University of Sussex

Small-scale probes of the early Universe

Philippa S. Cole

Submitted for the degree of Doctor of Philosophy

Department of Astronomy

University of Sussex

March 2020

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

The work in this thesis has been completed in collaboration with Christian Byrnes, Subodh Patil and Joseph Silk, and is comprised of the following papers:

- P. S. Cole and C. T. Byrnes, Extreme Scenarios: the tightest possible constraints on the power spectrum due to primordial black holes, (2017), arXiv:1706.10288 [astro-ph.CO], published in JCAP 1802 (2018) no.02, 019.

  Christian Byrnes was responsible for the original idea of this paper. The paper and calculations were completed by myself under the supervision of Christian Byrnes, who made some minor adjustments and additions.

- C. T. Byrnes, P. S. Cole and S. P. Patil, Steepest growth of the power spectrum and primordial black holes, (2018), arXiv: 1811.11158 [astro-ph.CO], published in JCAP 1906 (2019) 028.

  Christian Byrnes and Subodh Patil came up with the initial idea for this paper, which we then all refined throughout the course of the project. Sections 3.2 and A were predominantly written by myself with additions from Subodh Patil, and the calculations were done by myself and Subodh Patil to cross-check results. Section 3.2.3 was predominantly worked on by Christian Byrnes. Section 3.3 was predominantly worked on by Christian Byrnes with minor contributions from myself. Section 3.4 was worked on by all three of us. Sections B and C were predominantly worked on by Subodh Patil. All of the numerical calculations of the power spectrum were completed by myself, with guidance from my second supervisor David Seery.

- P. S. Cole and J. Silk, Small scale primordial fluctuations in the 21cm Dark Ages signal, arXiv:1912.02171.

  Joseph Silk came up with the initial idea for this paper, which I then refined and

steered for the remainder of the project. I wrote the paper and completed all of the calculations myself with advice from Joseph Silk.

Signature:

Philippa Cole

UNIVERSITY OF SUSSEX

Philippa Cole, Doctor of Philosophy

Small-scale probes of the early universe

Summary

Large-scale observations of the Cosmic Microwave Background have provided the most substantial evidence for how the Universe expanded rapidly in the first fraction of a second after the Big Bang. This phase of rapid expansion known as inflation provided the seeds for all of the structure that we see in the Universe today, and therefore understanding its dynamics will provide explanations for why the Universe is the way it is, in both the visible and dark sectors.

Constraining the dynamics of inflation is best done with the primordial power spectrum, which measures the overdensities and underdensities left over at the end of inflation. This thesis investigates different probes for measuring and constraining the primordial power spectrum on small scales, where constraints are currently much weaker than those deduced from large-scale measurements.

The existence or lack of primordial black holes is primarily investigated as a means for both constraining the primordial power spectrum and hence inflation, as well as due to their interest as a dark matter candidate. Large-amplitude scalar perturbations are usually required for the production of primordial black holes, and therefore observational signatures of such perturbations including spectral distortions, stochastic gravitational waves, and the 21cm signal will be explored as means of detecting them. Constraints and signatures from these observations will be compared with classes of inflationary models so as to understand the inflationary dynamics that would be necessary to produce the results and predictions of multiple current and future experiments.

Acknowledgements

I would like to thank my supervisor, Chris Byrnes, for supporting me through what has been an amazing three and a half years. He has always encouraged me to travel, meet new people and talk about my work, and these experiences have enriched my time as a PhD student no-end. I've felt incredibly lucky to have him as my supervisor.

I would also like to thank David Seery and Antony Lewis for good advice and support during my time at Sussex, and for creating such a friendly atmosphere in the department.

Furthermore, it was a pleasure to work with Subodh Patil and Joe Silk during my PhD, and I especially enjoyed my time as a Balzan fellow at Johns Hopkins University with Joe. Thank you also to Yacine Ali-Haïmoud for a fabulous experience at New York University during my long term attachment, and for the ongoing collaboration that I hope will continue long after my PhD.

To the PhD cohort at Sussex who were always ready for a cold beer at Falmer - Steve Duivenvoorden, Lucas Porth, Dimitrios Irotodou, Aswin Vijayan, Jussi Kuusisto, Luke Conaboy, Kareem Marzouk, Maria Campos Varillas, Dan Pryer, Azizah Hosein and Carlos Vergara - you've all been the best bunch to share this experience with.

I'd also particularly like to thank Guilio Fabbian and his high standards of Italian food and wine which kept us well fed and watered since his arrival last year; Ciaran Fairhurst for not wringing my neck as I gave him daily percentage updates of how far through the PhD we were (and for being a rock during the last few months – we made it!); Jessie Hislop for being the wonder woman that she is, I've missed you hugely since you left; and to Chris Lovell, who has been both an incredible support and a terrible influence all at the same time. Finally, to Sunayana Bhargava, who has become a best friend through this process. Thank you for living with me in the highs, and scooping me out of the lows – I'm so proud of us.

I'd also like to thank Niamh Kerr, Jade Orekoya, Hannah Picton, Fab Trovato and everyone at Marina Studios for the best memories of Brighton.

Thank you also to all of my friends not in Brighton who've supported me immeasurably, there really are too many to name, but I hope that you know who you are, and that you've

learnt that I'm not studying star signs.

Finally, thank you to my family. You've inspired me to go out there and do whatever I want to do with the utmost support, and unwavering belief that I am capable of success, even when I didn't believe in myself. I hope I've made you proud. I'm looking forward to Grandma's review of this thesis, as she once memorably put it, "the Universe is just too big to comprehend".

I would like to dedicate this thesis to the best teacher I ever had, Reverend Tim Hardingham. I am so sorry that I didn't get a chance to tell him that I became a doctor in the subject he inspired me to study.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Reader's guide

Everything we know about inflation, the phase of rapid expansion immediately after the Big Bang, is based on evidence from large-scale observations that can only tell us about a tenth of its total duration. There is a wealth of information hidden on small scales that can be unlocked by comparing theoretical predictions for various relics of inflation including primordial black holes, gravitational waves and spectral distortions with observations from upcoming experiments.

For an introduction to inflationary dynamics and constraining the primordial power spectrum, see section 1.2 of the introduction. For an overview of primordial black hole motivations and production, see section 1.3. For brief introductions to spectral distortions and gravitational waves as probes of the primordial power spectrum see sections 1.4 and 1.5 respectively, and for an introduction to 21cm observations as a probe for primordial fluctuations, see section 1.6.

Chapter 2 will then make the connection between primordial black holes and the primordial power spectrum, whilst chapter 3 will make the connection between the primordial power spectrum required for primordial black hole production and inflationary dynamics. Constraining the primordial power spectrum with current observations will also be explored in chapter 3, including via the cosmic microwave background, spectral distortions and gravitational waves. Finally, chapter 4 will investigate using 21cm measurements as a probe for the matter power spectrum during the Dark Ages, which is a tracer for the primordial power spectrum. This will be used to look for signatures of primordial black holes.

1.2 Cosmological inflation

1.2.1 Motivation

During the first second after the Big Bang, the Universe underwent a period of rapid expansion, where it grew by over twenty orders of magnitude in size. This phase is known as inflation, and it is currently the best explanation we have for the dawn of the Universe's existence. Tiny quantum fluctuations present during the expansion were blown up and stretched out, providing seeds for the growth of all of the structure that we see today. The need for this rapid period of expansion is due to the fact that the Universe we observe today is very flat, no exotic relics have yet been observed, and central to this work, is very homogeneous.

The cosmological principle states that on large enough scales, an observer should view the Universe to look the same in every direction, no matter their location. This has been evidenced by the fact that patches of the sky separated by the largest distances that we are able to probe have been measured to have almost exactly the same temperatures. This suggests that they were at one point in physical contact, so as to exchange heat energy and equilibrate. However, if the Universe has expanded according to the Hot Big Bang model, namely at a rate governed by a radiation-dominated epoch followed by a matter-dominated one, those regions could never have been causally connected, as there hasn't been enough time for them to get so far apart. A period of rapid expansion, therefore, explains how regions that were initially in causal contact and hence equilibrium, could have been blown apart and end up seemingly too far away from each other to share the same properties.

More quantitatively, a photon can travel the distance

$$d_p = \int_0^t \frac{dt'}{a(t')} \tag{1.1}$$

in time t, where $a(t)$ is the scale factor which is governed by the Friedmann equation

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3} - \frac{K}{a^2} \tag{1.2}$$

with the Hubble factor H, G the gravitational constant, ρ the energy density of the Universe and K the curvature. Throughout this thesis, derivatives with respect to time t will be denoted with dots. According to the conservation of energy, and assuming adiabatic expansion, we can write down a relationship between the energy density and the pressure p:

$$\dot{\rho} + 3H(\rho + p) = 0. \tag{1.3}$$

Then for $p = \omega\rho$ where ω is the equation of state we find

$$\rho \propto a^{-3(1+\omega)} \propto \begin{cases} a^{-3} & \text{for pressureless matter, } \omega = 0 \\ a^{-4} & \text{for radiation, } \omega = 1/3. \end{cases} \tag{1.4}$$

Matter dilutes like a^3, i.e. the rate at which the Universe is expanding, whilst the wavelength of the radiation also gets redshifted to lower energies, and therefore dilutes faster like a^4. This means that the early Universe will be radiation-dominated, before becoming matter-dominated at matter-radiation equality $t_{\rm eq}$ when $\rho_r = \rho_m$. By integrating equation (1.2) (with $K = 0$), with these relations, we find that the scale factor behaves as $a \propto t^{1/2}$ in radiation-domination and $a \propto t^{2/3}$ in matter-domination. Returning to equation (1.1), we can then see that the maximum comoving distance a photon can have travelled by recombination $t_{\rm rec}$ (more on this later) is

$$d_p = \int_0^{t_{\rm eq}} \left(\frac{t_0}{t}\right)^{1/2} dt + \int_{t_{\rm eq}}^{t_{\rm rec}} \left(\frac{t_0}{t}\right)^{2/3} dt \tag{1.5}$$

if the Universe was just made up of matter and radiation. Plugging in the values $t_{\rm rec} \sim 10^{13}\,$s, $t_{\rm eq} \sim 10^{12}\,$s and $t_0 \sim 4 \times 10^{17}\,$s gives a maximum distance of $d_p \sim 300\,$Mpc [19], meaning that we shouldn't expect to see correlations in temperature between patches of the sky larger than this. However, the comoving distance to recombination is 14000 Mpc, and we indeed observe correlations in the temperature of the sky on these scales. Regions this large can't have been in causal contact if the Universe had only experienced periods of radiation and matter-domination, and instead there needs to have been a phase of accelerated expansion so that previously causally connected regions can now be much further apart than causal processes would allow. Such a phase can be realised if the equation of state is sufficiently negative, and in particular, accelerated expansion occurs when $\omega < -1/3$. See e.g. [20] for a review. During the rapid expansion, the comoving Hubble horizon scale, $(aH)^{-1}$, will decrease, while the comoving scale of perturbations will remain constant. This means that the largest modes will exit the Hubble horizon, going out of causal contact, first. After inflation ends, the smallest scale modes will reenter, as the comoving Hubble horizon scale starts to grow during radiation-domination. This is illustrated in figure 1.1.

Exactly how this phase of expansion occurred therefore determines the initial conditions for the dark matter distribution, and hence the gravitational potential wells into which matter eventually collapses, forming stars, galaxies and black holes. There are two ways of approaching the problem of how exactly this phase of rapid expansion took place. Either, we can observe structures today that allow us to put constraints on the initial

overdensities and underdensities left over at the end of inflation and therefore infer how inflation worked. Conversely, given a model or class of inflationary models we can put restrictions on the type of objects, structures or signals that should be detectable today. If those observables are there, then the features of those models become a necessary ingredient for inflationary dynamics. If they're not, we can rule out that model or class of inflationary models.

Quantifying the dynamics of inflation will require a prescription for tracking the field that drives the expansion, as well as for measuring the overdensities and underdensities left over at the end. We will now see how this is most commonly approached, using cosmological perturbation theory.

1.2.2 The primordial perturbation

In cosmological perturbation theory, overdensities are quantified as a small perturbation to the background energy density

$$\rho(\mathbf{x}, t) = \rho_{\rm bg}(t) + \delta\rho(\mathbf{x}, t) \tag{1.6}$$

and only the lowest-order (linear) terms in $\delta\rho$ are kept[1]. The background energy density is assumed to be flat and homogeneous (and hence only a function of time), whilst the perturbation is a function of time and space. An isotropic and homogeneous universe is described by the Friedmann-Robertson-Walker (FRW) metric

$$\begin{aligned} ds^2 =& a^2(\tau)(d\tau^2 - d\mathbf{x}^2) \\ =& a^2(\tau)\left(d\tau^2 - \frac{dr^2}{1-Kr^2} - r^2 d\Omega^2\right) \end{aligned} \tag{1.7}$$

with conformal time τ defined by $dt = ad\tau$, K is the curvature, r is the radial coordinate and Ω is the angular coordinate. On large scales, the Universe will locally appear smooth and homogeneous, and such regions can therefore be treated as separate FRW universes [21]. Due to the initial presence of quantum fluctuations, some of these regions will expand more or less than others, and we can quantify the difference between regions either by the relative change in energy density, or instead slice the space-time such that each region has the same energy density, $\delta\rho = 0$, and define the different regions by their relative curvature. This relative curvature can then be used to describe the perturbation to the background.

[1] However, note that this linear approximation starts to become far less accurate in the context of primordial black holes which form from large overdensities, as will be discussed later.

The simplest scalar quantity that represents the curvature of this metric is the Ricci scalar of constant-time hypersurfaces, $R = 6K/a^2$. We want to define a scalar to describe the primordial perturbations which is proportional to this curvature, with a gauge choice such that $\delta\rho = 0$. The quantity which satisfies these constraints is the comoving curvature perturbation:

$$\mathcal{R} = \frac{a^2 R}{4k^2}. \tag{1.8}$$

This definition of the comoving curvature perturbation is just a particularly convenient way of describing a general scalar perturbation. There are, however, other parametrisation choices that can be made. Intuitively, we can also think of a perturbation being a difference in how much the Universe has expanded at a given position relative to the background, i.e. a change in the scale factor at a given position relative to the background. We could instead define the primordial perturbation directly as a change in the scale factor, $\zeta = \delta a/a = \delta \log a$, and then under a change of coordinates we can see that [20]

$$\begin{aligned} &\frac{\delta a}{a} \to \frac{\delta a}{a} - \frac{\dot{a}}{a}\delta t \\ \implies &\zeta \to \zeta - \frac{\dot{a}}{a}\delta t = \zeta - H\delta t. \end{aligned} \tag{1.9}$$

Parametrising the time step in terms of a change in energy density $\delta\rho$,

$$\zeta - H\frac{\delta\rho}{\dot{\rho}} = \zeta + \frac{\delta\rho}{3(\rho + p)} = \mathcal{R} \tag{1.10}$$

where the second equality is due to the conservation of energy given by equation (1.3), and then it is obvious that this 'general' perturbation is equal to the comoving curvature perturbation $\mathcal{R}$ in the uniform density gauge where $\delta\rho = 0$. In the separate universe approach, $\delta\rho = 0$ outside of the horizon, and therefore $\mathcal{R} = \zeta$ outside of the horizon too. This means that these two parameters are often used interchangeably in the literature. However, inside the horizon these two quantities are not equivalent, and in regimes where the perturbation is not conserved outside of the horizon, these quantities are also distinct and must not be used interchangeably.

1.2.3 The primordial power spectrum

The distribution of overdensities and underdensities left over at the end of inflation is quantified by the primordial power spectrum as a function of scale. This is the Fourier transform of the 2-point correlation function $\langle\delta(\mathbf{x}_1)\delta(\mathbf{x}_2)\rangle$ which gives the probability of a region with physical size $r \propto 1/k$ having a density δ greater than the average. The power

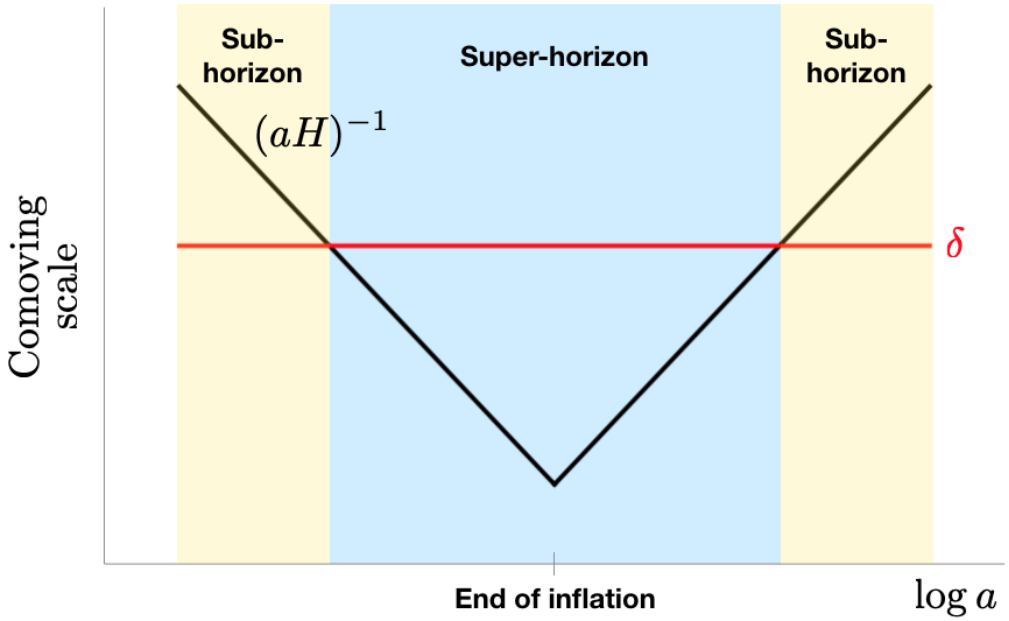

Figure 1.1: The comoving scale (wavelength) of the Hubble horizon in black, and of a perturbation in red, as a function of the scale factor. All modes begin inside the horizon, and are then blown up larger than the Hubble horizon during inflation. After inflation they re-enter the Hubble horizon, smallest scale first.

spectrum is then

$$\mathcal{P}(\mathbf{k}) = \frac{1}{(2\pi)^{3/2}} \int \langle \delta(\mathbf{x}_1)\delta(\mathbf{x}_2)\rangle \, e^{i\mathbf{k}\cdot\mathbf{x}} d^3\mathbf{k} \tag{1.11}$$

and is a measure of how over- or underdense regions of a particular size $r \propto 1/k$ are on average.

Throughout this thesis, the primordial power spectrum will represent the density in terms of the comoving curvature perturbation, $\mathcal{R}$. Using the separate universe approach, it can be shown that this quantity is conserved on superhorizon scales because each region is chosen such that it has constant curvature [21]. This means that inflationary dynamics can be inferred from the primordial power spectrum of the comoving curvature perturbation measured at the time that the perturbations exit the horizon and freeze-out. However, we will see in chapter 3 that there are circumstances where the comoving curvature perturbation does not remain constant on superhorizon scales, and instead needs to be tracked until the end of inflation before the primordial power spectrum can be evaluated. The primordial power spectrum acts as a bridge between inflationary dynamics and observations today, and this thesis will focus on comparing predictions of inflationary models with constraints from various observational probes across a wide range of scales.

1.2.4 Cosmic Microwave Background

To date, we have obtained precise measurements of the primordial power spectrum on large comoving scales $k \sim 10^{-3} - 0.1\,\mathrm{Mpc}^{-1}$ via measurements of the Cosmic Microwave Background [22]. At early times, the temperature of the Universe was still high enough that the hydrogen (the predominant constituent of the baryonic matter in the Universe) was completely ionized. Photons weren't able to travel great distances, because they were scattered frequently off the sea of free protons and electrons. However, by around redshift $z = 1100$, the Universe had cooled enough so that the protons and electrons could combine to form neutral hydrogen and the photons were able to decouple from the baryons and free-stream. This process is known as recombination, and the photons that were released at that time are known as the Cosmic Microwave Background. This radiation was first observed in 1965 by Arno Penzias and Robert Wilson at the Bell Telephone Laboratories in Murray Hill, New Jersey, and since then NASA's Cosmic Background Explorer (COBE), the Wilkinson Microwave Anisotropy Probe (WMAP) and the Planck Satellite have zeroed in on a precise measurement of the average CMB temperature today: $T_{\mathrm{CMB},0} = 2.72548 \pm 0.00057\,\mathrm{K}$ [23]. Since the temperature of radiation redshifts as $T_{\mathrm{CMB}} = T_{\mathrm{CMB},0}\,(1+z)$, the CMB temperature is then known at all times, and it was the measurement of this temperature across large distances that provided one of the main motivations for inflation, as discussed in section 1.2.1.

The Planck satellite has measured spatial fluctuations in the temperature of the CMB to deviate from the average by around 1 part in 100,000 across the range of scales $k \sim 10^{-3}$–$0.1\,\mathrm{Mpc}^{-1}$. Since the CMB fluctuations are small, there is a direct link between these fluctuations in CMB temperature and the initial primordial fluctuations of the comoving curvature perturbation via linear perturbation theory. This means that constraints on the primordial power spectrum can be derived from measurements of the CMB temperature fluctuations with the CMB transfer functions $T(k,t)$. These encode the linear evolution of various quantities such that $\mathcal{P}_{\mathcal{R}} T_x^2(k,t) = \mathcal{P}_{x,t}$ where x is the quantity of interest, for example the CMB temperature fluctuations or cold dark matter density fluctuations. Various codes including CAMB [24] and CLASS [25] are able to compute these transfer functions by solving the linear Boltzmann equations.

The latest results from the 2018 Planck release parametrise the primordial power spectrum with a power law,

$$\mathcal{P}_{\mathcal{R}} = A_s \left(\frac{k}{k_*}\right)^{n_s - 1} \tag{1.12}$$

with $A_s = 2 \times 10^{-9}$, $k_* = 0.05\,\mathrm{Mpc}^{-1}$ and $n_s = 0.965 \pm 0.004$ [26]. This shows that

the primordial power spectrum is nearly scale-invariant over this range, and that the overdensities and underdensities do not differ from the average by much. Using this tracer for the initial densities, we can then ask what kind of inflationary models are able to produce this distribution.

1.2.5 Inflationary model-building

Recalling the equation of state required for accelerated expansion, $\omega < -1/3$, we can see that a scalar field ϕ which has density and pressure given by

$$\begin{aligned} \rho_\phi &= \frac{1}{2}\dot{\phi}^2 + V(\phi) \\ p_\phi &= \frac{1}{2}\dot{\phi}^2 - V(\phi), \end{aligned} \tag{1.13}$$

where $V(\phi)$ is the potential of the scalar field, will provide a negative equation of state, and $\omega < -1/3$ given constraints on the potential. Inflation can be driven by one, usually referred to as the inflaton, or many, scalar fields. These classes of models are known as single-field and multifield inflation respectively.

The way that these scalar fields spend their potential energy determines the dynamics of inflation, and hence how much the initial fluctuations are blown up and stretched out. We track these fluctuations as modes with a given wavenumber k. All modes start off inside the horizon before exiting, largest first, as the Universe expands. After inflation ends, these modes reenter the horizon and come back into causal contact, see figure 1.1. This process is usually modelled as a scalar field rolling down a potential. In single-field inflation, there is just one degree of freedom for the field to roll, but in multifield inflation, there could be many.

Slow-roll inflation

Within the class of single-field models of inflation, we can further categorise a subset of those with the simplest dynamics. If a single scalar field rolls slowly, and at a nearly constant speed down it's potential, then it is said to be a slow-roll model of inflation. This distinction is quantified by the two slow-roll parameters

$$\epsilon = -\frac{\dot{H}}{H^2} = \frac{\dot{\phi}^2}{2M_p^2 H^2} \tag{1.14}$$

$$\eta = \frac{\dot{\epsilon}}{\epsilon H}, \tag{1.15}$$

where ϕ is the field value and $H = \dot{a}/a$ is the Hubble parameter. The first slow-roll parameter, ϵ, tracks how quickly the scalar field is moving down its potential, whilst the

second slow-roll parameter, η, tracks the acceleration. This means that if ϵ is small and constant, then $\eta \approx 0$, and the model is classed as slow-roll, meaning that the slow-roll approximation can be applied as follows. Inserting equation (1.13) into equation (1.2), the Friedmann equation for when the Universe is dominated by a scalar field reads

$$H^2 = \frac{1}{3M_p^2}\left(\frac{1}{2}\dot{\phi}^2 + V(\phi)\right), \tag{1.16}$$

and if $\epsilon \ll 1$, then

$$H^2 \approx \frac{V}{3M_p^2}. \tag{1.17}$$

This approximation states that ϵ and η can be recast in terms of the potential and its derivatives only:

$$\epsilon \approx \frac{M_p^2}{2}\left(\frac{V'(\phi)}{V(\phi)}\right)^2 \tag{1.18}$$

$$\eta \approx M_p^2 \frac{V''(\phi)}{V(\phi)}, \tag{1.19}$$

where $' = d/d\phi$. The power spectrum can be written in terms of ϵ as

$$\mathcal{P}_\mathcal{R} = \frac{H^2}{8\pi M_p^2 \epsilon}, \tag{1.20}$$

evaluated at the end of inflation, see e.g. [27] for a review of this formalism. And hence, in the slow-roll approximation, the primordial power spectrum becomes just a function of the potential V. It is not possible to write the slow-roll parameters, and hence the power spectrum, as a function of just the potential if $\eta \not\approx 0$. To understand the physical interpretation of being able to use the slow-roll approximation or not, we need to think about how the perturbations are behaving in each case.

1.2.6 Inflationary perturbation evolution

The joint evolution of the metric and the scalar field driving inflation, which we will assume is a minimally coupled scalar field ϕ, is given by the Friedmann equation as in (1.16) for the background, while the equation of motion for the scalar field perturbations is given by the Klein-Gordon equation:

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\mathrm{d}V}{\mathrm{d}\phi} = 0, \tag{1.21}$$

which is derived from the action for a minimally coupled scalar field. In the slow-roll approximation, the first term $\ddot{\phi}$ is neglected because $\eta \approx 0$, and hence $\epsilon \propto \dot{\phi}^2$ can be written in terms of $V'(\phi)$. In the case that $\eta \not\approx 0$, then $\ddot{\phi}$ is non-negligible in comparison to $V'(\phi)$.

If instead we wrote the action in terms of the comoving curvature perturbation $\mathcal{R}$ and then transformed to the canonically-normalised Mukhanov variable $\upsilon_k = z\mathcal{R}$ where $z = \sqrt{2\epsilon}a$ we would find the Mukhanov-Sasaki (MS) equation [28, 29]:

$$\upsilon_k'' + (k^2 - \frac{z''}{z})\upsilon_k = 0, \tag{1.22}$$

where primes denote differentiation with respect to conformal time τ, with $dt = ad\tau$. By directly differentiating z, it can be shown that

$$\frac{z''}{z} = (aH)^2\left(2 - \epsilon + \frac{3}{2}\eta + \frac{1}{4}\eta^2 - \frac{1}{2}\epsilon\eta + \frac{1}{2}\frac{\dot{\eta}}{H}\right). \tag{1.23}$$

This is exact - no approximations or higher order cut-offs have been introduced here.

To see what is happening to modes deep within the horizon, we can take the limit of equation (1.22) as τ approaches $-\infty$. In this case the k^2 term is dominant and the MS equation becomes

$$\upsilon_k'' + k^2\upsilon_k = 0, \tag{1.24}$$

which admits oscillating solutions $\upsilon_k \propto e^{\pm ik\tau}$.

The behaviour of the modes once they've left the horizon must be determined including the z''/z factor, since the wavenumber k associated with a particular mode is comparable to $1/\tau$ at the time of horizon crossing. We can't solve the MS equation easily with the form of (1.23) as it is, however if we assume that $\epsilon \ll 1$, then we can rewrite (1.23) as

$$\frac{z''}{z} = (aH)^2\left(2 + \frac{3}{2}\eta + \frac{1}{4}\eta^2 + \frac{1}{2}\frac{\dot{\eta}}{H}\right), \tag{1.25}$$

where we've just dropped all terms that involved ϵ. In fact, the condition for inflation to be occurring is only $\epsilon < 1$, however for inflation to be sustained, $\epsilon \ll 1$ is usually required. If we now define a new parameter ν by

$$\nu^2 = \frac{9}{4} + \frac{3}{2}\eta + \frac{1}{4}\eta^2 + \frac{\dot{\eta}}{2H}, \tag{1.26}$$

the MS equation can be rewritten as

$$\upsilon_k'' + (k^2 - \frac{\nu^2 - \frac{1}{4}}{\tau^2})\upsilon_k = 0 \tag{1.27}$$

where again, the only approximation here is that ϵ is small compared to 1 and η.

It is now clear that the evolution of υ and hence $\mathcal{R}$ is very dependent on the value of η, which tracks the rate of change in ϵ.

We will explore this behaviour fully in chapter 3, however for now the important thing to note is that if we write the equation of motion for $\mathcal{R}$ which is found by substituting $v_k = z\mathcal{R}$ into equation (1.22)

$$\mathcal{R}'' + \left(\frac{2a'}{a} + \frac{\epsilon'}{\epsilon}\right)\mathcal{R}' + k^2\mathcal{R} = 0 \tag{1.28}$$

in the superhorizon limit $k \to 0$, then we have

$$\mathcal{R}_{k\to 0} = C_k + D_k \int \frac{dt}{a^3\epsilon}, \tag{1.29}$$

and we can see that there is always a constant mode C_k (the usual conserved quantity in single-field slow-roll), but that there is also a time-dependent mode. If ϵ is small and constant, as is the case in slow-roll, then the whole second term will be decaying as the scale factor dominates and there is no additional contribution to $\mathcal{R}$ other than the value that is constant at horizon exit. This means that in the slow-roll approximation, the power spectrum can be evaluated at horizon exit. However, if ϵ is not approximately constant, $\eta \not\approx 0$ and the slow-roll approximation is not valid due to the contribution of the second term in equation (1.29). If the approximation is used to evaluate the primordial power spectrum, the effect of η which encapsulates the rate of decrease of ϵ will be disregarded, and the resulting power spectrum will not be an accurate measure of the perturbations left over at the end of inflation.

Beyond slow-roll

If $\eta < 0$ at some point during inflation, this signifies a decrease in ϵ caused by the field slowing down on its potential. Looking again at equation (1.29), we can see that if ϵ is decreasing faster than a^3 (which is equivalent to $\eta = -3$), the second term starts to grow and there will be an additional contribution to $\mathcal{R}$ beyond that of the constant value at horizon exit. In these cases, it is vital to track both the constant mode and the decaying or growing mode, depending on the behaviour of ϵ, throughout inflation and to evaluate the primordial power spectrum only at the end of inflation, or when both modes have clearly stopped evolving. These models that do not qualify for the slow-roll approximation are called 'beyond slow-roll' models of inflation. The deceleration of the field is often caused by a flattening of the potential, and if the potential becomes completely flat, the scalar field enters a phase known as ultra-slow-roll (see e.g. [30]). See figure 1.2 for an illustration of an example potential. The value of η which signals ultra-slow-roll can be found from equation (1.21), where if $dV/d\phi = 0$, we have

$$-\frac{\ddot{\phi}}{\dot{\phi}H} = \epsilon - \frac{\eta}{2} = 3 \tag{1.30}$$

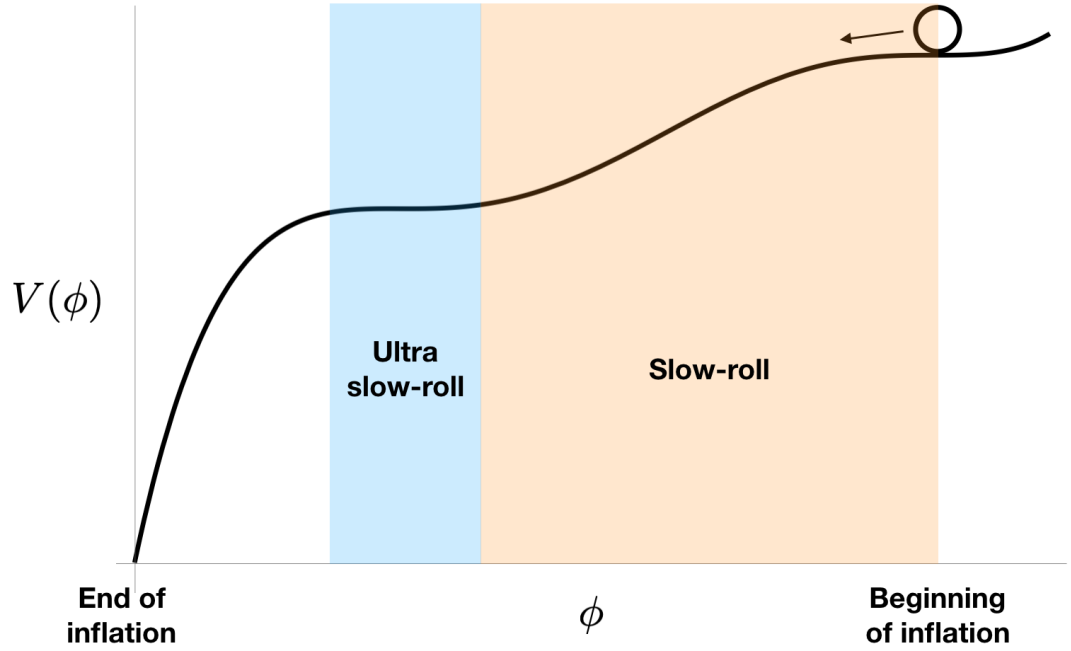

Figure 1.2: An illustration of an inflationary potential where the inflaton rolls from right to left in ϕ-space, down its potential $V(\phi)$. The gentle slope of the potential in the slow-roll phase flattens out to lead into a phase of ultra-slow roll. The inflaton will then need to speed up again so as to end inflation and for the Universe to reheat.

and hence $\eta = -6$.

The resulting power spectra from beyond slow-roll models are expected to deviate from scale-invariance due to ϵ not being constant as it is for slow-roll, and the extra boost in perturbations caused by the superhorizon growth quantified by the second term in equation (1.29). Instead, the power spectrum usually becomes scale-dependent and increases (or decreases if $\eta > 0$), forming some sort of feature or peak. Various models have been constructed that aim to produce a period of ultra-slow-roll in order to maximise the growth of superhorizon perturbations and hence the size of the peak in the power spectrum. For example, inflection points [11] and local maxima [31] in the potential have been explored as mechanisms for this to succeed.

Since the Planck measurements have shown that the primordial power spectrum is almost scale-invariant on large scales, ϵ must have been nearly constant at the time during inflation when those modes exited the horizon. This means that the simplest class of inflationary models, namely single-field, slow-roll models, are able to explain the measurements of the CMB. However, more complicated models could still reproduce these measurements even on just this range of scales, and narrowing down precisely which model is correct,

and ruling out others, will require extra constraints from different observables.

1.2.7 Accessing smaller scales

In terms of single-field inflation, distinctions between models can be inferred by looking at the tensor-to-scalar ratio, r, which quantifies the perturbations to the metric, i.e. the gravitational degrees of freedom, produced in a given model. These can be generated by quantum fluctuations of the gravitational field, analogous to the quantum fluctuations that produce the scalar perturbations, or else via mechanisms including particle production during inflation [32] or the presence of extra scalar fields [33]. Single-field models are capable of producing a wide range of values for r on the scales that Planck has explored, and there is a direct link, in the slow-roll approximation, between r and the energy scale of inflation at the time that a given mode leaves the horizon [20, 34]:

$$V = \left(1.88 \times 10^{16}\,\mathrm{GeV}\right)^4 \frac{r}{0.1}. \tag{1.31}$$

Currently, CMB polarisation experiments (combined with Planck data and Baryon Acoustic Oscillation results) have placed constraints in the range $r < 0.07 - 0.09$ at the Planck pivot scale $k_* = 0.05\,\mathrm{Mpc}^{-1}$ [35]. Indeed, a combination of the tensor-to-scalar ratio and the spectral index n_s at the scale k_* is the most common classification of viable inflationary models, for example see the 'zoo' plot in figure 1.3 which shows the Planck constraints in the blue, red and grey shaded regions overlaid with various models of inflation as functions of their predicted r and n_s. As constraints on r improve with data from upcoming experiments such as the Laser Interferometer Space Antenna (LISA), we should get a handle on the amount of tensor perturbations produced during inflation, and hence narrow down the class of models that can be responsible for the inflationary relics we have observed on large scales.

There is, however, still more that can be done by studying just the density perturbations left over at the end of inflation, if we can access smaller scales. CMB anisotropy measurements cannot help constrain smaller scales because fluctuations in temperature are washed out by Silk-damping [36]. As baryons began to fall into gravitational potential wells on smaller scales after recombination, photons (which are decoupled from the baryons at this point) diffused from hot regions to cold regions, evening out the temperature of the CMB and erasing perturbations on small scales.

Whilst we know that η can't deviate far from zero on large scales from CMB temperature anisotropies, there are very few constraints on smaller scales that restrict this possibility. Such a deviation from $\eta \approx 0$ would produce a feature or peak in the primordial

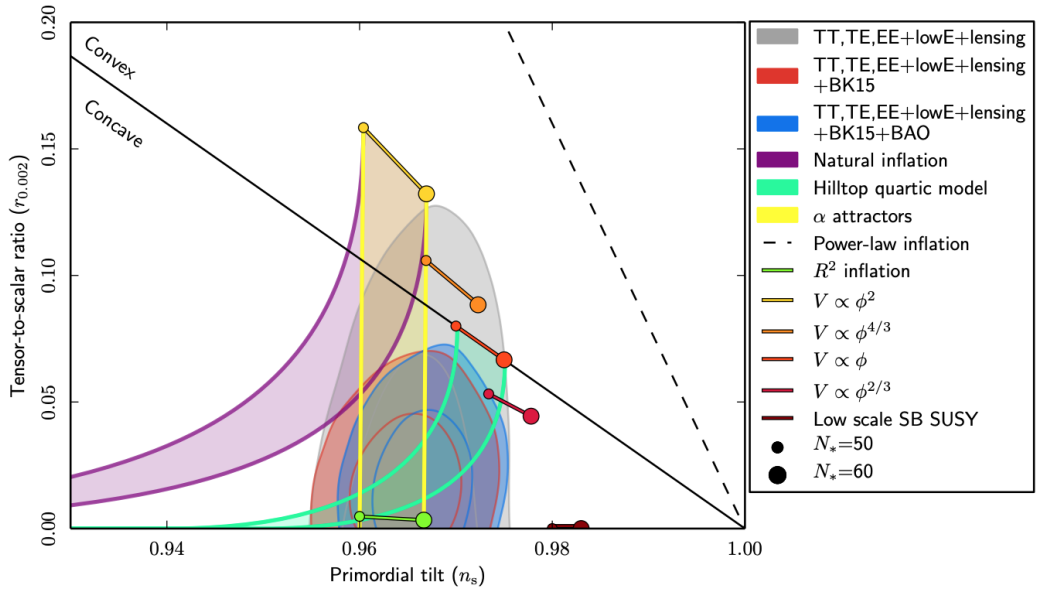

Figure 1.3: Planck 2018 constraints on the tensor-to-scalar ratio and the spectral index n_s at a scale of $k = 0.002\,\mathrm{Mpc}^{-1}$ are shown in the grey, red and blue contour regions. Various inflationary models' predicted range of values for r and n_s are overlaid. Figure credit: [1].

power spectrum, which would not only be very prescriptive for inflationary models, but would also imply the existence of interesting relics that could be observable today. In fact, the existence or lack of such observables are the best hope for constraining the primordial power spectrum, and hence inflationary dynamics, on small scales.

Non-Gaussianity and multifield inflation

Throughout this thesis we will assume that the probability density function of the initial perturbations can be described by a Gaussian distribution, which is a generic prediction of canonical slow-roll single-field inflation. Since a Gaussian distribution is entirely characterised by just two quantities, its mean and its variance, the statistics of the initial density field can be entirely described by a power spectrum. However, if the probability density function was non-Gaussian, higher-order statistics would be required to describe the initial density field. The next-simplest statistic to characterise deviations from Gaussianity is the bispectrum, which is the Fourier transform of the three point function. It is a function of three wavenumbers k_1, k_2, k_3, and the triangular configuration of these wavenumbers is known as the 'shape' of the bispectrum, whilst the amplitude is denoted by the quantity f_{NL}. Different early universe scenarios predict different shapes and amplitudes of the bispectrum, however the two most commonly considered shapes are:

- Local (also known as squeezed), which is maximised when $k_1 << k_2 \sim k_3$

- Equilateral, which is maximised when $k_1 \sim k_2 \sim k_3$.

There are two different amplitudes corresponding to these configurations, $f_{NL}^{\rm local}$ and $f_{NL}^{\rm equil}$. Of particular relevance for testing single-field inflation is $f_{NL}^{\rm local}$, since single-field models of inflation predict small ($<< 1$) values for this parameter according to the Maldacena consistency relation [37–39]. A detection of $f_{NL}^{\rm local} \sim \mathcal{O}(1)$ would strongly suggest that there was more than one scalar field present during inflation.

Furthermore, $f_{NL}^{\rm equil}$ is also usually expected to be small in canonical slow-roll single-field inflation according to the Maldacena consistency relation. However, beyond slow-roll and/or ultra-slow-roll regimes break this relation [40].

$f_{NL}^{\rm equil}$ is also able to test the presence of a non-canonical kinetic field term[2] c_s in the description of the inflaton's dynamics. Such a term can affect the primordial power spectrum as $\mathcal{P}_\mathcal{R} \propto 1/\epsilon c_s$ in single-field inflation, and can also produce a large bispectrum in the equilateral limit where $k_1 \sim k_2 \sim k_3$, i.e. in the case when all three modes exit the horizon at a similar time [41].

On large scales, the Planck 2018 analysis has set the most stringent constraints on these amplitudes: $f_{NL}^{\rm local} = -0.9 \pm 5.1$ and $f_{NL}^{\rm equil} = -26 \pm 47$ [42]. As the error bars are still large with respect to unity, these constraints are unable to confirm or rule out the simplest single-field models of inflation at this stage, but upcoming experiments such as improved CMB polarisation data from Simons Observatory [43], as well as large-scale structure probes including SPHEREX [44], will aim to improve constraints on the shape and amplitude of the bispectrum, targeting the threshold of $f_{NL}^{\rm local} < 1$. If a detection of $f_{NL}^{\rm local} > 1$ is made, this will confirm that multifield models are preferred over single-field models.

Finally, it is possible that the bispectrum is scale-dependent, and therefore there could be more information on small scales that is yet to be probed. However, this is even more experimentally difficult than probing the small-scale power spectrum, due to the non-linear physics that introduces late-time non-Gaussianities which can be difficult to separate from those of primordial origin [39].

[2] A non-canonical kinetic field term usually means that there are higher powers of the kinetic term $(\partial_\mu \phi)^2$ in the action for the inflaton which are dominant over the leading term.

1.3 Primordial black holes

There are multiple possibilities for inflationary relics, including cosmic strings or domain walls [45], that may be observable today due to, for example, their gravitational wave signatures or effects on the CMB. However, we will focus on relics that would be produced by a boost in the power spectrum on small scales. As detailed in the previous subsection, this could be due to a flattening of the inflationary potential in single-field inflation, or due to, for example, a second sufficiently light field that produces a blue-tilted power spectrum on small scales [46].

If there is a large peak in the primordial power spectrum at a given scale, perturbations of that scale will be particularly overdense or underdense. If they're overdense enough, they could form structures such as ultra compact mini haloes (UCMHs) [47] or primordial black holes (PBHs) immediately after inflation. The former are able to probe the primordial power spectrum and have interesting consequences for structure formation [48] and dark matter annihilation signals [49], however this thesis will focus on the latter.

1.3.1 Threshold for collapse

PBHs form when a critical threshold for collapse is reached by an overdensity re-entering the horizon post-inflation. The amplitude of the perturbation, and hence the value of the critical density threshold for collapse, needs to be defined in terms of the density contrast $\delta = \delta\rho/\rho_{bg}$ as opposed to the curvature perturbation defined as either $\mathcal{R}$ or ζ, because the amplitude of a given perturbation around horizon reentry is affected by superhorizon modes, the effects of which are not captured by $\mathcal{R}$ or ζ [50]. Furthermore, there is a non-linear relationship between δ and ζ which can have an order 1 effect on the threshold for collapse [50, 51], but we will not consider that here.

In order to consistently measure the scale and amplitude of a perturbation, the compaction function is used:

$$\mathcal{C}(r,t) = \frac{2(M(r,t) - M_{bg}(r,t))}{R(r,t)} \tag{1.32}$$

where $M(r,t)$ is the Misner-Sharp mass [52], i.e. the amount of energy within a sphere of areal radius $R(r,t) = a(t)r$, assuming spherical symmetry. Assuming spherical symmetry in this context is a good approximation because only the rarest peaks will be overdense enough to collapse, and rare peaks have been shown to be more spherical [53]. $M_{bg}(r,t)$ is the Misner-Sharp mass of the flat FRW background.

