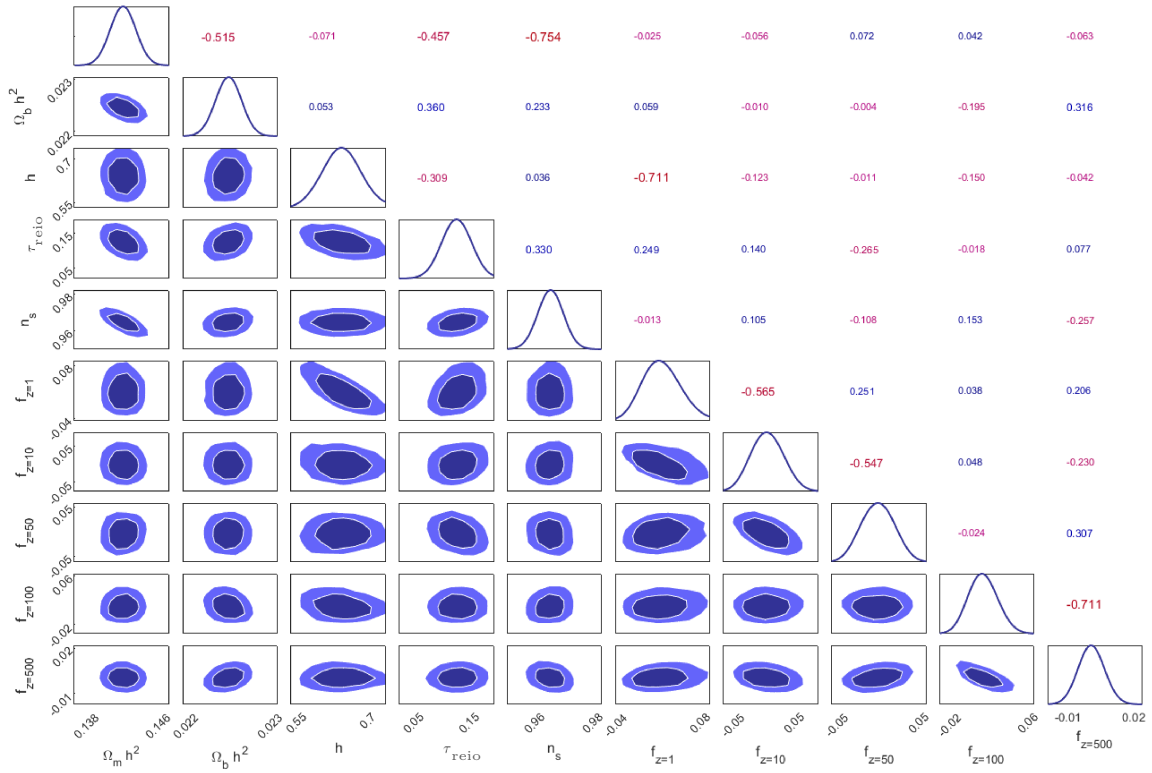


Figure C.2: Results from an 11 parameter MCMC analysis constraining the expansion history after the last scattering surface, are shown. A_s, which just gives a scaling on the CMB power spectrum, is omitted to make the plot less cluttered. The lower triangle panels show plots of the 68% and 95% confidence contours for pairs of parameters. The average, standard deviation and the best fit values are stated in Table 4.2.

This method of calculating the posterior distribution is faster than Methods 1 and 2 because we can calculate $\mathbb{P}(\theta|s, D)$ for different seeds in parallel. Therefore, this is the method we have used in Section 4.1.

C.4 Full posteriors for fluctuating dark energy

We present the 68% and 95% confidence levels of marginalised posteriors for the main parameters, A_s excluded, used in our MCMC analysis of Planck2018 data. Figures C.2 and C.3 are for the fluctuating dark energy models (4.17) and (4.18), respectively. The upper triangles show the correlation-coefficient between the pair of parameters. Negative correlation-coefficients are shown in red and positive correlation-coefficients are shown in blue. The diagonal plots are the 1-dimensional marginalised distribution of the parameters.