The maximum of the compaction function represents the scale at which the perturb-

ation differs most from the background value, and hence the scale r_m is defined as the scale of the perturbation which satisfies $\mathcal{C}'(r_m) = 0$. The horizon crossing time for the perturbation is given by the time at which

$$r_m(t_H)a(t_H)H(t_H) = 1. \tag{1.33}$$

The physical scale of the perturbation is therefore $R_m = r_m a(t)$, and the amplitude of the perturbation is defined as the excess density within the radius R_m at horizon crossing time t_H. This is calculated by finding the amount of excess mass over the spherical volume $V_m = 4\pi R_M^3/3$ by integrating the density contrast between 0 and R_m

$$\begin{aligned}\delta(r_m, t_H) &= \frac{1}{V_m}\int_0^{\pi}\int_0^{2\pi}\int_0^{R_m}\frac{\delta\rho}{\rho_{bg}}R^2 dR d\phi d\theta \\ &= \frac{3}{R_m^3}\int_0^{R_m}\frac{\delta\rho}{\rho_{bg}}R^2 dR.\end{aligned} \tag{1.34}$$

Simulations of the collapse of perturbations have been conducted to find the value, δ_c, which this average mass excess needs to reach in order for the perturbation to collapse and form a black hole. Therefore it is the averaged value $\delta(r_m, t_H)$ which should be compared directly with δ_c to determine whether that particular overdensity will collapse to form a black hole.

So far, we have discussed the value of the overdensity in terms of a general spherically symmetric curvature perturbation, however in reality, the perturbations will have distinct and possibly varying profiles which are related to the curvature perturbation $\mathcal{R}$ and can be characterised by a curvature profile $K(r)$ [54]. It is then not so trivial how to take the average of the energy density, and the choice of window function can affect the result [55, 56]. The integral between 0 and R_m in equation (1.34) is equivalent to applying a real-space top-hat window function to the curvature perturbation profile, since it just represents a hard cut-off at $r = R_m$. However, Gaussian window functions have also been suggested as appropriate, so as not to disregard overdensities just outside of the horizon which may too fall into the resultant black hole by the time it has actually formed.

The threshold for collapse depends on both the shape of the profile and the window function used to determine the mass excess associated with a given scale R_m. Latest results have shown that this value can vary between $\sim 0.45 - 0.65$ for collapse during a radiation-dominated era [57, 58].

If an overdense region meets these criteria, it will collapse within around a Hubble time to form a PBH. Once formed, PBHs decouple from the Hubble flow and evolve like matter. For the majority of this thesis, we refer to PBHs forming in a radiation-dominated era,

however there are other possibilities. In particular, a phase of early matter-domination will be explored in chapter 2. There are also more exotic PBH formation mechanisms, like the collapse of cosmic strings or the collision of domain walls, which we will not explore but have received attention in the literature, e.g. [59, 60] and references within [61].

Given that we don't know each of the individual profiles of the curvature perturbations, the power spectrum is usually used to quantify the average amplitude of the perturbations as a function of scale. Whilst the above formalisms inform the value of the critical threshold, as well as the appropriate window function according to the expected profile of the perturbations, below we will lay out some simplifications which will suffice given that we are working with the average value of the perturbations according to a given power spectrum. Work has been done to connect the form of the power spectrum with the curvature perturbation profile [62], moving towards a unified approach that takes into account the curvature profile, the window function, and a critical threshold calculated with consistent numerical simulations. These considerations in turn should predict an accurate distribution of PBH masses.

For simplicity, as will be followed in chapter 2, we will assume that the mass, and hence scale of the PBH formed, is simply related to the mass of the horizon at the time of reentry, given by [9, 10]

$$M_{\rm PBH} = \gamma M_{\rm hor} = \gamma \frac{3c^3 t_i}{4G}, \tag{1.35}$$

where γ is the fraction of the horizon mass that falls into the PBH found numerically to be ~ 0.2 and t_i is the cosmic time that the mode enters the horizon. i.e. the later the reentry, the larger the mode reentering, and hence the larger the black hole which is formed. The comoving scale that corresponds to each black hole mass is then found by scaling the particle horizon

$$d_{p,i} = \int_0^{t_i} \frac{a_i}{a} dt = \int_0^{t_i} \left(\frac{t_i}{t}\right)^{\frac{1}{2}} = 2t_i \tag{1.36}$$

at the time of reentry

$$t_i = \frac{4GM_{\rm PBH}}{3c^3} \tag{1.37}$$

by the change in the scale factor over the subsequent epochs. If the PBH formed in radiation-domination, as is the standard case, then the comoving scale today corresponding to a PBH that formed shortly after a horizon entry time of t_i is

$$k_i = \pi \left(\frac{t_{\rm eq}}{t_0}\right)^{\frac{2}{3}} \left(\frac{3c}{4GM_{\rm PBH} t_{\rm eq}}\right)^{\frac{1}{2}} \tag{1.38}$$

where $t_{\rm eq}$ is the cosmic time of matter-radiation equality.

1.3.2 Abundance of PBHs

According to the Press-Schechter formalism[3] [64], the number of black holes formed of a given mass is determined by integrating the distribution of density between δ_c and 1. If the overdensities are assumed to follow a Gaussian distribution, then this results in a mass fraction [9]

$$\begin{aligned}\beta(M_{\rm PBH}) =& \frac{2}{\sqrt{2\pi}\sigma(R)}\int_{\delta_c}^{\infty}\exp\left(\frac{-\delta^2(R)}{2\sigma^2(R)}\right)d(\delta(R)) \\ =& \text{erfc}\left(\frac{\delta_{\rm c}}{\sqrt{2}\sigma(R)}\right),\end{aligned} \tag{1.39}$$

where erfc is the complementary error function, $\sigma(R)$ is the variance of the density field at a given scale and we have included the Press-Schechter[4] factor of 2. The mass fraction is the proportion of the Universe that is in regions dense enough to form PBHs, and is defined at the time of formation. $\sigma(R)$ is related to the primordial power spectrum via

$$\sigma^2(R) = \int_0^{\infty}\widetilde{W}^2(kR)\mathcal{P}_{\delta}(k)\frac{dk}{k}, \tag{1.40}$$

where $\tilde{W}$ is the window function which smooths the perturbations over a given scale R. We will discuss the consequences of this relationship in chapter one.

1.3.3 Dark matter candidate

PBHs are interesting potential relics of inflation, but they were also first postulated as a dark matter candidate in the 1970s by George Chapline and Stephen Hawking among others [65, 66]. There are various motivations for a dark sector, the first evidence of which came in 1933 when Fritz Zwicky observed that galaxies in the Coma Cluster were moving too fast to be held together by the gravitational force accounted for by just the visible mass [67]. Furthermore, the stars and gas towards the outskirts of a disk galaxy should rotate around the centre at a slower rate than those nearer the centre, because the gravitational force decreases with distance from the centre. However, observations (notably first by Vera Rubin [68]) have shown (see [69] for an early review), that the velocity does not decrease with distance from the centre but instead levels off. This suggests that there must be more gravitational energy than the visible matter would imply, which in turn points to

[3]Note that there has been a lot work done on the correct formalism to use for calculating PBH abundance, as the Press-Schechter formalism was developed in the context of galaxy and cluster formation, not PBH formation. Alternatives include Peaks Theory and modifications thereof, see e.g. [63].

[4]Since the factor of 2 ends up in the argument of the complementary error function in the expression for the power spectrum, including it or not only changes the results for the power spectrum by 1-2%.

the fact that there must be some form of dark matter which only interacts gravitationally with the visible matter in these galaxies.

Over the last half-decade, this evidence has been corroborated by numerous observations, and the current observational value for the energy density of dark matter today is $\Omega_{\rm DM} = \rho_{\rm DM}/\rho_{\rm c} = 0.26$ [26]. However, whilst the existence of dark matter is widely agreed upon, its origin and form is yet to be determined.

Largely due to cosmological numerical simulations, 'cold' particulate dark matter models where the particle decouples from thermal equilibrium at a temperature below its mass ($m = k_B T/c^2$) and is therefore non-relativistic during structure formation have become favoured over 'hot' dark matter models where the particles instead exhibit relativistic velocities. This is because non-relativistic particles have a much shorter free-streaming length, and therefore structures form on much smaller scales. These can then merge and form larger structures in a hierarchical way, which matches much more closely sub-cluster structure and cosmic web formation that has been seen in galaxy surveys [70]. Hot dark matter, in contrast, has a much larger free-streaming length, and hence small-scale structure is washed out by the thermal motion of the particles [71].

Direct detection searches have therefore predominantly focused on particulate cold dark matter candidates. For example firstly, the Weakly Interacting Massive Particle (WIMP), which self-annihilates and is capable of producing the correct dark matter density today based on a mass-dependent cross-section of the order $10^{-26}{\rm cm}^3{\rm s}^{-1}$ [72]. Alternatively, the axion, which was originally motivated by wanting to solve the 'fine-tuning' problem in quantum chromodynamics which describes the strong force that affects quarks and gluons known as the strong CP problem e.g. [73]. However there have been a multitude of candidates put forward, such as sterile neutrinos [74], supersymmetric particles [75], topological defects [76], or even that modified gravity could be responsible [77]. Despite extensive searches for evidence of dark matter particles, there have been no confirmed detections as of yet.

PBHs also fulfil all of the necessary properties of dark matter, and unlike all of the examples just mentioned, require no new physics beyond the standard model. As we will see, they have had a resurgence in popularity in recent years due to possible signatures of their existence, but still various observational constraints threaten to rule them out as making up all of the dark matter. The proportion of the dark matter energy density that could be accounted for by PBHs today is described by the quantity $f_{\rm PBH} = \rho_{\rm PBH}/\rho_{\rm DM}$

which is related to the abundance at the time of formation β (equation (1.39)) by [9]

$$f_{\rm PBH} = \left(\frac{M_{\rm PBH}}{M_{\rm eq}}\right)^{-\frac{1}{2}} \frac{\beta(M_{\rm PBH})}{\Omega_{\rm DM}}, \tag{1.41}$$

assuming they formed during radiation-domination.

1.3.4 Observational direct detection constraints

Due to the fact that PBHs can in theory form with any mass and corresponding scale there need to be a variety of observational strategies for detecting them. Four of the main classes of constraints are plotted in figure 1.4 and described below.

They evaporate according to Hawking radiation, for which the timescale of evaporation depends on the mass [10]

$$t_{\rm ev} \sim \frac{G^2 M^3}{\hbar c^4}, \tag{1.42}$$

where G is Newton's gravitational constant, $\hbar$ is the reduced Planck's constant and c is the speed of light. Only PBHs lighter than $\sim 10^{15}\,$g would have totally evaporated by today. This means that on the low-mass end, tight constraints come from the lack of evidence for energy injection in the CMB due to their evaporation.

There are various microlensing constraints which are based on searches for cases of PBHs passing between ground-based telescopes and far-off stars. The PBH should act as a lens and cause the star to exhibit a distinctive feature in the light-curve. The HSC Subaru telescope has put constraints on PBH masses between $\sim 10^{-11} - 10^{-5}\,M_\odot$ that reach $f_{\rm PBH} \lesssim 2 \times 10^{-3}$ at their most stringent (for a monochromatic mass function) [2], whilst the Optical Gravitational Lensing Experiment (OGLE) [78], Expérience pour la Recherche d'Objets Sombres' (EROS) [79] and MACHO [80] have put constraints on PBH masses between $\sim 10^{-7} - 10^{-1}\,M_\odot$ with $f_{\rm PBH} \lesssim 10^{-2}$ for the mass where the constraints are strongest [81].

There are various constraints for high-mass PBHs, for example via supernovae lensing [82], heating of the stellar content in ultra-faint dwarf galaxies [83], or disruption of wide halo binaries [84]. However, the most stringent constraints on high-mass PBHs are due to accretion effects [85, 86]. PBHs are expected to accrete gas in the early Universe, and when radiation pressure overcomes gravity, some of the mass will be re-radiated. This radiation can then heat and ionise the intergalactic medium, which could change the temperature and the spectrum of the CMB locally. Depending on the number of PBHs, this would show up in CMB anisotropy and spectral distortion measurements from, for example, the

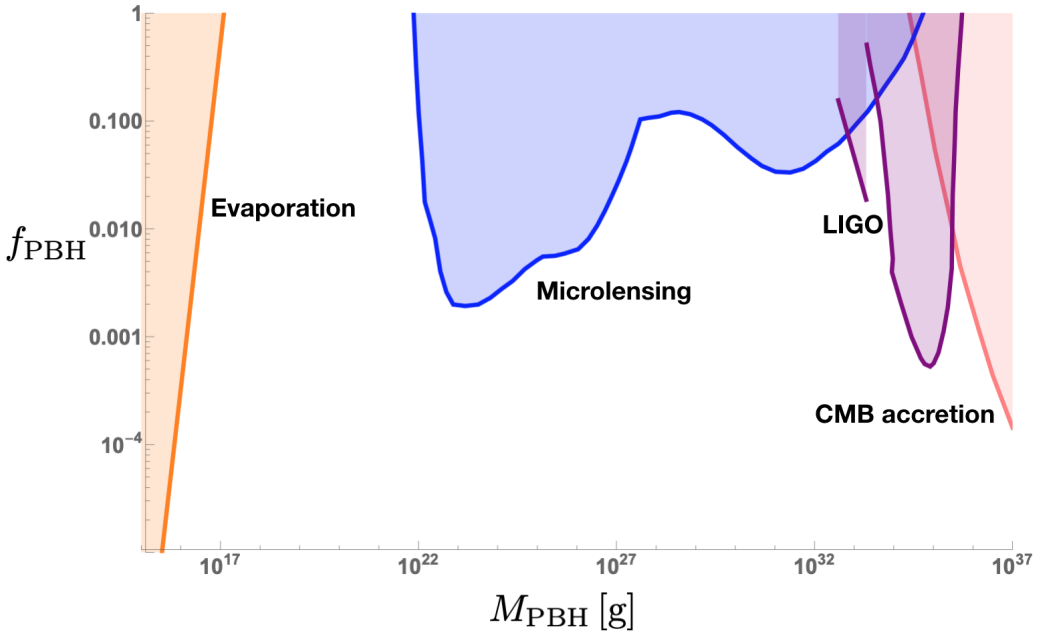

Figure 1.4: Direct constraints on $f_{\rm PBH}$ as a function of PBH mass due to the non-detection of PBHs from four of the main probes. On the low-mass end, evaporation constraints are shown in orange, the envelope of microlensing constraints from HSC Subaru as well as EROS/MACHO are in blue, sub-solar mass gravitational wave constraints from O2 LIGO and O1 LIGO constraints in the $1 - 10 M_\odot$ range are in purple, and constraints from CMB accretion are in pink. There are two different LIGO constraints due to separate analyses on the different mass ranges - a consistent analysis over the full range would result in a smooth constraint. Data taken from [2–4].

Planck Satellite and COBE/FIRAS (see section 1.4). Conservative constraints have been placed on the fraction of dark matter that could be made up of PBHs, with tighter limits possible as uncertainties in, for example, the radial profile of the accretion are better understood. For masses $\sim 1000\,M_\odot$, $f_{\rm PBH} \leq 10^{-2}$, and for $\sim 10^4\,M_\odot$, $f_{\rm PBH} \leq 10^{-4}$ [87].

Finally, there has been a resurgence in popularity of PBHs in recent years due to the detection of gravitational waves. LIGO/Virgo have detected more than 10 black hole merger events since the detectors came online in 2015 [88]. All of the black holes detected during the first two observing runs have been constrained to be in the mass-range $\sim 10 - 50\,M_\odot$. Black holes of these masses can be explained with astrophysical formation mechanisms, i.e. due to the collapse of stars. However, they could also be explained by PBH formation mechanisms, due to the fact that they can form at any mass in principle. Definitively distinguishing between the possible origins of these black holes will rely on theoretical advances including a better understanding of the mass function of PBHs [89–91], as well as the expected spins for primordial and astrophysical black holes, and the effects and/or likelihood of a merger history for primordial and astrophysical black holes.

Future 'extreme' detections will also help to distinguish between explanations for the black holes' origins. For example, if a black hole larger than $100\,M_\odot$ is observed, it will be very difficult to explain its origin with stellar collapse mechanisms, and subsequent mergers would be relied upon to produce a black hole with such a a large mass [92]. Similarly, black holes of masses below the Chandrasekhar limit of $\sim 1.4\,\mathrm{M}_\odot$ are difficult to produce via astrophysical mechanisms, since this limit represents the maximum mass that a white dwarf star can have before gravity will overcome pressure and initiate collapse to form a neutron star or black hole. Conversely this represents the minimum mass that a neutron star or astrophysical black hole can have. On the other hand, depending on the spread of the mass-distribution of a larger population of detections, it may become difficult to account for the mass-distribution through standard PBH formation mechanisms.

1.3.5 PBH mass function

Often PBHs are assumed, as we did above in section 1.3.1 for simplicity, to form with just a single mass, especially when calculating their constraints as a fraction of the dark matter budget. However, this is unphysical. As we will show in chapter 3, even for a monochromatic peak in the power spectrum, a spread of PBH masses are produced

according to the critical collapse formalism:

$$\frac{M_{\rm PBH}}{M_{\rm hor}} = \mathcal{K}\left(\delta - \delta_c\right)^{\gamma} \tag{1.43}$$

where $\mathcal{K} \approx 3.3$ which encodes the profile of the overdensity and $\gamma \approx 0.36$ are both calculated with numerical simulations. Note that this γ is not the same as defined in equation (1.35). If we then calculate the abundance of PBHs at the time of formation, β, for this distribution of masses, the integral in (1.39) becomes [93, 94]

$$\begin{aligned}\beta(M_{\rm PBH}) =& \frac{2}{\sqrt{2\pi}\sigma(R)}\int_{\delta_c}^{\infty}\frac{M_{\rm PBH}}{M_{\rm hor}}\exp\left(\frac{-\delta^2(R)}{2\sigma^2(R)}\right)d(\delta(R)) \\ =& \frac{2}{\sqrt{2\pi}\sigma(R)}\int_{\delta_c}^{\infty}\mathcal{K}\left(\delta-\delta_c\right)^{\gamma}\exp\left(\frac{-\delta^2(R)}{2\sigma^2(R)}\right)d(\delta(R)).\end{aligned} \tag{1.44}$$

Combining this with the definition of $f_{\rm PBH}$ in equation (1.41), we can write the PBH mass function as

$$\begin{aligned}f(M_{\rm PBH}) \equiv& \frac{1}{\Omega_{\rm CDM}}\frac{d\Omega_{\rm PBH}}{d\ln M_{\rm PBH}} \\ =& \frac{1}{\Omega_{\rm CDM}}\int_{-\infty}^{\infty}\frac{2}{\sqrt{2\pi\sigma^2(M_{\rm hor})}}\exp\left[-\frac{(\mu^{1/\gamma}+\delta_c(M_{\rm hor}))^2}{2\sigma^2(M_{\rm hor})}\right]\frac{M_{\rm PBH}}{\gamma M_{\rm hor}}\mu^{1/\gamma}\sqrt{\frac{M_{\rm eq}}{M_{\rm hor}}}d\ln M_{\rm hor},\end{aligned} \tag{1.45}$$

where $\mu \equiv \frac{M_{\rm PBH}}{\mathcal{K}M_{\rm hor}}$, and σ is the mass variance which is related to the power spectrum according to equation (1.40), which should be smoothed over the scale R_m, and $\Omega_{\rm CDM} = 0.26$ is the dark matter density today.

Furthermore, the primordial power spectrum itself cannot have a monochromatic peak [95, 96], and therefore there are observational consequences for both primordial power spectrum constraints and direct detection constraints on PBHs for the spread of overdensities at a particular scale. The non-detection constraints are shown in figure 1.4 for monochromatic mass functions. However these should really be recalculated individually for a realistic extended mass function [97, 98]. It has been shown that log-normal distributions fit the predictions of PBH-motivated inflationary models well, but other forms have also been postulated, including multiple peak forms [99].

Whilst most PBH mass ranges are now ruled out as being able to explain the entire dark matter budget, there is one remaining low-mass window, $10^{20} - 10^{24}\,$g, where this could still be possible. This window is observationally very difficult to probe due to the finite source effect and the wave optics effect which are both problematic for microlensing. The former is a problem when the source star is of a comparable apparent size to the Einstein radius of the lens, and hence assuming the source is much smaller than the lens for the point-source approximation is no longer valid. Only a small part of the star will

be magnified when the finite size of the source is taken into account, and this makes the constraints weaker [2]. The latter effect is important when the Schwarzschild radius of the lens (in the case that it's a PBH) becomes comparable to the wavelength of the light being observed by the microlensing experiment. The geometric optics approximation is valid when $r_s > \lambda$, but otherwise interference and diffraction effects occur due to the fact that the lens does not 'see' the full waveform of the light if it is too small, again resulting in weaker constraints as the magnification is generally lower [100].

However, indirect constraints due to constraints on the primordial power spectrum could rule out this window more easily.

If constraints on the primordial power spectrum can be placed such that the critical density threshold for collapse cannot be reached, then not only will PBHs making up all of the dark matter at that scale be ruled out, but PBHs of that mass existing at all will be disfavoured, and a more exotic mechanism for their formation would be required. This indirect method of constraining PBH abundance is therefore less model-independent than direct detection, but potentially more powerful.

1.4 Spectral distortions

There are various observational probes able to constrain different scales of the primordial power spectrum. As already discussed, Planck's measurements of the temperature (and polarisation) anistropies in the CMB have already tightly constrained the largest scales. Spectral distortions of the energy spectrum of the CMB are able to constrain smaller scales. They quantify deviations from the black-body temperature distribution of the CMB, caused by energy injection and removal from the plasma in the early Universe.

At high redshift, if some process changes the energy density ρ_γ of the photon-baryon plasma, the blackbody distribution of the CMB will also change momentarily. This is because the number density of photons N_γ would need to change according to $\Delta N_\gamma/N_\gamma \approx (3/4)\Delta\rho_\gamma/\rho_\gamma$ in order for the relations $\rho_\gamma \propto T^4$ and $N_\gamma \propto T^3$ to hold, and therefore for the distribution to be a blackbody. In addition, the shift in temperature required to maintain the shape of the blackbody distribution as well as the amplitude is given by $T' \approx T + (1/4)\Delta\rho_\gamma/\rho_\gamma$. This means that unless all three quantities shift simultaneously to maintain these relations, the blackbody distribution will be distorted [101].

Above redshift $z \sim 2 \times 10^6$, thermalisation processes were extremely efficient. Double Compton scattering, where two photons (and an electron) are emitted from an interaction between one photon and one electron, and Bremsstrahlung processes, caused by decelera-

tion of electrons, mediate the number of photons N_γ, whilst Compton scattering processes mediate ρ_γ by enabling the photons to diffuse their energy. Whilst these processes occur efficiently, any momentary distortion to the CMB blackbody is soon erased, and no memory of the energy injection or removal process is stored. However, as these processes becomes less efficient at lower redshifts, the blackbody distribution may not recover and a distortion could become detectable by experiments today.

There are two types of distortion that could contribute to $\Delta\rho_\gamma/\rho_\gamma$. Between redshifts $3\times10^5 < z < 2\times10^6$, double Compton scattering and Bremsstrahlung processes become less efficient, whilst Compton scattering remains active. This means that the energy of the photons is still being redistributed efficiently, and it's the number density of photons which can shift from equilibrium. This results in a chemical potential, or pure μ-distortion (and a temperature shift) of the CMB.

At lower redshifts, $z < 10^4$, Compton scattering becomes inefficient too and photons are unable to diffuse their energy. This causes a shift in the average energy density of the photons ρ_γ, and a Compton-y distortion. At intermediate redshifts, $10^4 < z < 3\times10^5$, a non-linear combination of the two distortions are produced. A thermalisation Green's function, G_ν, can be calculated by numerically solving the linear Boltzmann equation [102, 103]

$$\frac{\partial n_\nu}{\partial t} - H\nu\frac{\partial n_\nu}{\partial \nu} = \frac{\mathrm{d}n_\nu}{\mathrm{d}t}\Big|_{\mathrm{C}} + \frac{\mathrm{d}n_\nu}{\mathrm{d}t}\Big|_{\mathrm{DC}} + \frac{\mathrm{d}n_\nu}{\mathrm{d}t}\Big|_{\mathrm{BR}} + \frac{\mathrm{d}n_\nu}{\mathrm{d}t}\Big|_{\mathrm{S}}, \tag{1.46}$$

where the left-hand side describes the redshifting of the photons due to the Hubble expansion, and the right-hand side encapsulates the change in photon number due to Compton scattering, Double Compton scattering, Bremsstrahlung radiation and direct photon production (S). Then, given an energy injection history $Q(z)$ for some process of interest, the change in the intensity of photons $I_\nu = (2h\nu^3/c^2)n_\nu$ is given by [103]

$$\Delta I_\nu(z=0) = \int G_\nu(\nu, z', 0)\frac{\mathrm{d}(Q/\rho_\gamma)}{\mathrm{d}z'}\mathrm{d}z'. \tag{1.47}$$

If the transition between producing μ-distortions and y-distortions is treated as instantaneous so that mixed distortions are neglected, the shape of the distortion can be estimated and I_ν can be split up as

$$\Delta I_\nu \approx Y_\nu \Delta\rho_\gamma/\rho_\gamma\Big|_y + M_\nu\Delta\rho_\gamma/\rho_\gamma\Big|_\mu + G_\nu\Delta\rho_\gamma/\rho_\gamma\Big|_T \tag{1.48}$$

with Y_ν and M_ν the appropriate Green's functions for the low and high redshift processes relevant for y and μ distortions respectively. The total temperature shift, as opposed to

the distortion of the shape of the blackbody distribution, is accounted for in the last term of (1.48).

There are various causes of energy injection and/or removal from the plasma, including particle production or decay and scattering between baryons and dark matter [104, 105]. However, we will focus on spectral distortions caused by large scalar perturbations produced during inflation. A large boost in the primordial power spectrum at a particular scale or over a range of scales will lead to fluctuations in the density of the baryons and photons as a function of scale after reheating. This will mean that the photon distributions on different scales will be described by different blackbodies, and as those photons mix via Thomson scattering, a spectral distortion will be induced if Compton scattering, Double Compton scattering and Bremsstrahlung processes aren't efficient enough to bring them into equilibrium. The energy injection history for adiabatic perturbations described by a smooth power spectrum $P_{\mathcal{R}}(k)$ can be approximated by [101]

$$\frac{\mathrm{d}(Q/\rho_\gamma)}{\mathrm{d}z} \approx \frac{A^2}{aH}\frac{32c^2}{45\dot{\tau}}\int \frac{\mathrm{d}k}{2\pi^2}k^4 P_{\mathcal{R}}(k)e^{-2k^2/k_D^2}, \tag{1.49}$$

where $A \simeq (1+(4/15)R_\nu)^{-1} \simeq 0.9$, $R_\nu = \rho_\nu/(\rho_\gamma+\rho_\nu)$, $\dot{\tau} = \sigma_T N_e c$ is the rate of Thomson scattering, and k_D is the damping scale at recombination, i.e. the typical distance (k_D^{-1}) a photon can travel before being scattered. Equation (1.49) can then be inserted into equation (1.47) and integrated numerically. The final μ and y-distortions induced by the scalar perturbations can then be approximated by [101]

$$\begin{aligned}\mu &\approx \int_{k_{\min}}^{\infty}\frac{dk}{k}\mathcal{P}_{\mathcal{R}}(k)\,W_\mu(k),\\ y &\approx \int_{k_{\min}}^{\infty}\frac{dk}{k}\mathcal{P}_{\mathcal{R}}(k)\,W_y(k),\end{aligned} \tag{1.50}$$

with k-space window functions of the form

$$W_\mu(k) \approx 2.27\left[\exp\left(-\left[\frac{\hat{k}}{1360}\right]^2\Big/\left[1+\left[\frac{\hat{k}}{260}\right]^{0.3}+\frac{\hat{k}}{340}\right]\right)-\exp\left(-\left[\frac{\hat{k}}{32}\right]^2\right)\right], \tag{1.51}$$

$$W_y(k) \approx \frac{A^2}{2}\exp\left(-\left[\frac{\hat{k}}{32}\right]^2\right),$$

where $\hat{k} = k/1\,\mathrm{Mpc}^{-1}$ and $k_{\min} \simeq 1\,\mathrm{Mpc}^{-1}$. These window functions are calculated by integrating the energy injection history analytically for the y-distortion, and numerically for the μ-distortion, for details, see [106]. They up- and down-weight contributions at different wavenumbers according to the denominators in equation (1.51). Given a particular form for the power spectrum, this can then be used to compute the total induced

μ or y-distortion. Comparing this with observations then results in constraints on the primordial power spectrum.

Due to the fact that they are produced at very different times, the μ and y-distortions probe different k-ranges of the power spectrum. y-distortions place constraints on larger modes $k < 3\,\mathrm{Mpc}^{-1}$, whilst μ-distortions constrain the smaller scales, down to $k \sim 10^4\,\mathrm{Mpc}^{-1}$.

The Far-InfraRed Absolute Spectrophotometer (FIRAS) instrument on board the COBE satellite measured spectral distortions to be smaller than $\Delta\rho_\gamma/\rho_\gamma < 6 \times 10^{-5}$ [14], and a proposed future detector such as PIXIE [107], or a more recent proposal [18] aims for constraints of $\Delta\rho_\gamma/\rho_\gamma < 8 \times 10^{-9}$.

In chapter 3, we will see that these observations place constraints on the primordial power spectrum amplitude of the order $\mathcal{P}_\mathcal{R} \lesssim 10^{-4} - 10^{-5}$.

1.5 Gravitational waves

Gravitational waves are tensor perturbations to the background metric, and can therefore be described by a modification to the Minkowski metric $\eta_{\mu\nu}$ with the spacetime metric given by

$$g_{\mu\nu} \approx \eta_{\mu\nu} + h_{\mu\nu}, \tag{1.52}$$

and $h_{\mu\nu}$ the metric perturbation. They are produced by the mass quadrupole moment changing in time, and can be sourced by, for example, black hole or neutron star mergers, supernova explosions, continuous waves from e.g. pulsars, or via early universe processes [108]. Individual events, such as the black hole and neutron star mergers already detected by LIGO/Virgo have distinctive signals that can be searched for in the data with templates. Alternatively, the accumulation of unresolved gravitational waves will contribute to a stochastic background of gravitational waves. This will include signals from astrophysical sources that are too faint to be resolved in the data, but it could also include a contribution from primordial gravitational waves. In this thesis, we will specifically focus on those induced by second-order scalar perturbations, since the large scalar perturbations required for primordial black hole production could produce gravitational waves that should be detectable with current and future experiments.

1.5.1 Stochastic background

A gravitational wave background could be sourced by large scalar perturbations reentering the horizon after inflation. Gravitational waves would be necessarily produced as a second-order effect and these would contribute to the stochastic gravitational wave signal. Measuring this signal in turn is then able to probe the scalar primordial power spectrum on very small scales.

Following [109], the equation of motion for the metric perturbation at a given scale $\mathbf{k}$ is

$$h''_{\mathbf{k}} + 2aHh'_{\mathbf{k}} + k^2 h_{\mathbf{k}} = 4S_{\mathbf{k}} \tag{1.53}$$

where $S_{\mathbf{k}}$ is the source term which is a function of the scalar perturbations. By solving for $h_{\mathbf{k}}$, the power spectrum of the tensor perturbations can be found;

$$P_h(\tau, k) = 4\int_0^\infty dv \int_{|1-v|}^{1+v} du \left(\frac{4v^2 - (1+v^2-u^2)^2}{4vu}\right)^2 I^2(v, u, k\tau)\mathcal{P}_{\mathcal{R}}(kv)\mathcal{P}_{\mathcal{R}}(ku), \tag{1.54}$$

where $u = |\mathbf{k} - \tilde{\mathbf{k}}|/k$, $v = \tilde{k}/k$ and $\tilde{k}$ is the wavenumber corresponding to the scalar source. $I(v, u, k\tau)$ is a highly oscillating function which contains the source information and the Green's function which solves for $h_{\mathbf{k}}$. The observational quantity related to this power spectrum is the energy density of gravitational waves given by

$$\Omega_{GW}(\tau, k) = \frac{\rho_{GW}(\tau, k)}{\rho_{tot}(\tau)} = \frac{1}{24}\left(\frac{k}{aH}\right)^2 \mathcal{P}_h(\tau, k). \tag{1.55}$$

If we assume that the entire contribution to any stochastic background detection is from the tensor power spectrum in equation (1.54), then constraints on the stochastic background can be translated to constraints on the scalar power spectrum. This is a conservative constraint, as there may be other unresolved astrophysical contributions to the signal. If a detection is made, as opposed to an upper limit on the amplitude from non-detection, spectral information of the signal will be required to distinguish between the possible sources.

1.5.2 Possibilities for detection

The stochastic background could either be detected directly by gravitational wave observatories, or by the monitoring of stable and rapidly rotating pulsars, namely millisecond pulsars. See e.g. [110] for a review.

So far, the best constraints on the primordial power spectrum due to secondarily produced gravitational waves have come from pulsar timing array (PTA) experiments. The

beams of radio waves given off by pulsars as they spin can be picked up by ground-based radio telescopes. Averaged over many rotations, beams arrive from individual pulsars with a very consistent time period. Whilst factors such as changes in the interstellar medium or orbital motion can cause individual pulsar's average time periods to vary, gravitational waves, as well as local gravitational effects from planets in the Solar System, will affect all pulsars in the same way. These 'global' effects are also expected to have a much longer time period than the time period of beams from millisecond pulsars, meaning that averaging over many rotations allows precise measurements of the longer time-scale effects. Signals from multiple pulsars can be cross-correlated so as to subtract astrophysical effects on individual pulsars and be left with just the effects that are common to all pulsars. Furthermore, each of the different possible sources affecting all of the pulsars should exhibit distinct signals in the residuals and can therefore be separated out. However, it was shown recently [111] that Solar System effects can mimic a gravitational wave background signal and therefore need to be treated carefully.

The North American Nanohertz Observatory for Gravitational Waves' (NANOGrav) 11 year observation of 45 pulsars [111], the European Pulsar Timing Array (EPTA) data from 18 years of observation of 6 pulsars [112], and the Parkes Pulsar Timing Array (PPTA) observations of 20 pulsars [113] have all put limits on the amplitude of the gravitational wave background at frequencies of the order of $10^{-8}\,\mathrm{Hz}$. This corresponds to probing scales around $k \sim 10^{6}\,\mathrm{Mpc}^{-1}$ of the primordial power spectrum. Furthermore, the upcoming Square Kilometre Array (SKA) radio experiment will monitor up to 100 pulsars [114] and improve constraints on a commensurate frequency range by around an order of magnitude.

Finally, looking to the future, the Laser Interferometer Space Antenna (LISA) has been designed to directly detect high-frequency gravitational wave signals, especially from the mergers of supermassive black holes. However, scalar perturbations on scales of the order $k \sim 10^{12}\,\mathrm{Mpc}^{-1}$ would contribute to a secondarily produced gravitational wave background on the same frequency range, meaning that LISA will be able to constrain the primordial power spectrum on very small scales. These constraints will be explored in chapter 3.

1.6 21cm

Finally, 21cm observations may also be able to constrain the primordial power spectrum on scales down to $k \sim 0.1\,\mathrm{Mpc}^{-1}$ from the ground, or down to $k \sim \mathcal{O}(10)\,\mathrm{Mpc}^{-1}$ from

space. After the photons decoupled from the baryons during recombination and were able to free-stream, the Universe went totally dark in an epoch known as the Dark Ages. The baryonic matter was made up almost entirely of neutral hydrogen, and therefore the only observational signature from this epoch is due to the fact that neutral hydrogen undergoes a spin-flip transition at a characteristic wavelength of $\lambda = 21$cm. This occurs in three circumstances:

- spontaneous emission; a random, quantum effect which causes the bound electron in a hydrogen atom to drop from its excited state to its lower energy state, emitting a photon with energy equal to the difference between the two states (in the case of neutral hydrogen, at a wavelength of 21cm). This process is independent of the radiation field which is incident upon the hydrogen.

- stimulated emission; when an incoming photon interacts with a hydrogen atom and causes another photon to be emitted at a wavelength of 21cm, and for the atom to drop from its excited to its ground state.

- stimulated absorption; when a hydrogen atom absorbs a photon at a wavelength of 21cm and this causes it to jump from its ground to its excited state.

The photons that are incident on the neutral hydrogen during the Dark Ages are the free-streaming CMB photons, which act as a backlight to the distribution of the hydrogen during this epoch. The distribution of the neutral hydrogen is in turn a tracer of the matter distribution, and hence a tracer of the underlying dark matter distribution. Since the densities remain linear during this epoch (redshift ~ 50), this distribution can therefore be related to the primordial density distribution via a linear transfer function. Thus, a connection is made between the 21cm signal during the Dark Ages and inflationary dynamics and relics, via the primordial power spectrum.

The distribution of neutral hydrogen is mapped out by measuring the ratio of neutral hydrogen in its ground versus its excited state. This is defined by

$$\frac{n_1}{n_0} = 3e^{-\frac{T_*}{T_s}} \tag{1.56}$$

where n_1 and n_0 are the number densities of neutral hydrogen in their excited and ground states respectively, $T_* = 0.068K$ is the temperature corresponding to the energy difference between the ground and excited states, and T_s is the spin temperature. See e.g. [115] for a review of 21cm cosmology.

The number densities of each state are balanced according to

$$n_0(n_H\kappa_{01} + B_{01}u_\nu) = n_1(A_{10} + B_{10}u_\nu + n_H\kappa_{10}), \tag{1.57}$$

where $n_H = n_0 + n_1$ is the total number density of neutral hydrogen, A_{10} is the Einstein coefficient which quantifies the probability of spontaneous emission, B_{10} is the probability of stimulated emission and B_{01} is the probability of stimulated absorption. The CMB photons are described by the radiation field $u_\nu = 2k_\mathrm{b}T_\gamma/\lambda^2$ which is the blackbody intensity in the Rayleigh-Jeans tail of the CMB with k_b the Boltzmann constant and $\lambda = 21\,\mathrm{cm}$ the wavelength of the radiation [116], and κ_{01}, κ_{10} are the collisional rate coefficients, which describe the exchange rates between states due to hydrogen-hydrogen collisions. Note that at lower redshifts, $z < 50$, hydrogen-electron collisions also become important [117].

Provided that the temperature of the neutral hydrogen, T_gas, and the CMB, T_CMB, are much greater than T_* such that $e^{-\frac{T_*}{T_s}} \approx 1 - T_*/T_s$, equation (1.57) can be rewritten in terms of these three temperatures:

$$T_s = T_\mathrm{CMB} + (T_\mathrm{gas} - T_\mathrm{CMB})\frac{n_H\kappa_{10}}{n_H\kappa_{10} + A_{10}\frac{T_\mathrm{gas}}{T_*}}. \tag{1.58}$$

Putting in the time-dependence of $T_\mathrm{CMB} \propto (1+z)$ and $T_\mathrm{gas} \propto (1+z)^2$, we can see that the spin temperature and the gas temperature drop below the CMB temperature after redshift $z \sim 150$, when there are no longer enough free electrons to scatter off the gas and CMB photons, which was previously keeping them in equilibrium. After around $z \sim 30$, the Hubble expansion becomes dominant over collisions and hence the spin temperature returns to align with the CMB temperature. This means that there is a window $z \sim 30 - 150$ where it is possible for a significant number of CMB photons to be absorbed by the neutral hydrogen, and hence a 21cm absorption signal should be present.

The spin temperature is not directly observable, but the contrast between it and the CMB temperature is, which is what we define as the brightness temperature:

$$T_{21} = \tau\frac{(T_s - T_\mathrm{CMB})}{1+z} \tag{1.59}$$

with τ the optical depth which quantifies how many photons will get scattered by the local environment and is given by

$$\tau = \frac{3c\lambda^2(z)hA_{10}n_H}{32\pi k_b T_s H(z)} \tag{1.60}$$

with λ the wavelength of the 21cm radiation. The redshift dependence of T_{21} is shown in figure 1.5 for four different models. The absorption signal from the Dark Ages as discussed above is shown by the trough at high redshift. Then, as the first stars start to form, Lyman-alpha photons are produced which cause the spin temperature to once again follow the gas temperature and hence there is another trough in the brightness temperature which represents the epoch of Cosmic Dawn. As the first sources start to heat up the gas, the

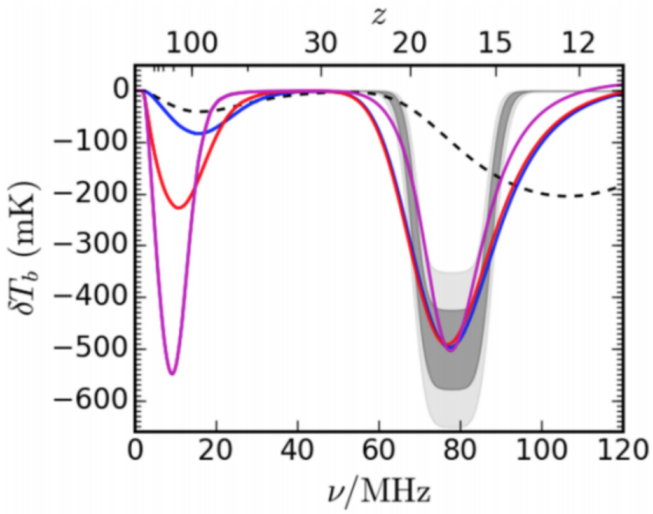

Figure 1.5: The 21cm brightness temperature as a function of redshift/frequency. The dashed black line is a standard model that fits with ΛCDM predictions, the high redshift trough being that of the Dark Ages absorption signal and the low redshift trough being that of the Cosmic Dawn. The three coloured lines represent different models that would fit the low redshift EDGES prediction, but could be distinguished by differing Dark Ages signals at high redshift, detectable from the Moon. Figure credit: [5].

spin temperature follows the gas temperature until it is much larger than $T_{\rm CMB}$ and there is a 21cm emission signal. As reionization subsequently sets in, the neutral hydrogen is ionized and the 21cm signal vanishes. The time evolution of the 21cm signal across these incredibly physics-rich epochs in cosmic history allows for tomographic analysis in order to learn about either the matter distribution during the Dark Ages, or about the first sources and epoch of reionization (EoR). Given that the 21cm radiation will simply be redshifted as $(1+z)$ as it travels towards us, it is possible to identify the redshift slice from which the signal originated. Adding up the contributions from 21cm photons with wavelength $\lambda_{21}(1+z)$ from all directions will probe the 'global' signal as plotted in figure 1.5. Alternatively, interferometry can be used to also measure the scale-dependence of the signal and build a 21cm power spectrum for each redshift slice. This is the observable which will map out the distribution of neutral hydrogen as a function of scale and lead to constraints on the primordial power spectrum.

The spatial fluctuations in the 21cm signal δ_{21} are a result of fluctuations in the baryons δ_b, the gas temperature δ_T, the free electron fraction δ_x (or the neutral hydrogen fraction), the Lyman-alpha photons δ_α, as well as the peculiar velocities of the hydrogen clouds δ_v.

The isotropic 3d 21cm power spectrum at a particular redshift is then defined by

$$\langle \delta_{21}(\mathbf{k_1})\delta_{21}(\mathbf{k_2})\rangle = (2\pi)^3\delta^D(\mathbf{k_1}-\mathbf{k_2})P_{21}(\mathbf{k_1}). \tag{1.61}$$

This is the quantity that we will be interested in for tracing the primordial power spectrum in chapter 4.

1.6.1 Observations

21cm photons emitted during the Dark Ages will reach ground-based detectors in the radio band, at a frequency of $f_{21} = c/(0.21(1+z))\,\mathrm{Hz}$. Various radio telescopes have been constructed in order to detect the 21cm signal, for example EDGES and its reported detection of the global signal from Cosmic Dawn [118]. In order to measure the spatial fluctuations however, radio interferometery is required so as to measure the signal at a given redshift as a function of scale.

In a radio array, each pair of antennae forms a baseline. The interference pattern of the signal that each pair of antennae detects results in a measurement of the intensity of the signal for one Fourier mode, corresponding to one baseline, after careful cross-correlation.

The Precision Array for Probing the Epoch of Reionization (PAPER) [119] the Murchison Widefield Array (MWA) [120] and the Low Frequency Array (LOFAR) are radio arrays that have started to place constraints on the 21cm power spectrum at redshifts up to $z \sim 12$ [121]. The Hydrogen Epoch of Reionization Array (HERA) is likely to be the first radio interferometer that may be able to make a detection of the 21cm power spectrum during the Epoch of Reionization. It is an array of 350 14m dishes, that will probe down to scales of $\sim 0.1\,\mathrm{Mpc}^{-1}$. HERA is thought of as a precursor to the Square Kilometer Array (SKA), which will aim to probe slightly smaller scales and similar redshifts but with greater sensitivity due to its larger collecting area.

Furthermore, in order to reach the higher redshifts for the Dark Ages cosmological signal, a space-based radio telescope will be required. This is because the 21cm signal from the Dark Ages will have been redshifted to frequencies below 50 MHz, and the Earth's ionosphere will reflect radio waves of frequencies lower than this back into space, meaning that the signal can never reach a ground-based detector. Various proposals for space-based interferometers have been made [122, 123], including a space-based fleet of satellites [124] or a lunar array on the far-side of the Moon [125]. The Netherlands-China Low-Frequency Explorer (NCLE) has already landed a probe on the lunar far-side with a prototype radio antenna on board that will be sensitive to $80\,\mathrm{kHz} - 80\,\mathrm{MHz}$ signals, and will provide

an excellent proof of concept that these types of observations are possible, and that the investment in such large-scale missions is worthwhile.

In addition, due to the fact that the smallest scale that an interferometer can be sensitive to is set by the largest separation between pairs of antennae $R_{\max}$,

$$k_{\max} \propto \frac{2\pi R_{\max}}{\lambda(z)}\,\mathrm{Mpc}^{-1}, \tag{1.62}$$

with $\lambda(z)$ the redshifted wavelength of the 21cm radiation, it will be possible to target much smaller scales with space-based interferometers, as they can have much longer baselines due to the lack of restrictions on their footprint which limit ground-based detectors.

However, there are many observational challenges that must be overcome in order for these detections to be realised. Foreground contamination is the primary obstacle for 21cm observations. The temperature of astrophysical sources including both unresolved and resolved extragalactic point sources, as well as galactic synchrotron radiation is of the order 100 K, whereas the 21cm brightness temperature signal is of the order 10^4 times smaller, at the mK level. Removing these foregrounds accurately and being sure that the remainder of the signal is purely the 21cm contribution is a very difficult task, to which many groups have concentrated a lot of effort [126–128].

It has been noted that the anisotropic power spectrum may enable a circumvention of foreground cleaning, as most of the foregrounds appear in a 'wedge' shape in the line-of-sight vs perpendicular ($k_\parallel$ vs $k_\perp$) plane [127]. This means that the foregrounds can be more simply avoided as opposed to subtracted from the anisotropic power spectrum signal, given in terms $k_\parallel$ and $k_\perp$. Whilst the signal-to-noise suffers dramatically at the loss of so many independent modes, the constraints and/or detection may be more robust [129].

Instrumental difficulties are hoped to be separate from the issue of foreground removal, since it should be possible to remove instrumental noise by cross-correlating detections at different times, whilst the foregrounds are largely time-independent [130].

If PBHs exist, there are observational signatures in the 21cm power spectrum that should be detectable with future experiments. For example, the accretion of matter onto PBHs would effect the spin temperature, and the Poisson distribution of the PBHs would also be present in the 21cm signal. In chapter 4, we will investigate the inclusion of the primordial fluctuations required to form the PBHs in the first place in the 21cm signal, and note that due to the small scales at which these signatures show up, a space-based radio telescope would be required for their detection.

Chapter 2

Extreme Scenarios: tightest possible constraint on the power spectrum due to primordial black holes

Philippa S. Cole[1] and Christian T. Byrnes[1]

[1] Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, United Kingdom

2.1 Introduction

Primordial black holes (PBHs) can form from the collapse of large density fluctuations in the early Universe. If an overdensity of order unity in a given region reenters the Hubble sphere[1] after inflation, then the region collapses to form a PBH with mass roughly equal to the mass within the Hubble sphere.

The abundance of PBHs in our observable Universe today can constrain the primordial power spectrum, and hence models of inflation, on scales much smaller than are accessible via observations of the CMB and LSS (which provide the tightest constraints on the primordial power spectrum for scales between $k \sim 10^{-3} - 1\,\mathrm{Mpc}^{-1}$). Constraints on the abundance of PBHs are very tight at the time of their formation, due to their gravitational effects and the consequences of their evaporation if they were sufficiently light [10]. See

[1] 'Hubble sphere' is used interchangeably with 'horizon' throughout.

[131] for some of the most up to date constraints on the abundance of PBHs[2]. Despite constraints on PBHs being tight, the fact that PBH abundance and the power spectrum are related logarithmically during radiation domination means that even massively tightening the PBH abundance constraints does not translate into a great improvement on the constraints on the primordial power spectrum. Here we derive the tightest constraints possible on the primordial power spectrum given the most extreme constraints on the abundance of PBHs, i.e. that there are none.

Apart from providing the tightest possible future constraints on the primordial power spectrum[3], this extreme assumption is motivated by an argument that the decay of a PBH would destabilise the Higgs vacuum; hence PBHs of masses small enough that they would have decayed by today can't ever have formed [6, 7, 132]. If they had they would have seeded the decay of the metastable Higgs vacuum, the Universe would have tunnelled to the true vacuum, and hence been destroyed.

Since the logarithmic relation between PBH abundance and the power spectrum is the main cause of the suppression of the constraint, the most effective way of improving the constraint is by looking at scenarios where the power spectrum amplitude is more sensitive to the PBH abundance. This is the case during an early matter-dominated phase (see [133] and [134] for motivations) where the relation between abundance and power spectrum is instead given by a power law [135, 136]. This means that for an observed abundance of a particular PBH mass, the constraint on the primordial power spectrum is tightened by many orders of magnitude, suggesting that the best constraint possible will come from such a scenario where the equation of state is at its minimum non-negative value, $\omega = 0$. Models with $-\frac{1}{3} < \omega < 0$ are rarely considered, and $\omega < -1/3$ corresponds to inflation. In terms of the worst constraints, according to [137], the value of ω for which it is most difficult to produce PBHs is $\omega = 1/3$ due to the value of critical overdensity being at its maximum. The threshold density then decreases as ω approaches 1. However, this is different to the results of, for example [10] and [138] whose values of critical overdensity increase with ω (for $0 < \omega < 0.6$ in the latter case - other values of ω were not simulated), suggesting that the larger ω is, the more difficult it is to produce PBHs.

Generally, calculations done in the radiation-dominated scenario assume spherical symmetry of the collapsing region (which makes collapse as likely as possible) but those done in the matter-dominated scenario do not [139]. Full numerical simulations of the collapse of density fluctuations to form PBHs are required to gain a complete understanding of the

[2] Presented in the context of whether PBHs can make up all of dark matter.

[3] Assuming radiation domination and Gaussian initial perturbations.

process. The critical overdensity required for a region to collapse and form a PBH also depends on the density profile. Additionally, phenonema such as non-Gaussianity could have a significant influence on the abundance of PBHs since they are formed from rare, large overdensities. Such large fluctuations are susceptible to changes in the tail of the fluctuation distribution caused by the amount of non-Gaussianity present [140]. We will not show explicitly the effects of non-Gaussianity on our results but its potential effect should be kept in mind. Furthermore for simplicity, we will assume that all PBHs form with the same mass for a given time. In reality, PBH constraints depend on the PBH mass function (see [94] for a recent update) which would affect the scale at which primordial power spectrum constraints are correlated to.

The paper is laid out as follows: in section 2.2 we will discuss the argument in [6, 7] for zero PBHs due to the Higgs instability and calculate constraints on the primordial power spectrum in this case. In section 2.3 we will look at motivations for a matter-dominated phase and find constraints on the primordial power spectrum for different durations of matter domination prior to BBN. Finally, in section 2.4 we will combine these two frameworks and see the result of their co-existence.

2.2 No PBH formation during radiation domination

It is believed that the electroweak vacuum is metastable, depending on the mass of the top quark, with a lifetime longer than the present age of the Universe [141, 142]. The notion that impurities initiate phase transitions gives rise to the idea that natural inhomogeneities such as PBHs may be capable of seeding rapid vacuum decay from the metastable vacuum to the true vacuum. If this had happened before today, the process would have had catastrophic consequences for the Universe. This is proposed and explored in [6, 7, 143, 144].

The conclusion drawn from this argument is that since the Universe has not been destroyed, no PBHs with masses small enough such that they would have already decayed can ever have formed. Assuming that the PBH mass, $M_{\rm PBH}$, is of the same order as the horizon mass, the relationship between PBH mass and horizon entry time, t_i, is given by

$$M_{\rm PBH} = \gamma 10^{15} \left(\frac{t_i}{10^{-23}\text{s}} \right) \text{g}, \tag{2.1}$$

where γ is the ratio between horizon mass and PBH mass [10]. The mass of a PBH which would just be decaying today is around 10^{15}g [10] so from equation (2.1) we can say that PBHs which could have formed at or before $\sim 10^{-23}$ seconds would have catalysed the

rapid vacuum decay, and therefore never existed. We will use this bound on abundance of PBHs being zero to find the tightest possible constraint on the primordial power spectrum originating from the non-detection of PBHs.

2.2.1 Constraint relations

The abundance of PBHs is usually described by the PBH mass fraction:

$$\beta(M_{\rm PBH}) = \frac{\rho_{\rm PBH}(M_{\rm PBH})}{\rho_{\rm tot}}, \tag{2.2}$$

which denotes what fraction of the total energy density of the Universe is contained in regions overdense enough to generate PBHs, measured at the time of their formation, where $\rho_{\rm PBH}$ is the energy density contained within PBHs, and $\rho_{\rm tot}$ is the total energy density of the Universe. During radiation domination, PBHs form shortly after horizon entry. As will be seen later, during matter domination PBH formation occurs a significant time after horizon entry.

Assuming radiation domination, in order for a region of space-time to collapse and form a PBH, the smoothed density contrast at horizon crossing, $\delta(R)$, needs to exceed some critical level of over-density, δ_c, which is of order 1. If the initial density perturbations have a Gaussian distribution then the pdf of the smoothed density contrast is given by:

$$P(\delta(R)) = \frac{1}{\sqrt{2\pi}\sigma(R)}\exp\left(\frac{-\delta^2(R)}{2\sigma^2(R)}\right) \tag{2.3}$$

where $\sigma(R)$ is the mass variance:

$$\sigma^2(R) = \int_0^\infty \widetilde{W}^2(kR)\mathcal{P}_\delta(k)\frac{dk}{k}. \tag{2.4}$$

$\mathcal{P}_\delta(k)$ is the primordial power spectrum of δ at horizon entry, describing how the overdensities and underdensities are distributed according to scale, $\delta = (\rho - \bar{\rho})/\bar{\rho}$ is the comoving density contrast and $\widetilde{W}$ is a smoothing function [145].

The PBH mass fraction in equation (2.2) is related to the probability density function as [9]

$$\begin{aligned}\beta(M_{\rm PBH}) &= \frac{2}{\sqrt{2\pi}\sigma(R)}\int_{\delta_c}^\infty \exp\left(\frac{-\delta^2(R)}{2\sigma^2(R)}\right) d(\delta(R)) \\ &= \text{erfc}\left(\frac{\delta_{\rm c}}{\sqrt{2}\sigma(R)}\right),\end{aligned} \tag{2.5}$$

where Erfc is the complementary error function and we have included the Press Schechter factor of 2. By inverting this expression, it is possible to find constraints on the mass

variance (and hence power spectrum via $\sigma^2 \sim \mathcal{P}$) given constraints on β. We will therefore be using this expression:

$$\mathcal{P}_\delta = \left(\frac{\delta_c}{\sqrt{2}\text{erfc}^{-1}(\beta)} \right)^2 \tag{2.6}$$

to plot our constraints on the power spectrum[4].

In order to construct a constraint on β that represents there being zero PBHs, or more precisely less than one PBH per current horizon volume, for a certain range of masses, we can model the observable universe as a cube of volume L^3, made up of N_l smaller cubes each with volume l^3. These small cubes represent the size of the patches that may have collapsed to form black holes in the early Universe. In our observable Universe today, if there is just one black hole that formed in the early Universe, then a patch of size l will have been overdense enough to collapse at the time that the patch reentered the horizon.

If less than one out of all of these patches (i.e. none of them) contain a black hole then

$$\beta < \frac{1}{N_l} = \left(\frac{l}{L} \right)^3 . \tag{2.7}$$

Our constraint on the primordial power spectum therefore becomes:

$$\mathcal{P}_\delta < \left(\frac{\delta_c}{\sqrt{2}\text{erfc}^{-1}(\frac{1}{N_l})} \right)^2 . \tag{2.8}$$

2.2.2 Relevant scales for Higgs stability argument

As we saw from equation (2.1), a PBH with mass 10^{15}g formed at 10^{-23}s is the largest and latest PBH that could have seeded rapid vacuum decay. We want to work out the physical size of the overdense region that would have needed to collapse to form a black hole of such mass, how large that region has expanded to today, and how many patches of that size there are in the Universe today. This will provide the threshold scale for the range of PBH masses capable of seeding rapid vacuum decay.

Assuming radiation domination, the physical size (length), l_{phys}, of the horizon at horizon entry time is given by

$$l_{\text{phys}} = \int_0^{t_i} \frac{a_i}{a} dt = \int_0^{t_i} \left(\frac{t_i}{t} \right)^{\frac{1}{2}} dt = 2t_i \tag{2.9}$$

setting the speed of light to $c = 1$. The scale that this has grown to today, i.e. the comoving scale, can be found by multiplying the physical scale by the ratio between the scale factor today (defined as $a_0 = 1$) and the scale factor at the time that the black hole formed, i.e.

$$\frac{a_0}{a_i} = \frac{a_{\text{eq}}}{a_i} \frac{a_0}{a_{\text{eq}}} = \left(\frac{t_{\text{eq}}}{t_i} \right)^{\frac{1}{2}} \left(\frac{t_0}{t_{\text{eq}}} \right)^{\frac{2}{3}} , \tag{2.10}$$

[4]In this paper we always plot the constraints assuming a monochromatic mass spectrum of PBHs.

where $a_{\rm eq}$ is the scale factor at matter-radiation equality[5]. We therefore have that the size of the horizon at a given horizon entry time has grown to a size today, $l_i|_{\rm t_0}$, and comoving scale, k_i, given by:

$$l_i|_{\rm t_0} = 2t_i c \left(\frac{t_{\rm eq}}{t_i}\right)^{\frac{1}{2}} \left(\frac{t_0}{t_{\rm eq}}\right)^{\frac{2}{3}} = \frac{2\pi}{k_i} \tag{2.11}$$

Inserting values for radiation-matter equality and the age of the Universe ($t_{\rm eq} \simeq 10^{12}\,$s, $t_0 \simeq 4 \times 10^{17}\,$s), and using $t_i \sim 10^{-23}$s, we find that the physical size converted to Mpc of the region today is $l_i|_{\rm t_0} \sim 10^{-15}\,$Mpc $\sim 3 \times 10^7\,$m, and the corresponding scale is $k_i \sim 10^{16}\,{\rm Mpc}^{-1}$. This represents the largest horizon scale of a PBH which would have decayed by today.

In order to determine the smallest scale that we can probe with PBHs in this scenario, we need to find the scale that left the horizon just before inflation ended and reentered immediately afterwards. If H can be approximated as being constant during inflation, which is typically the case for small-field inflation, then the number of e-folds that occur between the time that today's horizon scale left the Hubble sphere during inflation and the end of inflation, ΔN is

$$\Delta N = \ln\left(\frac{k_{\rm end}}{k_0}\right) \qquad \Rightarrow \quad k_{\rm end} = k_0 e^{\Delta N}. \tag{2.12}$$

Taking a value of ΔN within the expected range [146], for example 60, and $k_0 = a_0 H_0 \simeq 2.3 \times 10^{-4}\,{\rm Mpc}^{-1}$ with $H_0 = 68\,{\rm km\,s^{-1}\,Mpc^{-1}}$ [147], the smallest scale to leave the horizon just before the end of inflation is $k_{\rm end} \simeq 2.6 \times 10^{22}\,{\rm Mpc}^{-1}$.

If H cannot be approximated as being constant during inflation, which is the case for large-field models of inflation, then the above expression becomes less accurate. For example in quadratic inflation, with $V = m^2\phi^2/2$, the smallest scale is instead given by

$$k_{\rm end} = \frac{\sqrt{2}k_0}{\sqrt{2\Delta N + 1}} e^{\Delta N}. \tag{2.13}$$

For $\Delta N = 60$, this decreases $k_{\rm end}$ by about a factor of 10.

2.2.3 Results

We can now plot the consequences for the primordial power spectrum of no PBHs forming on all scales and compare this to the current constraints. For scales smaller than $k \sim 10^{16}\,{\rm Mpc}^{-1}$, indicated by everything to the right of the vertical red line in figure 2.1, the

[5] We do not account for any late time dark energy domination in the evolution of scales nor β throughout our work. The duration of this era is too short to have a dominant effect over some of the other uncertainties in the calculations.

constraint plotted is a consequence of the claim in [6, 7]. For wavenumbers larger than $k \sim 10^{16}\,\mathrm{Mpc}^{-1}$, we do not have any evidence to suggest that no PBHs can have formed, but the plot demonstrates the effect on the constraints of the power spectrum if this were to be the case. It therefore also provides the tightest possible future constraint from PBHs, assuming there are none.

Using (2.7) and $k_0 = 2\pi/L$, the constraint on β as a function of scale is

$$\beta < 1.2 \times 10^{-11} (k\,\mathrm{Mpc})^{-3}. \tag{2.14}$$

Then with equation (2.6) we calculate the constraint on the power spectrum $\mathcal{P}_\delta(k)$ against scale k, measured in Mpc^{-1}. Converting from the comoving density contrast δ to the comoving curvature perturbation $\mathcal{R}$ with

$$\delta = \frac{2(1+\omega)}{5+3\omega} \left(\frac{k}{aH}\right)^2 \mathcal{R} \tag{2.15}$$

evaluated at horizon crossing so $k = aH$ [148] implies

$$\mathcal{P}_\mathcal{R} = \left(\frac{5+3\omega}{2(1+\omega)}\right)^2 \mathcal{P}_\delta. \tag{2.16}$$

During radiation domination ($\omega = 1/3$), $\mathcal{P}_\mathcal{R} = (81/16)\mathcal{P}_\delta$. The resulting plot of $\mathcal{P}_\mathcal{R}$ against k is shown in figure 2.1. Plotting $\mathcal{P}_\mathcal{R}$ instead of $\mathcal{P}_\delta$ allows a comparison to be made with Planck's observed value of the amplitude of the primordial power spectrum on large scales, $A_s = 2 \times 10^{-9}$ [147]. Figure 2.1 shows that despite the constraints on β being as extreme as they can be, the constraint on the primordial power spectrum only improves by around half an order of magnitude in comparison to the current constraints given in [10]. The value of δ_c chosen is important since the constraint varies with the square of this value, and this has much more of an effect than any variation in β. However, the effect of δ_c is the same for both the current constraints and our new constraints, so the improvement from the current constraints to the constraints based on there being no PBHs is the same for whichever value of δ_c is chosen. For our choice of $\delta_c = 0.42$ [149] (see also [150, 151] which derive a similar value of δ_c), the tightest constraint $\mathcal{P}_\mathcal{R} \simeq 2.5 \times 10^{-3}$ is reached at the scale $k \simeq 2.6 \times 10^{22}\,\mathrm{Mpc}^{-1}$, which is the smallest scale to reenter the horizon post-inflation (found in section 2.2.2 with $\Delta N = 60$) and the largest value of k plotted in figure 2.1. This shows that we cannot do any better than $\mathcal{P}_\mathcal{R} \simeq 2.5 \times 10^{-3}$ from only knowing the constraint on β and taking the smallest scale to reenter the horizon post-inflation to be the value found in section 2.2.2 with $\Delta N = 60$. These calculations assume spherical collapse, and as pointed out by [152], constraints on the power spectrum from non-detection of PBHs are very uncertain due to effects such as non-spherical collapse and critical collapse.

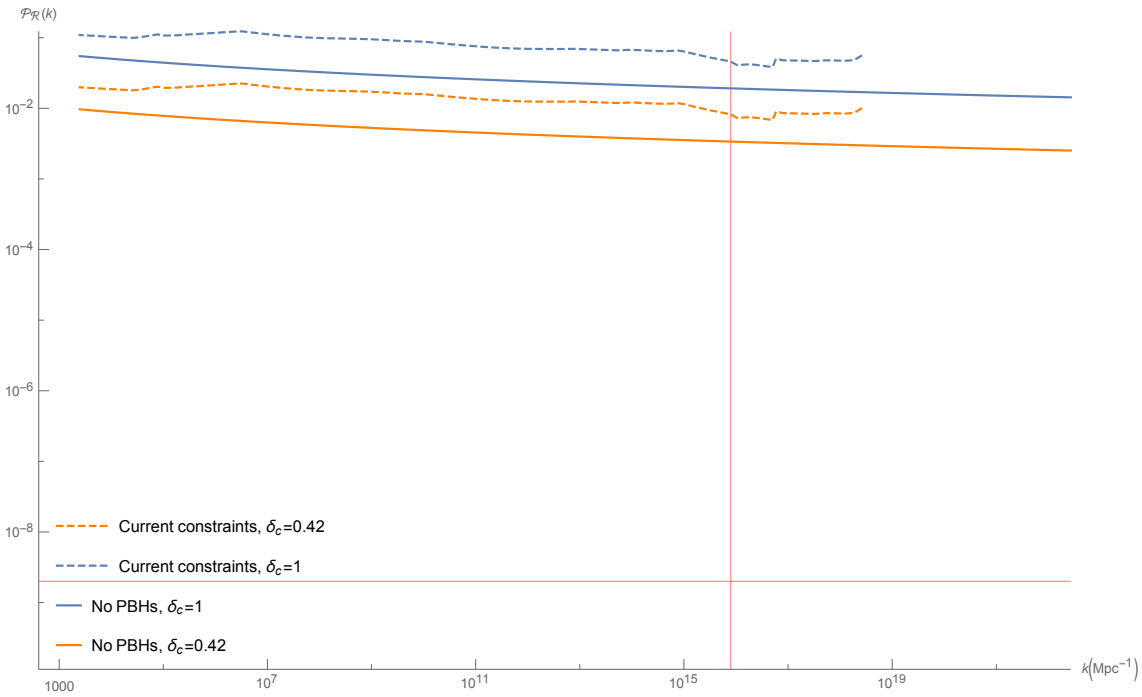

Figure 2.1: Power spectrum constraints from PBH formation in radiation domination. The dashed lines represent the current constraints with two different values of critical level of overdensity. The solid lines represent the constraints if no PBHs form. The red vertical line represents the scale of a PBH that would just have decayed by today, so everything to the right of this line is the constraint due to the argument in [6, 7]. The horizontal red line is drawn at $\mathcal{P}_\mathcal{R}(k) = 2 \times 10^{-9}$. The smallest value of k plotted corresponds to $M_{\text{PBH}} = 10^{40}$g, whilst the largest value of k plotted corresponds to the smallest scale to reenter the horizon post-inflation as found in section 2.2.2 with $\Delta N = 60$. We used $\gamma = 0.2$ to plot these results [8, 9].

Finally, we note that because PBHs form deep into the tail of the probability distribution, the effect of non-Gaussianity can have a much larger effect on the constraints, potentially changing the constraints by two orders of magnitude [145, 153–157].

2.3 Early matter-dominated phase

2.3.1 Motivations

In order to achieve tighter constraints on the primordial power spectrum using PBH abundance constraints, we require a scenario where the power spectrum depends on β more sensitively than logarithmically. This is the case during an early matter-dominated phase which can be caused by a scalar field which dominates the background energy density (e.g. the inflaton or curvaton) oscillating in a quadratic potential [133, 134]. PBH formation in an early matter-dominated phase has been studied in various previous works [135, 136, 158–163], where it was shown that the relationship is governed by a power law instead of a logarithmic function. The exponentially enhanced probability of formation is due to the fact that the Jeans pressure which would normally halt PBHs from forming on sub-Hubble scales during radiation domination vanishes in matter domination, and so PBHs are able to form more easily.

2.3.2 PBH formation likelihood and power spectrum constraints

Based on the results of [136], the expression relating PBH abundance and the mass variance σ (and hence the power spectrum via $\mathcal{P}_\delta \sim \sigma^2$) is

$$\beta_0 \simeq 0.056\sigma^5, \tag{2.17}$$

where β_0 is the PBH abundance fraction defined at the time of formation. This expression does not assume spherical symmetry in the initial density profile. Using data from the plot of $\rho_{\rm PBH}/\rho_{\rm DM}$ against $M_{\rm PBH}$ from [163] (which is in turn collated from constraints due to evaporation [10], femto-lensing of gamma-ray bursts [164], neutron star capture[6] [165], white dwarf explosions [166], microlensing [79, 131, 167], Planck results [87], survival of stars in Segue I [168] and Eridanus II [169], and distribution of wide binaries [84]), it is possible to scale the observed constraints on PBH abundance such that they include a period of evolution in the matter-dominated phase. Taking constraints on $\rho_{\rm PBH}/\rho_{\rm DM}$

[6] Note that since this paper was published, both the neutron star and white dwarf constraints have been dismissed.

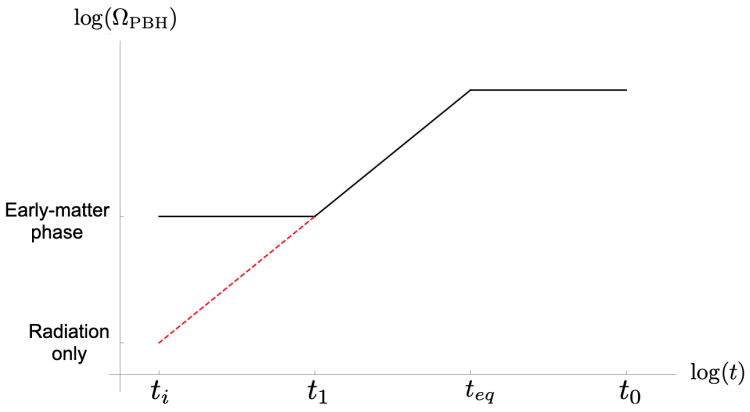

Figure 2.2: Sketch of how $\Omega_{\rm PBH} = \rho_{\rm PBH}/\rho_{\rm tot}$ scales with time depending on the periods of matter and radiation-domination, where β is equal to $\Omega_{\rm PBH}$ at horizon entry. The solid black line shows how $\Omega_{\rm PBH}$ evolves given an early matter-dominated phase that lasts until t_1, followed by a radiation-dominated phase where $\Omega_{\rm PBH}$ grows as a and then the later matter-dominated phase. The red dashed line shows where β from [10] is evaluated with no early matter-dominated phase.

from [163], we can relate them to constraints on $\beta(M)$ via

$$\beta(M) = \frac{\rho_{\rm PBH}}{\rho_{\rm DM}}\Omega_{\rm DM}\left(\frac{M_i}{M_\odot}\right)^{\frac{1}{2}}\left(\frac{M_\odot}{M_{\rm eq}}\right)^{\frac{1}{2}} \tag{2.18}$$

where $\Omega_{\rm DM} \simeq 0.26$ [147] and the horizon mass at matter-radiation equality will be taken as 7×10^{50}g [148].

The values of β found via equation (2.18) are calculated under the assumption that PBHs were forming in a radiation dominated phase. During radiation domination $\rho_{\rm PBH}/\rho_{\rm tot}$ grows like the scale factor a, whereas during matter domination $\rho_{\rm PBH}/\rho_{\rm tot}$ stays constant. In order to account for the period of time between formation and the end of the early matter dominated phase where $\rho_{\rm PBH}/\rho_{\rm tot}$ will remain constant instead of growing as a we need to scale the observed value for β by a. Figure 2.2 gives a graphical depiction of this difference in scaling. The relation between $\beta(M)$ from equation (2.18) and β_0 from equation (2.17) is then

$$\beta_0 = \frac{a_1}{a_i}\beta(M) = \left(\frac{t_1}{t_i}\right)^{\frac{2}{3}}\beta(M), \tag{2.19}$$

where the subscript '1' refers to the end of the early matter-dominated phase (and hence the beginning of the radiation dominated phase), and the subscript 'i' refers to the time that the overdensity enters the Hubble sphere. In order to relate the PBH masses to different scales, we will take $\gamma = 1$ to remain agnostic while the precise value is uncertain

for matter domination and assume that the mass of the resultant black hole is of the order of the mass of the horizon as the overdensity enters:

$$M_{\rm PBH} \sim M_{\rm hor} = \frac{c^3}{2GH(t_i)} = \frac{3c^3 t_i}{4G}, \tag{2.20}$$

where we have used $H(t_i) = 2/(3t_i)$ during matter domination. Improved simulations of the collapse of PBHs in matter domination will provide a better estimate for the correct value of γ. The physical radius of the horizon at the time of horizon crossing is given by the gravitational radius,

$$r_{\rm phys} = \frac{2GM_{\rm hor}}{c^2} = \frac{2GM_{\rm PBH}}{c^2}, \tag{2.21}$$

which has expanded to a radius today given by

$$r_i|_{t_0} = r_{\rm phys}\left(\frac{a_0}{a_i}\right) = r_{\rm phys}\left(\frac{t_1}{t_i}\right)^{\frac{2}{3}}\left(\frac{t_{\rm eq}}{t_1}\right)^{\frac{1}{2}}\left(\frac{t_0}{t_{\rm eq}}\right)^{\frac{2}{3}}, \tag{2.22}$$

where t_1 is the time that the early matter-dominated phase ends. The scale today of PBHs with a particular mass at formation is:

$$k_i(M_{\rm PBH}) = \frac{2\pi}{r_i|_{t_0}} = \frac{2\pi c^2}{2GM_{\rm PBH}}\left(\frac{t_1}{t_i}\right)^{-\frac{2}{3}}\left(\frac{t_{\rm eq}}{t_1}\right)^{-\frac{1}{2}}\left(\frac{t_0}{t_{\rm eq}}\right)^{-\frac{2}{3}}. \tag{2.23}$$

Comparing this to the scale we found for a PBH that would have just decayed by today with mass 10^{15}g, we find that

$$k_i(10^{15}{\rm g}) \simeq \frac{(1\,{\rm s})^{\frac{1}{6}}}{t_1^{\frac{1}{6}}} \times 5 \times 10^{12}\,{\rm Mpc}^{-1}. \tag{2.24}$$

Taking the longest possible duration for the matter phase so that it lasts right up until BBN at $t_1 = 1$ second, we find $k_i(10^{15}{\rm g}) \simeq 5 \times 10^{12}\,{\rm Mpc}^{-1}$ which is around 4 orders of magnitude larger than if there was no early matter era. Choosing $t_1 = 10^{-23}$ seconds instead so that the evolution of the PBH is solely within the radiation dominated phase, we see that $t_1 = t_i$ and $k_i \sim 10^{16}\,{\rm Mpc}^{-1}$ which matches the value found in section 2.2. This consistency check assumes that the matter phase ends and transitions to a radiation phase instantaneously, with the PBH forming as this happens - we will discuss the validity of this assumption as well as the collapse time of the PBHs shortly.

We will plot power spectrum constraint against lengthscale using the inverse of equation (2.17):

$$\mathcal{P}_\delta = \left(\frac{\beta_0}{0.056}\right)^{\frac{2}{5}}, \tag{2.25}$$

and (2.23). Both of these quantities depend on the PBH mass so we take data from the plot in [163] at PBH masses from 10^{10}g to 10^{40}g. Each mass has a corresponding value

for β which we find from equation (2.18), and then scale to β_0 via equation (2.19). This is the value that is substituted into equation (2.25) to give values of the power spectrum corresponding to the scale for each PBH mass from equation (2.23).

Since the power spectrum and the scale both depend on when the early matter dominated era ends, each chosen value of t_1 will result in a different constraint on the power spectrum. Additionally, each value of t_1 will determine the largest and smallest scales with observable consequences for PBH masses $10^{10} - 10^{40}$g that can enter the horizon during the matter-dominated phase. We always assume that the early matter era begins before the horizon mass has grown to 10^{10}g. There will need to be enough time between horizon entry and turnaround time for initial overdensities δ_i to grow to order 1 (or the chosen δ_c) if they are to collapse to form PBHs, so the constraints will weaken on larger scales that enter the horizon close to the end of the matter dominated phase as there won't be enough time for small initial density perturbations to grow before the matter-dominated phase ends.

Looking first at the largest and smallest scales that can enter the horizon before the end of the matter-dominated phase, we see that the smallest mass we use from [163] is 10^{10} g, which corresponds to a horizon entry time of $t_i \simeq 3 \times 10^{-29}$ seconds. This means that the earliest time t_1 to the nearest power of 10 that matter domination can end and there still have been time for PBHs within the mass ranges we have data for to form is $t_1 = 10^{-28}$ seconds.

The smallest scale that enters the horizon post-inflation will be given by $k_i(M_{\rm PBH})$ evaluated at $M_{\rm PBH} = 10^{10}$ g for each value of t_1. The largest scale to enter the horizon before the phase transition from matter domination to radiation domination is determined by the value of $k_i(M_1)$ with M_1 given by

$$M_1 = \frac{3c^3 t_1}{4G}. \tag{2.26}$$

If the overdensities were to collapse instantaneously after horizon entry then $k_i(10^{10}\,\mathrm{g})$ and $k_i(M_1)$ would determine the range of scales to be plotted for each value of t_1. However, overdensities do not collapse instantaneously to form PBHs after crossing the horizon. Instead, there needs to be enough time between horizon crossing and the end of the matter dominated phase for the overdensity to grow to order 1 and begin to collapse at the time of the 'turnaround' [162]. We impose for simplicity the requirement that the overdensity must reach turnaround before the end of the early matter dominated phase if it is to collapse, so the scale factor will need to grow by a factor of δ_i^{-1} between t_i and t_1

for the density fluctuation entering the horizon at t_i to have had time to grow to order 1:

$$\delta(M_i)|_{\mathrm{t=t_1}} \simeq 1 = \delta_i^{-1}\delta_i = \left(\frac{a_1}{a_i}\right)\delta_i. \tag{2.27}$$

During matter domination $k = aH \propto a^{-1/2}$ so the ratio between the scale at horizon entry, k_i, and the scale at the end of matter domination, k_1, goes as

$$\frac{k_i}{k_1} = \left(\frac{a_1}{a_i}\right)^{\frac{1}{2}}. \tag{2.28}$$

Therefore, the scale of the horizon needs to grow by at least $\delta_i^{-\frac{1}{2}}$ if that density fluctuation is to go on to successfully collapse. Only the most extreme fluctuations need to be given time to collapse in order to achieve the observed constraint on β for each PBH mass. How far into the tail of the distribution we must go for each β corresponding to a PBH mass entering the horizon at t_i is given by solving for x_{tail}

$$\beta = \mathrm{Erfc}\left(\frac{x_{\mathrm{tail}}}{\sqrt{2}}\right). \tag{2.29}$$

For example, if $x_{\mathrm{tail}} = 5$, then only 5-sigma fluctuations (i.e. those with $\delta_i > 5 \times \sigma$) need to collapse in order to achieve the observed constraint on β (appropriately scaled via equation (2.19)). We therefore require

$$\sigma > \frac{\left(\frac{k_i}{k_1}\right)^{-2}}{x_{\mathrm{tail}}}. \tag{2.30}$$

If the value of σ from equation (2.25) satisfies equation (2.30), then there is enough time for the overdensity corresponding to that σ to grow to order 1 and collapse before the end of matter domination. If a value of σ does not satisfy equation (2.30), then the constraint on the power spectrum must weaken to the minimum value of σ that allows enough time for the growth of the density fluctuation to order 1 from the time of horizon crossing to the end of matter domination. This minimum value for the power spectrum of δ measured at $k = k_i$ is

$$\sigma^2_{\mathrm{min}} = \left(\frac{\left(\frac{k_i}{k_1}\right)^{-2}}{x_{\mathrm{tail}}}\right)^2. \tag{2.31}$$

For each scale, we will choose the maximum of σ and σ_{min}. The results are plotted in figure 2.3, where we have again converted from $\mathcal{P}_\delta$ to $\mathcal{P}_\mathcal{R}$ via equation (2.16) with $\omega = 0$, hence $\mathcal{P}_\mathcal{R} = (25/4)\mathcal{P}_\delta$.

For the shortest periods of matter domination, $t_1 \sim 10^{-28} - 10^{-25}\,\mathrm{s}$, the value of σ found from equation (2.25) is too small for the corresponding overdensity to have time to

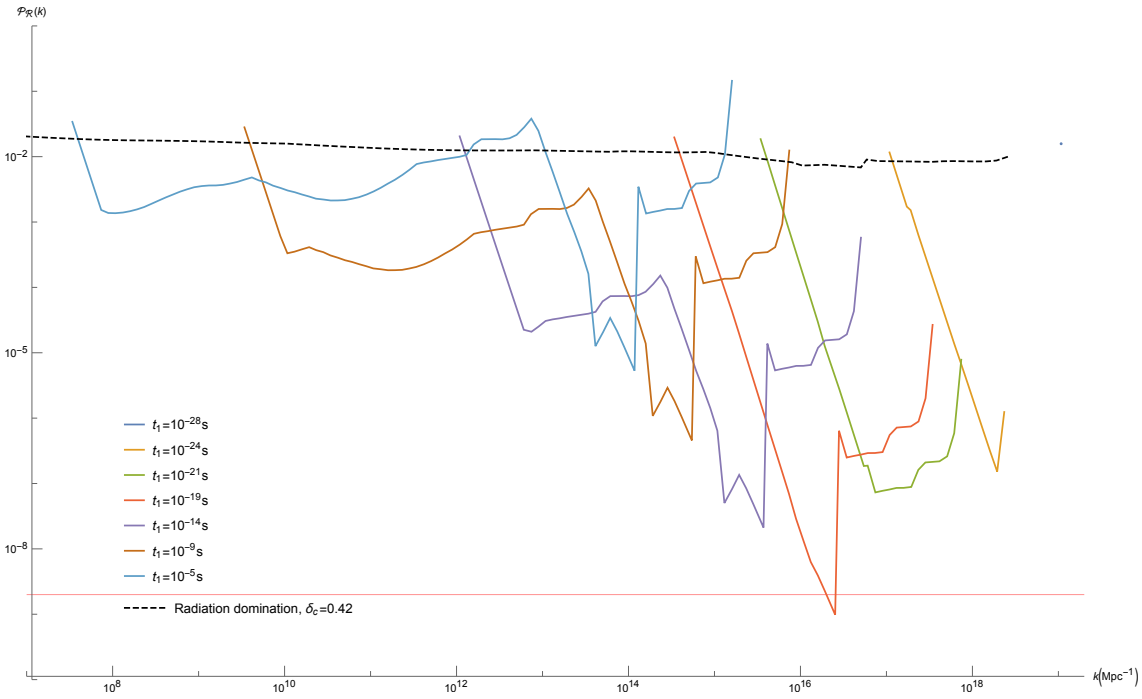

Figure 2.3: Power spectrum constraints from PBH formation in matter domination for values of t_1 from 10^{-28} seconds, represented by the single rightmost point, to 10^{-5} s, represented by the pale blue line that reaches the largest scales. The horizontal red line is drawn at $\mathcal{P}_\mathcal{R}(k) = 2 \times 10^{-9}$. The dashed black line shows the constraints from PBHs which formed during radiation domination.

grow to order 1 before radiation domination on any scale that we have an observed value of β for, and so the constraints are weakened. The rightmost point of each line represents the smallest PBH mass that we constrain with data from [163], 10^{10}g. For $t_1 > 10^{-25}\,$s, there begins to be sufficient time for the initial density fluctuation corresponding to σ on some scales to grow to order unity before the end of matter domination. The tightest constraint is achieved for $t_1 \sim 10^{-19}\,$s at a scale of $k \simeq 5 \times 10^{16}\,\mathrm{Mpc}^{-1}$, surpassing $\mathcal{P}_\mathcal{R} = 2 \times 10^{-9}$, Planck's measurement of the amplitude of the primordial power spectrum. Constraints on the power spectrum improve for all values of t_1 between $10^{-28}\,$s and $10^{-6}\,$s in comparison to the constraints from PBH formation in radiation domination, with constraints for $t_1 \sim 10^{-5}\,$s only just overlapping with the constraint from radiation domination on some scales. The constraints due to values of $\sigma_{\min}$ join up with those from radiation domination as expected, since PBHs forming at the end of the early matter era will predominantly feel the effects of radiation domination if the transition occurs very soon after their formation. Note that uncertainties are introduced in the comparison between matter domination constraints and radiation domination constraints because those from radiation domination assume spherical symmetry, whereas those from matter domination do not. Additionally, for radiation domination we take $\gamma = 0.2$ and $\delta_c = 0.42$, but for matter domination we approximate $\gamma = 1$. This explains why the radiation constraint is slightly stronger than the matter constraint on the left-hand end of each line.

We choose the latest termination of matter domination to be $t_1 \sim 10^{-5}\,$s for two reasons. The expression in equation (2.25) is only valid for $\sigma < 0.05$ [136], so we cannot trust the relation between σ and β for $\mathcal{P}_\delta \gtrsim 10^{-3}$. Secondly, the QCD phase transition occurs around $10^{-5}\,$s at an energy scale which can be probed in the laboratory, and it is generally expected that the hot big bang will be complete by this time, with the Universe dominated by radiation (although counter examples exist, see e.g. [170]).

We assume an instantaneous phase transition from matter domination to radiation domination, but of course the true dynamics of this transition would affect the constraints. A smooth weakening of the constraints from matter domination to match those from radiation domination is most likely. In addition, when cutting our plots at the end of matter domination, we require that the overdensities must have reached turnaround by the time t_1. We expect that some overdensities will have grown considerably but not quite reached order 1, however it is possible that they would still collapse to form PBHs at some point during or after the transition from matter to radiation domination - these cases have been disregarded in our constraints. Simulating the growth and subsequent collapse of the

overdensity during the phase transition, as well as the dynamics of the phase transition itself, would be necessary to gain a full understanding of the effect of these cases on the constraints.

2.3.3 Inhomogeneous effects

The relationship (2.17) was derived by considering departures from spherical symmetry but neglecting potential effects of inhomogeneities in the collapsing region. If these effects were to be included, they would account for a scenario where a caustic could form in the centre of the region, and the increase in pressure could prevent a PBH from forming. The probability of this happening adds an additional factor of $\sigma^{\frac{3}{2}}$ in the relationship between β and σ, calculated using the Lemaitre-Tolman-Bondi dust solution [135] (see also the more recent review [171]). With the most conservative reasoning, this effect is considered to be independent of the probability that the region is spherical enough[7] to collapse into a PBH rather than a pancake or cigar shape, which accounts for the factor of σ^5 that we have been using [135, 172]. Multiplying these gives the minimum probability of PBH formation to be

$$\beta_0 \simeq 2 \times 10^{-2} \sigma^{\frac{13}{2}}. \tag{2.32}$$

Using this relationship originally from [135, 172] instead of (2.17), we can plot the power spectrum vs. scale again, shown in figure 2.4. The strongest constraint on β now produces a constraint on the power spectrum of order $\mathcal{P}_\mathcal{R} \simeq 10^{-7}$, two orders of magnitude weaker than if the inhomogeneous effects are neglected. Additionally, values of σ become larger than 0.05 which were considered not valid for the equation (2.17) in [136], however [135] does not cite this as a limiting factor of the equation that accounts for inhomogeneous effects given in equation (2.32). We caution against concluding that the constraints in a matter dominated era may be weaker than those in radiation domination (which would be very surprising since pressure can only act against gravitational collapse); the radiation era constraints are derived assuming spherical symmetry which maximises the probability of PBHs forming.

Arguments for neglecting the effect of inhomogeneities include that they are very dependent on the matter model, and since we have not specified what has caused the early matter dominated phase, these effects are quite uncertain. Additionally, it was argued in [136] that pressure arising in the central region could just slow down the collapse as opposed to halting it completely.

[7] Note that the more overdense a region is, the more spherical it is according to [172].

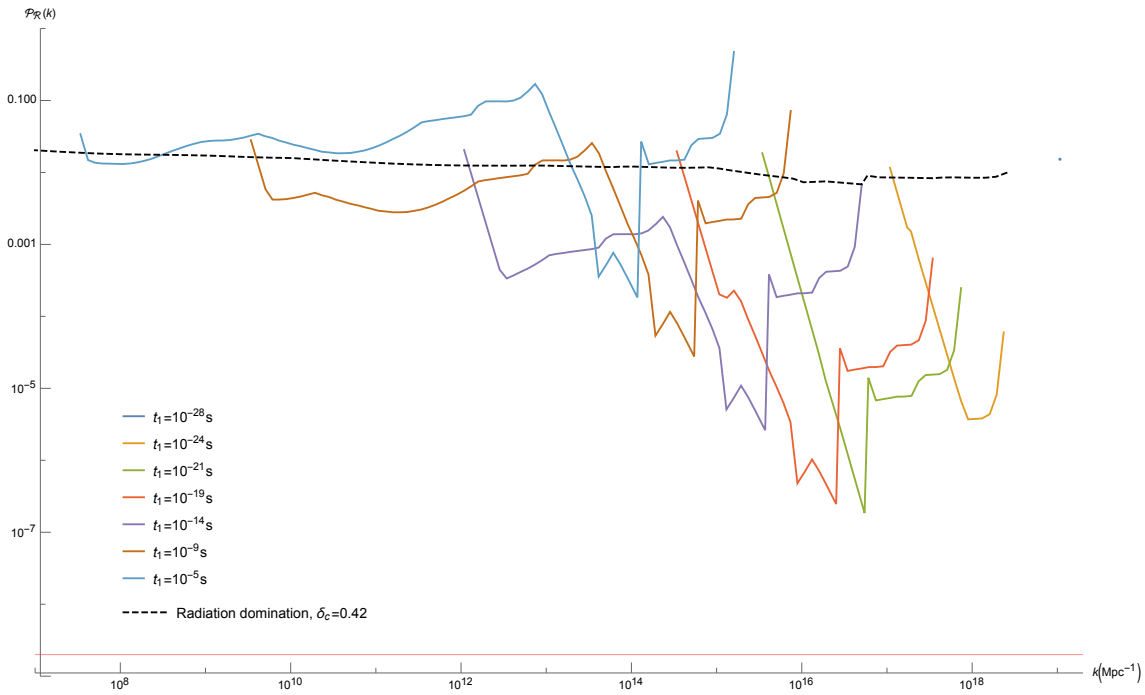

Figure 2.4: Inhomogeneous effects on PBH formation during matter domination translated to the power spectrum. Constraints are weakened in comparison to figure 2.3 for all values of t_1 from 10^{-28} seconds, represented by the single rightmost point, to 10^{-5} s, represented by the pale blue line that reaches the largest scales. The horizontal red line is drawn at $\mathcal{P}_\mathcal{R}(k) = 2 \times 10^{-9}$. The dashed black line shows the constraints from PBHs formed in radiation domination.