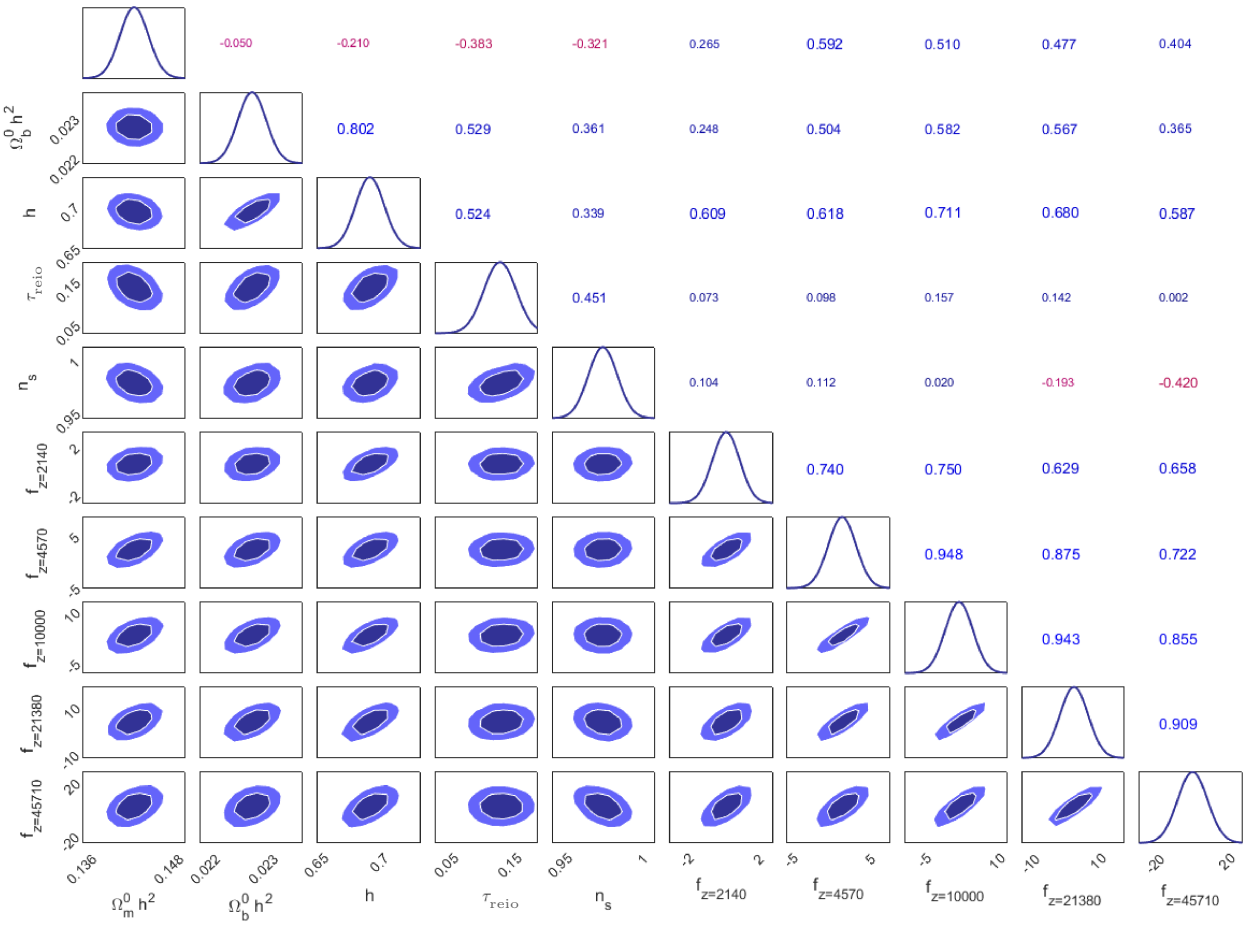


Figure C.3: Results from an 11 parameter MCMC analysis constraining the expansion history before the last scattering surface, are shown. A_s, which just gives a scaling on the CMB power spectrum, is omitted to make the plot less cluttered. The lower triangle panels show plots of the 68% and 95% confidence contours for pairs of parameters. The average standard deviation and the best fit values are stated in Table 4.2. In this plot the f(z) values have been multiplied by 1000 to avoid the zeros after the decimal point.