2.4 No PBHs and an early-matter phase

When the two scenarios explored so far in sections 2.2 and 2.3 are combined such that there is presumed to be a period of early matter domination during which no PBHs of masses up to 10^{15} g form, because they would have otherwise seeded rapid vacuum decay, constraints on the power spectrum tighten by many orders of magnitude. Using the same argument from equation (2.7) that the Universe can be split up into regions of the scale of a potential PBH, we can reformulate our expression relating the power spectrum to the PBH mass fraction so that it reads

$$\mathcal{P}_\delta \simeq \left(\frac{1}{0.056 N_l}\right)^{\frac{2}{5}}, \tag{2.33}$$

using the relationship from [136] that disregards inhomogeneous effects. If we were instead to account for inhomogeneous effects and use the more conservative expression (2.32) the constraints would weaken similarly to the case in section 2.3.3.

Figure 2.5 demonstrates how the constraints on the primordial power spectrum tighten by many orders of magnitude when the two scenarios are combined. The right hand side of each line is the smallest scale plotted for each value of t_1, corresponding to the scale when the horizon mass is 10^3g. This is the lightest PBH mass that could form after inflation assuming that approximately 60 efolds occur between today's horizon scale exiting the Hubble sphere during inflation and the end of inflation. Horizon entry time for a PBH mass of 10^3g is $t_i \simeq 3 \times 10^{-36}$s, so the earliest value of t_1 we use is $t_1 = 10^{-35}$s. Only constraints due to PBHs of masses small enough that they would have decayed by today are plotted in figure 2.5, so the left hand side of each line corresponds to the scale when the horizon mass is 10^{15}g.

For $t_1 < 10^{-16}$s, initial fluctuations corresponding to values of σ from equation (2.33) are too small on all scales that enter the horizon before the end of the matter phase to grow to order 1 before radiation domination, so $\sigma_{\min}$ is always chosen. For t_1 greater than this, constraints from equation (2.33) start to become valid. The best constraint is achieved at a scale of $k \simeq 5 \times 10^{19}\,\mathrm{Mpc}^{-1}$ with $t_1 \sim 10^{-17}\,$s, and reaches $\mathcal{P}_\mathcal{R} \simeq 5 \times 10^{-27}$. Models which generate a nearly scale-invariant spectrum are ruled out for values of t_1 between 10^{-5}s and $10^{-30}\,$s as constraints on the power spectrum tighten further than an amplitude of $\mathcal{P}_\mathcal{R} = 2 \times 10^{-9}$ (as measured by Planck) on some portion of scales depending on the value of t_1. If constraints were required for the situation where less than one PBH of any mass formed, the lines plotted in figure 2.5 can be extrapolated from the left-most point of each line to values of around $\mathcal{P}_\mathcal{R} \sim 10^{-2}$, depending on the value of t_1. PBHs forming right

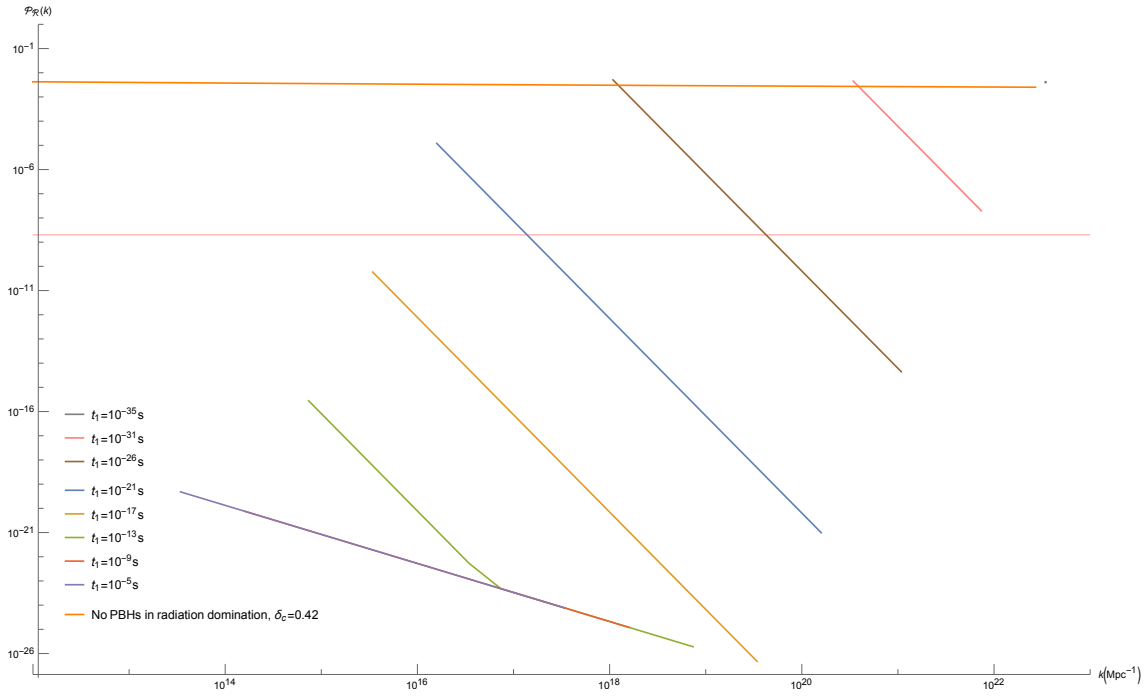

Figure 2.5: The upper bound on the power spectrum assuming that no PBHs have ever decayed. The horizontal red line is the power spectrum amplitude measured on CMB scales. For the lines in matter domination, the left hand side of the line corresponds to the scale when the horizon mass is 10^{15}g (which corresponds to the heaviest PBH which would have decayed by today) and the right hand side corresponds to the scale when the horizon mass is 10^{3}g (which is the lightest PBH mass that could form after inflation assuming that approximately 60 efolds occur between today's horizon scale exiting the Hubble sphere during inflation and the end of inflation). The nearly horizontal orange line is the constraint from no PBHs forming in radiation domination for $\delta_c = 0.42$. Notice that a much smaller range of scales is being plotted here compared to all previous plots showing the constraints on the power spectrum.

at the end of a matter-dominated phase will essentially only be affected by the radiation dominated phase that begins soon after, so it is unsurprising that the constraints become similar to those from 'zero' PBHs forming in a radiation dominated background (shown by the nearly horizontal orange line). The constraints from radiation domination are slightly stronger because they assume spherical collapse, whereas the constraints from matter domination do not, and similarly we choose $\delta_c = 0.42$ for radiation domination but only that the perturbation must grow to order 1 for matter domination.

Such extreme constraints on the power spectrum suggest that an early-matter phase is incompatible with no light PBHs forming during this phase. If the power spectrum is quasi-scale invariant over all scales, then only a sufficiently brief matter dominated period of less than about 7 efoldings is allowed, as the density perturbations would not have time to collapse into a PBH before radiation domination. As presented recently in [162], it is possible that the probability of PBHs nucleating rapid vacuum decay has been overestimated, or otherwise constraints on the energy level at the end of inflation must be enforced to avoid light PBHs forming if an early matter dominated phase occurred. They show that if the energy density at the end of inflation is less than $(2 \times 10^9 \mathrm{GeV})^4$, then the first and hence lightest PBHs produced after inflation will be sufficiently massive to not have decayed. Even more recently, [173] show that a nonminimal, but renormalizable coupling between the Standard Model Higgs field and gravity can have a large effect on the decay rate of the vacuum, possibly allowing light PBHs to form without nucleating decay to the true vacuum.

2.5 Conclusion

PBHs constrain the primordial power spectrum over an extremely broad range of scales, covering over 20 orders of magnitude. However, the constraints are much weaker than the observed amplitude on CMB scales, due to the high pressure forces during radiation domination, which mean only very large amplitude perturbations can collapse. In this paper, we have considered two possible scenarios to tighten the constraints. Firstly we consider the ultimate observational constraint that no PBHs formed during a standard radiation era and secondly the softest possible equation of state - an early matter dominated era which makes PBH formation much easier. Finally we combine the two scenarios.

Although improving the observational constraints is very important for potentially ruling out PBHs as a dark matter candidate (see for example [8, 131, 161, 174–177]), we show in Fig. 2.1 that the consequent improvement in the constraints on the primordial

power spectrum is modest, being less than an order of magnitude. The only exception is on very small scales corresponding to PBHs which form with such small masses that they decay before big bang nucleosynthesis (corresponding to $k \gtrsim 10^{18}$/Mpc), on these small scales there are no standard observational constraints. However, the argument that the evaporation of any black hole could destroy the stability of the Universe [6, 7] suggests that arbitrarily small scales (right down to the horizon scale at the end of inflation) can be constrained. The power spectrum constraints are so insensitive to changes in β of even 50 orders of magnitude, that a better understanding of the formation of PBHs (e.g. simulations of non-spherical initial conditions or reaching a better understanding of the expected initial density profile of PBHs from inflation) will have a larger effect than improving observational constraints. Non-Gaussianity of the primordial density perturbation has already been shown to effect the constraints by up to two orders of magnitude.

However, we have shown that in the matter domination case, due to the enhanced probability of PBH formation, the constraints on the power spectrum can be improved by many orders of magnitude and they become much more sensitive to the observational constraints on β. The constraints also depend on the length of the early matter dominated phase, as shown in Fig. 2.3. For an early matter dominated phase lasting $10^{-19}\,$s, the constraint on the primordial power spectrum surpasses Planck's measurement of $\mathcal{P}_{\mathcal{R}} = 2 \times 10^{-9}$ at a scale of $k \simeq 5 \times 10^{16}\,\mathrm{Mpc}^{-1}$. We further include full constraints for scales entering the horizon right up until the end of matter domination, which consistently match with those from radiation domination to within the uncertainties between the two calculations. If the constraints were to tighten further than this on a wider range of scales, models of inflation that generate a quasi scale-invariant spectrum could be ruled out when combined with an early-matter dominated phase.

In combining these two scenarios, that no PBHs that would have evaporated by today can ever have formed, and that there was a phase of early-matter domination post inflation, we show that incredibly tight constraints on the power spectrum can be achieved for some lengths of a matter dominated phase. If both scenarios can be shown to be realised in nature, then this provides an excellent direction for highlighting the most probable models of inflation, the power spectrum of which would need to decrease by many orders of magnitude on small scales. This may suggest that one or both of these scenarios are not realised, or that the energy scale at the end of inflation must be so low such that sufficiently light PBHs that would have decayed by today could not have formed [162].

Note added: Whilst our paper was being written, [163] produced a similar plot to Fig. 2.3, showing the constraints on $\mathcal{P}_\mathcal{R}$ for a radiation and matter dominated era. Our constraints are an order of magnitude weaker than theirs, which we believe is primarily due to a factor of 9 difference in the conversion between δ to $\mathcal{R}$. We use eq. (2.16) evaluated at horizon entry, while they relate δ and $\mathcal{R}$ on super-horizon scales. They also do not consider constraints due to $\sigma_{\rm min}$ (see eq. (2.31)), which was not important for their purposes. We note that our constraint on scales which enter the horizon during radiation domination of $\mathcal{P}_\mathcal{R} \lesssim 10^{-2}$ agrees with [178].

Two papers relating the vacuum stability argument that no PBHs may have evaporated during a matter dominated phase also recently appeared [162, 173], we comment on these papers in Sec. 2.3.3.

Acknowledgements

The authors thank Tommi Tenkanen for helpful discussions. PC acknowledges support from the UK Science and Technology Facilities Council via Research Training Grant ST/N504452/1, CB is supported by a Royal Society University Research Fellowship.

Chapter 3

Steepest growth of the power spectrum and primordial black holes

Christian T. Byrnes[1], Philippa S. Cole[1] and Subodh P. Patil[2]

[1] Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, United Kingdom

[2] Niels Bohr International Academy and Discovery Center, Niels Bohr Institute, Blegdamsvej 17, Copenhagen, DK 2100, Denmark

3.1 Introduction

There has been a recent surge of interest in the possibility that primordial black holes (PBHs) might constitute a non-negligible fraction of the dark matter in the Universe. This was largely sparked by the question posed in [179] (see also [180, 181]) – whether the order ten solar mass black holes observed by LIGO [182] could be primordial[1]. This is motivated by the fact that there are several hints for the existence of PBHs [186], for example the progenitor BH spins of the LIGO detections being consistent with zero in most cases [187–190], which is expected for PBHs formed during radiation-domination [191] (but not matter-domination [192]) and arguably unexpected for astrophysical BHs [193, 194]. This begs the follow-up question – if these observed black holes are of primordial origin[2], how were they produced and what are the implications for inflationary model building?

[1] See however [183–185] for an alternative interpretation of the LIGO data.

[2] For more about a stellar origin of the detected BHs see e.g. [92, 195, 196].

The idea that black holes could be primordial relics (albeit of non-thermal origin) was first discussed in [158, 197]. Since then, the possibility that they could be produced through inflationary dynamics has been vigorously investigated, see e.g. [11, 176, 198–214], and [61] for a recent review of PBHs in the context of the recent observations by LIGO. In order for primordial black holes to form, the primordial power spectrum has to grow by about seven orders of magnitude above the amplitude of $\mathcal{P}_{\mathcal{R}} \sim 2 \times 10^{-9}$ observed on CMB scales (we discuss the uncertainties in this estimate in Sec. 3.4.3). Growing to such an amplitude on smaller scales takes time during inflation due to the causality of the underlying background field dynamics, but to date no one has quantified just how quickly the power spectrum can grow. We show that at least assuming canonical single-field inflation, neglecting transients, the power spectrum cannot grow faster than $n_s - 1 = 4$, even allowing for arbitrary and instantaneous changes in the derivatives of the inflaton field potential. That is, the (inverse) length-scale k must change by at least an order of magnitude in order for the power spectrum to grow by four orders of magnitude. This implies that any observational constraint on the allowed amplitude of the power spectrum which is tighter than the required amplitude to generate PBHs on a particular scale can be extended over a broader range of scales than directly implied by the observations, due to the restriction on how quickly the power spectrum can grow. We also discuss how observational constraints from the CMB, large-scale structure, spectral distortions and Pulsar Timing Arrays (PTA) all provide constraints on the allowed masses of PBHs which could have formed.

In the context of PBHs as dark matter (DM), the PBH mass function is important for determining the fraction of the energy density in DM that could be made up of PBHs given current constraints on their detection. We therefore investigate whether restrictions on the primordial power spectrum growth rate have an effect on the PBH mass function and find that vastly different power spectra produce very similar mass functions. This means that if one is interested in producing PBHs within a particular range of masses, observational constraints on the power spectrum will need to be avoided without the slope increasing faster than k^4, and the resulting mass function - which will be largely independent of the power spectrum - must then also avoid constraints on the allowed fraction of PBHs in dark matter. Placing analytic bounds on the steepest growth of the power spectrum defines the largest windows possible for PBH production, and targets for future experiments to aim for.

In section 3.2 we define the slow-roll approximation and deviations from it, and use

analytical approximations to find the steepest growth of the power spectrum, and discuss its possible physical basis. We also discuss a dip in the power spectrum that is common to both numerical and analytical results. Fig. 3.2 shows the steepest growth. In section 3.3 we look at the dependence of the mass function on the shape of the power spectrum. In section 3.4 we review the relevant observational constraints on the power spectrum and discuss the effect our bound on the power spectrum has for model-building. Fig. 3.8 shows our "master" plot of the constraints on the power spectrum across a huge range of scales, and future forecasts are shown in Fig. 3.9. Finally, we present our conclusions in section 3.5 with various details deferred to the appendices.

3.2 The fastest possible growth in power

The simplest models of single-field inflation can be described by the so-called slow-roll approximation. This assumes that the inflaton field's kinetic energy is very sub-dominant compared to the potential as it descends it, and is described by the slow-roll parameters:

$$\epsilon = -\frac{\dot{H}}{H^2} = \frac{\dot{\phi}^2}{2H^2 M_{\rm pl}^2}, \tag{3.1}$$

$$\eta = \frac{\dot{\epsilon}}{\epsilon H}. \tag{3.2}$$

These are the first two terms in the so-called Hubble hierarchy, defined (for $i \geq 1$) as

$$\epsilon_{i+1} := \frac{\dot{\epsilon}_i}{H\epsilon_i}. \tag{3.3}$$

For the background to be inflating, ϵ must be less than unity, and provided it varies slowly as inflation progresses, the resulting primordial power spectrum is nearly scale-invariant. The Planck collaboration [22] have measured the amplitude of the primordial power spectrum at scales sampled by the CMB ($k \sim 10^{-3} - 10^{-1}$ Mpc^{-1}) to be of the order 10^{-9} and nearly scale-invariant – consistent with the simplest models of slow-roll inflation. However, CMB measurements tell us nothing about the power spectrum at scales $k \gg 1\,\text{Mpc}^{-1}$. The tightest constraints for $k \gtrsim 1$ Mpc^{-1} are disputed, but are certainly orders of magnitudes weaker than those on CMB scales (however, see [215] for recent claims to the contrary). For $k \gtrsim 10^7$ Mpc^{-1}, the constraint is $\mathcal{P}_\mathcal{R}(k) \lesssim 10^{-2}$ [10], where we use $\mathcal{P}_\mathcal{R}(k)$ to denote the dimensionless power spectrum of the comoving curvature perturbation. This means that the power spectrum is free to grow to around 10^{-2} on small scales, and such growth would indicate strong deviations from the standard slow-roll regime. Any peak-like features in the power spectrum are of topical interest since primordial black holes are produced if the power spectrum is of order 10^{-2} [10]. In what

follows, we analytically derive a steepest growth index of k^4 for the power spectrum in the context of single-field inflation.

3.2.1 Slow-Roll, Beyond Slow-Roll, Ultra Slow-Roll Inflation

In order for the power spectrum to grow during single-field inflation, the potential must become very flat, meaning that ϵ must decrease rapidly. The quantity that governs this is η, as can be seen from the Klein-Gordon equation for a minimally coupled scalar field:

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\mathrm{d}V}{\mathrm{d}\phi} = 0. \tag{3.4}$$

When the potential is exactly flat, $\mathrm{d}V/\mathrm{d}\phi = 0$, so that

$$-\frac{\ddot{\phi}}{\dot{\phi}H} = \epsilon - \frac{\eta}{2} = 3. \tag{3.5}$$

Hence the smallest value of η attainable for a monotonically decreasing potential is $\eta = -6$. Through the defining equation (3.2), we see that the fastest ϵ can therefore decrease is

$$\epsilon \propto e^{-6\mathcal{N}}, \tag{3.6}$$

where $\mathcal{N}$ is the number of e-folds, and we have used $\frac{d}{dt} = H\frac{d}{d\mathcal{N}}$. The limiting case for a monotonic potential is an inflection point or an extended period of $V' = 0$. As verified explicitly through a potential reconstruction exercise in App. C, we indeed see that a phase of Ultra Slow-Roll (USR) inflation [216, 217], defined as a phase of constant $\eta = -6$, is attained as one approaches an inflection point, which can also be reasonably well approximated by a small enough first derivative [11, 30, 176, 208, 218].

If the potential is non-monotonic and the inflaton field rolls uphill then an arbitrarily negative value of η is possible, but the potential needs to be extremely tuned and will have many transients associated with it (which could end up dominating) [40]. For this reason we will mainly focus on regimes of $\eta \geq -6$, however, we will show that our result for the steepest growth of the power spectrum also holds for non-monotonic potentials with $\eta < -6$. Any deviation from $\eta \simeq 0$ goes beyond what has typically come to be known as the standard slow-roll (SR) approximation, which consists of neglecting the acceleration term in (3.4), a valid approximation only when $\eta \simeq 0$. However, qualitatively different behaviour for the mode functions can result from different regimes of $\eta < 0$ even as the background remains approximately de Sitter, with $a(t) \sim e^{Ht}$. We see this by first recalling the equation of motion for the curvature perturbation in conformal time

$$\mathcal{R}_k'' + 2\frac{z'}{z}\mathcal{R}_k' + k^2\mathcal{R}_k = 0 \tag{3.7}$$

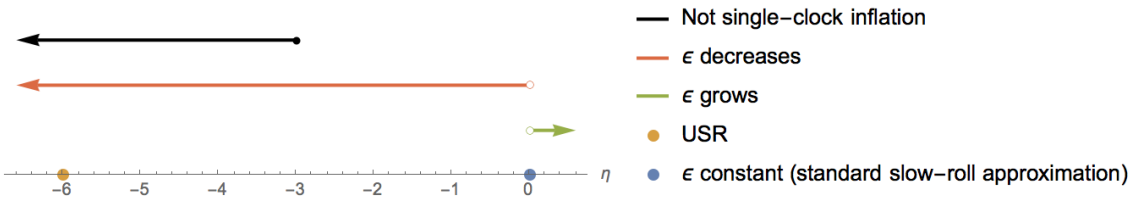

Figure 3.1: The behaviour of the backgrounds and perturbations as a function of the slow-roll parameter η.

where $z^2 = 2a^2 M_{\rm pl}^2 \epsilon$. In the long wavelength limit we find the general solution

$$\mathcal{R}_{k\to 0} = C_k + D_k \int^\tau \frac{d\tau'}{a^2\epsilon}. \tag{3.8}$$

The first term is the usual constant super-horizon mode, and the second term ordinarily decays. However, when ϵ decays at least as fast as $\epsilon \propto a^{-3}$ in cosmological time (i.e. a^{-2} in conformal time), the second term no longer decays. That is, for $\eta < -3$, one has a growing super-horizon mode in addition to the usual constant mode, whereas for $\eta > -3$ one has the customary constant mode and decaying mode of standard slow-roll inflation. This implies that our inflationary background is not an attractor whenever $\eta \leq -3$, and we are in the peculiar regime of single-field, but non-single-clock inflation[3]. This is because on an attractor, we only have one linearly independent perturbation that can persist – a local reparameterisation of the background solution, with the other linearly independent perturbation decaying exponentially. When this is no longer the case, an arbitrary perturbation can no longer be described as just a local time reparameterisation of the background – the defining characteristic of the single-clock regime[4]. For this reason, we find it useful for the purposes of the following discussion to classify different phases of η as standard slow-roll approximation ($\eta \approx 0$), beyond the slow-roll approximation ($-3 < \eta < 0$), and non-single-clock inflation ($\eta \leq -3$) with the limiting case of ultra slow-roll at $\eta = -6$. See Fig. 3.1 for a visual representation of these regimes.

Inflationary potentials which have inflection points or sufficiently flat sections have been studied in e.g. [11, 176, 208, 209, 212, 213], and are generally found to be severely tuned if one stipulates that a peak be produced in the power spectrum with amplitude of order 10^{-2}. In what follows, we will show that on the way to such a peak, one cannot increase primordial power arbitrarily fast in k-space. Were we to consider a phase of strictly constant η evolution, one can straightforwardly derive a steepest possible growth

[3]Non-single-clock inflation means that there is more than one mode (usually the constant mode in slow-roll inflation) that needs to be kept track of in order to understand the evolution of the perturbations.

[4]See [38] for a detailed discussion of this point.

of $\mathcal{P}_\mathcal{R} \propto k^3$ (cf. App. B). However this is too simple an approximation, as any realistic inflationary background must eventually exit such a phase. A more careful multi-phase matching calculation demonstrates a steepest growth of $\mathcal{P}_\mathcal{R} \propto k^4$. This implies that the generation of primordial black holes due to peaks in the power spectrum is subject to further model-independent integral constraints from CMB spectral distortions and pulsar-timing array bounds.

3.2.2 Regimes of constant η

As noted above, the power spectrum grows quickly if ϵ decreases quickly – a process tracked by the second slow-roll parameter η. In order to determine the fastest possible growth, we consider regimes where η decreases monotonically from 0 to different negative values. Finding the behaviour of the power spectrum for instant transitions between different phases is possible analytically via a matching calculation [219], and stitching together sufficiently many phases of constant η evolution can approximate a smooth transition. For the purposes of deriving a steepest growth index, we note that the growth produced by an instant transition between different phases of constant η will be steeper than the growth produced in a smooth transition.

As elaborated upon in App. A, because there are no jumps in the energy momentum tensor of the background between phases of different η, the Israel junction conditions (see (A.10)) require us to match the curvature perturbation and its first derivatives across the matching surface [220, 221]. All modes begin in the Bunch-Davies vacuum in the initial phase of $\eta = 0$, after which we match to a phase of constant $\eta < 0$, and then again to a terminal phase of $\eta = 0$. We derive analytic expressions for the power spectrum for instant transitions between 0 and the 6 integer values of η up to the ultra slow-roll regime in (A.26) - (A.28), the results of which are plotted in Fig. 3.2. The duration of each $\eta < 0$ phase is chosen such that the same growth in amplitude is achieved in all six cases, facilitating a straightforward comparison of the spectral index.

From Fig. 3.2, we see that the slowest growth occurs for the integer values $\eta = -1$, with $n_s - 1 = 1$ once transients have died down (due to the unrealistic instantaneous transitions[5]), and for $\eta = -2$, where the growth is given by $n_s - 1 = 2$. Since we only expect the previously decaying mode to start growing once $\eta \leq -3$, the slope of the power spectrum for $\eta > -3$ is determined by the value of η only, and matches the expression given in (B.2). For $\eta \leq -3$, a qualitatively different behaviour emerges. Since we are no

[5]Note that even smooth but rapid transitions in the inflationary potential can lead to oscillation in the power spectrum, see e.g. [222].

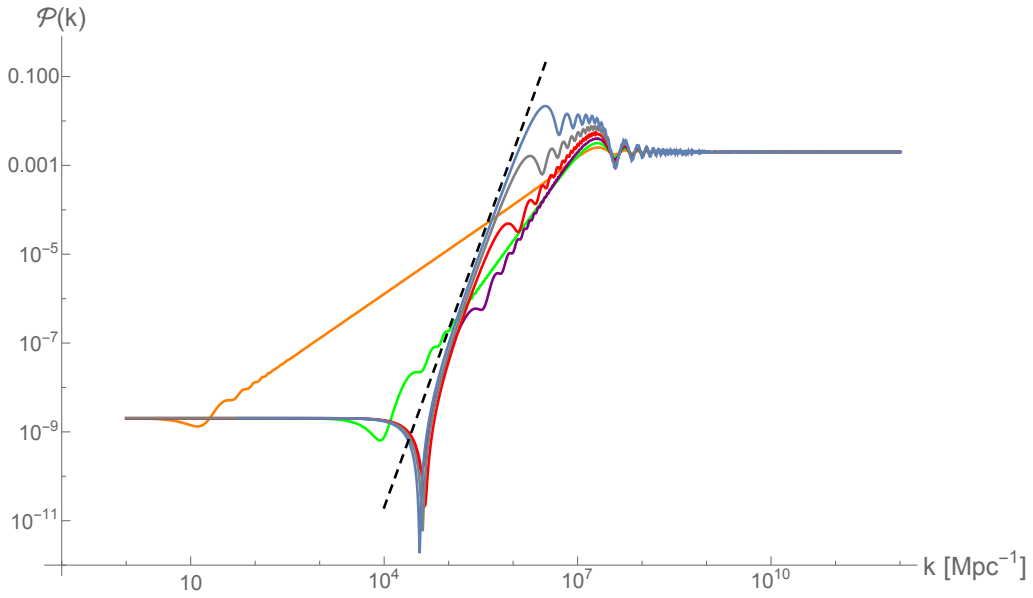

Figure 3.2: Analytical matching from $\eta = 0$ to values of constant η between -1 and -6, to a final phase of $\eta = 0$ slow roll. For $\eta \leq -3$ (purple, red, grey and blue lines in decreasing value of η), the steepest slope of $n_s - 1 = 4$ is achieved after the dip, before relaxing to a shallower slope decided by the value of η, see (B.2). For $\eta > -3$ (green and orange lines), the slope is constant for the whole range of k that the power spectrum increases for. See the online version for colour figures.

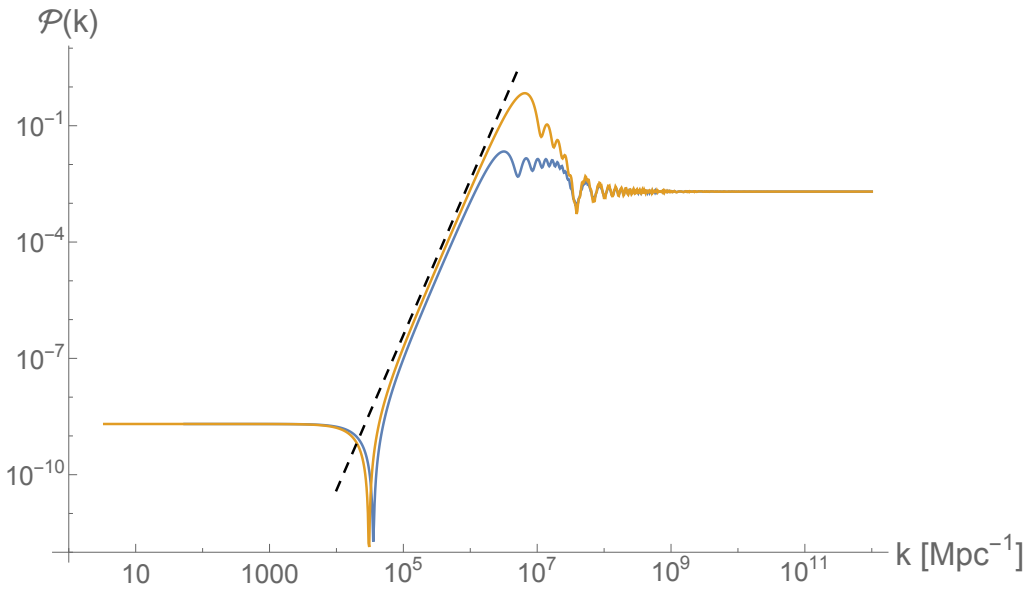

Figure 3.3: The blue line is the same power spectrum as plotted in Fig. 3.2 for 2.3 e-folds of USR ($\eta = -6$). The yellow line is a matching from $\eta = 0$ to $\eta = -8$ for 1.725 e-folds and back to $\eta = 0$. The spectral index of the power spectrum is $n_s = 4$ initially (after the dip), followed by a brief period of negative spectral index $n_s \simeq -2$, before scale-invariance for the final $\eta = 0$ phase. The dashed black line has a k^4 slope.

longer in the single-clock regime, the previously decaying mode starts to grow, and with it, superhorizon perturbations. Here, the power spectrum has a pronounced dip occurring at scales that exit the horizon a few e-folds before the time of the first transition, followed by an initial growth index proportional to k^4, after which it settles to the constant-roll growth given in (B.2). The initial phase of k^4 growth is the steepest possible. In all cases of $\eta \leq -3$, the power spectrum begins to grow before the transition time, which is evidence for superhorizon growth[6]. Providing one adjusts the duration of the $\eta \neq 0$ phase such that the final amplitude of the power spectrum is always the same, the rapid growth lasts longest for $\eta = -6$ before reaching a scale-invariant spectrum. Evidently, the steepest growth is characteristic of the non-single clock phase. For an inflationary potential where the inflaton field transiently rolls uphill, η can become arbitrarily negative. However one finds a steepest growth of k^4 in this case as well, demonstrated in Fig. 3.3 with a matching from $\eta = 0$ to $\eta = -8$ and back. We offer an analytic understanding of this steepest growth in the next subsection.

A more realistic treatment would model the evolution of η as a series of non-zero

[6]In Fig. 3.2 the horizon exit scales at the transition times from $\eta = 0$ to constant η are (for decreasing values of η from -1 to -6) $k \sim 10, 10^4, 10^5, 3 \times 10^5, 6 \times 10^5$ and 10^6 Mpc^{-1}.

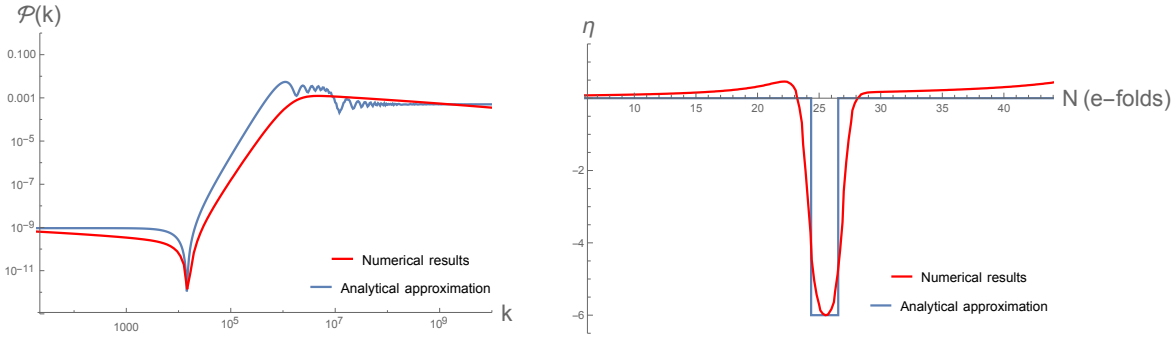

Figure 3.4: Left-hand plot: Numerical results for the potential in [11] are plotted in red and our analytical approximation is plotted in blue. The analytical approximation involves 3 constant phases of η from 0 to -6 and back to 0. The right-hand plot shows the piecewise form for η used for the analytical approximation in blue, with 2.2 e-folds of $\eta = -6$. The full numerical evolution of η for the potential in [11] is shown in red. Note that the units in e-folds have been defined arbitrarily, and we have chosen to centre the phase of $\eta = -6$ in our analytical approximation at the time N when the numerical η reaches -6 instantaneously.

constant phases of η, with instant transitions between each phase to approximate a smooth transition between slow roll and ultra slow roll. As expected, we again find a steepest growth of k^4, illustrated in Fig. A.1. We can test a 'realistic' example of a smooth transition by numerically calculating the final power spectrum for the inflection point model for the potential given in [11, 176], with the choice of parameters given in section 4 of [11]. We use CPPTransport [223, 224] to perform the numerical calculation. The red line in Fig. 3.4 is the resulting numerical power spectrum, and the blue line is the analytical result from approximating the evolution of η in the way shown. As expected, the analytic approximation following from an instant matching between phases of inflation grows more steeply than the more realistic (and smooth) potential of [11].

In App. C, we show how these analytical power spectra might be realised by constructing an example potential that smoothly traverses between $\eta = 0$ and $\eta = -6$ and back.

3.2.3 Steepest growth and the prior dip in the power spectrum

Inflationary models which include a phase of non-single-clock evolution (i.e. with $\eta < -3$) manifest a significant dip in the power spectrum before a steep rise caused by the growth of the perturbations on super-horizon scales, see Fig. 3.2. This has been observed in many

recent studies [11, 176, 208, 213, 218, 225] which have numerically computed the primordial power spectrum for inflationary models with deviations from the slow-roll approximation. It might be assumed that this is caused solely by an increase in ϵ before the rapid decrease, for example see ϵ plotted in Fig. 2 of [213]. However, on comparing numerical results with the analytical approximations from section 3.2.2 for which ϵ never increases, we show that the dip cannot be caused by this alone, and that it is actually a generic feature of transitioning from $\eta = 0$ to a regime where the decaying mode starts to grow. Perhaps surprisingly, the dip is located on scales which exit the horizon while normal slow roll is still taking place. We explain these features for the particular case of a transition from $\eta = 0$ to $\eta = -6$ by making an analytical matching between the two periods. We assume USR lasts for well over an e-folding and neglect the effects of transitioning out of USR, which is a subdominant effect (see Fig. A.2).

The expansion of the power spectrum in terms of $k/k_{\rm u}$, where $k_{\rm u}$ is the horizon scale at the time when USR begins (see eq. (A.26) for the full expression), is

$$\frac{\mathcal{P}_{\mathcal{R}}(k)}{\mathcal{P}_{\mathcal{R}}(0)} = \left[1 - \frac{4}{5}\left(\frac{k}{k_{\rm u}}\right)^2 e^{3N_{\rm USR}}\right]^2 + 2\left(\frac{k}{k_{\rm u}}\right)^2 - 0.10\left(\frac{k}{k_{\rm u}}\right)^6 e^{6N_{\rm USR}} + 0.0075\left(\frac{k}{k_{\rm u}}\right)^8 e^{6N_{\rm USR}} + \cdots \tag{3.9}$$

where $N_{\rm USR}$ is the number of e-folds during which $\eta = -6$ and we have dropped terms subleading in $e^{N_{\rm USR}}$ [7]. All higher-order terms, which come in even-powers of $k/k_{\rm u}$ also come dressed with pre-factors of $e^{6N_{\rm USR}}$ but with numerical coefficients that are down by an order of magnitude for each even order. Therefore, for $k \lesssim k_{\rm u}$, terms up to quartic order are an accurate approximation to the power spectrum. Once $k^2/k_{\rm u}^2 \sim \mathcal{O}(10)$, all of the terms in the alternating series are as important as each other and this is when the series begins to conditionally converge to an oscillating function. Solving for k such that the term in square brackets is zero gives the position of the dip, $k_{\rm dip}$, as

$$\frac{k_{\rm dip}}{k_{\rm u}} = \sqrt{\frac{5}{4}} e^{-\frac{3}{2}N_{\rm USR}} \tag{3.10}$$

and hence the dip occurs approximately $\frac{3}{2}N_{\rm USR}$ e-folds before USR begins. The amount by which the power spectrum is suppressed at this point is

$$\frac{\mathcal{P}_{\mathcal{R}}(k_{\rm dip})}{\mathcal{P}_{\mathcal{R}}(0)} \simeq 2.5 e^{-3N_{\rm USR}}. \tag{3.11}$$

[7] We note that if USR ends in a different way than an instant transition to constant ϵ then the numerical coefficients in the equation above change slightly, but the qualitative picture remains the same.

Finally the rapid k^4 growth during the transition to USR will end when the k^4 and k^6 terms in (3.9) become comparable, which happens when

$$k \simeq 2.5 k_{\rm u}, \tag{3.12}$$

and hence occurs about one e-folding after USR has begun, independently of the duration of USR (provided $N_{\rm USR} \gtrsim 1$).

Thus far, we've arrived at an analytic understanding of the shape of the primordial power spectrum via a matching calculation, and in particular, its steepest possible growth over an intermediate range of scales. This begs the immediate question – what is the underlying physical mechanism responsible for this steepest growth? Several independent arguments demonstrate a steepest growth of $\mathcal{P}_\mathcal{R} \propto k^3$ under the assumption that the large scale power spectrum is a strict power law over all relevant scales. Peebles showed that if the matter power spectrum is to accurately describe particulate matter over scales of cosmological interest, then the two-point function for the density contrast $\delta := \delta\rho/\rho$ can grow no faster than k^4 [226]. This implies that the dimensionless power spectrum for the curvature perturbations can grow no faster than k^3 since $\partial^2\mathcal{R} \propto \delta$. As shown in App. B, one can also derive a similar strongest possible scaling for the two point function of the curvature perturbation from the asymptotics of the mode functions. However, none of these arguments apply in the present context, where we do not assume constant power law behaviour for the primordial power spectrum, and the steepest growth is only over a limited (and in principle tunable) interval[8]. Although one might suspect causality or unitarity arguments to be at play – or perhaps conformal symmetry as the system tends towards $\epsilon \to 0$ – it seems that the bound may be due to an interplay of causality arguments and energy-momentum conservation, something we're currently investigating with a particular view to generalising to the multi-field context.

3.3 The PBH mass function

Having shown that there are limits to how quickly the power spectrum can grow, one may expect that this also places a sharp limit on how narrow the mass function of PBHs can be. In practice this is not the case, for (at least) three reasons: 1) for any given horizon mass, PBHs form with a spread of comparable masses; 2) the matter power spectrum is

[8] We also note in App. B that in assuming a (possibly distributional) power spectrum of the form $\mathcal{P} \propto k^n$ at all scales, one can show that it is not possible to regulate the short distance divergence of a spectrum with index $n > 4$ in four dimensions.

less 'sharp' than the primordial power spectrum because of the window function relating the two; and 3) PBH formation is exponentially sensitive to the amplitude of the power spectrum, so only perturbations comparable to the peak amplitude are important.

The phenomena of critical collapse [93, 227, 228] describes how PBHs of mass M can form with a variety of masses for any given horizon mass M_H according to the relation

$$M = kM_{\rm H}(\delta - \delta_c)^{\gamma}, \tag{3.13}$$

where during radiation-domination the constants have been numerically estimated as $k = 3.3,\ \gamma = 0.36,\ \delta_c = 0.45$ (the exact values depend upon the type of perturbations being considered but we use the values given here in order to be concrete which are for purely growing-mode perturbations [150]), see also [138, 229, 230] for details. From the expression given in [231], the mass function of PBHs, $f(M)$, is

$$\begin{aligned} f(M) &\equiv \frac{1}{\Omega_{\rm CDM}}\frac{d\Omega_{\rm PBH}}{d\ln M} \\ &= \frac{1}{\Omega_{\rm CDM}}\int_{-\infty}^{\infty}\frac{2}{\sqrt{2\pi\sigma^2(M_{\rm H})}}\exp\left[-\frac{(\mu^{1/\gamma}+\delta_c(M_{\rm H}))^2}{2\sigma^2(M_{\rm H})}\right]\frac{M}{\gamma M_{\rm H}}\mu^{1/\gamma}\sqrt{\frac{M_{\rm eq}}{M_{\rm H}}}d\ln M_{\rm H}, \end{aligned} \tag{3.14}$$

where $\mu \equiv \frac{M}{kM_{\rm H}}$.

Inspired by the observation that we cannot have an arbitrarily rapidly growing power spectrum, we calculate the resulting mass function $f(M)$ from 4 different power spectra. The first three grow at different rates towards smaller scales ($n_s - 1 = 0.1,\ 0.2$ and 4 but then drop to zero), the fourth is a Dirac delta function and we chose an overall normalisation of $A_s = 0.15$

$$\mathcal{P}_{\mathcal{R}}(k) = 0.242{\rm A_s(k/k_*)^{0.1}}\text{ for k} < 1.5{\rm k_*},\ 0\text{ otherwise;} \tag{3.15}$$

$$\mathcal{P}_{\mathcal{R}}(k) = 0.256{\rm A_s(k/k_*)^{0.2}}\text{ for k} < 1.35{\rm k_*},\ 0\text{ otherwise;} \tag{3.16}$$

$$\mathcal{P}_{\mathcal{R}}(k) = {\rm A_s(k/k_*)^4}\text{ for k} < {\rm k_*},\ 0\text{ otherwise;} \tag{3.17}$$

$$\mathcal{P}_{\mathcal{R}}(k) = 0.182A_s\delta(k - 0.83k_*). \tag{3.18}$$

The prefactors to the power spectra and the scale at which they drop to zero has been tuned in order to make the position and amplitude of the peak of the mass function as similar as possible, in order to easily compare the width of the mass function. The power spectrum dropping instantaneously to zero is unrealistic but unlike an increasing power spectrum, which cannot grow arbitrarily quickly, there is no theoretical limit to how rapidly the spectrum can decay. If the potential is discontinuous and drops to zero instantaneously, then once it goes to zero ϵ becomes 3 instantly, inflation ends and the Universe enters

kination-domination. For potentials which rapidly switch to a steep negative gradient an arbitrarily rapid transition to a rapidly growing ϵ can be engineered, which can make the power spectrum as blue as required without ending inflation quickly. For more details, see App. C.2.

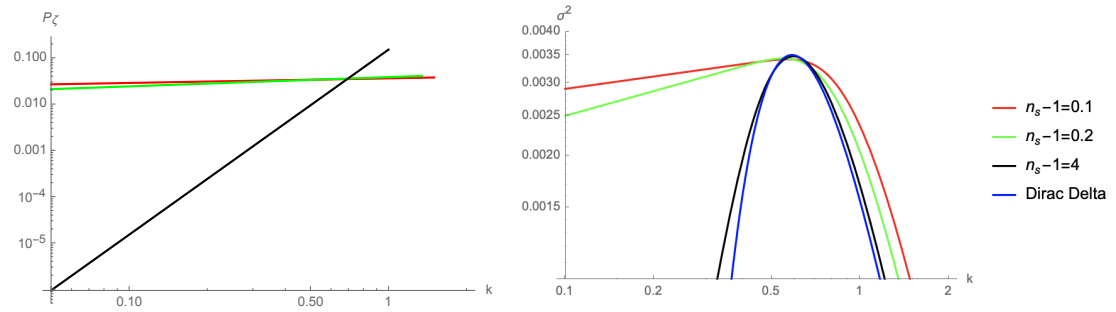

Figure 3.5: The power spectrum on the left and the smoothed variance of the density contrast on the right for the four models described in the text (the Dirac delta model is not plotted on the left plot). The power spectra are zero where no line is shown. The x–axis units are arbitrary.

The variance of the comoving density contrast at horizon entry (smoothed on a scale $R = 1/k$) is related to the power spectrum by

$$\sigma_R^2 = \int_0^\infty \frac{\mathrm{d}q}{q} \frac{16}{81} \left(qk^{-1}\right)^4 \mathcal{P}_\mathcal{R}(q) W_R(q)^2 \tag{3.19}$$

and we use a Gaussian[9] window function, $W_R(q)^2 = e^{-(qR)^2}$. The results are shown in Fig. 3.5. Fig. 3.6 demonstrates that the mass function is not very sensitive to how steep or spiked the primordial power spectrum was unless it varied very slowly with scale. There is almost no visible difference in $f(M)$ between a spike which is modelled by a Dirac delta power spectrum or one growing

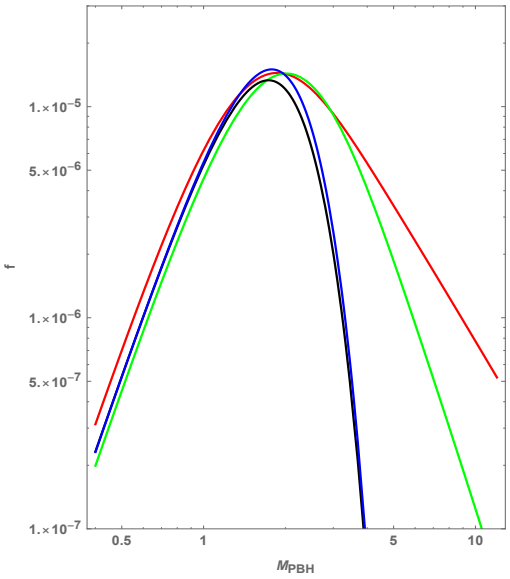

Figure 3.6: The mass function for the 4 power spectra plotted in Fig. 3.5, plotted using the same colour scheme. The arbitrary x–axis units are chosen such that the horizon mass is unity for $k = 1$ in the units used for Fig. 3.5.

[9]The influence of the choice of the window function is discussed in [55]. We neglect a transfer function which suppresses sub-horizon perturbations, because it has a negligible impact when using a Gaussian smoothing function.

like k^4. More surprisingly, even a slowly changing power spectrum with spectral index $n_s - 1 = 0.1$ and a cut off generates a mass function which is not substantially broader than the tightest possible mass function near the peak; compare the black and red lines in Fig. 3.6. Note that the mass functions agree extremely well for small masses because the power spectra all have a cut-off scale which has been chosen to align the peaks of the mass functions. The insensitivity of the PBH mass function to the shape of the power spectrum, due to the degeneracies between the effect of the amplitude and shape of the power spectrum, mean that the PBH mass function would have to be detected with very high precision in order to reconstruct the shape of the primordial power spectrum near the corresponding peak.

3.4 Observational constraints

Although the primordial power spectrum is tightly constrained on CMB scales (from roughly $k \sim 10^{-4} - 10^{-1}$ Mpc^{-1}) to be of the order $\mathcal{P}_\mathcal{R} \sim 10^{-9}$, there are upper bounds on scales far beyond those accessible in the CMB through a variety of tracers and indirect probes. Of these, the constraints most relevant to the present discussion are distortions of the CMB spectrum from the dissipation of acoustic modes, and bounds from gravitational wave backgrounds produced by scalar perturbations at second order. At the end of this section we produce a "master plot" of the key constraints on the power spectrum.

In order to quantify the effect of observational constraints on the allowed number of PBHs generated, which is quantified by $\beta = \rho_{\rm PBH}/\rho_{\rm tot}$ at the time of formation, we need to know the relationship between the amplitude of the power spectrum and β. In this paper we have neglected the impact of quantum diffusion of the inflaton field and non-Gaussianity of the primordial perturbations. In the context of PBH formation, quantum diffusion during inflation has been discussed subject to the slow-roll approximation by [232] and during USR with conflicting conclusions about its importance in [218, 233–235].

We have also neglected the impact of any non-Gaussianity of the primordial perturbations. This is despite the fact that one of the main reasons that USR inflation was initially considered interesting was that it appeared to generate local non-Gaussianity of order-unity amplitude [236], providing an exception to the statement that single-field inflation generates negligible local non-Gaussianity [37, 38, 237]. However, Cai et al. have recently shown that ending USR with a smooth transition tends

to erase the local non-Gaussianity [238]. If the local non-Gaussianity is not erased or modified by the way USR ends, it has a value $f_{\rm NL} = 5/2$ for modes which exited the horizon long after USR begins (see [239, 240] for coordinate-choice issues). $f_{\rm NL} \sim 1$ was shown to have a significant impact on the power spectrum constraints in [145], while the higher-order non-linearity parameters and mode coupling are also important [140, 156, 241]. The impact of non-Gaussianity on PBH formation during inflection point inflation was recently considered in [242].

Other uncertainties in relating the amplitude of the power spectrum to the number of black holes include non-sphericity of the initial density profile [152], the window function used to smooth the density contrast [55], the background equation of state when modes re-enter the horizon [138] (the QCD transition can motivate a population of solar mass PBHs [231, 243]) and the shape and sphericity of the initial energy-density profile [58, 139, 149, 244]. More broadly, the general calculation has been questioned recently in [57, 245], with particular uncertainty on the critical density threshold δ_c and more than an order of magnitude uncertainty in the relation between the horizon and PBH mass. There is not yet any consensus on how to calculate $\beta(M)$, given a particular primordial power spectrum.

Conditional on all of the aforementioned caveats, for a non-negligible number of PBHs to be generated, the amplitude of the power spectrum needs to be of order 10^{-2} depending on the mass of the PBH [246]. We will now see the relevant constraints.

3.4.1 CMB Spectral distortions

What we see in the CMB is a snapshot of acoustic excitations in the primordial plasma around the time of last scattering. Sound waves dissipate energy as they propagate, transferring energy into heating the ambient medium. Any heat dissipated into the primordial plasma has the possibility of showing up in the form of μ and y type distortions of the CMB [17, 101], provided they occur in the redshift window between $z \sim 10^6$ and last scattering at $z \sim 10^3$. The reason for this is that at sufficiently early times, Compton scattering is efficient enough to rapidly restore thermal equilibrium after any energy injection process. At around $z \sim 10^6$ its efficiency starts to drop, and distortions of the blackbody spectrum of the CMB can start to persist if they were initially large enough. The greater the power spectrum is at small scales, the greater the amount of energy that gets dissipated into the primordial plasma, hence spectral distortions offer a powerful probe of the power spectrum at scales beyond those accessible

in CMB anisotropies. Of the two varieties of distortions generated by dissipation, μ-type distortions are sensitive to power at smaller scales. It corresponds to a distortion of the black body spectrum mimicked by an effective chemical potential μ, given by [106]

$$\mu_{\mathcal{R}} \approx \int_{k_{\rm min}}^{\infty} \frac{dk}{k} \mathcal{P}_{\mathcal{R}}(k)\, W_{\mathcal{R},\mu}(k), \tag{3.20}$$

where the window function W_μ is given by

$$W_{\mathcal{R},\mu}(k) \approx 2.27 \left[\exp\left(-\left[\frac{\hat{k}}{1360}\right]^2 \Big/ \left[1 + \left[\frac{\hat{k}}{260}\right]^{0.3} + \frac{\hat{k}}{340}\right]\right) - \exp\left(-\left[\frac{\hat{k}}{32}\right]^2\right)\right], \tag{3.21}$$

with $\hat{k} = k/[1\,\mathrm{Mpc}^{-1}]$ and $k_{\rm min} \simeq 1\,\mathrm{Mpc}^{-1}$, below which the power spectrum is tightly constrained by large scale observations of the CMB. This window function is calculated by integrating the spectral distortion visibility function numerically, for details, see [106]. It behaves as a window function because it up- and down-weights contributions at different wavenumbers according to the numerical factors in equation (3.21). Measurements from COBE/FIRAS require the μ-distortion to be no greater than 9×10^{-5} [14, 15]. For a reasonably broad peak with approximately k^4 slope, centred on $k \sim 10^5 \mathrm{Mpc}^{-1}$, the resulting μ-distortion is $\mu \approx 9 \times 10^{-7}$. This scale corresponds to larger black hole masses than those detected by LIGO and even then, the constraint is not under pressure. The largest possible PBH that can be produced and be consistent with the $\mu-$distortion constraint has mass $\sim 4 \times 10^4 M_\odot$ assuming that the PBH mass equals the horizon mass at the time of horizon entry. This is calculated assuming that the amplitude of the power spectrum is required to reach the current constraints from PBHs, shown by the orange line in Fig. 3.8, which already rules out $f = \Omega_{\rm PBH}/\Omega_{\rm DM} = 1$ on a large range of scales. See Fig. 3.8 for a plot of the full μ-distortion constraints. Note that each point of the blue and purple lines represents the maximum allowed value of A_s at the scale of the peak, k_p, for a delta function power spectrum $\mathcal{P}_{\mathcal{R}} = A_s \delta(\log(k/k_p))$ (blue line) and a k^4 power spectrum $\mathcal{P} = 4A_s(k/k_p)^4$ cut off to zero for $k > k_p$ (purple line), after having integrated over each to find the total contribution to the μ-distortion value. We note that (3.20) captures only the μ-distortion induced by dissipation of scalar modes. Tensor modes can also produce dissipation distortions, with a resulting μ-distortion given by $\mu_h \approx \int_{k_{\rm min}}^{\infty} \frac{dk}{k} \mathcal{P}_h(k)\, W_{h,\mu}(k)$, [247]. However the corresponding window function $W_{h,\mu}(k)$ is such that the overall distortion is some six orders of magnitude smaller for a nearly scale-invariant spectrum (independent of the amplitude), although it has much broader support, and is sensitive to power up to scales approaching

$k \sim 10^5\,\mathrm{Mpc}^{-1}$. As we review in the next subsection, scalar perturbations can source tensor perturbations at second order, and any enhancement of the primordial power spectrum at small scales will source enhanced tensor perturbations at commensurate scales. Although these offer no meaningful constraints with present day sensitivities, a PIXIE-like survey [248] (with sensitivity to μ-distortions as small as $\mu \sim 10^{-8}$) could be sensitive to primordial power spectrum enhancements of up to $\mathcal{P}_\mathcal{R} \sim 10^{-2}$ at $k \sim 10^5\mathrm{Mpc}^{-1}$ due to the dissipation from secondarily produced tensors.

3.4.2 Pulsar Timing Arrays

Although scalar and tensor perturbations decouple at linear order, if the power spectrum is sufficiently boosted to generate PBHs a potentially observable amplitude of second-order gravitational waves (GWs) will be generated [109, 209, 249–259]. This can be intuited as arising from interactions of the form $h_{ij}\partial_i\mathcal{R}\partial_j\mathcal{R}$. Specifically, the transverse traceless projection of the spatial part of the 'stress-tensor'[10] of the curvature perturbation can source tensor perturbations at second order, resulting in an induced contribution to the tensor power of the form

$$\mathcal{P}_h(\tau, k) = \int_0^\infty dv \int_{|1-v|}^{1+v} du\, K(\tau, u, v)\mathcal{P}_\mathcal{R}(ku)\mathcal{P}_\mathcal{R}(kv) \tag{3.22}$$

where $K(\tau, u, v)$ is a rapidly oscillating kernel whose precise form can be found in e.g. [109, 252]. Given the convoluted nature of the integrand, many papers in the literature consider PTA constraints arising from a simple (though unphysical) delta function power spectrum but we also consider a more physical k^4 spectrum with a sharp cut-off. In Fig. 3.7 we plot the GW amplitude for the two power spectra[11]

$$\mathcal{P}_\mathcal{R}(k) = A\delta(\log(k/k_p)), \tag{3.23}$$

$$\mathcal{P}_\mathcal{R}(k) = 4A(k/k_p)^4 \quad \text{for} \quad k < k_p, \quad 0 \ \text{ otherwise.} \tag{3.24}$$

The factor of 4 is included in the latter power spectrum in order that both power spectra are normalised as $\int \mathcal{P}_\mathcal{R}(k)\mathrm{d}\ln k = A$. We also plot the GW spectrum averaged over an efolding (the dashed lines), because the gravitational wave energy of the delta function power spectrum diverges at $k = 2k_p/\sqrt{3}$. The smoothed spectrum is defined by

$$\Omega_{\mathrm{GW}}^{\mathrm{smooth}}(k) = \int_{ke^{-1/2}}^{ke^{1/2}} \Omega_{\mathrm{GW}}(k)\mathrm{d}\ln k. \tag{3.25}$$

[10] By this, we simply mean the symmetric tensor obtained by varying the cubic interaction terms of the form $\mathcal{R}\mathcal{R}h$ in the perturbed action w.r.t. h_{ij}.

[11] We computed this integral in Mathematica using the standard NIntegrate function and converged on the best error settings we could find. A better-purposed integration method should be used to eradicate the noise in Fig. 3.7, which we improved upon in future work [260].

The amplitude of the smoothed spectra is similar near the peaks, and the key difference is the different scaling behaviour at small k. The delta function scalar power spectrum produces a gravitational wave spectrum which scales like k^2 while the k^4 scalar power spectrum produces a gravitational wave spectrum which scales like k^3 at small k. This means that the constraints for values of k_p larger than the scales which PTA best constrain will differ markedly for the two power spectra. The counterintuitive result that a delta function power spectrum does not give rise to the narrowest possible GW spectrum has been observed in numerous papers, e.g. [261, 262]. Since the scalar power spectrum cannot grow faster than k^4, a k^2 tail in $\Omega_{\rm GW}$ cannot be produced by a narrowly peaked scalar power spectrum. It has been suggested that including non-Gaussianity in the calculation can mean that a delta-function or very narrow power spectrum will also induce a k^3 tail in $\Omega_{\rm GW}$ [263]. Other effects of including non-Gaussianities are discussed in [264]. To make Fig. 3.7, we have used $\Omega_{\rm rad,0}h^2 = 4 \times 10^{-5}$ in order to evolve the GW amplitude from horizon entry during radiation-domination until today.

There is, however, a discrepancy in the normalisation between various references in the literature that is most apparent when one tries to calculate the secondary tensor spectrum produced by a scale-invariant scalar power spectrum. In particular, the results of [249] (quoted in [258]), [251] and [109] differ, with the latter reference finding a normalisation that is order 10^{-2} less than the prior references. The precise source of this discrepancy is not immediately apparent to us. However, it is apparent that the numerical integrations necessary to arrive at the final answer are sufficiently involved as to make any analytic simplifications (such as those provided by [109, 265]) advantageous. For this reason, we utilise the explicit analytic form of the kernel $K(\tau, u, v)$ detailed in [109] in what follows – for the purposes of placing observational bounds, it is also the more conservative choice because it leads to a lower normalisation of $\Omega_{\rm GW}$ than several other references. Using the simplifications provided by [109, 265][12], we can also evaluate the secondary tensor perturbations produced by a k^4 steepest growth spectrum.

For each frequency, we use the tightest constraint for a stochastic GW background from various PTA experiments [112, 266, 267] in order to plot the power spectrum constraint in Fig. 3.8. We use the unsmoothed GW spectrum induced by the k^4 scalar power spectrum (in order to avoid needing to choose a smoothing scale which depends on the

[12]These references have resulting analytic forms for the kernel $K(\tau, u, v)$ that agree up to having taken different lower limits in eq. (15) in [109] and the corresponding eq. (33) in [265]. The resulting difference will be negligible whenever the source scalar modes are sub-Hubble. We thank Davide Racco for correspondence on this matter.

experiment) and convert from k to frequency space using $f = kc_s/(2\pi)$ where $c_s = 9.7 \times 10^{-15}Mpc/s$. Over the range of almost two orders of magnitude in k (over 3 orders of magnitude in horizon mass from $\sim 0.1M_\odot - 200M_\odot$, see equation (3.26)) the constraint on the power spectrum is stronger than the constraint from the non-detection of PBHs, meaning that PBHs will not form in significant numbers over this range of scales. We have discussed the many caveats at the beginning of this section. Notice that because the delta function power spectrum has a slower decay of $\Omega_{\rm GW}$ towards small k, the constraints on the power spectrum would become stronger than the k^4 spectrum constraints for scales sensitive to the low-frequency tail of $\Omega_{\rm GW}$.

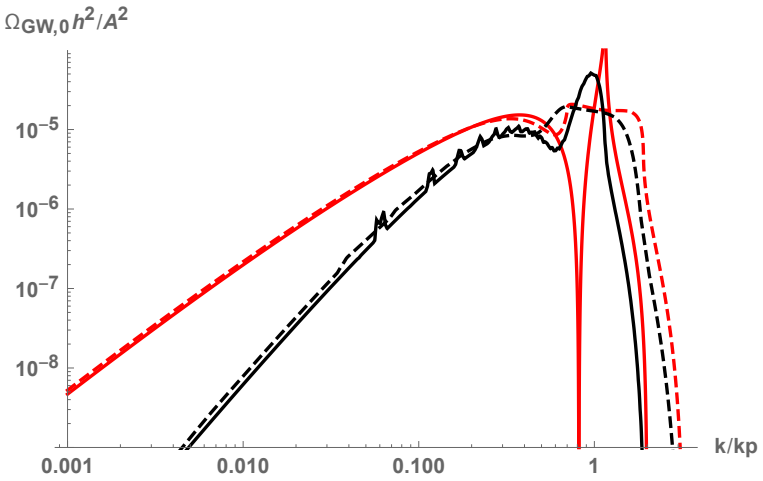

Figure 3.7: The GW amplitude today for a delta function power spectrum and one growing like k^4 in red and black respectively, as described in the main text. The dashed lines are the values smoothed over 1 efolding. The small 'teeth'-like features of the black line are due to numerical noise and they don't affect the power spectrum constraints derived from these curves.

3.4.3 Implications for model building

In Fig. 3.8 we produce a new plot of the constraints on the primordial power spectrum on all scales similar to Bringmann et al. [268], but unlike [268] we do not include constraints from ultracompact minihalos (UCMHs) since they rely on a WIMP DM scenario. For a discussion of recent UCMH constraints see [269–272], and mixed scenarios with both WIMPs and PBHs are discussed in [273–275].

On the largest scales we plot the Planck measurements of the power spectrum [22]. The next relevant constraints on smaller scales come from μ-distortions. Since the power spectrum cannot grow arbitrarily quickly, it is clear that the power spectrum cannot

become large enough to generate PBHs on scales $k < 10^4 \text{Mpc}^{-1}$, subject to the aforementioned assumptions, the most relevant being that the perturbations are Gaussian [276]. Hence there is no need to also show the y-distortion constraints which affect larger scales. The blue line is the upper bound on the amplitude for a monochromatic power spectrum, whilst the dashed purple line is the upper bound on the amplitude for a power spectrum with k^4 slope and immediate drop off. For a constraint on slightly smaller scales than spectral distortions, see [277].

The black line on scales $k \sim 10^7 \text{Mpc}^{-1}$ represents the upper bound due to the PTA constraints, while the relatively flat orange line represents the PBH constraint. The PBH constraints are calculated using values of $f = \Omega_{\text{PBH}}/\Omega_{\text{DM}}$ from [278] and [163] for PBH masses between $\sim 10^{-24}\text{M}_\odot$ and $10^7\text{M}_\odot$. These combine various constraints from e.g. their evaporation, femto-lensing of gamma-ray bursts, neutron-star capture, white dwarf explosions, and microlensing. We use $\delta_c = 0.45$ for definiteness, and we effectively use a delta function for the window function in translating the variance of the perturbations to the amplitude of the power spectrum (in place of $W(q)$ in equation (3.19) which we used for studying the mass function). We also use a linear relation between the density contrast δ and the comoving curvature perturbation $\mathcal{R}$, which is not realistic. We do not include constraints from microlensing for masses $\lesssim 10^{-10} M_\odot$ due to uncertainties concerning the wave effect [278]. The slight dip at $k \sim 10^7$ is caused by including the effect of the change in the equation of state around the time of the QCD transition [231]. Also note that we use

$$\frac{k}{3 \times 10^{22}\,\text{Mpc}^{-1}} = \frac{\pi}{t_i c}\left(\frac{t_i}{t_{\text{eq}}}\right)^{\frac{1}{2}}\left(\frac{t_{\text{eq}}}{t_0}\right)^{\frac{2}{3}}\left(\frac{g_{*,i}}{g_{*,0}}\right)^{-\frac{1}{12}}, \tag{3.26}$$

$$t_i = \frac{10^{-38}(M/1\text{g})}{\gamma}\text{s} \tag{3.27}$$

$$\implies \frac{k}{3 \times 10^{22}\,\text{Mpc}^{-1}} = 10^{19}\frac{\pi}{c}\gamma^{\frac{1}{2}}\left(\frac{M}{1\text{g}}\right)^{-\frac{1}{2}}\left(\frac{1\,\text{s}}{t_{\text{eq}}}\right)^{\frac{1}{2}}\left(\frac{t_{\text{eq}}}{t_0}\right)^{\frac{2}{3}}\left(\frac{g_{*,i}}{g_{*,0}}\right)^{-\frac{1}{12}} \tag{3.28}$$

to convert between k and PBH mass, where t_i is horizon entry time of the overdensity, M is the PBH mass, γ is the fraction of the horizon mass that will collapse to form the black hole which we take to be 1 given the uncertainty in the literature, $t_{\text{eq}} \approx 2 \times 10^{12}\,\text{s}$ is the cosmic time of matter-radiation equality and $t_0 \approx 4 \times 10^{17}\,\text{s}$ is the cosmic time today. We take the effective degrees of freedom today to be $g_{*,0} \approx 3.36$ and $g_{*,i}$ is the effective degrees of freedom at the time of horizon entry.

In order to reach the current constraint on the number of PBHs (orange line) from the amplitude of the power spectrum at CMB scales, the growth must begin at $k \gtrsim 10^3\,\text{Mpc}^{-1}$ in order to avoid constraints from the μ-distortions, since it can only

grow as fast as k^4. This is shown by the left-most dotted black line in Fig. 3.8. This implies a maximum PBH mass which can be generated of $4 \times 10^4 M_\odot$ corresponding to $k \sim 7 \times 10^4\,\mathrm{Mpc}^{-1}$. This point is where the left dashed black line (with k^4 slope) crosses the PBH constraint line, and it is also where the dashed purple line which marks the distortion constraints for a k^4 growth crosses the PBH constraint line. Notice that the blue line (for a delta function power spectrum) crosses the PBH constraint line at a larger scale $k \sim 4 \times 10^4\,\mathrm{Mpc}^{-1}$, corresponding to a PBH mass of $2 \times 10^5\,M_\odot$. The difference between these two masses demonstrates the additional restriction on PBHs caused by the restriction on the steepest possible growth of the power spectrum.
Similarly, in order to avoid PTA constraints, the power spectrum growth must begin at $k \lesssim 10^4\,\mathrm{Mpc}^{-1}$ as shown by the right-most dotted black line in Fig. 3.8. This assumes that the power spectrum can drop off instantaneously to $\mathcal{P}_\mathcal{R} \lesssim 10^{-3}$ after the peak to avoid the PTA constraints – see App. C.2 for a discussion on this point. For PBHs with masses larger than those constrained by PTA experiments, the power spectrum is free to grow from $k \gtrsim 6 \times 10^5\,\mathrm{Mpc}^{-1}$, as a k^4 slope will clear the PTA constraints from this scale, and there are no severe constraints on the power spectrum on smaller scales. In order to produce PBHs on a scale $k \sim 10^6\,\mathrm{Mpc}^{-1}$ and avoid the μ-distortion constraint, the power spectrum needs to grow at least as steeply as $k^{1.2}$ on the scales between those two constraints.
Early matter-dominated scenarios are of interest because the lack of pressure means that PBHs are able to form much more easily and have been considered recently in e.g. [136, 160, 163, 246]. This means that the amplitude of the power spectrum is related to the number of PBHs by a power law instead of logarithmically as is the case in radiation-domination that we have assumed to plot the orange line in Fig. 3.8. One could then question whether the constraints on the power spectrum change more quickly than the k^4 limit. Using constraints on the power spectrum from [246], we have verified that they do not change more quickly than k^4, and therefore that PBHs of every possible mass can still be generated while respecting this bound on the power spectrum growth.

3.5 Conclusions

We have shown that the steepest possible growth of the primordial power spectrum is given by $n_s - 1 = 4$ during canonical single-field inflation, independent of the shape of the inflaton potential. Such a rapid growth is only possible when the inflaton makes a rapid transition from "slow-roll" inflation to non-attractor inflation, characterised by an

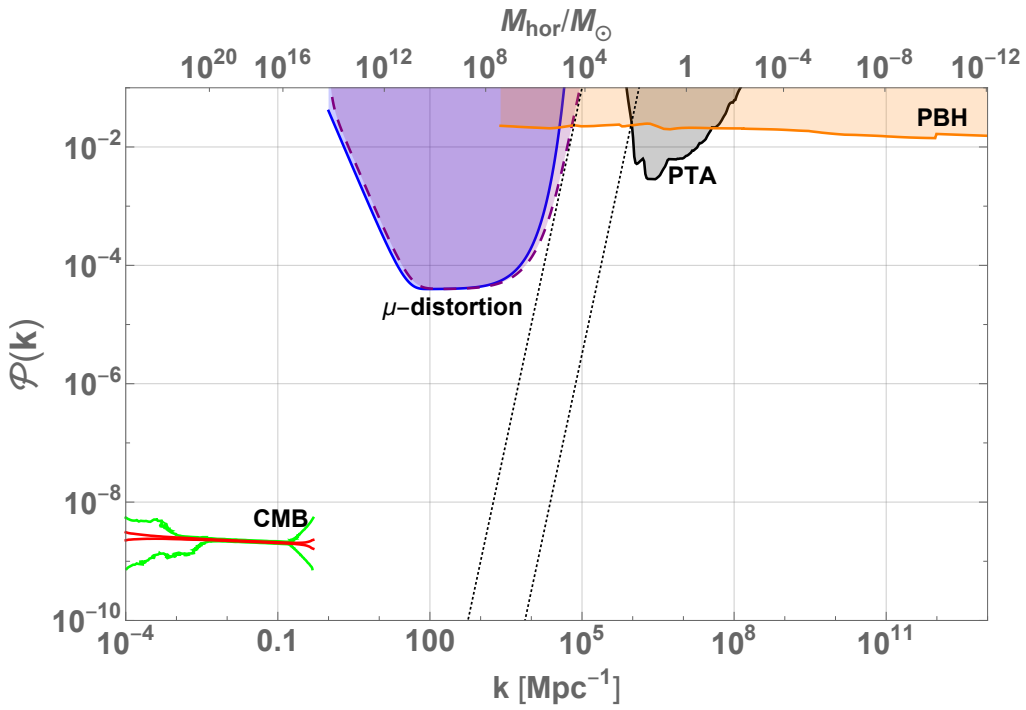

Figure 3.8: Observational constraints on the power spectrum. The lines at small k are the Planck 1σ and 3σ measurements. On much smaller scales there are only upper bounds; shaded regions are disallowed. The solid blue line shows the upper bound from μ-distortions for a delta function power spectrum, $\mathcal{P}_\mathcal{R} = A_s\delta(\log(k/k_p))$, as a function of k_p, and the solid orange line shows the PBH upper bounds, subject to the uncertainties discussed in the main text. The dashed purple line shows the upper bound from μ-distortions for the steepest growth power spectrum $\mathcal{P}_\mathcal{R} = 4A_s(k/k_p)^4$ which drops to zero for $k > k_p$, and the solid black line shows the PTA upper bounds for the same power spectrum. The factor of 4 is included so that it has the same normalisation as the delta function power spectrum when integrating with respect to $\ln k$. The dashed black lines have a k^4 slope.

almost exactly flat potential, and remains true even if the potential is not always decreasing[13]. In the standard case of single-clock inflation - implying the curvature perturbation freezes out shortly after horizon crossing - the power spectrum grows less steeply, and is bounded by $n_s - 1 < 3$. It would be of great interest to understand whether our bound can be violated in more complicated models of inflation. For example, see [55], which in some cases requires an ad hoc power spectrum with steepness up to k^8 in order to evade power spectrum constraints while generating PBHs, which our bound implies is not possible in the context of single-field inflation.

We have calculated analytic expressions for the most rapid growth of the power spectrum possible, by matching the curvature perturbation between various phases of inflation, characterised by different rates at which the slow-roll parameter ϵ decreases. The steep k^4 growth arises during times when modes exiting the horizon are affected by both periods of inflation. We have also provided a way to reconstruct the inflaton potential given an arbitrary time evolution of the expansion rate during inflation specified by $\epsilon(t)$ in App. C.

Due to the phenomena of critical collapse to form PBHs, the PBH mass spectrum cannot be arbitrarily close to monochromatic. We have shown that the mass spectrum is remarkably insensitive to the shape of the power spectrum close to its peak amplitude, with everything from a gentle growth, $n_s - 1 = 0.1$ to the extreme (and impossible) case of a delta function power spectrum producing a comparable width for the mass function of PBHs. This distribution approximately resembles a log-normal distribution and is often parametrised as such when studying extended mass functions, e.g. [94, 98].

In Fig. 3.8 we have combined the key measurements and constraints on the primordial power spectrum, showing how on various scales, CMB measurements, CMB spectral distortion constraints or PTA constraints all force the power spectrum to be too small to generate PBHs. There does however remain a window between the latter two constraints which is sufficiently broad such that the power spectrum can grow and produce large PBHs without conflicting either of those constraints and without requiring the perturbations to be non-Gaussian.

[13] Note added: As we were preparing this paper, ref. [279] appeared, aiming to derive a lower bound of $\eta > -6$ from causality arguments. However, the matching calculation on which this is based neglected to impose both Israel junction conditions (A.10), imposing only the continuity of $\mathcal{R}_k$. Furthermore, the correct causality criteria one should impose is that the commutator of the curvature perturbation at two points should vanish at space-like separation, trivially satisfied even when matching with an intermediate phase of $\eta < -6$.

We plot forecasted constraints on the power spectrum in Fig. 3.9. The sensitivity curves for SKA and LISA are extracted from [280] and do not include the possible degradation due to astrophysically generated gravitational waves. LISA covers the scales corresponding to a possible window where PBHs could consist of all the DM, with masses in the range $M_{\rm pbh} \sim 10^{-13} - 10^{-7} M_{\odot}$ [255]. Of particular interest is how the gap between future μ-distortion constraints assuming a PIXIE-like experiment which can probe $\mu = 2 \times 10^{-8}$ and the existing PTA constraints becomes a factor of 2 in k-space, corresponding to less than an e-folding of inflation, which in practice means constructing a model which grows at the maximum rate and then decreases again is unrealistic. A more detailed treatment of the PTA constraint at low frequency would probably completely close the gap, and the addition of SKA constraints from pulsars does close the gap. The difference of a factor of 8 in k between where the two different forecasted μ-distortion lines cross the PBH line show how much more powerful the PIXIE constraint on PBHs becomes once including the maximum growth rate of the power spectrum. Therefore, PIXIE combined with PTA constraints and the steepest growth rate that we have derived would be able to rule out the generation of LIGO mass PBHs, unless the initial perturbations are sufficiently non-Gaussian on the relevant range of scales. Finally, the combined constraints from the CMB, a PIXIE-like experiment, SKA and LISA will almost completely rule out Gaussian perturbations being able to generate any PBHs with masses greater than $10^{-15} M_{\odot}$.

Note added: after submitting the first version of our paper, [281] appeared which deals with current and future constraints on induced gravitational waves. We would like to thank the authors for helpful discussions and comments on both of our papers.

3.6 Acknowledgements

We thank David Seery for extensive help with his CppTransportCode [223]. CB is a Royal Society university research fellow. SP is supported by funds from Danmarks Grundforskningsfond under grant no. 1041811001. PC acknowledges support from the UK Science and Technology Facilities Council via Research Training Grant ST/N504452/1. We are grateful to Cliff Burgess, Jens Chluba, Dani Figueroa, Kazunori Kohri, Ilia Musco, Davide Racco, David Seery, Tommi Tenkanen, Takahiro Terada, Jennie Traschen, Michael Trott, and Sam Young for valuable discussions over the course of this investigation. We acknowledge the Centro de Ciencias de Benasque Pedro Pascual for providing an excellent working environment during the 2017 Understanding

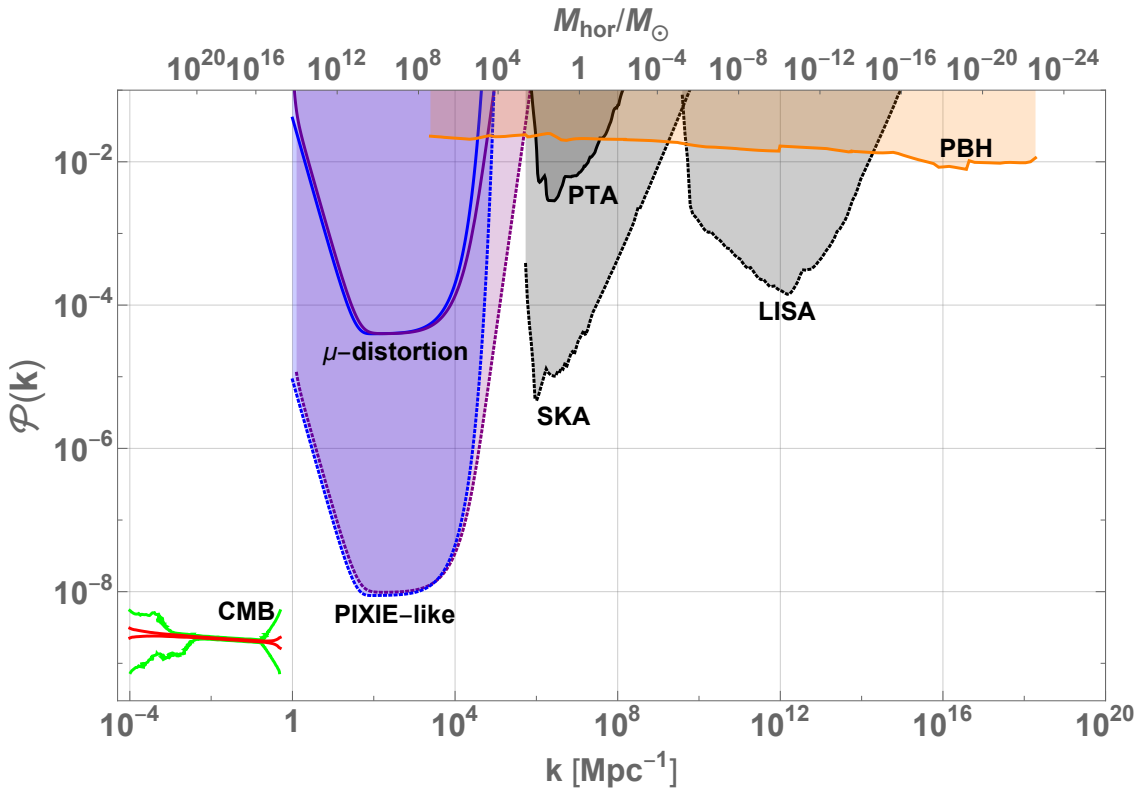

Figure 3.9: Current and forecasted constraints on the amplitude of the power spectrum. The solid lines are the same as in Fig. 3.8, apart from the x-axis which is extended to the smallest scales that PBHs constrain, corresponding to the horizon scale which generates a PBH that decays during big bang nucleosynthesis. The dashed lines show forecasted future constraints from a PIXIE-like satellite for μ-distortions (the dashed blue line assumes a delta function power spectrum while the purple line has a power spectrum growing at the maximum rate of k^4), and the dashed black lines are induced gravitational wave forecasts for a k^4 scalar power spectrum with a cut off, using PTA constraints from SKA and from the LISA satellite on smaller scales. Shaded regions are disallowed.

Cosmological Observations meeting where this work was initiated.

Chapter 4

Small-scale primordial fluctuations in the 21cm Dark Ages signal

Philippa S. Cole[1] and Joseph Silk[2,3,4]

[1] Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, United Kingdom

[2]Institut d'Astrophysique de Paris, UMR 7095 CNRS, Sorbonne University, 75014 Paris, France

[3]Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

[4]Beecroft Institute of Particle Astrophysics and Cosmology, Department of Physics, University of Oxford, Oxford OX1 3RH, United Kingdom

4.1 Introduction

The fluctuations in density left over at the end of inflation are the best probe for how inflation itself happened. Since these aren't observable directly, we must rely on mapping the evolution of these overdensities and underdensities which eventually gravitationally collapsed to form the structures that we see today. Measuring the late-time matter power spectrum will enable us to track back and predict how the fluctuations were distributed immediately after inflation, which is quantified with the primordial power spectrum. The 21cm power spectrum is a tracer of the underlying matter power spectrum, and there are two redshift windows where this signal appears in absorption: during the Dark Ages between redshifts $z \sim 30 - 200$ and during Cosmic Dawn around $z \sim 15 - 20$. In order to capture the matter distribution before it is complicated by the astrophysical

processes involved in reionization and galaxy formation, it is best to look at redshifts above ~ 30. Above redshift 30, the matter in the Universe was predominantly made up of neutral hydrogen and was therefore totally dark. However, due to neutral hydrogen's spin-flip transition, the distribution of hydrogen can be detected with 21cm observations. After recombination, when the photons decoupled from the newly formed neutral hydrogen and began free-streaming towards us as the Cosmic Microwave Background (CMB), the Universe continued to cool, but Compton scattering maintained the temperature of the CMB and the gas in equilibrium. By around $z \sim 200$, the Universe had cooled enough such that Compton scattering was no longer efficient enough to keep the gas and the CMB in equilibrium, and so the gas began to cool faster than the CMB. This meant that most of the neutral hydrogen was in its unexcited state, and therefore able to absorb CMB photons at the characteristic wavelength of 21cm. It is this difference in temperature of the CMB that is observable. The absorption line of the photons is redshifted from the initial wavelength of 21cm, and therefore the frequency of the radiation that arrives at detectors determines the redshift slice from which the signal originated.

Current CMB measurements [1] constrain the primordial power spectrum very tightly on scales $k \sim 10^{-4} - 0.1\,\mathrm{Mpc}^{-1}$ to be of amplitude 2×10^{-9}. This means that, while a detection of the 21cm signal from the Dark Ages on any scale would be a huge achievement, it is unlikely that anything new will be uncovered about the primordial power spectrum unless smaller scales are probed. This should in theory be possible with 21cm observations if high enough redshifts can be targeted. For the best hope of a detection of the Dark Ages 21cm power spectrum, an interferometer on the Moon [5, 123, 282] (or beyond [122, 283]) would be required to reach the small scales, that remain linear, at and above redshift 50. The constraining power of 3d 21cm power spectra measurements are illustrated in figure 4.1. Compared to 2d CMB measurements on large scales, and 3d large-scale structure probes on intermediate scales, both ground and space-based 21cm interferometers have the potential to access an unprecedented number of modes.

If Planck's measurements of the primordial power spectrum on large scales extrapolate to smaller scales, then the current most-favoured inflationary models (single-field, slow-roll) will continue to be preferred. However, any deviation from the low-amplitude, scale-invariant primordial power spectrum on small scales will point towards a different inflationary scenario, as well as lead to other potential observables [211]. For example,

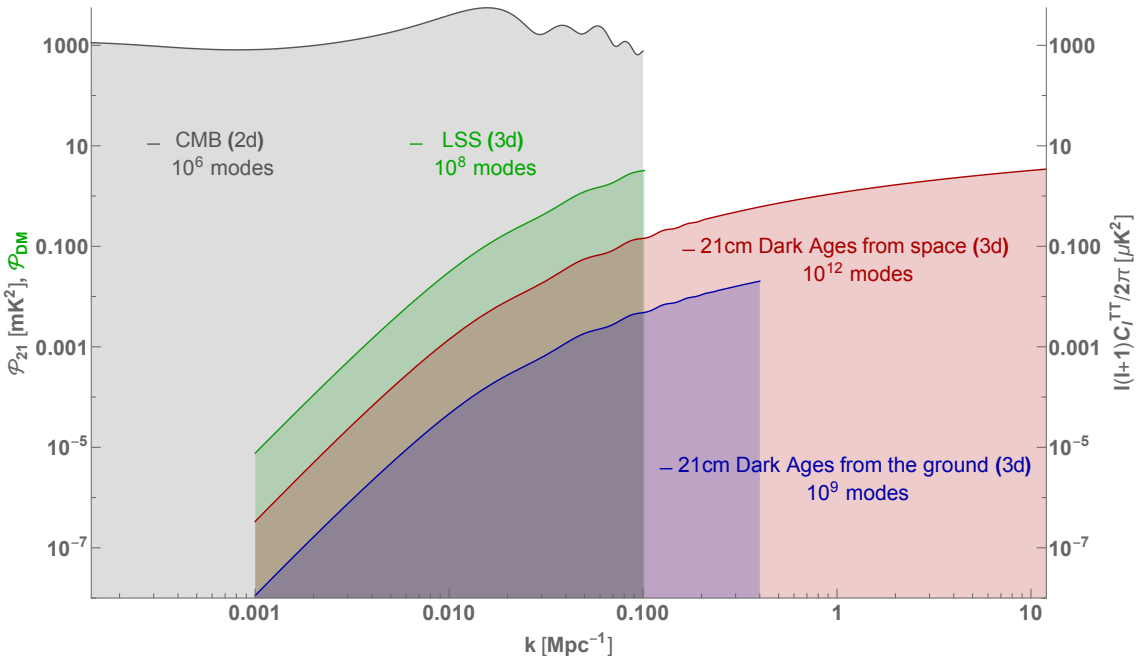

Figure 4.1: An illustration of the scope of different cosmological probes for accessing large numbers of modes. Note that the y-axes are different for each probe as described here. In grey is the TT angular power spectrum in units of μK^2 as shown on the right-hand axis, with multipoles roughly mapped to wavenumbers by $l \sim 14000k/\mathrm{Mpc}^{-1}$ [12]. In green is the dimensionless 3d matter power spectrum $\mathcal{P}_{\mathrm{DM}}$ computed with CAMB at redshift 1, which large-scale structure probes such as LSST and EUCLID will be sensitive to on scales between $k \sim 0.001 - 0.1\,\mathrm{Mpc}^{-1}$ [13] up to around redshift 2.5. In blue is the 3d 21cm power spectrum $\mathcal{P}_{21}$ in units of mK^2 at redshift 27, which is the highest redshift accessible from ground-based experiments such as HERA and SKA. In red is the 3d 21cm power spectrum in units of mK^2 at redshift 50, which would be accessible from the Moon. Note that the maximum k for 21cm experiments is solely based on the angular resolution for maximum baselines given in table 4.1.

an enhancement in small-scale power could lead to the production of primordial black holes or ultra-compact mini haloes, which could in turn provide the seeds for supermassive black holes and the most massive galaxies [284–286].
21cm observations can therefore teach us about both inflation and current observables at the same time. This could be complemented by a measurement of the integrated small-scale power via spectral distortions of the CMB, or by the detection of second order gravitational waves which would imply large primordial scalar perturbations, or by the detection of primordial non-Gaussianity.
This paper is laid out as follows. In section 4.2 we outline the basics of 21cm Cosmology that will enable us to produce the 21cm power spectra in section 4.3 from inflation-motivated primordial power spectra. In section 4.4 we demonstrate the play-off between density fluctuations produced during inflation and Poisson fluctuations in the 21cm power spectrum for different masses and abundances of PBHs and comment on their relevance with respect to accretion effects. Finally we discuss possibilities for detection in section 4.5 and then conclude.

4.2 21cm basics

The spin temperature T_s of neutral hydrogen is defined as

$$\frac{n_1}{n_0} = 3e^{-\frac{T_*}{T_s}} \tag{4.1}$$

where n_1 and n_0 are the number densities of neutral hydrogen in excited and ground states respectively with $n_H = n_0 + n_1$, and $T_* = 0.068$K is the temperature corresponding to the energy difference between the ground and excited states.
In order to see how the spin temperature evolves in time we can write down the rate equations for the hydrogen atoms

$$n_0(n_H\kappa_{01} + B_{01}u_\nu) = n_1(A_{10} + B_{10}u_\nu + n_H\kappa_{10}) \tag{4.2}$$

where A_{10} is the probability of spontaneous emission known as the Einstein A coefficient, B_{10} is the probability of stimulated emission (when an incoming CMB photon causes another photon to be emitted and for the atom to drop from its excited to its ground state), and B_{01} is the probability of stimulated absorption (when the atom absorbs a CMB photon and it jumps from its ground to its excited state). The blackbody CMB photons which mediate this process are described by the radiation field u_ν. κ_{10} and κ_{01} are the collisional rate coefficients for which we use the values tabulated in [287] - these

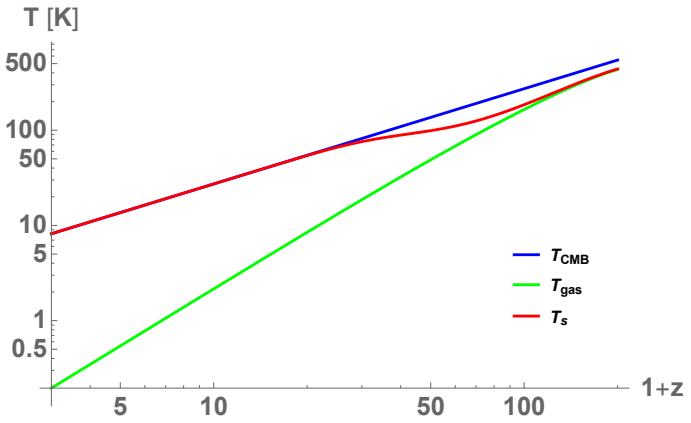

Figure 4.2: The evolution of the CMB temperature, gas temperature and spin temperature as a function of redshift.

describe the rate at which the atoms change states when they collide. In the limit of $T_* \ll T_{\rm CMB}, T_s$, (4.2) can be solved and the spin temperature can be written in terms of the gas temperature, the CMB temperature, the collisional rate coefficients and the Einstein A coefficient:

$$T_s = T_{\rm CMB} + (T_{\rm gas} - T_{\rm CMB})\frac{C_{10}}{C_{10} + A_{10}\frac{T_{\rm gas}}{T_*}} \tag{4.3}$$

with $C_{10} = n_H \kappa_{10}$. Figure 4.2 shows the evolution of the gas temperature, CMB temperature and spin temperature as a function of redshift. All three are in equilibrium until a redshift of $z \sim 200$ due to residual free electrons Thomson scattering off the gas and the CMB photons. The gas then begins to cool as $T_{\rm gas} \propto (1+z)^2$ while the CMB cools as $T_{\rm CMB} \propto (1+z)$ and the spin temperature therefore deviates from both. By around $z \sim 30$, the collision rate becomes subdominant to the Hubble expansion and the spin temperature couples to the CMB temperature once more. This redshift window $z \sim 30 - 200$ is therefore the window where a 21cm signal could be observable, in absorption relative to the CMB. The quantities $T_{\rm gas}$, n_H, $\overline{x}$ and $T_{\rm CMB}$ are computed using RECFAST [288]. Note that there is also an absorption signal during Cosmic Dawn and the Epoch of Reionization at lower redshifts, which result in another more prominent dip in T_s which we have not shown here.

The observable is not the spin temperature, but the brightness temperature T_{21} which describes the contrast between the spin temperature and the CMB

$$T_{21} = \tau \frac{T_s - T_{\rm CMB}}{z+1} \tag{4.4}$$

where the optical depth $\tau \ll 1$ depends on the neutral hydrogen density local to the

absorption

$$\tau = \frac{3c\lambda^2 h A_{10} n_H}{32\pi k_b T_s H(z)} \tag{4.5}$$

which can be approximated as [289]

$$\tau_a = 0.025 \frac{T_{\rm CMB}}{T_s} \left(\frac{1+z}{51}\right)^{\frac{1}{2}} \left(\frac{\Omega_m}{0.27}\right)^{-\frac{1}{2}} \left(\frac{\Omega_b h}{0.035}\right). \tag{4.6}$$

The sky-averaged brightness temperature can shed light on Cosmic Dawn and the Epoch of Reionization around redshift 10, but for the purposes of probing the scale-dependence of the 21cm signal at different redshift slices (and hence the primordial power spectrum), we will be interested in the 21cm fluctuations which track the matter fluctuations.
We will compute the 3d isotropic 21cm monopole transfer functions numerically using CAMB [24], which includes fluctuations in the density of the baryons, gas temperature, ionization fraction, radial peculiar velocities and Lyman-alpha pumping efficiency. However, fluctuations in the baryons will be largely dominant during the Dark Ages, before luminous sources have formed. The linear Boltzmann equations used in CAMB to calculate the 21cm monopole transfer functions are laid out in [290]. The non-linear effect of the relative velocity between dark matter and baryons is not captured by CAMB. This would enhance the 21cm power spectrum on large scales $k < 1\,{\rm Mpc}^{-1}$, suppress it on small scales $k > 200\,{\rm Mpc}^{-1}$ and enhance it again on very small scales $k > 2000\,{\rm Mpc}^{-1}$ by order unity [291]. Since we are interested in boosts in power beyond $k > 1\,{\rm Mpc}^{-1}$, and do not expect to be sensitive to scales smaller than $k \sim \mathcal{O}(10)\,{\rm Mpc}^{-1}$ even with futuristic radio interferometers, we do not include their effects here.

4.3 Predictions for 21cm power spectra given different primordial models

If the measurement of the primordial power spectrum on large scales by Planck extrapolates to small scales, it will be of the form

$$\mathcal{P}_{\mathcal{R}} = A_s \left(\frac{k}{k_p}\right)^{n_s - 1} \tag{4.7}$$

with $k_p = 0.05\,{\rm Mpc}^{-1}$ and $n_s \approx 0.965$ [1]. However, there may be an increase in power on small scales, which is theoretically motivated by the potential need to explain the seeds of supermassive black holes and the most massive galaxies, as well as the possible existence of primordial black holes or ultra compact minihaloes.

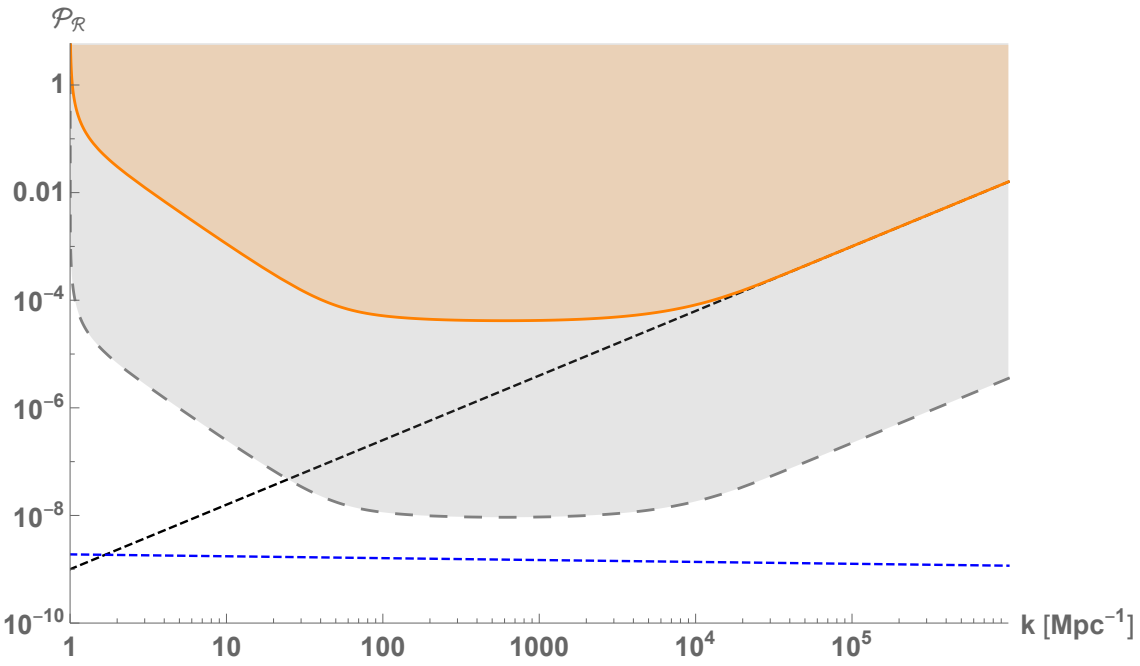

Figure 4.3: Constraints from COBE/FIRAS [14, 15] on the primordial power spectrum [16, 17] due to μ-distortions in orange - the shaded region is disallowed. Future constraints from a PIXIE-like [18] probe in grey. Constraints are calculated with an input primordial power spectrum that grows as $k^{1.2}$ with a sharp cut-off. The value of each point on the constraint curve represents the maximum amplitude that the peak of such a primordial power spectrum can be without conflicting with the spectral distortion constraints. The black dashed line grows as $k^{1.2}$ from $k = 1\,\mathrm{Mpc}^{-1}$ and $\mathcal{P}_\mathcal{R} = 10^{-9}$, i.e. the steepest that the power spectrum can be if it starts to grow at $k = 1\,\mathrm{Mpc}^{-1}$. The blue dashed line is the canonical CDM parameterisation of the primordial power spectrum with $A_s = 2.09 \times 10^{-9}$ and $n_s = 0.965$ [1].

There are various constraints on the primordial power spectrum which must be respected. The most relevant on the scales that 21cm observations may be able to probe are those from spectral distortions of the CMB [14, 15], in particular y-distortions constrain scales up to $k \sim 10\,\mathrm{Mpc}^{-1}$, and μ-distortions constrain scales between $k \sim 1 - 10^5\,\mathrm{Mpc}^{-1}$ [16, 17]. In order to avoid these, the fastest that the power spectrum can grow from $k = 1\,\mathrm{Mpc}^{-1}$ (where CMB constraints finish) is at a rate of $k^{1.2}$. See figure 4.3, where the constraint plotted at a given value of k represents the maximum amplitude allowed by μ-distortion constraints for a power spectrum that grows like $k^{1.2}$, peaks at that value of k, and then falls to 0 for larger k. This sharp cut-off is a conservative choice [95], however, if the power spectrum can't decrease that quickly [96] then the constraints will be tighter. For single-field models of inflation with canonical kinetic terms, the fastest that the power spectrum can grow is $k^5 \log k^2$ [96]. However, when limited observationally by a maximum growth of seven orders of magnitude between $\mathcal{P}_\mathcal{R} \sim 10^{-9}$ and 10^{-2}, the fastest growth can be approximated by k^4 [95], which also requires less restrictions on the evolution of the slow-roll parameters (see also [292]). The largest scale where such a fast boost can occur whilst still avoiding μ-type spectral distortion constraints is $k \sim 10^3\,\mathrm{Mpc}^{-1}$. 21cm observations offer a complimentary probe of the primordial power spectrum on these scales to spectral distortions, because they can probe the scale-dependence, whereas spectral distortion constraints are only sensitive to the integrated contribution of power across the range of scales. If an experiment such as PIXIE [248] (see [18] for a recent proposal) detected a larger signal than expected from a Planck extrapolated power spectrum, the 21cm Dark Ages signal could identify which scales are contributing to the surplus.

We now find the predicted 21cm signal for 4 different primordial power spectra at redshift 27, the largest redshift accessible from the ground, and at redshift 50 when the signal is largest and would be accessible by a future lunar array. We compute the 21cm transfer functions with CAMB [24], which we combine with the four different primordial power spectra to produce the 3d 21cm power spectra.

The four primordial power spectra chosen are shown in figure 4.4; in black is the spectrum extrapolated from the CMB measurements of equation (4.7), in orange is the piecewise primordial power spectrum that matches CMB measurements until $k = 1\,\mathrm{Mpc}^{-1}$ and then grows like $k^{1.2}$ representing the maximum growth possible whilst evading spectral distortion constraints, in grey is the primordial power spectrum that matches CMB measurements until $k = 1000\,\mathrm{Mpc}^{-1}$ and then grows like k^4 and in purple

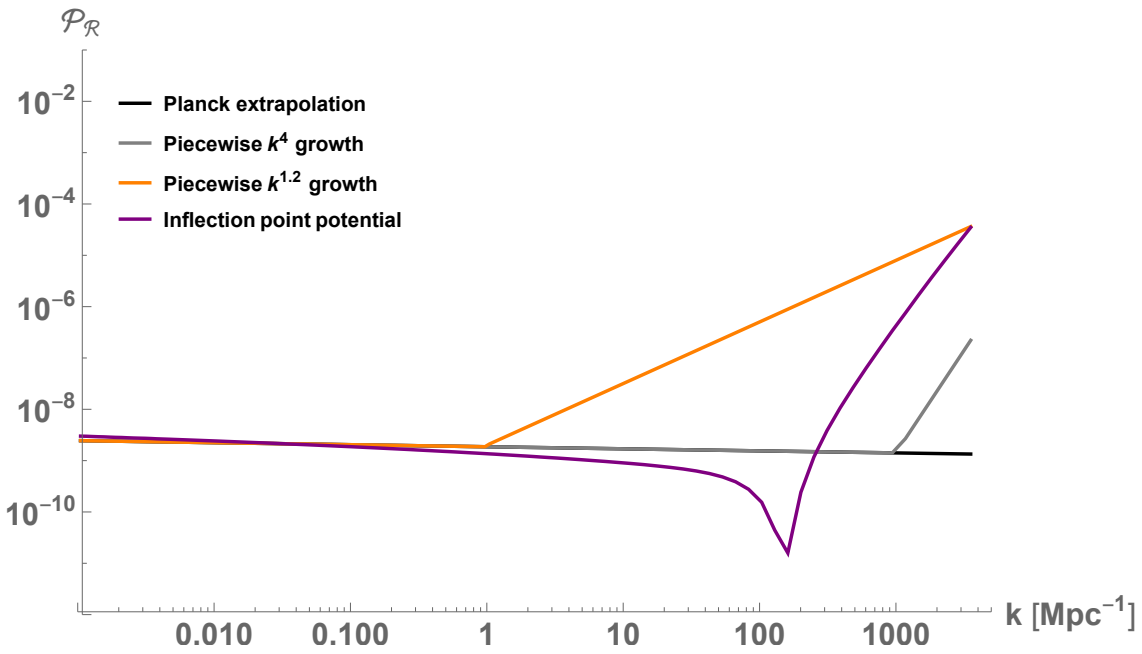

Figure 4.4: The primordial power spectra corresponding to the 21cm power spectra in figure 4.5.

is the primordial power spectrum produced from a 'realistic' inflationary potential [11] that grows steeply on small scales before flattening off. The corresponding 21cm power spectra are plotted at redshift 50 in figure 4.5.

An excess in power can be seen for the piecewise $k^{1.2}$ growth which begins at $k = 1\,\mathrm{Mpc}^{-1}$ because the 21cm signal has a chance to grow before it is damped at large k. The piecewise k^4 growth is just visible in comparison to the Planck-extrapolated spectrum at a scale of $k = 1000\,\mathrm{Mpc}^{-1}$. The realistic and smooth model of [11] shows a significant decrease in power which is common to inflection-point models of inflation [95, 293], and the subsequent growth also produces a signal in excess of the Planck model beyond $k \sim 300\,\mathrm{Mpc}^{-1}$, although note that this could be suppressed by relative velocity effects [291].

If the primordial power spectrum is boosted on scales beyond $k \sim 0.1\,\mathrm{Mpc}^{-1}$, it is plausible that 21cm interferometers will be sensitive to the signal (as well as possibly inferred from detections of the global signal [294, 295]), and be able to distinguish it from the signal expected from the simplest extrapolation of the Planck measurements to small scales. This would test whether a more complicated inflationary scenario that goes beyond the slow-roll approximation is required. We will discuss possibilities for detection in section 4.5.

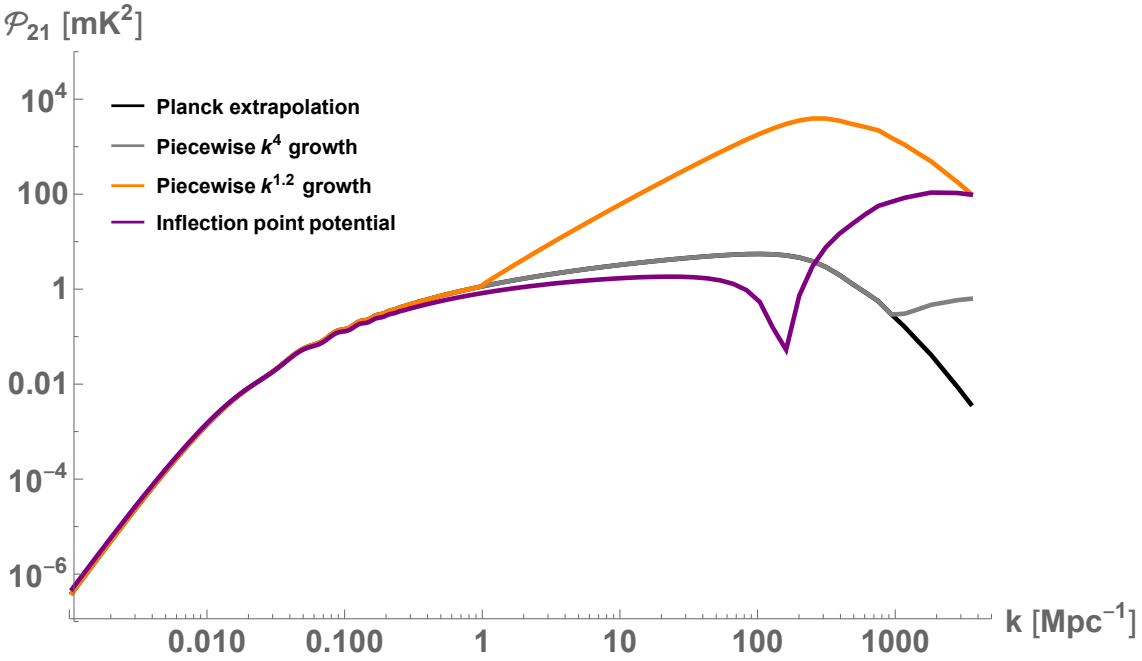

Figure 4.5: 21cm power spectrum predictions at redshift 50 for 4 different primordial power spectra as described in the text.

4.4 Primordial black hole production

If the primordial power spectrum continues to grow on small scales beyond those plotted in figure 4.4 until it reaches amplitudes of order $10^{-3} - 10^{-2}$ [50, 56, 57, 62, 296], primordial black holes would be necessarily formed on that scale (corresponding to a mass via $M/M_\odot \approx (k/k_\odot)^{-2}$ [50]) due to the collapse of large density perturbations reentering the horizon. The imprint on the 21cm signal of PBHs has been investigated before by, for example, [285, 297–299], where the effect of accretion onto the PBHs is taken into account, as well as the Poisson fluctuations sourced by the discrete distribution of PBHs [300]. However, the primordial fluctuations that are necessary for the PBHs to form in the first place have been previously neglected. In this work, we investigate the interplay between the Poisson fluctuations and the initial fluctuations generated during inflation, as plotted in the previous section, see figure 4.4. We focus on regimes where accretion effects are likely to be small, and show that the primordial power spectrum cannot be neglected when calculating the 21cm signal in these cases.

Whilst PBHs will begin accreting matter at the beginning of the matter-dominated epoch, they only have an effect on the 21cm signal once the temperature of the CMB is low enough so that deviations to the spin temperature and hence the brightness temperature are noticeable. Deviations are caused by the heating and ionisation of the

IGM due to the matter falling onto the PBHs. The energy is radiated by either x-ray emission or advection-dominated accretion flow and can have both local and global effects [299].

In this paper, our focus is on the signal at redshift 50, since it is with very low frequency radio interferometers that the smallest scales will be detectable. At redshift 50, effects of accretion on the brightness temperature are expected to be small [299], except for in the case of very large PBH masses and/or abundances. Since a boost in small-scale primordial power would be necessary for even just one PBH to be formed [246], we consider small $f_{\rm PBH}$ so as to emphasize the importance of including the primordial power spectrum contribution, when there would be no PBH signature in the 21cm signal from accretion effects. In order to explain the seeds of supermassive black holes with PBHs, only small abundances of PBHs would be required, meaning that quantifying the 21cm power spectrum in these cases is well-motivated. There may still be small effects on the 21cm power spectrum due to accretion, for example [299] show that the 21cm power spectrum differs by a factor of ~ 4 between the cases of no PBHs and $f_{\rm PBH} = 10^{-4}$ with $M_{\rm PBH} = 100\,{\rm M}_\odot$ at redshifts between $z \sim 15 - 20$, i.e. for the Cosmic Dawn absorption signal. The effect at higher redshifts during the Dark Ages should be even smaller given that the signal itself is expected to be around five times smaller than during Cosmic Dawn. Given also the uncertainties in the modelling of the accretion mechanism, for example the fact that spherical accretion is assumed [299], we will not include them here and leave a detailed analysis to future work.

If one was instead interested in constraining PBHs as a dark matter candidate with $f_{\rm PBH}$ as close to 1 as possible, accretion effects would be imperative to understand fully and include in the calculation. In addition, at redshifts below $z \sim 30$, the effects of accretion are much more pronounced, although still heavily dependent on the mass and abundance of the PBH population. They may directly compete with the contribution from primordial fluctuations, and we leave a full investigation to future work. Given that ground-based interferometers which would be sensitive to redshifts up to $z \sim 27$ cannot reach small enough scales to be sensitive to a boost in the primordial power spectrum, accretion effects may be the only way of detecting PBH signatures in the 21cm signal, as has been previously investigated [285, 297–299].

Even if they made up all of the dark matter, the typical separation between PBHs is much larger than the comoving horizon size at the time of formation, and therefore their distribution can be described by a Poisson distribution (unlike particulate dark matter).

The Poisson fluctuations are sourced by the already-formed PBHs, and the power spectrum of the Poisson fluctuations is [300, 301]

$$P_{\rm Poisson}(z) = \frac{9}{4}(1+z_{\rm eq})^2 D^2(z) \frac{f_{\rm PBH}^2}{n_{\rm PBH}} \tag{4.8}$$

where $\frac{9}{4}(1+z_{\rm eq})^2$ is the transfer function for isocurvature perturbations, since they are only coupled to the dark matter content, and $z_{\rm eq}$ is the redshift of matter-radiation equality which we take to be 3449. $D(z)$ is the growth factor normalised to unity today, which we calculate using CAMB to be approximately 0.025 at redshift 50 and approximately 0.05 at redshift 27. $n_{\rm PBH}$ is the comoving number density of PBHs, and the factor $f_{\rm PBH}^2/n_{\rm PBH}$ can be rewritten as $f_{\rm PBH}M_{\rm PBH}/\Omega_{\rm DM}\rho_{\rm crit}$ which will be of importance later when we discuss the degeneracy of the mass and abundance of PBHs in the 21cm signal. The combined contribution to the matter power spectrum is then given by [300]

$$\mathcal{P}_{21,\rm combined} = \mathcal{P}_{21,\rm adiabatic} + \frac{T_{21}^2}{T_{\rm DM}^2}\frac{k^3}{2\pi^2}P_{\rm Poisson}, \tag{4.9}$$

where $\mathcal{P} = k^3/2\pi^2 P$ for all quantities. The combined 21cm power spectrum is then calculated by using the 21cm transfer functions, T_{21}^2, and the CDM transfer functions, $T_{\rm DM}$, from CAMB at a given redshift.

We focus on 2 different masses of PBHs, $100M_\odot$ and $10^4 M_\odot$, as these are the largest and smallest mass PBHs that can be produced without conflicting with either spectral distortion or pulsar timing array constraints (e.g. [95, 281]), and could be respectively produced from the primordial power spectra growing like $k^{1.2}$ and k^4 plotted in figure 4.4. A non-monochromatic PBH mass function is inherent in the non-monochromatic power spectra, however we assume a monochromatic mass spectrum for the PBH population in the Poisson contribution. It was shown in [95] that the mass function of PBHs produced from even very shallow primordial power spectra on the low-mass end drops off quickly. We therefore expect Poisson contributions due to extended mass functions to affect a very small range of scales larger than the peak of the PBH distribution and do not include them here. For a fuller discussion of PBHs with extended mass distributions and the 21cm signal, see [302]. Note that PBHs with masses $< 0.1M_\odot$ can be readily produced without conflicting with any additional power spectrum constraints - their abundance is only limited by various constraints of their 'direct' non-detection which vary between $f_{\rm PBH} = 1$ and $f_{\rm PBH} \sim 10^{-5}$ depending on the mass. PBHs with masses this low, however, would form on scales too small to be detectable with 21cm experiments. For $M_{\rm PBH} = 100M_\odot$ the Poisson contribution is by far a sub-dominant effect for all $f_{\rm PBH}$

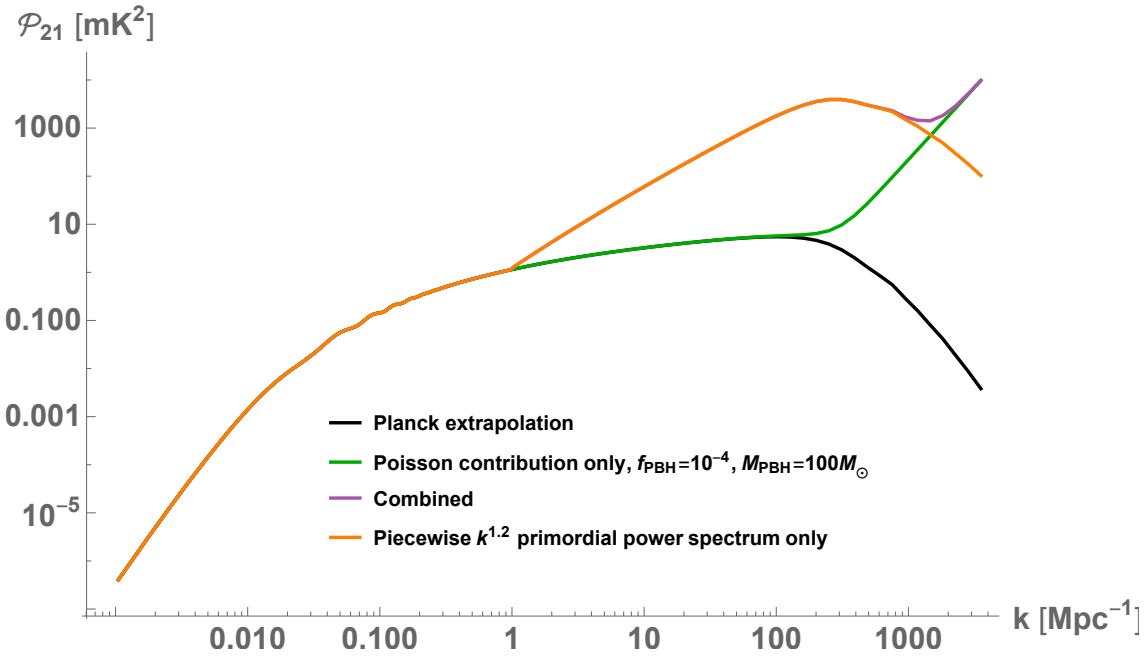

Figure 4.6: The 21cm power spectrum at redshift 50 for the scenario where $100M_\odot$ PBHs are produced with abundance $f_{\rm PBH} = 10^{-4}$. In orange is the 21cm signal prediction taking into account just the boost in the primordial power spectrum, in green is just the Poisson contribution, and in purple is the combined result. In black is the 21cm power spectrum produced by extrapolating the primordial power spectrum measured by Planck to small scales. This demonstrates that it is important to include the primordial power spectrum boost, so as not to underestimate the 21cm power spectrum signal.

in comparison with the primordial power spectrum growing like $k^{1.2}$ because it occurs on much smaller scales. This means that neglecting the primordial power spectrum would predict a 21cm signal that is too small. This is shown in figure 4.6, where the 21cm power spectrum at redshift 50 is plotted. For $100M_\odot$ PBHs, the orange line is just the primordial power spectrum contribution, the green line is just the Poisson contribution, and the purple line is the combined result. On scales beyond $k \sim 1\,\mathrm{Mpc}^{-1}$, the primordial signal is much larger than the Poisson contribution, showing that only including the Poisson fluctuations underestimates the 21cm signal if the primordial power spectrum is boosted on larger scales than the Poisson fluctuations affect. Any boost in the primordial power spectrum that occurs in the range $k \sim 0.1 - 100\,\mathrm{Mpc}^{-1}$ should therefore be included in 21cm signal predictions.

The k^4 primordial power spectrum (grey line in figures 4.4 and 4.5) would produce PBHs with masses around $M_{\mathrm{PBH}} = 10^4 M_\odot$. In this scenario, since the the primordial fluctuations grow very steeply, the boost only needs to occur on very small scales and the Poisson fluctuations generally dominate. We show this in figure 4.7. The orange line is just the primordial power spectrum contribution, whilst the green and purple lines show the signal including the Poisson fluctuations for the combinations $f_{\mathrm{PBH}} M_{\mathrm{PBH}}/M_\odot = 100, 1$ respectively. Whilst the boosted primordial fluctuations can be extrapolated to infer a most likely PBH mass produced (up to uncertainties in the mass function and horizon mass relationship), due to the degeneracy between f_{PBH} and M_{PBH} in the Poisson power spectrum, if the Poisson fluctuations dominate, the information about the mass and abundance individually is lost. In this situation, accretion effects may be able to distinguish between the two, however at redshift 50 they are likely to be small and therefore need to be accounted for very accurately. We leave an investigation of the interplay between all three effects at high redshift for future work.

Note again with relevance to all of the plots in this section, that relative velocity effects may boost the large-scale signal ($k < 1\,\mathrm{Mpc}^{-1}$) and suppress the small-scale signal ($k > 200\,\mathrm{Mpc}^{-1}$) [291]. The large-scale signal is unimportant for distinguishing between the usual Planck-extrapolated power spectrum and a boost in primordial fluctuations, and as will be shown in section 4.5, despite the small-scale signal being important in the context of primordial black hole production and the associated primordial fluctuations, even space-based detectors will not be able to reach small enough scales where this effect would be important.

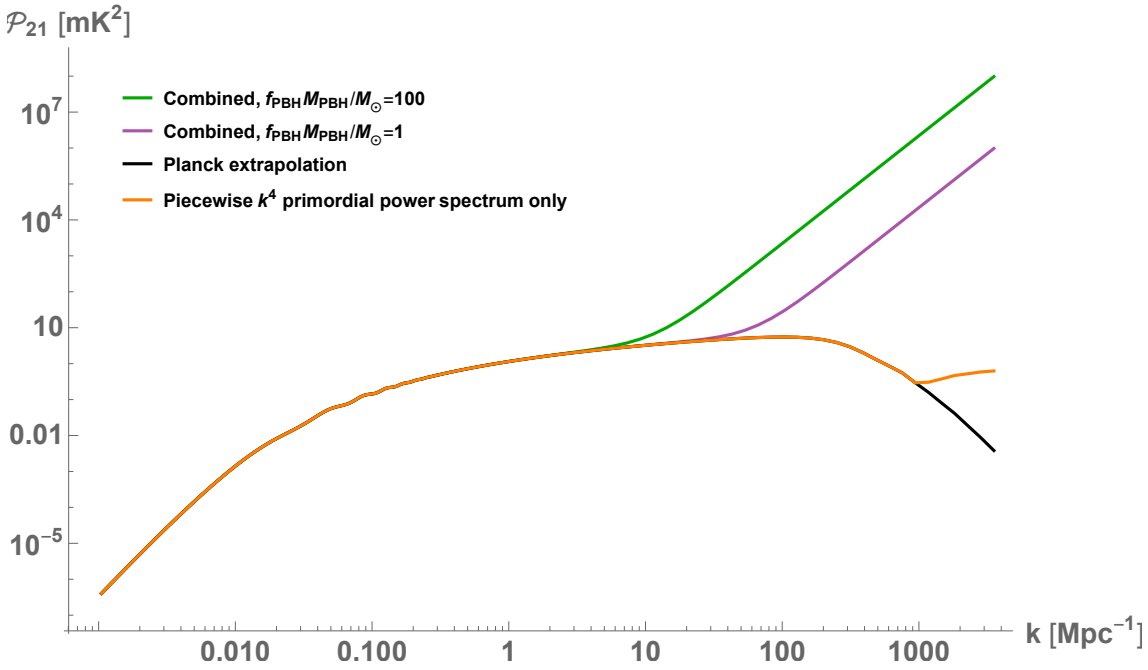

Figure 4.7: The 21cm power spectrum at redshift 50. The orange line only includes the primordial fluctuations contribution, for the primordial power spectrum that grows like k^4 and would produce $10^4 M_\odot$ PBHs if extrapolated. The green and purple lines include the Poisson fluctuations for $f_{\rm PBH} M_{\rm MPBH}/M_\odot = 100, 1$ respectively. Since the primordial boost happens on very small scales, the Poisson contribution is dominant.

	HERA	SKA	Lunar
$R_{\max}$	0.876km	5km	300km
$f_{\rm cov}$	0.08	0.01	0.75
Bandwidth	100MHz	2GHz	50, 100 MHz
$T_{\rm sky}$	2000 K	2000 K	10^4K
$t_{\rm int}$	1000 hours	1000 hours	1000 hours
ϵ	1	1	1

Table 4.1: Parameters describing HERA, SKA and lunar arrays.

4.5 Possibilities for detection

For a rough estimate on the sensitivity of SKA to the Dark Ages 21cm signal, the scaling relation derived in [303] from the prescription in [304] can be used which gives an approximation for the $1-\sigma$ error on the 21cm power spectrum at a given k and z:

$$\sqrt{\frac{k^3\delta P_{21}}{2\pi^2}} \sim \frac{0.1\,\mathrm{mK}}{\epsilon^{1/4} f_{\rm cov}} \left(\frac{k}{0.04\,\mathrm{Mpc}^{-1}}\right)^{3/4} \left(\frac{T_{\rm sky}}{10^4\,K}\frac{2\,\mathrm{km}}{R_{\max}}\right) \left(\frac{10\,\mathrm{MHz}}{B}\right)^{1/4} \left(\frac{1000\,\mathrm{hr}}{t_{\rm int}}\right)^{1/2} \left(\frac{1+z}{50}\right) \tag{4.10}$$

where ϵ is the frequency in k that the data is binned, $f_{\rm cov}$ is the array covering factor, $T_{\rm sky}$ is the temperature of the Galactic synchrotron foreground at the frequency of the observation, $R_{\max}$ is the radius of the (circular) array, B is the bandwidth, and $t_{\rm int}$ is the number of hours of integration. Whilst this scaling relation only gives a rough estimate on the detectability of the signal, it does capture the sensitivity of the errors to the various interferometer design parameters. Since SKA-Low is yet to be built, and a lunar interferometer yet to be funded, this rough estimate suffices in our case to get a guide on the sensitivity required.

For the best hope of observing low frequencies, i.e. high redshifts and smaller scales, it will be necessary to go to the Moon. Using the parameters proposed by [285] for a lunar radio interferometer in table 4.1, according to equation (4.10), the sensitivity is shown by the red region in figure 4.8.

Assuming perfect foreground removal, the lunar array should be sensitive enough to measure the 21cm power spectrum at redshift 50 up to $k \sim 12\,\mathrm{Mpc}^{-1}$. This would enable a clean distinction between the expected matter power spectrum from an extrapolation of the Planck measurements on large scales, and any deviations. Extra power, or a lack

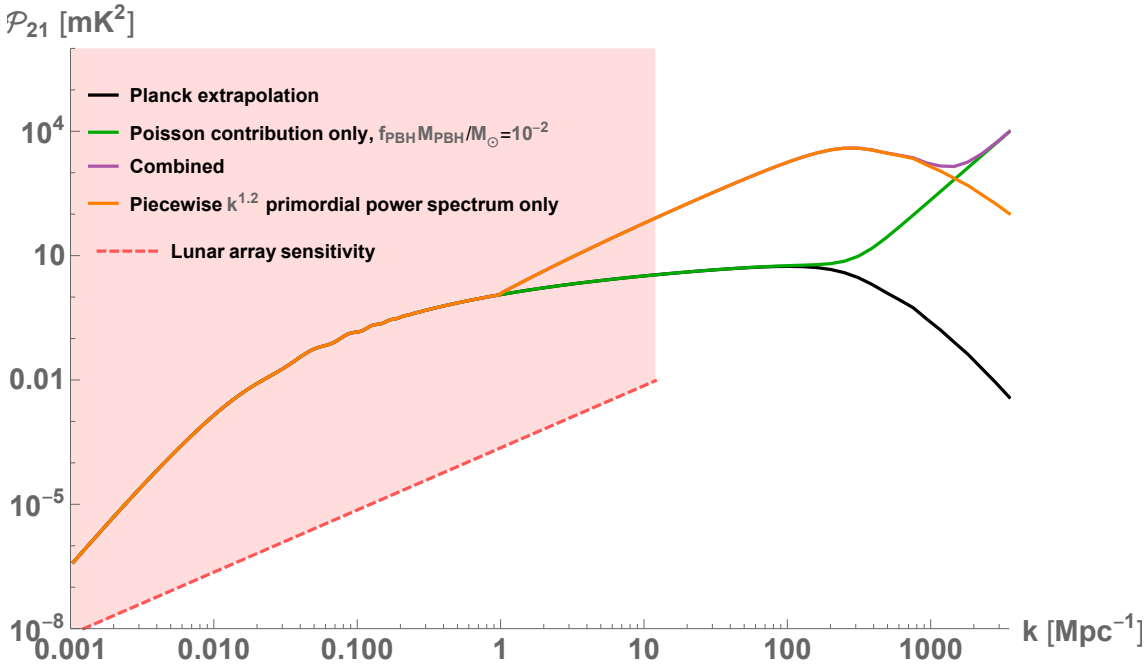

Figure 4.8: The same 21cm power spectrum as in figure 4.6 at redshift 50. A rough estimate of the sensitivity of a possible configuration for a radio interferometer on the far side of the moon is shown by the red dashed line.

of power, on the as of yet unexplored small scales beyond $k \sim 0.1\,\mathrm{Mpc}^{-1}$ should be observable. In addition, multiple 3d power spectra at several redshift slices could be stacked in order to increase the signal-to-noise of the detection. We demonstrate this below with a Fisher forecast for three parameters that describe a small-scale boost in power.

We have defined the smallest scale detectable as determined entirely by the angular resolution of the detector, given by $k_{\mathrm{max}} \sim 2\pi R_{\mathrm{max}}/14000\lambda(z)\,\mathrm{Mpc}^{-1}$, and we have focused on the isotropic power spectrum for which the signal should be largest. However, given that foregrounds are expected to especially dominate the Fourier modes in the angular direction, $k_\perp$, (see, for example, [128]) it might be possible to reach smaller scales in the line-of-sight direction $k_\parallel$. Whilst the non-isotropic power spectrum would exhibit a smaller signal, better spectral resolution of the detector might be possible and therefore a larger $k_\parallel$ could be reached than the k_{max} defined by the angular resolution. A signal in the parameter space away from the foreground 'wedge' would simplify the foreground removal task somewhat, however the number of independent modes lost to the wedge would decrease the signal-to-noise of any detection. See [129] for a recent investigation of using the anisotropic power spectrum to extract more cosmological information from line-intensity mapping.

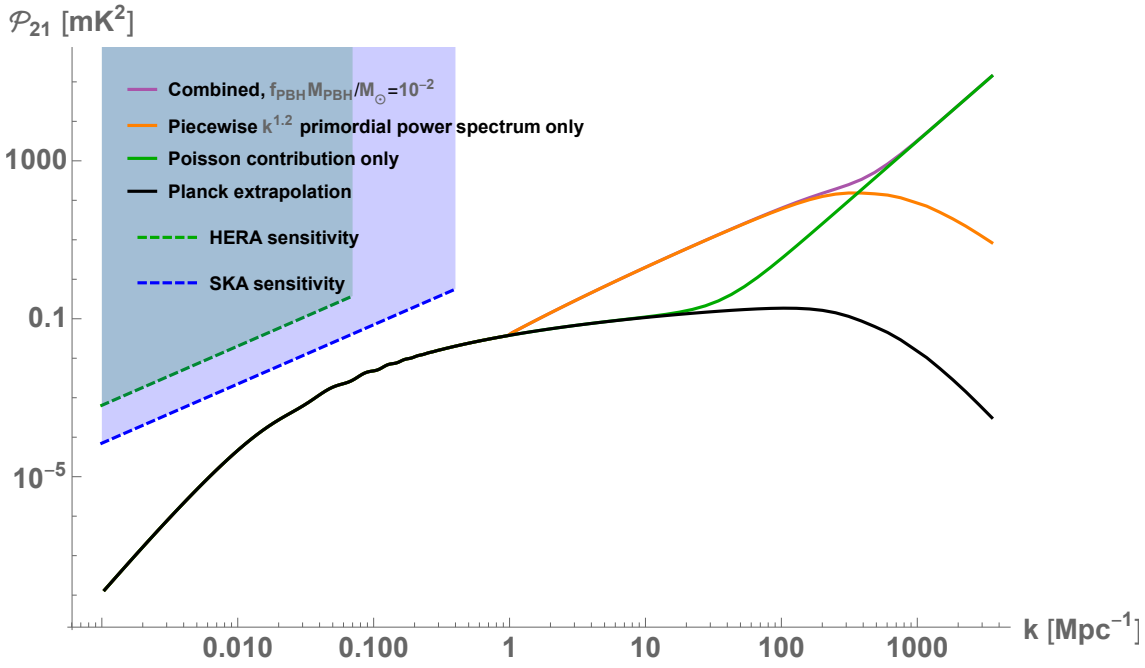

Figure 4.9: The 21cm power spectrum at redshift 27 for the scenario where $100M_\odot$ PBHs are produced with abundance $f_{\rm PBH} = 10^{-4}$. In orange is the 21cm signal prediction taking into account just the boost in the primordial power spectrum, in green is just the Poisson contribution, and in purple is the combined result. In black is the 21cm power spectrum produced by extrapolating the primordial power spectrum measured by Planck to small scales. A rough estimate of the sensitivity of HERA and SKA are shown by the green and blue regions.

To put into context the foreground removal challenge, the temperature of the galactic synchrotron foregrounds are of the order $T_{\rm sync} \sim (1+z)^{2.6}$. During the Cosmic Dawn/Epoch of Reionization era, the foregrounds therefore have a temperature of between $500 - 1000\,\mathrm{K}$ whilst the global 21cm signal is of the order $\mathcal{O}(100)\,\mathrm{mK}$. During the Dark Ages, the temperature of the foregrounds (not including extragalactic radio sources) are around an order of magnitude larger, $T_{\rm sync} \sim 10^5\,\mathrm{K}$, and the global 21cm signal is around five times smaller at the $T_{21} \sim 20\,\mathrm{mK}$ level. Whilst this is a big challenge to remove foregrounds 4-6 orders of magnitude larger than the signal and be sure that everything left-over is pure signal alone, there is hope in that the foregrounds are expected to be spectrally smooth, and in foreground avoidance methods as opposed to subtraction as mentioned above [128, 129].

Using parameters that emulate the HERA configuration and a possible SKA-Low configuration given in table 4.1, it is possible to put a rough estimate on the sensitivity to the 21cm signals predicted in the previous section at $z \sim 27$. This is shown by the green and blue regions in figure 4.9. The angular resolution means that neither HERA nor SKA-Low will be sensitive to small enough scales to go beyond the tightly constrained Planck measurements of the isotropic power spectrum up to $k \sim 0.1\,\mathrm{Mpc}^{-1}$. This means that PBH signatures will only be detectable if the PBH masses and abundances are large so that accretion effects dominate [285, 298, 299, 302]. Note that the scaling relation (4.10) does not take into account sample variance. This suffices in our case because we are interested in the sensitivity at small scales where the noise dominates, however it would be important for an accurate SKA error estimate on large scales. Furthermore, astrophysical sources would contaminate the signal at these redshifts [305], and would need to be taken into account for an accurate prediction.

We perform a Fisher forecast for three parameters that describe a boost in the power spectrum which could be detected in the 21cm signal. We parametrise the 21cm power spectrum as

$$\mathcal{P}_{21} = T_{21}^2 \left(A_s \left(\frac{k}{k_*} \right)^{n_s - 1} + B_s \left(\frac{k}{k_{inc}} \right)^{n_b} \right) + \frac{T_{21}^2}{T_{\rm DM}^2} \frac{9}{4} (1+z_{\rm eq})^2 D^2(z) \frac{k^3}{2\pi^2} \frac{f_{\rm PBH} M_{\rm PBH}}{\Omega_{\rm DM} \rho_c} \tag{4.11}$$

with T_{21} the 21cm monopole transfer function and $T_{\rm DM}$ the cold dark matter transfer function at a given redshift, k_{inc} the scale at which the primordial power spectrum is boosted from near scale-invariance, n_b is the spectral index of the boosted part of the spectrum, and $B_s = A_s \left(k_{inc}/k_*\right)^{n_s - 1}$. We will use k_{inc}, n_b and $f_{\rm PBH} M_{\rm PBH}$ as the three parameters for our Fisher forecast. The Fisher matrix for the 21cm power spectrum is

	$k_{\rm inc}$ [Mpc^{-1}]	n_b
$R_{\rm max} = 300$km	1 ± 0.0037	1.2 ± 0.0063
	5 ± 0.22	1.2 ± 0.046
	5 ± 0.18	2 ± 0.044
$R_{\rm max} = 500$km	1 ± 0.0022	1.2 ± 0.0038
	5 ± 0.13	1.2 ± 0.028
	5 ± 0.11	2 ± 0.026
$R_{\rm max} = 500$km	1 ± 0.0017	1.2 ± 0.0028
$f_{\rm cov} = 1$	5 ± 0.098	1.2 ± 0.021
	5 ± 0.081	2 ± 0.020

Table 4.2: $1-\sigma$ errors on fiducial values of the parameters $k_{\rm inc}$ and n_b for the lunar array as described in table 4.1.

defined as [306]

$$F_{\alpha\beta} = \sum_{k,z} \frac{1}{\varepsilon^2(k,z)} \frac{\partial \mathcal{P}_{21}(k,z)}{\partial \theta_\alpha} \frac{\partial \mathcal{P}_{21}(k,z)}{\partial \theta_\beta} \tag{4.12}$$

with θ representing the three parameters we have chosen, ε is the error on the 21cm signal given in equation (4.10) and the $1-\sigma$ error bars on a single parameter we calculate with $\sigma = \sqrt{F^{-1}_{\alpha\alpha}}$. We bin the 21cm transfer functions in increments of $\Delta k = k$ to be consistent with $\epsilon = 1$ in equation (4.10), and we sum over three redshift slices at $z = 49, 50, 51$. We assume these slices are independent based on a frequency resolution of $\Delta\nu = 0.1\,$MHz, whilst the redshifts $z = 49, 50, 51$ correspond to frequencies of $\nu = 28.6, 28.0, 27.5\,$MHz respectively. A calculation of the correlation length between redshift slices shows that this is a reasonable assumption in [307]. As HERA and SKA are only likely to be sensitive to the 21cm power spectrum up to $k \sim 0.07\,$Mpc^{-1} and $k \sim 0.4\,$Mpc^{-1} respectively, the lunar array is the only experiment that would be sensitive to $k_{\rm inc} \geq 1$. We find that the proposed specifications for the lunar array will be very sensitive to $k_{\rm inc}$ and n_b but will not be able to constrain the parameter $f_{\rm PBH}M_{\rm PBH}$ at all, given that it becomes important at much smaller scales. We therefore just report the resulting $1-\sigma$ error bars for fiducial values of $k_{\rm inc}$ and n_b in table 4.2, and show the effect of varying $R_{\rm max}$ and $f_{\rm cov}$.

4.6 Conclusions

Dark ages exploration has unexcelled reach in probing excess power in the primordial spectrum on scales far smaller than those probed by the CMB or LSS, see the illustration in figure 4.1. Not only is this range of parameter space uniquely accessibly via 21cm spectroscopy at high $z \sim 30 - 80$, without any contamination from the first stars, but the huge number of modes available, further boosted by 21cm tomography, in exploiting power down to $k > 10$ Mpc^{-1} makes this potentially the most sensitive cosmological probe possible. Of course this is a futuristic view as the foregrounds are many orders of magnitude larger at such low frequencies, ideally ~ 30 MHz, amounting to a brightness temperature thousands of times larger than the elusive signal at the tens of mK level. However, CMB primordial B-mode detection faces a comparable foreground challenge, where the current CMB-S4 goal of B-mode sensitivity at the few mK level may not be insurmountable. We hence consider that it is worthwhile to develop predictions in this paper without entering into the details of the foreground limitations.

Identification of the nature of dark matter remains the highest priority in particle astrophysics and cosmology. The primordial black hole is the principal weakly interacting candidate for non-baryonic dark matter that does not require the existence of a new particle beyond the standard model. The challenge is to develop initial conditions in the post-inflationary universe that can produce PBHs in the empirically allowed mass range. Whilst the mass window for all of the non-baryonic dark matter to be made up of PBHs is limited to the sub-lunar range, bounded by Hawking evaporation limits from the diffuse gamma ray background at the lower end and gravitational microlensing of M31 at the upper end, more specifically to the mass range 10^{-17} - $10^{-9}\,\mathrm{M}_\odot$, a population of larger PBHs could still be astrophysically significant even with small abundances, which is what we consider here. This is due to the fact that they could account for the population of seed black holes required to account for the presence of supermassive black holes at $z \gtrsim 6$, namely $f_{\rm PBH} \sim 10^{-4}$. Furthermore they could account for some or all of the LIGO detections of unexpectedly massive black holes, possible if the PBH mass fraction satisfies $f_{\rm PBH} \sim 0.01$. In addition, standard PBH production scenarios require a deviance from scale-invariance in the primordial power spectrum, and therefore a detection of small-scale power would also be informative for understanding inflationary dynamics.

We have found that PBH production in the observationally motivated range of $10 - 10^4\mathrm{M}_\odot$, requires the power spectrum to be sufficiently boosted by primordial

fluctuations. If this boost occurs on larger scales, $k \sim 0.1 - 100\mathrm{Mpc}^{-1}$, this contribution must be accounted for in the 21cm power spectrum so as not to underestimate the signal. Depending on the mass and abundance, Poisson fluctuations can also become important, and in that case accurate modelling of accretion effects at high redshifts will be vital to identify the underlying PBH population producing the signal. These signatures could become potentially observable in the 21cm power spectrum with the new generation of filled low frequency interferometers. Evidently our predictions, which lack any modelling of foregrounds, are unrealistic, but we hope that they will motivate improved cleaning algorithms that can enable us to access this intriguing corner of PBH-motivated parameter space.

4.7 Acknowledgements

PC acknowledges support from the UK Science and Technology Facilities Council via Research Training Grant ST/N504452/1. We are grateful to Nicola Bellomo, Christian Byrnes, Joshua Dillon, Guilio Fabbian, Anastasia Fialkov, Danny Jacobs, Antony Lewis, Jose Luis Bernal, Julian Munoz and Pablo Villanueva-Domingo for valuable discussions over the course of this investigation.

Chapter 5

Conclusions

Deciphering how inflation happened is one of the biggest challenges in cosmology today. Our best hope for answering this question is by studying the overdensities and underdensities left over at the end of inflation. Large-scale measurements of the primordial power spectrum have shown that the simplest models of inflation, namely single-field slow-roll, are able to produce the almost scale-invariant distribution of overdensities and underdensities observed today. More complicated models, for example multifield models or those with non-canonical kinetic terms could also produce this distribution, and higher-order statistics such as the bispectrum or a detection of primordial gravitational waves would be required to make the distinction. However, on smaller scales, there are far fewer constraints on the primordial power spectrum which means that there could be some sort of feature or peak that can't be explained by a single-field slow-roll model of inflation. Just by studying the two-point statistics of the smaller scales, it is possible to make connections between inflationary models and observations today.

We have seen that primordial black holes are a probe of primordial fluctuations, because they should generically form from the collapse of very large overdensities re-entering the horizon immediately after inflation. This means that if they are observed today, a strong scale-dependence in the primordial power spectrum would be required for their production in order to boost the amplitude of the fluctuations from the small value that has been measured on large scales via anisotropies of the Cosmic Microwave Background [26].

The investigation of this probe is further motivated by the fact that primordial black holes are a dark matter candidate that do not require the existence of any new particles beyond the standard model. Their recent resurgence in popularity is predominantly due

to the gravitational wave detections of order 10 solar mass black holes by the LIGO/Virgo collaboration [88]. Even though it is looking unlikely that primordial black holes of this mass-range can make up all of the dark matter, there is a lower mass-window of the order $10^{-13}\,M_\odot$ where this could still be possible. Furthermore, even a very small fraction of the dark matter being accounted for by primordial black holes would have strong implications for inflationary dynamics, and could also provide the seeds of supermassive black holes [284, 285].

In chapter 2, we explored the requirements on the primordial power spectrum for even just one primordial black hole to form. Due to the logarithmic sensitivity of the primordial power spectrum to the abundance of primordial black holes, assuming Gaussian initial fluctuations, the requirement on the amplitude of the primordial power spectrum to form a single primordial black hole versus enough to make up all of the dark matter only differs by around between a factor of 2 and 5 depending on the scale. Whilst this means that improving the direct detection constraints on primordial black holes will not help to constrain the primordial power spectrum considerably more than they already have, it does mean that the detection of just one primordial black hole would be extremely prescriptive for inflationary model-building, since it would imply that a very large peak in the power spectrum at a particular scale must have been present (or else some other primordial black hole production mechanism would be required).

Non-Gaussian initial conditions may soften this requirement by an order of magnitude or so, whilst a phase of early matter-domination after inflation could circumvent the need for a boost in the primordial fluctuations at all in order to produce a considerable number of primordial black holes. We showed in chapter 2 that for certain durations of an early matter-dominated phase, even with a scale-invariant power spectrum extrapolated from the large-scale Planck measurements, PBHs could actually be overproduced, i.e. this scenario would be in conflict with current observational constraints on the abundance of primordial black holes.

Having determined that there are tough requirements on the primordial power spectrum for primordial black hole production in the standard, radiation-dominated scenario, it then becomes a question of how the large overdensities can be produced from a single-field inflationary model. It is necessary to go beyond the slow-roll regime in order to achieve a 7 order of magnitude growth in the power spectrum from single-field inflation, which is usually realised by slowing down the inflaton as it rolls down its potential. An inflection point, flat section, or local maximum in the potential is capable

of doing this, although 'realistic' inflationary models with this feature are usually very delicately designed with fine-tuning of the order 1 part in 10^7 in order to achieve a large enough boost in the fluctuations [11, 208, 212].

Not only must the primordial power spectrum's boost in amplitude be large enough, but the width of the feature must be able to thread itself through the small-scale constraints, although they are currently quite weak. In chapter 3, we determined an upper bound on the spectral index of the power spectrum that can be realised by single-field models of inflation which exhibit a phase of slow-roll, followed by some beyond slow-roll regime. By breaking the slow-roll conditions with increasing degrees of severity, we found that the power spectrum can't grow faster than k^4. This has an effect on power spectrum constraints from probes such as spectral distortions and primordial gravitational waves because they should be calculated assuming that, at best, the power spectrum can grow only as fast as k^4, rather than assuming that it can be boosted instantaneously at a given scale. In terms of implications for primordial black hole production, this means that the viable mass ranges for primordial black holes to have been produced are narrower than previously calculated, because the power spectrum can't always grow fast enough to avoid constraints at a given scale. We recalculated the current constraints on the primordial power spectrum from spectral distortions and secondarily produced primordial gravitational waves assuming a k^4 growth for the power spectrum instead of a monochromatic boost. We found that an extra half an order of magnitude in k can be constrained, which corresponds to an order of magnitude in primordial black hole mass. With future observations, it could be possible to entirely close mass-ranges of primordial black hole production that would naively be expected to still be viable if the constraints are calculated for monochromatic power spectra. Furthermore, due to the results of chapter 2, if constraints on the primordial power spectrum rule out amplitudes of $10^{-2} - 10^{-3}$ across a range of scales, not only will primordial black holes be ruled out as making up the dark matter on those mass-ranges, they will be ruled out as existing at all, assuming Gaussian initial fluctuations and that they formed in radiation-domination.

After the publication of the work that makes up chapter 3, reference [96] appeared. They showed that it is in fact possible to obtain a slightly steeper growth of the power spectrum, if a period of ultra-slow-roll is preceded by a period of gently broken slow-roll (a phase of $\eta = -1$). The slope can reach a limit of $k^5(\log k)^2$ in this case. However, if the large-scale measurements from Planck are to be respected, a period of slow-roll is required to reproduce the scale-invariance. This means that fluctuations over only a

small range of scales can experience a faster than k^4 growth in the power spectrum, and hence the implications found for the observational constraints in chapter 3 do not differ significantly. Furthermore, an inflationary potential that can exhibit this evolution of the second slow-roll parameter has not been found.

In chapter 4, we investigated a 'late-time' probe of the primordial fluctuations in the way of the 21cm Dark Ages signal. Between redshifts $z \sim 30 - 150$ the distribution of neutral hydrogen can be mapped out via the power spectrum of the 21cm absorption signal, which is a consequence of neutral hydrogen's spin-flip transition. The existence of primordial black holes may alter the 'standard' predicted 21cm signal in three ways: due to accretion of matter onto the primordial black holes, the Poissonian contribution due to the discrete nature of the primordial black hole distribution, and due to the fact that large primordial fluctuations must have been present to produce the primordial black holes in the first place as demonstrated in chapters 2 and 3. The first two effects have been studied previously [285, 297, 299, 302], whilst chapter 4 investigates the contribution of the third effect to the signal, especially in cases where the effects of accretion are expected to be small, for example if there is only a small population of primordial black holes.

We found that a boost in small-scale primordial fluctuations would be present in the 21cm power spectrum, but that in order to detect such a signal, a space-based detector would be required. This is due to the fact that ground-based detectors can't reach low enough frequencies because the Earth's atmosphere reflects radio waves below around 50 MHz, as well as that a very large array is required to reach small scales. We showed that it is vital to include the contribution from primordial fluctuations so as not to underestimate the signal, however there is a lot of work to be done in terms of both modelling and detecting the 21cm signal. Concerning primordial black holes, an accurate prescription for the accretion effects on the 21cm signal is required, and this should be included with the contributions from the primordial fluctuations. In general, precise and reliable foreground modelling is going to be paramount for successfully detecting the 21cm power spectrum at high redshift [128].

5.1 Future outlook

Arguably, the detection of a primordial black hole is the most exciting prospect for the future of this field. The distinction from those of astrophysical origin would be most readily provided if its mass was less than the Chandrasekhar limit ($\lesssim 1.3\,\mathrm{M}_\odot$). Not only

would this explain at least a (possibly very small) fraction of the dark matter budget, the required inflationary potential for the production of the overdensities that could collapse to form the primordial black hole would be very restricted if single-field models are still preferred. Otherwise, a more complicated production mechanism will need to be specified.

In terms of direct detection, microlensing experiments have the best sensitivity to sub-solar mass black holes, although the LIGO/Virgo collaboration does have some sensitivity and the results from the third observing run will provide more constraints. Future proposed detectors such as the Einstein Telescope [308] will also improve chances of a detection and/or constraints. Theoretically, a better understanding of black hole spins, mass functions and merger rates in both the primordial and astrophysical cases will facilitate more robust analyses of the observational data.

In terms of indirectly constraining primordial black holes via the primordial power spectrum, and hence constraining inflation directly at the same time, 21cm experiments, future spectral distortion probes and gravitational wave searches provide the most likely chances of success. A space-based radio interferometer will be able to map out the 21cm power spectrum at redshifts $z \sim 30 - 50$ down to a scale of $k \sim 12\,\mathrm{Mpc}^{-1}$. A spectral distortion probe similar to the PIXIE proposal [107] would constrain the primordial power spectrum to an amplitude of $\mathcal{P}_{\mathcal{R}} \sim 10^{-8} - 10^{-9}$ across scales down to $k \sim 10^{5}\,\mathrm{Mpc}^{-1}$. Pulsar timing array constraints from SKA will constrain even smaller scales, $k \sim 10^{6} - 10^{8}\,\mathrm{Mpc}^{-1}$, to the level of $\mathcal{P}_{\mathcal{R}} \sim 10^{-5}$, and finally LISA will constrain the primordial power spectrum via secondarily produced gravitational waves to the $\mathcal{P}_{\mathcal{R}} \sim 10^{-4}$ level on scales around $k \sim 10^{12}\,\mathrm{Mpc}^{-1}$. This latter constraint will also rule out primordial black holes as forming from the standard mechanism in the last remaining window where they could make up all of the dark matter.

There are various theoretical uncertainties on the amplitude of the primordial power spectrum required for primordial black hole production which need further investigation. Further to the numerical work in chapter 3, we have begun to investigate the amount of primordial non-Gaussianity produced from beyond-slow-roll single-field models of inflation. If it is non-negligible, then the calculation of the abundance of primordial black holes will need to be revised so as to account for this effect. There is also ongoing work, see for example [232, 233], on the effect of quantum fluctuations on inflationary dynamics in beyond-slow-roll models of inflation that could effect primordial black hole production. Finally, uncertainties on the relationship between the horizon mass and the

PBH mass produced, and the effect of the window function used to smooth perturbations in the calculation of the primordial black hole abundance also need to be clarified.

If primordial black holes are ruled out as forming from overdensities collapsing after inflation via either direct detection constraints[1] or indirectly via the primordial power spectrum, the focus will need to shift to more exotic production mechanisms such as the collapse of topological defects or considering a different thermal history such as an early matter-dominated phase for them to survive as either a dark matter candidate or a relic of inflation.

More broadly, the distinction between single-field and multifield models of inflation will be probed further with upcoming CMB and large-scale structure experiments such as Simons Observatory [43], SPHEREX [44] and the proposed PICO mission [309], which will all improve constraints on primordial non-Gaussianity on large scales to beyond the threshold of $f_{NL}^{local} \lesssim 1$.

Primordial black holes have the ability to both explain dark matter and provide specific requirements on the inflationary potential. Confirming or falsifying their existence has drawn together theorists and observers from many different fields, and as a concept that has endured since the 1970s, its increased recent interest will hopefully facilitate a conclusion in the near future. Advances in the understanding of small-scale probes including spectral distortions, gravitational waves and the Dark Ages 21cm signal will complement the primordial black hole discussion, as well as concurrently exploring varied elements of early universe cosmology.

[1] Note that it will be very difficult to rule them out entirely via direct detection, but constraining them to make up a negligible fraction of the dark matter budget would be possible.

Bibliography

[1] Planck Collaboration et al. Planck 2018 results. X. Constraints on inflation. *arXiv e-prints*, art. arXiv:1807.06211, July 2018. xii, xvii, 14, 85, 89, 90

[2] Hiroko Niikura, Masahiro Takada, Naoki Yasuda, Robert H. Lupton, Takahiro Sumi, Surhud More, Toshiki Kurita, Sunao Sugiyama, Anupreeta More, Masamune Oguri, and et al. Microlensing constraints on primordial black holes with subaru/hsc andromeda observations. *Nature Astronomy*, 3(6):524–534, Apr 2019. ISSN 2397-3366. doi: 10.1038/s41550-019-0723-1. URL http://dx.doi.org/10.1038/s41550-019-0723-1. xii, 21, 22, 25

[3] Ville Vaskonen and Hardi Veermäe. A lower bound on the primordial black hole merger rate. *Phys. Rev.*, D101(4):043015, 2020. doi: 10.1103/PhysRevD.101.043015.

[4] B. P. Abbott et al. Search for Subsolar Mass Ultracompact Binaries in Advanced LIGO's Second Observing Run. *Phys. Rev. Lett.*, 123(16):161102, 2019. doi: 10.1103/PhysRevLett.123.161102. xii, 22

[5] Steven Furlanetto et al. Astro 2020 Science White Paper: Fundamental Cosmology in the Dark Ages with 21-cm Line Fluctuations. *arXiv e-prints*, art. arXiv:1903.06212, Mar 2019. xiii, 33, 85

[6] Philipp Burda, Ruth Gregory, and Ian Moss. Gravity and the stability of the Higgs vacuum. *Phys. Rev. Lett.*, 115:071303, 2015. doi: 10.1103/PhysRevLett.115.071303. xiii, 37, 38, 42, 43, 56

[7] Philipp Burda, Ruth Gregory, and Ian Moss. The fate of the Higgs vacuum. *JHEP*, 06:025, 2016. doi: 10.1007/JHEP06(2016)025. xiii, 37, 38, 42, 43, 56

[8] Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, Yuichiro Tada, and

Tsutomu T. Yanagida. Inflationary Primordial Black Holes as All Dark Matter. 2017. xiii, 43, 55

[9] Bernard J. Carr. The Primordial black hole mass spectrum. *Astrophys. J.*, 201: 1–19, 1975. doi: 10.1086/153853. xiii, 18, 19, 21, 39, 43

[10] B. J. Carr, Kazunori Kohri, Yuuiti Sendouda, and Jun'ichi Yokoyama. New cosmological constraints on primordial black holes. *Phys. Rev.*, D81:104019, 2010. doi: 10.1103/PhysRevD.81.104019. xiii, 18, 21, 36, 37, 38, 42, 44, 45, 60

[11] Cristiano Germani and Tomislav Prokopec. On primordial black holes from an inflection point. *Phys. Dark Univ.*, 18:6–10, 2017. doi: 10.1016/j.dark.2017.09.001. xv, 12, 59, 61, 62, 66, 67, 92, 108

[12] Andrew R Liddle, David Parkinson, Samuel M Leach, and Pia Mukherjee. The WMAP normalization of inflationary cosmologies. *Phys. Rev.*, D74:083512, 2006. doi: 10.1103/PhysRevD.74.083512. xvi, 86

[13] Hu Zhan, Lloyd Knox, Anthony Tyson, and Vera Margoniner. Exploring large-scale structure with billions of galaxies. *Astrophys. J.*, 640:8–17, 2006. doi: 10.1086/500077. xvi, 86

[14] D. J. Fixsen, E. S. Cheng, J. M. Gales, John C. Mather, R. A. Shafer, and E. L. Wright. The Cosmic Microwave Background spectrum from the full COBE FIRAS data set. *Astrophys. J.*, 473:576, 1996. doi: 10.1086/178173. xvii, 28, 73, 90, 91

[15] John C. Mather et al. Measurement of the Cosmic Microwave Background spectrum by the COBE FIRAS instrument. *Astrophys. J.*, 420:439–444, 1994. doi: 10.1086/173574. xvii, 73, 90, 91

[16] J. Chluba, R. Khatri, and R. A. Sunyaev. CMB at 2 2 order: the dissipation of primordial acoustic waves and the observable part of the associated energy release. *Monthly Notices of the Royal Astronomical Society*, 425(2):1129–1169, Sep 2012. doi: 10.1111/j.1365-2966.2012.21474.x. xvii, 90, 91

[17] Jens Chluba, Adrienne L. Erickcek, and Ido Ben-Dayan. Probing the inflaton: Small-scale power spectrum constraints from measurements of the CMB energy spectrum. *Astrophys. J.*, 758:76, 2012. doi: 10.1088/0004-637X/758/2/76. xvii, 72, 90, 91

[18] Jens Chluba et al. Spectral Distortions of the CMB as a Probe of Inflation, Recombination, Structure Formation and Particle Physics. *Bulletin of the AAS*, 51 (3):184, May 2019. xvii, 28, 90, 91

[19] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, and et al. Nine-year wilkinson microwave anisotropy probe (wmap) observations: Final maps and results. *The Astrophysical Journal Supplement Series*, 208(2):20, Sep 2013. ISSN 1538-4365. doi: 10.1088/0067-0049/208/2/20. URL http://dx.doi.org/10.1088/0067-0049/208/2/20. 3

[20] Antony Lewis. *Advanced Cosmology: Inflation and the generation of fluctuations*. https://cosmologist.info/teaching/EU/notes_inflation.pdf [Accessed: March 2020]. 3, 5, 13

[21] David Wands, Karim A. Malik, David H. Lyth, and Andrew R. Liddle. New approach to the evolution of cosmological perturbations on large scales. *Physical Review D*, 62(4), Jul 2000. ISSN 1089-4918. doi: 10.1103/physrevd.62.043527. URL http://dx.doi.org/10.1103/PhysRevD.62.043527. 4, 6

[22] Y. Akrami et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. 2018. 7, 60, 76

[23] D. J. Fixsen. The temperature of the cosmic microwave background. *The Astrophysical Journal*, 707(2):916–920, Nov 2009. ISSN 1538-4357. doi: 10.1088/0004-637x/707/2/916. URL http://dx.doi.org/10.1088/0004-637X/707/2/916. 7

[24] Antony Lewis, Anthony Challinor, and Anthony Lasenby. Efficient computation of CMB anisotropies in closed FRW models. *Astrophysical Journal*, 538:473–476, 2000. doi: 10.1086/309179. 7, 89, 91

[25] Julien Lesgourgues. The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview. *arXiv e-prints*, art. arXiv:1104.2932, April 2011. 7

[26] Planck Collaboration et al. Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints*, art. arXiv:1807.06209, Jul 2018. 7, 20, 106

[27] Daniel Baumann. TASI Lectures on Inflation. *arXiv e-prints*, art. arXiv:0907.5424, July 2009. 9

[28] V. F. Mukhanov. Gravitational instability of the universe filled with a scalar field. *Soviet Journal of Experimental and Theoretical Physics Letters*, 41:493, May 1985. 10

[29] Misao Sasaki. Large Scale Quantum Fluctuations in the Inflationary Universe. *Progress of Theoretical Physics*, 76(5):1036–1046, November 1986. doi: 10.1143/PTP.76.1036. 10

[30] Konstantinos Dimopoulos. Ultra slow-roll inflation demystified. *Phys. Lett.*, B775: 262–265, 2017. doi: 10.1016/j.physletb.2017.10.066. 11, 61

[31] Swagat S. Mishra and Varun Sahni. Primordial Black Holes from a tiny bump in the Inflaton potential. *arXiv e-prints*, art. arXiv:1911.00057, Oct 2019. 12

[32] Jessica L. Cook and Lorenzo Sorbo. Particle production during inflation and gravitational waves detectable by ground-based interferometers. *Physical Review D*, 85(2), Jan 2012. ISSN 1550-2368. doi: 10.1103/physrevd.85.023534. URL http://dx.doi.org/10.1103/PhysRevD.85.023534. 13

[33] M. C. Guzzetti, N. Bartolo, M. Liguori, and S. Matarrese. Gravitational waves from inflation. *Riv. Nuovo Cim.*, 39(9):399–495, 2016. doi: 10.1393/ncr/i2016-10127-1. 13

[34] David H. Lyth. What would we learn by detecting a gravitational wave signal in the cosmic microwave background anisotropy? *Physical Review Letters*, 78(10): 1861–1863, Mar 1997. ISSN 1079-7114. doi: 10.1103/physrevlett.78.1861. URL http://dx.doi.org/10.1103/PhysRevLett.78.1861. 13

[35] P.A.R. Ade, Z. Ahmed, R.W. Aikin, K.D. Alexander, D. Barkats, S.J. Benton, C.A. Bischoff, J.J. Bock, R. Bowens-Rubin, J.A. Brevik, and et al. Improved constraints on cosmology and foregrounds from bicep2 and keck array cosmic microwave background data with inclusion of 95 ghz band. *Physical Review Letters*, 116(3), Jan 2016. ISSN 1079-7114. doi: 10.1103/physrevlett.116.031302. URL http://dx.doi.org/10.1103/PhysRevLett.116.031302. 13

[36] Joseph Silk. Cosmic Black-Body Radiation and Galaxy Formation. *Astrophysical Journal*, 151:459, Feb 1968. doi: 10.1086/149449. 13

[37] Juan Martin Maldacena. Non-Gaussian features of primordial fluctuations in single

field inflationary models. *JHEP*, 05:013, 2003. doi: 10.1088/1126-6708/2003/05/013. 15, 71

[38] Paolo Creminelli and Matias Zaldarriaga. A single-field consistency relation for the three-point function. *Journal of Cosmology and Astroparticle Physics*, 2004(10): 006–006, Oct 2004. ISSN 1475-7516. doi: 10.1088/1475-7516/2004/10/006. URL http://dx.doi.org/10.1088/1475-7516/2004/10/006. 62, 71

[39] Pieter Daniel Meeburg, Daniel Green, Raphael Flauger, Benjamin Wallisch, M. C. David Marsh, Enrico Pajer, Garret Goon, Cora Dvorkin, Azadeh Moradinezhad Dizgah, Daniel Baumann, Guilherme L. Pimentel, Simon Foreman, Eva Silverstein, Elisa Chisari, Benjamin Wandelt, Marilena Loverde, and Anze Slosar. Primordial Non-Gaussianity. *Bulletin of the AAS*, 51(3):107, May 2019. 15

[40] Jerome Martin, Hayato Motohashi, and Teruaki Suyama. Ultra Slow-Roll Inflation and the non-Gaussianity Consistency Relation. *Phys. Rev.*, D87(2):023514, 2013. doi: 10.1103/PhysRevD.87.023514. 15, 61, 154

[41] Júlio Fabris, Oliver Piattella, Davi Rodrigues, H. Velten, and Winfried Zimdahl. *The Cosmic Microwave Background: Proceedings of the II José Plínio Baptista School of Cosmology*, volume 45. 01 2016. ISBN 978-3-319-44768-1. doi: 10.1007/978-3-319-44769-8. 15

[42] Planck Collaboration et al. Planck 2018 results. IX. Constraints on primordial non-Gaussianity. *arXiv e-prints*, art. arXiv:1905.05697, May 2019. 15

[43] Peter Ade, James Aguirre, Zeeshan Ahmed, Simone Aiola, Aamir Ali, David Alonso, Marcelo A. Alvarez, Kam Arnold, Peter Ashton, Jason Austermann, and et al. The simons observatory: science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056–056, Feb 2019. ISSN 1475-7516. doi: 10.1088/1475-7516/2019/02/056. URL http://dx.doi.org/10.1088/1475-7516/2019/02/056. 15, 111

[44] Olivier Doré et al. Cosmology with the SPHEREX All-Sky Spectral Survey. *arXiv e-prints*, art. arXiv:1412.4872, Dec 2014. 15, 111

[45] Edmund J. Copeland and T. W. B. Kibble. Cosmic strings and superstrings. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering*

Sciences, 466(2115):623–657, Jan 2010. ISSN 1471-2946. doi: 10.1098/rspa.2009.0591. URL http://dx.doi.org/10.1098/rspa.2009.0591. 16

[46] Masahiro Kawasaki, Naoya Kitajima, and Tsutomu T. Yanagida. Primordial black hole formation from an axionlike curvaton model. *Phys. Rev. D*, 87:063519, Mar 2013. doi: 10.1103/PhysRevD.87.063519. URL https://link.aps.org/doi/10.1103/PhysRevD.87.063519. 16

[47] Veniamin Berezinsky, Vyacheslav Dokuchaev, and Yury Eroshenko. Small-scale clumps in the galactic halo and dark matter annihilation. *Physical Review D*, 68 (10), Nov 2003. ISSN 1089-4918. doi: 10.1103/physrevd.68.103003. URL http://dx.doi.org/10.1103/PhysRevD.68.103003. 16

[48] Dong Zhang. Impact of primordial ultracompact minihaloes on the intergalactic medium and first structure formation. *Monthly Notices of the Royal Astronomical Society*, 418(3):1850–1872, Nov 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.19602.x. URL http://dx.doi.org/10.1111/j.1365-2966.2011.19602.x. 16

[49] Julian Adamek, Christian T. Byrnes, Mateja Gosenca, and Shaun Hotchkiss. Wimps and stellar-mass primordial black holes are incompatible. *Physical Review D*, 100(2), Jul 2019. ISSN 2470-0029. doi: 10.1103/physrevd.100.023506. URL http://dx.doi.org/10.1103/PhysRevD.100.023506. 16

[50] Sam Young, Ilia Musco, and Christian T. Byrnes. Primordial black hole formation and abundance: contribution from the non-linear relation between the density and curvature perturbation. *JCAP*, 2019(11):012, Nov 2019. doi: 10.1088/1475-7516/2019/11/012. 16, 93

[51] V. De Luca, G. Franciolini, A. Kehagias, M. Peloso, A. Riotto, and C. Ünal. The ineludible non-Gaussianity of the primordial black hole abundance. *JCAP*, 2019 (7):048, Jul 2019. doi: 10.1088/1475-7516/2019/07/048. 16

[52] Charles W. Misner and David H. Sharp. Relativistic equations for adiabatic, spherically symmetric gravitational collapse. *Phys. Rev.*, 136:B571–B576, Oct 1964. doi: 10.1103/PhysRev.136.B571. URL https://link.aps.org/doi/10.1103/PhysRev.136.B571. 16

[53] J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay. The Statistics of Peaks of

Gaussian Random Fields. *Astrophysical Journal*, 304:15, May 1986. doi: 10.1086/164143. 16

[54] Alexander G. Polnarev and Ilia Musco. Curvature profiles as initial conditions for primordial black hole formation. *Classical and Quantum Gravity*, 24(6):1405–1431, March 2007. doi: 10.1088/0264-9381/24/6/003. 17

[55] Kenta Ando, Keisuke Inomata, and Masahiro Kawasaki. Primordial black holes and uncertainties in the choice of the window function. *Phys. Rev.*, D97(10): 103528, 2018. doi: 10.1103/PhysRevD.97.103528. 17, 70, 72, 80

[56] Sam Young. The primordial black hole formation criterion re-examined: Parametrisation, timing and the choice of window function. *International Journal of Modern Physics D*, page 2030002, Oct 2019. ISSN 1793-6594. doi: 10.1142/s0218271820300025. URL http://dx.doi.org/10.1142/S0218271820300025. 17, 93

[57] Cristiano Germani and Ilia Musco. Abundance of primordial black holes depends on the shape of the inflationary power spectrum. *Physical Review Letters*, 122(14), Apr 2019. ISSN 1079-7114. doi: 10.1103/physrevlett.122.141302. URL http://dx.doi.org/10.1103/PhysRevLett.122.141302. 17, 72, 93

[58] Ilia Musco. Threshold for primordial black holes: Dependence on the shape of the cosmological perturbations. *Physical Review D*, 100(12):123524, December 2019. doi: 10.1103/PhysRevD.100.123524. 17, 72

[59] Alexander Polnarev and Robert Zembowicz. Formation of Primordial Black Holes by Cosmic Strings. *Phys. Rev.*, D43:1106–1109, 1991. doi: 10.1103/PhysRevD.43.1106. 18

[60] Jaume Garriga, Alexander Vilenkin, and Jun Zhang. Black holes and the multiverse. *JCAP*, 2016(2):064, February 2016. doi: 10.1088/1475-7516/2016/02/064. 18

[61] Misao Sasaki, Teruaki Suyama, Takahiro Tanaka, and Shuichiro Yokoyama. Primordial black holes–perspectives in gravitational wave astronomy. *Class. Quant. Grav.*, 35(6):063001, 2018. doi: 10.1088/1361-6382/aaa7b4. 18, 59

[62] Alba Kalaja, Nicola Bellomo, Nicola Bartolo, Daniele Bertacca, Sabino Matarrese, Ilia Musco, Alvise Raccanelli, and Licia Verde. From primordial black holes

abundance to primordial curvature power spectrum (and back). *Journal of Cosmology and Astroparticle Physics*, 2019(10):031–031, Oct 2019. ISSN 1475-7516. doi: 10.1088/1475-7516/2019/10/031. URL http://dx.doi.org/10.1088/1475-7516/2019/10/031. 18, 93

[63] Sam Young and Marcello Musso. Application of peaks theory to the abundance of primordial black holes. *arXiv e-prints*, art. arXiv:2001.06469, January 2020. 19

[64] William H. Press and Paul Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. , 187:425–438, February 1974. doi: 10.1086/152650. 19

[65] G. F. Chapline. Cosmological effects of primordial black holes. , 253(5489): 251–252, January 1975. doi: 10.1038/253251a0. 19

[66] Stephen Hawking. Gravitationally collapsed objects of very low mass. *Mon. Not. Roy. Astron. Soc.*, 152:75, January 1971. doi: 10.1093/mnras/152.1.75. 19

[67] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, Jan 1933. 19

[68] Vera C. Rubin. The rotation of spiral galaxies. *Science*, 220(4604):1339–1344, 1983. ISSN 0036-8075. doi: 10.1126/science.220.4604.1339. URL https://science.sciencemag.org/content/220/4604/1339. 19

[69] M. S. Roberts. The Rotation Curve of Galaxies. In Avram Hayli, editor, *Dynamics of the Solar Systems*, volume 69 of *IAU Symposium*, page 331, Jan 1975. 19

[70] M. Davis, J. Huchra, D. W. Latham, and J. Tonry. A survey of galaxy redshifts. II. The large scale space distribution. *Astrophysical Journal*, 253:423–445, Feb 1982. doi: 10.1086/159646. 20

[71] Gianfranco Bertone and Dan Hooper. History of dark matter. *Reviews of Modern Physics*, 90(4):045002, Oct 2018. doi: 10.1103/RevModPhys.90.045002. 20

[72] Gary Steigman, Basudeb Dasgupta, and John F. Beacom. Precise relic WIMP abundance and its impact on searches for dark matter annihilation. *Physical Review D*, 86(2):023506, Jul 2012. doi: 10.1103/PhysRevD.86.023506. 20

[73] Roberto D. Peccei. The strong cp problem and axions. *Axions*, page 3–17, 2008. ISSN 0075-8450. doi: 10.1007/978-3-540-73518-2_1. URL http://dx.doi.org/10.1007/978-3-540-73518-2_1. 20

[74] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens, and O. Ruchayskiy. Sterile neutrino Dark Matter. *Progress in Particle and Nuclear Physics*, 104:1–45, January 2019. doi: 10.1016/j.ppnp.2018.07.004. 20

[75] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Physics Reports*, 267:195–373, March 1996. doi: 10.1016/0370-1573(95)00058-5. 20

[76] Mark Hindmarsh, Russell Kirk, and Stephen M. West. Dark matter from decaying topological defects. *Journal of Cosmology and Astroparticle Physics*, 2014(03): 037–037, Mar 2014. ISSN 1475-7516. doi: 10.1088/1475-7516/2014/03/037. URL http://dx.doi.org/10.1088/1475-7516/2014/03/037. 20

[77] Clare Burrage, Edmund J. Copeland, Christian Käding, and Peter Millington. Symmetron scalar fields: Modified gravity, dark matter, or both? *Phys. Rev.*, D99 (4):043539, 2019. doi: 10.1103/PhysRevD.99.043539. 20

[78] Łukasz Wyrzykowski, Alicja E. Rynkiewicz, Jan Skowron, Szymon Kozłowski, Andrzej Udalski, Michał K. Szymański, Marcin Kubiak, Igor Soszyński, Grzegorz Pietrzyński, Radosław Poleski, Paweł Pietrukowicz, and Michał Pawlak. OGLE-III Microlensing Events and the Structure of the Galactic Bulge. *Astrophysical Journal, Supplement*, 216(1):12, January 2015. doi: 10.1088/0067-0049/216/1/12. 21

[79] P. Tisserand et al. Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *Astron. Astrophys.*, 469:387–404, 2007. doi: 10.1051/0004-6361:20066017. 21, 44

[80] C. Alcock, R. A. Allsman, D. R. Alves, T. S. Axelrod, A. C. Becker, D. P. Bennett, K. H. Cook, N. Dalal, A. J. Drake, K. C. Freeman, and et al. The macho project: Microlensing results from 5.7 years of large magellanic cloud observations. *The Astrophysical Journal*, 542(1):281–307, Oct 2000. ISSN 1538-4357. doi: 10.1086/309512. URL http://dx.doi.org/10.1086/309512. 21

[81] Hiroko Niikura, Masahiro Takada, Shuichiro Yokoyama, Takahiro Sumi, and Shogo Masaki. Constraints on earth-mass primordial black holes from ogle 5-year microlensing events. *Physical Review D*, 99(8), Apr 2019. ISSN 2470-0029. doi: 10.1103/physrevd.99.083503. URL http://dx.doi.org/10.1103/PhysRevD.99.083503. 21

[82] Miguel Zumalacárregui and Uroš Seljak. Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae. *Physical Review Letters*, 121(14):141101, Oct 2018. doi: 10.1103/PhysRevLett.121.141101. 21

[83] Jakob Stegmann, Pedro R. Capelo, Elisa Bortolas, and Lucio Mayer. Improved constraints from ultra-faint dwarf galaxies on primordial black holes as dark matter. *Mon. Not. Roy. Astron. Soc.*, 492(4):5247–5260, Mar 2020. doi: 10.1093/mnras/staa170. 21

[84] Miguel A. Monroy-Rodríguez and Christine Allen. The end of the MACHO era-revisited: new limits on MACHO masses from halo wide binaries. *Astrophys. J.*, 790(2):159, 2014. doi: 10.1088/0004-637X/790/2/159. 21, 44

[85] B. J. Carr. Pregalactic black hole accretion and the thermal history of the Universe. *Monthly Notices of the Royal Astronomical Society*, 194(3):639–668, 03 1981. ISSN 0035-8711. doi: 10.1093/mnras/194.3.639. URL https://doi.org/10.1093/mnras/194.3.639. 21

[86] Massimo Ricotti, Jeremiah P. Ostriker, and Katherine J. Mack. Effect of primordial black holes on the cosmic microwave background and cosmological parameter estimates. *The Astrophysical Journal*, 680(2):829–845, Jun 2008. ISSN 1538-4357. doi: 10.1086/587831. URL http://dx.doi.org/10.1086/587831. 21

[87] Yacine Ali-Haïmoud and Marc Kamionkowski. Cosmic microwave background limits on accreting primordial black holes. *Phys. Rev.*, D95(4):043534, 2017. doi: 10.1103/PhysRevD.95.043534. 23, 44

[88] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, and et al. Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Physical Review X*, 9(3), Sep 2019. ISSN 2160-3308. doi: 10.1103/physrevx.9.031040. URL http://dx.doi.org/10.1103/PhysRevX.9.031040. 23, 107

[89] V. De Luca, G. Franciolini, and A. Riotto. On the Primordial Black Hole Mass Function for Broad Spectra. *arXiv e-prints*, art. arXiv:2001.04371, Jan 2020. 23

[90] Andrew D. Gow, Christian T. Byrnes, Alex Hall, and John A. Peacock. Primordial

black hole merger rates: distributions for multiple LIGO observables. *JCAP*, 2020 (1):031, Jan 2020. doi: 10.1088/1475-7516/2020/01/031.

[91] Anne M. Green and Andrew R. Liddle. Critical collapse and the primordial black hole initial mass function. *Physical Review D*, 60(6):063509, Sep 1999. doi: 10.1103/PhysRevD.60.063509. 23

[92] S. E. de Mink and I. Mandel. The chemically homogeneous evolutionary channel for binary black hole mergers: rates and properties of gravitational-wave events detectable by advanced LIGO. *Mon. Not. Roy. Astron. Soc.*, 460(4):3545–3553, 2016. doi: 10.1093/mnras/stw1219. 23, 58

[93] Jun'ichi Yokoyama. Cosmological constraints on primordial black holes produced in the near critical gravitational collapse. *Phys. Rev.*, D58:107502, 1998. doi: 10.1103/PhysRevD.58.107502. 24, 69

[94] Bernard Carr, Martti Raidal, Tommi Tenkanen, Ville Vaskonen, and Hardi Veermae. Primordial black hole constraints for extended mass functions. 2017. 24, 38, 80

[95] Christian T. Byrnes, Philippa S. Cole, and Subodh P. Patil. Steepest growth of the power spectrum and primordial black holes. *JCAP*, 1906:028, 2019. doi: 10.1088/1475-7516/2019/06/028. 24, 91, 92, 95

[96] Pedro Carrilho, Karim A. Malik, and David J. Mulryne. Dissecting the growth of the power spectrum for primordial black holes. *Phys. Rev.*, D100(10):103529, 2019. doi: 10.1103/PhysRevD.100.103529. 24, 91, 108

[97] Anne M. Green. Microlensing and dynamical constraints on primordial black hole dark matter with an extended mass function. *Physical Review D*, 94(6), Sep 2016. ISSN 2470-0029. doi: 10.1103/physrevd.94.063530. URL http://dx.doi.org/10.1103/PhysRevD.94.063530. 24

[98] Nicola Bellomo, José Luis Bernal, Alvise Raccanelli, and Licia Verde. Primordial black holes as dark matter: converting constraints from monochromatic to extended mass distributions. *Journal of Cosmology and Astroparticle Physics*, 2018 (01):004–004, Jan 2018. ISSN 1475-7516. doi: 10.1088/1475-7516/2018/01/004. URL http://dx.doi.org/10.1088/1475-7516/2018/01/004. 24, 80

[99] Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, and Tsutomu T. Yanagida. Double inflation as a single origin of primordial black holes for all dark matter and ligo observations. *Physical Review D*, 97(4), Feb 2018. ISSN 2470-0029. doi: 10.1103/physrevd.97.043514. URL http://dx.doi.org/10.1103/PhysRevD.97.043514. 24

[100] Sunao Sugiyama, Toshiki Kurita, and Masahiro Takada. Revisiting the wave optics effect on primordial black hole constraints from optical microlensing search. *arXiv e-prints*, art. arXiv:1905.06066, May 2019. 25

[101] Jens Chluba, Jan Hamann, and Subodh P. Patil. Features and New Physical Scales in Primordial Observables: Theory and Observation. *Int. J. Mod. Phys.*, D24(10):1530023, 2015. doi: 10.1142/S0218271815300232. 25, 27, 72

[102] J. Chluba and R. A. Sunyaev. The evolution of CMB spectral distortions in the early Universe. , 419(2):1294–1314, January 2012. doi: 10.1111/j.1365-2966.2011.19786.x. 26

[103] J. Chluba. Green's function of the cosmological thermalization problem. *Mon. Not. Roy. Astron. Soc.*, 434(1):352–357, Sep 2013. doi: 10.1093/mnras/stt1025. 26

[104] Tracy R. Slatyer and Chih-Liang Wu. Early-Universe constraints on dark matter-baryon scattering and their implications for a global 21 cm signal. *Physical Review D*, 98(2):023013, Jul 2018. doi: 10.1103/PhysRevD.98.023013. 27

[105] Yacine Ali-Haïmoud, Jens Chluba, and Marc Kamionkowski. Constraints on Dark Matter Interactions with Standard Model Particles from Cosmic Microwave Background Spectral Distortions. *Physical Review Letters*, 115(7):071304, Aug 2015. doi: 10.1103/PhysRevLett.115.071304. 27

[106] Jens Chluba and Daniel Grin. CMB spectral distortions from small-scale isocurvature fluctuations. *Mon. Not. Roy. Astron. Soc.*, 434:1619–1635, 2013. doi: 10.1093/mnras/stt1129. 27, 73

[107] Alan Kogut, Jens Chluba, Dale J. Fixsen, Stephan Meyer, and David Spergel. *The Primordial Inflation Explorer (PIXIE)*, volume 9904 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99040W. 2016. doi: 10.1117/12.2231090. 28, 110

[108] Nelson Christensen. Stochastic gravitational wave backgrounds. *Reports on Progress in Physics*, 82(1):016903, Jan 2019. doi: 10.1088/1361-6633/aae6b5. 28

[109] Kazunori Kohri and Takahiro Terada. Semianalytic calculation of gravitational wave spectrum nonlinearly induced from primordial curvature perturbations. *Phys. Rev.*, D97(12):123532, 2018. doi: 10.1103/PhysRevD.97.123532. 29, 74, 75

[110] Pravin Kumar Dahal. Review of Pulsar Timing Array for Gravitational Wave Research. *arXiv e-prints*, art. arXiv:2002.01954, February 2020. 29

[111] Z. Arzoumanian et al. The NANOGrav 11-year Data Set: Pulsar-timing Constraints On The Stochastic Gravitational-wave Background. *Astrophys. J.*, 859 (1):47, 2018. doi: 10.3847/1538-4357/aabd3b. 30

[112] L. Lentati et al. European Pulsar Timing Array Limits On An Isotropic Stochastic Gravitational-Wave Background. *Mon. Not. Roy. Astron. Soc.*, 453(3):2576–2598, 2015. doi: 10.1093/mnras/stv1538. 30, 75

[113] R. N. Manchester. Detection of Gravitational Waves using Pulsar Timing. *arXiv e-prints*, art. arXiv:1004.3602, Apr 2010. 30

[114] T J W Lazio. The square kilometre array pulsar timing array. *Classical and Quantum Gravity*, 30(22):224011, nov 2013. doi: 10.1088/0264-9381/30/22/224011. URL https://doi.org/10.1088%2F0264-9381%2F30%2F22%2F224011. 30

[115] Jonathan R. Pritchard and Abraham Loeb. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, August 2012. doi: 10.1088/0034-4885/75/8/086901. 31

[116] George B. Rybicki and Alan P. Lightman. *Radiative processes in astrophysics*. 1979. 32

[117] Steven R. Furlanetto and Michael R. Furlanetto. Spin-exchange rates in electron-hydrogen collisions. *Mon. Not. Roy. Astron. Soc.*, 374(2):547–555, Jan 2007. doi: 10.1111/j.1365-2966.2006.11169.x. 32

[118] Judd D. Bowman, Alan E. E. Rogers, Raul A. Monsalve, Thomas J. Mozdzen, and Nivedita Mahesh. An absorption profile centred at 78 megahertz in the sky-averaged spectrum. *Nature*, 555(7694):67–70, 2018. doi: 10.1038/nature25792. URL https://doi.org/10.1038/nature25792. 34

[119] Zaki S. Ali et al. PAPER-64 Constraints on Reionization: The 21 cm Power Spectrum at z = 8.4. *Astrophysical Journal*, 809(1):61, August 2015. doi: 10.1088/0004-637X/809/1/61. 34

[120] A. Ewall-Wice, Joshua S. Dillon, J. N. Hewitt, A. Loeb, A. Mesinger, A. R. Neben, A. R. Offringa, M. Tegmark, N. Barry, A. P. Beardsley, and et al. First limits on the 21cm power spectrum during the epoch of x-ray heating. *Monthly Notices of the Royal Astronomical Society*, 460(4):4320–4347, May 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw1022. URL http://dx.doi.org/10.1093/mnras/stw1022. 34

[121] F. G. Mertens et al. Improved upper limits on the 21 cm signal power spectrum of neutral hydrogen at z ≈ 9.1 from LOFAR. *Mon. Not. Roy. Astron. Soc.*, 493(2): 1662–1685, April 2020. doi: 10.1093/mnras/staa327. 34

[122] Leon Koopmans et al. Peering into the Dark (Ages) with Low-Frequency Space Interferometers. *arXiv e-prints*, art. arXiv:1908.04296, Aug 2019. 34, 85

[123] Jack Burns, Stuart Bale, Neil Bassett, Judd Bowman, Richard Bradley, Anastasia Fialkov, Steven Furlanetto, Michael Hecht, Marc Klein-Wolt, Colin Lonsdale, Robert MacDowall, Jordan Mirocha, Julian Muñoz, Bang Nhan, Jonathan Pober, David Rapetti, Alan Rogers, and Keith Tauscher. Dark Cosmology: Investigating Dark Matter and Exotic Physics in the Dark Ages using the Redshifted 21-cm Global Spectrum. *Bulletin of the AAS*, 51(3):6, May 2019. 34, 85

[124] Xuelei Chen et al. Discovering the Sky at the Longest Wavelengths with Small Satellite Constellations. *arXiv e-prints*, art. arXiv:1907.10853, Jul 2019. 34

[125] Jack Burns, Gregg Hallinan, Jim Lux, Andres Romero-Wolf, Lawrence Teitelbaum, Tzu-Ching Chang, Jonathan Kocz, Judd Bowman, Robert MacDowall, Justin Kasper, Richard Bradley, Marin Anderson, and David Rapetti. FARSIDE: A Low Radio Frequency Interferometric Array on the Lunar Farside. In *Bulletin of the AAS*, volume 51, page 178, Sep 2019. 34

[126] Adrian Liu, Max Tegmark, Judd Bowman, Jacqueline Hewitt, and Matias Zaldarriaga. An improved method for 21-cm foreground removal. *Monthly Notices of the Royal Astronomical Society*, 398(1):401–406, Sep 2009. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2009.15156.x. URL http://dx.doi.org/10.1111/j.1365-2966.2009.15156.x. 35

[127] Adrian Liu, Aaron R. Parsons, and Cathryn M. Trott. Epoch of reionization window. ii. statistical methods for foreground wedge reduction. *Physical Review D*, 90(2), Jul 2014. ISSN 1550-2368. doi: 10.1103/physrevd.90.023019. URL http://dx.doi.org/10.1103/PhysRevD.90.023019. 35

[128] Miguel F. Morales, Bryna Hazelton, Ian Sullivan, and Adam Beardsley. Four fundamental foreground power spectrum shapes for 21 cm cosmology observations. *The Astrophysical Journal*, 752(2):137, Jun 2012. ISSN 1538-4357. doi: 10.1088/0004-637x/752/2/137. URL http://dx.doi.org/10.1088/0004-637X/752/2/137. 35, 100, 102, 109

[129] José Luis Bernal, Patrick C. Breysse, Héctor Gil-Marín, and Ely D. Kovetz. User's guide to extracting cosmological information from line-intensity maps. *Physical Review D*, 100(12), Dec 2019. ISSN 2470-0029. doi: 10.1103/physrevd.100.123522. URL http://dx.doi.org/10.1103/PhysRevD.100.123522. 35, 100, 102

[130] Adrian Liu, Max Tegmark, and Matias Zaldarriaga. Will point sources spoil 21-cm tomography? *Mon. Not. Roy. Astron. Soc.*, 394(3):1575–1587, Apr 2009. doi: 10.1111/j.1365-2966.2009.14426.x. 35

[131] Hiroko Niikura, Masahiro Takada, Naoki Yasuda, Robert H. Lupton, Takahiro Sumi, Surhud More, Anupreeta More, Masamune Oguri, and Masashi Chiba. Microlensing constraints on $10^{-10}M_{\odot}$-scale primordial black holes from high-cadence observation of M31 with Hyper Suprime-Cam. 2017. 37, 44, 55

[132] Nikolaos Tetradis. Black holes and Higgs stability. *JCAP*, 1609(09):036, 2016. doi: 10.1088/1475-7516/2016/09/036. 37

[133] Gordon Kane, Kuver Sinha, and Scott Watson. Cosmological Moduli and the Post-Inflationary Universe: A Critical Review. *Int. J. Mod. Phys.*, D24(08): 1530022, 2015. doi: 10.1142/S0218271815300220. 37, 44

[134] Bobby Samir Acharya, Gordon Kane, Scott Watson, and Piyush Kumar. A Non-thermal WIMP Miracle. *Phys. Rev.*, D80:083529, 2009. doi: 10.1103/PhysRevD.80.083529. 37, 44

[135] A. G. Polnarev and M. Y. Khlopov. Primordial Black Holes and the ERA of Superheavy Particle Dominance in the Early Universe. *Soviet Journal of Astrophysics*, 25:406, 1981. 37, 44, 51

[136] Tomohiro Harada, Chul-Moon Yoo, Kazunori Kohri, Ken-ichi Nakao, and Sanjay Jhingan. Primordial black hole formation in the matter-dominated phase of the Universe. *Astrophys. J.*, 833(1):61, 2016. doi: 10.3847/1538-4357/833/1/61. 37, 44, 50, 51, 53, 78

[137] Tomohiro Harada, Chul-Moon Yoo, and Kazunori Kohri. Threshold of primordial black hole formation. *Phys. Rev.*, D88(8):084051, 2013. doi: 10.1103/PhysRevD.88.084051,10.1103/PhysRevD.89.029903. [Erratum: Phys. Rev.D89,no.2,029903(2014)]. 37

[138] Ilia Musco and John C. Miller. Primordial black hole formation in the early universe: critical behaviour and self-similarity. *Class. Quant. Grav.*, 30:145009, 2013. doi: 10.1088/0264-9381/30/14/145009. 37, 69, 72

[139] Florian Kühnel and Marit Sandstad. Ellipsoidal collapse and primordial black hole formation. *Phys. Rev.*, D94(6):063514, 2016. doi: 10.1103/PhysRevD.94.063514. 37, 72

[140] Sam Young and Christian T. Byrnes. Primordial black holes in non-Gaussian regimes. *JCAP*, 1308:052, 2013. doi: 10.1088/1475-7516/2013/08/052. 38, 72

[141] Greg W. Anderson. New Cosmological Constraints on the Higgs Boson and Top Quark Masses. *Phys. Lett.*, B243:265–270, 1990. doi: 10.1016/0370-2693(90)90849-2. 38

[142] Peter Brockway Arnold and Stamatis Vokos. Instability of hot electroweak theory: bounds on m(H) and M(t). *Phys. Rev.*, D44:3620–3627, 1991. doi: 10.1103/PhysRevD.44.3620. 38

[143] Ruth Gregory, Ian G. Moss, and Benjamin Withers. Black holes as bubble nucleation sites. *JHEP*, 03:081, 2014. doi: 10.1007/JHEP03(2014)081. 38

[144] Philipp Burda, Ruth Gregory, and Ian Moss. Vacuum metastability with black holes. *JHEP*, 08:114, 2015. doi: 10.1007/JHEP08(2015)114. 38

[145] Christian T. Byrnes, Edmund J. Copeland, and Anne M. Green. Primordial black holes as a tool for constraining non-Gaussianity. *Phys. Rev.*, D86:043512, 2012. doi: 10.1103/PhysRevD.86.043512. 39, 44, 72

[146] Andrew R Liddle and Samuel M Leach. How long before the end of inflation were observable perturbations produced? *Phys. Rev.*, D68:103503, 2003. doi: 10.1103/PhysRevD.68.103503. 41

[147] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016. doi: 10.1051/0004-6361/201525830. 41, 42, 45

[148] Anne M. Green, Andrew R. Liddle, Karim A. Malik, and Misao Sasaki. A New calculation of the mass fraction of primordial black holes. *Phys. Rev.*, D70:041502, 2004. doi: 10.1103/PhysRevD.70.041502. 42, 45

[149] Tomohiro Harada, Chul-Moon Yoo, Tomohiro Nakama, and Yasutaka Koga. Cosmological long-wavelength solutions and primordial black hole formation. *Phys. Rev.*, D91(8):084057, 2015. doi: 10.1103/PhysRevD.91.084057. 42, 72

[150] Ilia Musco, John C. Miller, and Luciano Rezzolla. Computations of primordial black hole formation. *Class. Quant. Grav.*, 22:1405–1424, 2005. doi: 10.1088/0264-9381/22/7/013. 42, 69

[151] Masaru Shibata and Misao Sasaki. Black hole formation in the Friedmann universe: Formulation and computation in numerical relativity. *Phys. Rev.*, D60: 084002, 1999. doi: 10.1103/PhysRevD.60.084002. 42

[152] Yashar Akrami, Florian Kuhnel, and Marit Sandstad. Uncertainties in primordial black-hole constraints on the primordial power spectrum. *Phys. Dark Univ.*, 19: 124–128, 2018. doi: 10.1016/j.dark.2018.01.002. 42, 72

[153] David H. Lyth. The hybrid inflation waterfall and the primordial curvature perturbation. *JCAP*, 1205:022, 2012. doi: 10.1088/1475-7516/2012/05/022. 44

[154] Sarah Shandera, Adrienne L. Erickcek, Pat Scott, and Jhon Yana Galarza. Number Counts and Non-Gaussianity. *Phys. Rev.*, D88(10):103506, 2013. doi: 10.1103/PhysRevD.88.103506.

[155] E. V. Bugaev and P. A. Klimai. Primordial black hole constraints for curvaton models with predicted large non-Gaussianity. *Int. J. Mod. Phys.*, D22:1350034, 2013. doi: 10.1142/S021827181350034X.

[156] Sam Young and Christian T. Byrnes. Long-short wavelength mode coupling tightens primordial black hole constraints. *Phys. Rev.*, D91(8):083521, 2015. doi: 10.1103/PhysRevD.91.083521. 72

[157] Sam Young, Donough Regan, and Christian T. Byrnes. Influence of large local and non-local bispectra on primordial black hole abundance. *JCAP*, 1602(02):029, 2016. doi: 10.1088/1475-7516/2016/02/029. 44

[158] B. J. Carr, J. H. Gilbert, and James E. Lidsey. Black hole relics and inflation: Limits on blue perturbation spectra. *Physical Review D*, 50(8):4853–4867, Oct 1994. ISSN 0556-2821. doi: 10.1103/physrevd.50.4853. URL http://dx.doi.org/10.1103/PhysRevD.50.4853. 44, 59

[159] A. G. Polnarev and M. Yu. Khlopov. COSMOLOGY, PRIMORDIAL BLACK HOLES, AND SUPERMASSIVE PARTICLES. *Sov. Phys. Usp.*, 28:213–232, 1985. doi: 10.1070/PU1985v028n03ABEH003858. [Usp. Fiz. Nauk145,369(1985)].

[160] Julian Georg, Gizem Sengor, and Scott Watson. Nonthermal WIMPs and primordial black holes. *Phys. Rev.*, D93(12):123523, 2016. doi: 10.1103/PhysRevD.93.123523. 78

[161] Julian Georg and Scott Watson. A Preferred Mass Range for Primordial Black Hole Formation and Black Holes as Dark Matter Revisited. 2017. 55

[162] Dmitry Gorbunov, Dmitry Levkov, and Alexander Panin. Fatal youth of the Universe: black hole threat for the electroweak vacuum during preheating. 2017. 47, 55, 56, 57

[163] Bernard Carr, Tommi Tenkanen, and Ville Vaskonen. Primordial black holes from inflaton and spectator field perturbations in a matter-dominated era. 2017. 44, 45, 46, 47, 50, 57, 77, 78

[164] A. Barnacka, J. F. Glicenstein, and R. Moderski. New constraints on primordial black holes abundance from femtolensing of gamma-ray bursts. *Phys. Rev.*, D86: 043001, 2012. doi: 10.1103/PhysRevD.86.043001. 44

[165] Fabio Capela, Maxim Pshirkov, and Peter Tinyakov. Constraints on primordial black holes as dark matter candidates from capture by neutron stars. *Phys. Rev.*, D87(12):123524, 2013. doi: 10.1103/PhysRevD.87.123524. 44

[166] Peter W. Graham, Surjeet Rajendran, and Jaime Varela. Dark Matter Triggers of Supernovae. *Phys. Rev.*, D92(6):063007, 2015. doi: 10.1103/PhysRevD.92.063007. 44

[167] R. A. Allsman et al. MACHO project limits on black hole dark matter in the 1-30 solar mass range. *Astrophys. J.*, 550:L169, 2001. doi: 10.1086/319636. 44

[168] Savvas M. Koushiappas and Abraham Loeb. Dynamics of dwarf galaxies disfavor stellar-mass black hole dark matter. 2017. 44

[169] Timothy D. Brandt. Constraints on MACHO Dark Matter from Compact Stellar Systems in Ultra-Faint Dwarf Galaxies. *Astrophys. J.*, 824(2):L31, 2016. doi: 10.3847/2041-8205/824/2/L31. 44

[170] Daniel G. Figueroa and Christian T. Byrnes. The Standard Model Higgs as the origin of the hot Big Bang. *Phys. Lett.*, B767:272–277, 2017. doi: 10.1016/j.physletb.2017.01.059. 50

[171] M. Yu. Khlopov. Primordial Black Holes. *Res. Astron. Astrophys.*, 10:495–528, 2010. doi: 10.1088/1674-4527/10/6/001. 51

[172] A. G. Doroshkevich. Spatial structure of perturbations and origin of galactic rotation in fluctuation theory . *Astrophysics*, 6:320, 1970. doi: 10.1007/BF01001625. 51

[173] D. Canko, I. Gialamas, G. Jelic-Cizmek, A. Riotto, and N. Tetradis. On the Catalysis of the Electroweak Vacuum Decay by Black Holes at High Temperature. 2017. 55, 57

[174] Ely D. Kovetz. Probing Primordial-Black-Hole Dark Matter with Gravitational Waves. 2017. 55

[175] Juan Garcia-Bellido. Massive Primordial Black Holes as Dark Matter and their detection with Gravitational Waves. *J. Phys. Conf. Ser.*, 840(1):012032, 2017. doi: 10.1088/1742-6596/840/1/012032.

[176] Juan Garcia-Bellido and Ester Ruiz Morales. Primordial black holes from single field models of inflation. 2017. 59, 61, 62, 66, 67

[177] Florian Kuhnel and Katherine Freese. Constraints on Primordial Black Holes with Extended Mass Functions. *Phys. Rev.*, D95(8):083508, 2017. doi: 10.1103/PhysRevD.95.083508. 55

[178] Amandeep S. Josan, Anne M. Green, and Karim A. Malik. Generalised constraints on the curvature perturbation from primordial black holes. *Phys. Rev.*, D79: 103520, 2009. doi: 10.1103/PhysRevD.79.103520. 57

[179] Simeon Bird, Ilias Cholis, Julian B. Munoz, Yacine Ali-Haimoud, Marc Kamionkowski, Ely D. Kovetz, Alvise Raccanelli, and Adam G. Riess. Did LIGO detect dark matter? *Phys. Rev. Lett.*, 116(20):201301, 2016. doi: 10.1103/PhysRevLett.116.201301. 58

[180] Sebastien Clesse and Juan Garcia-Bellido. The clustering of massive Primordial Black Holes as Dark Matter: measuring their mass distribution with Advanced LIGO. *Phys. Dark Univ.*, 15:142–147, 2017. doi: 10.1016/j.dark.2016.10.002. 58

[181] Misao Sasaki, Teruaki Suyama, Takahiro Tanaka, and Shuichiro Yokoyama. Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914. *Phys. Rev. Lett.*, 117(6):061101, 2016. doi: 10.1103/PhysRevLett.121.059901,10.1103/PhysRevLett.117.061101. [erratum: Phys. Rev. Lett.121,no.5,059901(2018)]. 58

[182] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016. doi: 10.1103/PhysRevLett.116.061102. 58

[183] P. Naselsky, A. D. Jackson, and Hao Liu. Understanding the LIGO GW150914 event. *JCAP*, 1608(08):029, 2016. doi: 10.1088/1475-7516/2016/08/029. 58

[184] James Creswell, Sebastian von Hausegger, Andrew D. Jackson, Hao Liu, and Pavel Naselsky. On the time lags of the LIGO signals. *JCAP*, 1708(08):013, 2017. doi: 10.1088/1475-7516/2017/08/013.

[185] James Creswell, Hao Liu, Andrew D. Jackson, Sebastian von Hausegger, and Pavel Naselsky. Degeneracy of gravitational waveforms in the context of GW150914. *JCAP*, 1803(03):007, 2018. doi: 10.1088/1475-7516/2018/03/007. 58

[186] Sebastien Clesse and Juan Garcia-Bellido. Seven Hints for Primordial Black Hole Dark Matter. 2017. 58

[187] B. P. Abbott et al. The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. *Astrophys. J.*, 833(1):L1, 2016. doi: 10.3847/2041-8205/833/1/L1. 58

[188] B. P. Abbott et al. GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence. *Astrophys. J.*, 851(2):L35, 2017. doi: 10.3847/2041-8213/aa9f0c.

[189] B. P. Abbott et al. GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 119(14):141101, 2017. doi: 10.1103/PhysRevLett.119.141101.

[190] Javier Roulet and Matias Zaldarriaga. Constraints on Binary Black Hole Populations from LIGO-Virgo Detections. 2018. 58

[191] Takeshi Chiba and Shuichiro Yokoyama. Spin Distribution of Primordial Black Holes. *PTEP*, 2017(8):083E01, 2017. doi: 10.1093/ptep/ptx087. 58

[192] Tomohiro Harada, Chul-Moon Yoo, Kazunori Kohri, and Ken-Ichi Nakao. Spins of primordial black holes formed in the matter-dominated phase of the Universe. *Phys. Rev.*, D96(8):083517, 2017. doi: 10.1103/PhysRevD.96.083517. 58

[193] K. Belczynski et al. The origin of low spin of black holes in LIGO/Virgo mergers. 2017. 58

[194] Tsvi Piran and Kenta Hotokezaka. Who Ordered That? On The Origin of LIGO's Merging Binary Black Holes. 2018. 58

[195] Krzysztof Belczynski, Daniel E. Holz, Tomasz Bulik, and Richard O'Shaughnessy. The first gravitational-wave source from the isolated evolution of two 40-100 Msun stars. *Nature*, 534:512, 2016. doi: 10.1038/nature18322. 58

[196] Carl L. Rodriguez, Carl-Johan Haster, Sourav Chatterjee, Vicky Kalogera, and Frederic A. Rasio. Dynamical Formation of the GW150914 Binary Black Hole. *Astrophys. J.*, 824(1):L8, 2016. doi: 10.3847/2041-8205/824/1/L8. 58

[197] B. J. Carr and James E. Lidsey. Primordial black holes and generalized constraints on chaotic inflation. *Phys. Rev. D*, 48:543–553, Jul 1993. doi: 10.1103/PhysRevD.48.543. URL https://link.aps.org/doi/10.1103/PhysRevD.48.543. 59

[198] P. Ivanov, P. Naselsky, and I. Novikov. Inflation and primordial black holes as dark matter. *Phys. Rev.*, D50:7173–7178, 1994. doi: 10.1103/PhysRevD.50.7173. 59

[199] Juan Garcia-Bellido, Andrei D. Linde, and David Wands. Density perturbations and black hole formation in hybrid inflation. *Phys. Rev.*, D54:6040–6058, 1996. doi: 10.1103/PhysRevD.54.6040.

[200] James S. Bullock and Joel R. Primack. NonGaussian fluctuations and primordial black holes from inflation. *Phys. Rev.*, D55:7423–7439, 1997. doi: 10.1103/PhysRevD.55.7423.

[201] Kazunori Kohri, David H. Lyth, and Alessandro Melchiorri. Black hole formation and slow-roll inflation. *JCAP*, 0804:038, 2008. doi: 10.1088/1475-7516/2008/04/038.

[202] Hiranya V. Peiris and Richard Easther. Primordial Black Holes, Eternal Inflation, and the Inflationary Parameter Space after WMAP5. *JCAP*, 0807:024, 2008. doi: 10.1088/1475-7516/2008/07/024.

[203] Amandeep S. Josan and Anne M. Green. Constraints from primordial black hole formation at the end of inflation. *Phys. Rev.*, D82:047303, 2010. doi: 10.1103/PhysRevD.82.047303.

[204] Manuel Drees and Encieh Erfani. Running-Mass Inflation Model and Primordial Black Holes. *JCAP*, 1104:005, 2011. doi: 10.1088/1475-7516/2011/04/005.

[205] Andrei Linde, Sander Mooij, and Enrico Pajer. Gauge field production in supergravity inflation: Local non-Gaussianity and primordial black holes. *Phys. Rev.*, D87(10):103506, 2013. doi: 10.1103/PhysRevD.87.103506.

[206] Edgar Bugaev and Peter Klimai. Axion inflation with gauge field production and primordial black holes. *Phys. Rev.*, D90(10):103501, 2014. doi: 10.1103/PhysRevD.90.103501.

[207] Masahiro Kawasaki, Alexander Kusenko, Yuichiro Tada, and Tsutomu T. Yanagida. Primordial black holes as dark matter in supergravity inflation models. *Phys. Rev.*, D94(8):083523, 2016. doi: 10.1103/PhysRevD.94.083523.

[208] Mark P. Hertzberg and Masaki Yamada. Primordial Black Holes from Polynomial Potentials in Single Field Inflation. *Phys. Rev.*, D97(8):083509, 2018. doi: 10.1103/PhysRevD.97.083509. 61, 62, 67, 108, 157

[209] Haoran Di and Yungui Gong. Primordial black holes and second order gravitational waves from ultra-slow-roll inflation. *Journal of Cosmology and Astroparticle Physics*, 2018(07):007, 2018. URL http://stacks.iop.org/1475-7516/2018/i=07/a=007. 62, 74

[210] Kristjan Kannike, Luca Marzola, Martti Raidal, and Hardi Veermae. Single Field Double Inflation and Primordial Black Holes. *JCAP*, 1709(09):020, 2017. doi: 10.1088/1475-7516/2017/09/020.

[211] Hayato Motohashi and Wayne Hu. Primordial Black Holes and Slow-Roll Violation. *Phys. Rev.*, D96(6):063503, 2017. doi: 10.1103/PhysRevD.96.063503. 85

[212] Guillermo Ballesteros and Marco Taoso. Primordial black hole dark matter from single field inflation. *Phys. Rev.*, D97(2):023501, 2018. doi: 10.1103/PhysRevD.97.023501. 62, 108

[213] Michele Cicoli, Victor A. Diaz, and Francisco G. Pedro. Primordial Black Holes from String Inflation. *JCAP*, 1806(06):034, 2018. doi: 10.1088/1475-7516/2018/06/034. 62, 67

[214] Ogan Özsoy, Susha Parameswaran, Gianmassimo Tasinato, and Ivonne Zavala. Mechanisms for Primordial Black Hole Production in String Theory. *JCAP*, 1807: 005, 2018. doi: 10.1088/1475-7516/2018/07/005. 59

[215] Mansour Karami, Niayesh Afshordi, and Jesus Zavala. Forward modelling of quasar light curves and the cosmological matter power spectrum on milliparsec scales. 2018. 60

[216] N. C. Tsamis and Richard P. Woodard. Improved estimates of cosmological perturbations. *Phys. Rev.*, D69:084005, 2004. doi: 10.1103/PhysRevD.69.084005. 61

[217] William H. Kinney. Horizon crossing and inflation with large eta. *Phys. Rev.*, D72: 023515, 2005. doi: 10.1103/PhysRevD.72.023515. 61

[218] Matteo Biagetti, Gabriele Franciolini, Alex Kehagias, and Antonio Riotto. Primordial Black Holes from Inflation and Quantum Diffusion. 2018. 61, 67, 71

[219] Syksy Rasanen and Eemeli Tomberg. Planck scale black hole dark matter from Higgs inflation. 2018. 63

[220] W. Israel. Singular hypersurfaces and thin shells in general relativity. *Nuovo Cim.*, B44S10:1, 1966. doi: 10.1007/BF02710419,10.1007/BF02712210. [Nuovo Cim.B44,1(1966)]. 63, 146

[221] Nathalie Deruelle and Viatcheslav F. Mukhanov. On matching conditions for cosmological perturbations. *Phys. Rev.*, D52:5549–5555, 1995. doi: 10.1103/PhysRevD.52.5549. 63, 146

[222] Jennifer A. Adams, Bevan Cresswell, and Richard Easther. Inflationary perturbations from a potential with a step. *Phys. Rev.*, D64:123514, 2001. doi: 10.1103/PhysRevD.64.123514. 63

[223] David Seery. CppTransport: a platform to automate calculation of inflationary correlation functions. 2016. doi: 10.5281/zenodo.61239. 66, 81

[224] Mafalda Dias, Jonathan Frazer, David J. Mulryne, and David Seery. Numerical evaluation of the bispectrum in multiple field inflation?the transport approach with code. *JCAP*, 1612(12):033, 2016. doi: 10.1088/1475-7516/2016/12/033. 66

[225] Juan Garcia-Bellido, Marco Peloso, and Caner Unal. Gravitational waves at interferometer scales and primordial black holes in axion inflation. *JCAP*, 1612 (12):031, 2016. doi: 10.1088/1475-7516/2016/12/031. 67

[226] P. J. E. Peebles. The Effect of a Lumpy Matter Distribution on the Growth of Irregularities in an Expanding Universe. *Astron. & Astrophys.*, 32:391, June 1974. 68

[227] Florian Kühnel, Cornelius Rampf, and Marit Sandstad. Effects of Critical Collapse on Primordial Black-Hole Mass Spectra. *Eur. Phys. J.*, C76(2):93, 2016. doi: 10.1140/epjc/s10052-016-3945-8. 69

[228] Bernard Carr, Florian Kuhnel, and Marit Sandstad. Primordial Black Holes as Dark Matter. *Phys. Rev.*, D94(8):083504, 2016. doi: 10.1103/PhysRevD.94.083504. 69

[229] Jens C. Niemeyer and K. Jedamzik. Near-critical gravitational collapse and the initial mass function of primordial black holes. *Phys. Rev. Lett.*, 80:5481–5484, 1998. doi: 10.1103/PhysRevLett.80.5481. 69

[230] Ilia Musco, John C. Miller, and Alexander G. Polnarev. Primordial black hole formation in the radiative era: Investigation of the critical nature of the collapse. *Class. Quant. Grav.*, 26:235001, 2009. doi: 10.1088/0264-9381/26/23/235001. 69

[231] Christian T. Byrnes, Mark Hindmarsh, Sam Young, and Michael R. S. Hawkins. Primordial black holes with an accurate QCD equation of state. *JCAP*, 1808(08): 041, 2018. doi: 10.1088/1475-7516/2018/08/041. 69, 72, 77

[232] Chris Pattison, Vincent Vennin, Hooshyar Assadullahi, and David Wands. Quantum diffusion during inflation and primordial black holes. *JCAP*, 1710(10): 046, 2017. doi: 10.1088/1475-7516/2017/10/046. 71, 110

[233] Jose Maria Ezquiaga and Juan Garcia-Bellido. Quantum diffusion beyond slow-roll: implications for primordial black-hole production. 2018. 71, 110

[234] Diego Cruces, Cristiano Germani, and Tomislav Prokopec. Failure of the stochastic approach to inflation in constant-roll and ultra-slow-roll. 2018.

[235] Hassan Firouzjahi, Amin Nassiri-Rad, and Mahdiyar Noorbala. Stochastic Ultra Slow Roll Inflation. 2018. 71

[236] Mohammad Hossein Namjoo, Hassan Firouzjahi, and Misao Sasaki. Violation of non-Gaussianity consistency relation in a single field inflationary model. *EPL*, 101 (3):39001, 2013. doi: 10.1209/0295-5075/101/39001. 71

[237] Sander Mooij and Gonzalo A. Palma. Consistently violating the non-Gaussian consistency relation. *JCAP*, 1511(11):025, 2015. doi: 10.1088/1475-7516/2015/11/025. 71

[238] Yi-Fu Cai, Xingang Chen, Mohammad Hossein Namjoo, Misao Sasaki, Dong-Gang Wang, and Ziwei Wang. Revisiting non-Gaussianity from non-attractor inflation models. *JCAP*, 1805(05):012, 2018. doi: 10.1088/1475-7516/2018/05/012. 72, 150

[239] Rafael Bravo, Sander Mooij, Gonzalo A. Palma, and Bastián Pradenas. A generalized non-Gaussian consistency relation for single field inflation. *JCAP*, 1805 (05):024, 2018. doi: 10.1088/1475-7516/2018/05/024. 72

[240] Rafael Bravo, Sander Mooij, Gonzalo A. Palma, and Bastián Pradenas. Vanishing of local non-Gaussianity in canonical single field inflation. *JCAP*, 1805(05):025, 2018. doi: 10.1088/1475-7516/2018/05/025. 72

[241] G. Franciolini, A. Kehagias, S. Matarrese, and A. Riotto. Primordial Black Holes from Inflation and non-Gaussianity. *JCAP*, 1803(03):016, 2018. doi: 10.1088/1475-7516/2018/03/016. 72

[242] Vicente Atal and Cristiano Germani. The role of non-gaussianities in Primordial Black Hole formation. 2018. 72

[243] K. Jedamzik and Jens C. Niemeyer. Primordial black hole formation during first order phase transitions. *Phys. Rev. D*, 59:124014, 1999. doi: 10.1103/PhysRevD.59.124014. 72

[244] Tomohiro Nakama, Tomohiro Harada, A. G. Polnarev, and Jun'ichi Yokoyama. Identifying the most crucial parameters of the initial curvature profile for primordial black hole formation. *JCAP*, 1401:037, 2014. doi: 10.1088/1475-7516/2014/01/037. 72

[245] Chul-Moon Yoo, Tomohiro Harada, Jaume Garriga, and Kazunori Kohri. PBH abundance from random Gaussian curvature perturbations and a local density threshold. 2018. 72

[246] Philippa S. Cole and Christian T. Byrnes. Extreme scenarios: the tightest possible constraints on the power spectrum due to primordial black holes. *JCAP*, 1802(02): 019, 2018. doi: 10.1088/1475-7516/2018/02/019. 72, 78, 94

[247] Jens Chluba, Liang Dai, Daniel Grin, Mustafa Amin, and Marc Kamionkowski. Spectral distortions from the dissipation of tensor perturbations. *Mon. Not. Roy. Astron. Soc.*, 446:2871–2886, 2015. doi: 10.1093/mnras/stu2277. 73

[248] A. Kogut et al. The Primordial Inflation Explorer (PIXIE): A Nulling Polarimeter for Cosmic Microwave Background Observations. *JCAP*, 1107:025, 2011. doi: 10.1088/1475-7516/2011/07/025. 74, 91

[249] Kishore N. Ananda, Chris Clarkson, and David Wands. The Cosmological gravitational wave background from primordial density perturbations. *Phys. Rev.*, D75:123518, 2007. doi: 10.1103/PhysRevD.75.123518. 74, 75

[250] Bob Osano, Cyril Pitrou, Peter Dunsby, Jean-Philippe Uzan, and Chris Clarkson. Gravitational waves generated by second order effects during inflation. *JCAP*, 0704:003, 2007. doi: 10.1088/1475-7516/2007/04/003.

[251] Daniel Baumann, Paul J. Steinhardt, Keitaro Takahashi, and Kiyotomo Ichiki. Gravitational Wave Spectrum Induced by Primordial Scalar Perturbations. *Phys. Rev.*, D76:084019, 2007. doi: 10.1103/PhysRevD.76.084019. 75

[252] Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, Yuichiro Tada, and Tsutomu T. Yanagida. Inflationary primordial black holes for the LIGO gravitational wave events and pulsar timing array experiments. *Phys. Rev.*, D95 (12):123510, 2017. doi: 10.1103/PhysRevD.95.123510. 74

[253] Juan Garcia-Bellido, Marco Peloso, and Caner Unal. Gravitational Wave signatures of inflationary models from Primordial Black Hole Dark Matter. *JCAP*, 1709(09):013, 2017. doi: 10.1088/1475-7516/2017/09/013.

[254] Kenta Ando, Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, and Tsutomu T. Yanagida. Primordial black holes for the LIGO events in the axionlike curvaton model. *Phys. Rev.*, D97(12):123512, 2018. doi: 10.1103/PhysRevD.97.123512.

[255] Ryo Saito and Jun'ichi Yokoyama. Gravitational wave background as a probe of the primordial black hole abundance. *Phys. Rev. Lett.*, 102:161101, 2009. doi: 10.1103/PhysRevLett.102.161101,10.1103/PhysRevLett.107.069901. [Erratum: Phys. Rev. Lett.107,069901(2011)]. 81

[256] Tomohiro Nakama and Teruaki Suyama. Primordial black holes as a novel probe of primordial gravitational waves. II: Detailed analysis. *Phys. Rev.*, D94(4):043507, 2016. doi: 10.1103/PhysRevD.94.043507.

[257] Nicholas Orlofsky, Aaron Pierce, and James D. Wells. Inflationary theory and pulsar timing investigations of primordial black holes and gravitational waves. *Phys. Rev.*, D95(6):063518, 2017. doi: 10.1103/PhysRevD.95.063518.

[258] Chiara Caprini and Daniel G. Figueroa. Cosmological Backgrounds of Gravitational Waves. *Class. Quant. Grav.*, 35(16):163001, 2018. doi: 10.1088/1361-6382/aac608. 75

[259] Kenta Ando, Masahiro Kawasaki, and Hiromasa Nakatsuka. Formation of primordial black holes as dark matter or LIGO black hole binaries in an axion-like curvaton model. 2018. 74

[260] Andrew D. Gow, Christian T. Byrnes, Philippa S. Cole, and Sam Young. The power spectrum on small scales: Robust constraints and comparing PBH methodologies. *arXiv e-prints*, art. arXiv:2008.03289, August 2020. 74

[261] Edgar Bugaev and Peter Klimai. Induced gravitational wave background and primordial black holes. *Phys. Rev.*, D81:023517, 2010. doi: 10.1103/PhysRevD.81.023517. 75

[262] N. Bartolo, V. De Luca, G. Franciolini, M. Peloso, D. Racco, and A. Riotto. Testing Primordial Black Holes as Dark Matter through LISA. 2018. 75

[263] Rong-gen Cai, Shi Pi, and Misao Sasaki. Gravitational Waves Induced by non-Gaussian Scalar Perturbations. 2018. 75

[264] Caner Unal. Imprints of Primordial Non-Gaussianity on Gravitational Wave Spectrum. 2018. 75

[265] Jose Ramon Espinosa, Davide Racco, and Antonio Riotto. A Cosmological Signature of the SM Higgs Instability: Gravitational Waves. *JCAP*, 1809(09):012, 2018. doi: 10.1088/1475-7516/2018/09/012. 75

[266] Z. Arzoumanian et al. The NANOGrav Nine-year Data Set: Limits on the Isotropic Stochastic Gravitational Wave Background. *Astrophys. J.*, 821(1):13, 2016. doi: 10.3847/0004-637X/821/1/13. 75

[267] R. M. Shannon et al. Gravitational waves from binary supermassive black holes missing in pulsar observations. *Science*, 349(6255):1522–1525, 2015. doi: 10.1126/science.aab1910. 75

[268] Torsten Bringmann, Pat Scott, and Yashar Akrami. Improved constraints on the primordial power spectrum at small scales from ultracompact minihalos. *Phys. Rev.*, D85:125027, 2012. doi: 10.1103/PhysRevD.85.125027. 76

[269] Mateja Gosenca, Julian Adamek, Christian T. Byrnes, and Shaun Hotchkiss. 3D simulations with boosted primordial power spectra and ultracompact minihalos. *Phys. Rev.*, D96(12):123519, 2017. doi: 10.1103/PhysRevD.96.123519. 76

[270] M. Sten Delos, Adrienne L. Erickcek, Avery P. Bailey, and Marcelo A. Alvarez. Are ultracompact minihalos really ultracompact? *Phys. Rev.*, D97(4):041303, 2018. doi: 10.1103/PhysRevD.97.041303.

[271] Tomohiro Nakama, Teruaki Suyama, Kazunori Kohri, and Nagisa Hiroshima. Constraints on small-scale primordial power by annihilation signals from extragalactic dark matter minihalos. *Phys. Rev.*, D97(2):023539, 2018. doi: 10.1103/PhysRevD.97.023539.

[272] M. Sten Delos, Adrienne L. Erickcek, Avery P. Bailey, and Marcelo A. Alvarez. The density profiles of ultracompact minihalos: implications for constraining the primordial power spectrum. 2018. 76

[273] Brian C. Lacki and John F. Beacom. Primordial Black Holes as Dark Matter: Almost All or Almost Nothing. *Astrophys. J.*, 720:L67–L71, 2010. doi: 10.1088/2041-8205/720/1/L67. 76

[274] Yu. N. Eroshenko. Dark matter density spikes around primordial black holes. *Astron. Lett.*, 42(6):347–356, 2016. doi: 10.1134/S1063773716060013. [Pisma Astron. Zh.42,no.6,359(2016)].

[275] Sofiane M. Boucenna, Florian Kuhnel, Tommy Ohlsson, and Luca Visinelli. Novel Constraints on Mixed Dark-Matter Scenarios of Primordial Black Holes and WIMPs. *JCAP*, 1807(07):003, 2018. doi: 10.1088/1475-7516/2018/07/003. 76

[276] Tomohiro Nakama, Bernard Carr, and Joseph Silk. Limits on primordial black holes from μ distortions in cosmic microwave background. *Phys. Rev.*, D97(4): 043525, 2018. doi: 10.1103/PhysRevD.97.043525. 77

[277] Donghui Jeong, Josef Pradler, Jens Chluba, and Marc Kamionkowski. Silk damping at a redshift of a billion: a new limit on small-scale adiabatic perturbations. *Phys. Rev. Lett.*, 113:061301, 2014. doi: 10.1103/PhysRevLett.113.061301. 77

[278] Keisuke Inomata, Masahiro Kawasaki, Kyohei Mukaida, and Tsutomu T. Yanagida. Double inflation as a single origin of primordial black holes for all dark matter and LIGO observations. *Phys. Rev.*, D97(4):043514, 2018. doi: 10.1103/PhysRevD.97.043514. 77

[279] Shu-Lin Cheng, Wolung Lee, and Kin-Wang Ng. Superhorizon curvature perturbation in ultra-slow-roll inflation. 2018. 80

[280] C. J. Moore, R. H. Cole, and C. P. L. Berry. Gravitational-wave sensitivity curves. *Class. Quant. Grav.*, 32(1):015014, 2015. doi: 10.1088/0264-9381/32/1/015014. 81

[281] Keisuke Inomata and Tomohiro Nakama. Gravitational waves induced by scalar perturbations as probes of the small-scale primordial spectrum. 2018. 81, 95

[282] J. O. Burns. A moon-earth radio interferometer. In W. W. Mendell, editor, *Lunar Bases and Space Activities of the 21st Century*, pages 293–300, Jan 1985. 85

[283] Baptiste Cecconi et al. NOIRE Study Report: Towards a Low Frequency Radio Interferometer in Space. In *EGU General Assembly Conference Abstracts*, volume 20 of *EGU General Assembly Conference Abstracts*, page 3648, Apr 2018. 85

[284] Bernard Carr and Joseph Silk. Primordial Black Holes as Generators of Cosmic Structures. *Mon. Not. Roy. Astron. Soc.*, 478(3):3756–3775, 2018. doi: 10.1093/mnras/sty1204. 87, 107

[285] JosÃ© Luis Bernal, Alvise Raccanelli, Licia Verde, and Joseph Silk. Signatures of primordial black holes as seeds of supermassive black holes. *JCAP*, 1805(05):017, 2018. doi: 10.1088/1475-7516/2018/05/017. 93, 94, 99, 102, 107, 109

[286] Muhammad A. Latif and Andrea Ferrara. Formation of supermassive black hole seeds. *Publ. Astron. Soc. Austral.*, 33:e051, 2016. doi: 10.1017/pasa.2016.41. 87

[287] B. Zygelman. Hyperfine Level-changing Collisions of Hydrogen Atoms and Tomography of the Dark Age Universe. *Astrophysical Journal*, 622:1356–1362, April 2005. doi: 10.1086/427682. 87

[288] S. Seager, D. D. Sasselov, and D. Scott. A New Calculation of the Recombination Epoch. *apjl*, 523:L1–L5, September 1999. doi: 10.1086/312250. 88

[289] Annalisa Pillepich, Cristiano Porciani, and Sabino Matarrese. The bispectrum of redshifted 21-cm fluctuations from the dark ages. *Astrophys. J.*, 662:1–14, 2007. doi: 10.1086/517963. 89

[290] Antony Lewis and Anthony Challinor. 21cm angular-power spectrum from the dark ages. *Physical Review D*, 76(8), Oct 2007. ISSN 1550-2368. doi: 10.1103/physrevd.76.083005. URL http://dx.doi.org/10.1103/PhysRevD.76.083005. 89

[291] Yacine Ali-Haïmoud, P. Daniel Meerburg, and Sihan Yuan. New light on 21 cm intensity fluctuations from the dark ages. *Phys. Rev.*, D89(8):083506, 2014. doi: 10.1103/PhysRevD.89.083506. 89, 92, 97

[292] Ogan Özsoy and Gianmassimo Tasinato. On the slope of curvature power spectrum in non-attractor inflation. *arXiv e-prints*, art. arXiv:1912.01061, Dec 2019. 91

[293] Samuel Passaglia, Wayne Hu, and Hayato Motohashi. Primordial black holes and local non-Gaussianity in canonical inflation. *Phys. Rev.*, D99(4):043536, 2019. doi: 10.1103/PhysRevD.99.043536. 92

[294] Shintaro Yoshiura, Keitaro Takahashi, and Tomo Takahashi. Probing Small Scale Primordial Power Spectrum with 21cm Line Global Signal. *arXiv e-prints*, art. arXiv:1911.07442, Nov 2019. 92

[295] Julian B. Muñoz, Cora Dvorkin, and Francis-Yan Cyr-Racine. Probing the Small-Scale Matter Power Spectrum with Large-Scale 21-cm Data. *arXiv e-prints*, art. arXiv:1911.11144, Nov 2019. 92

[296] A. Kehagias, I. Musco, and A. Riotto. Non-Gaussian formation of primordial black holes: effects on the threshold. *JCAP*, 2019(12):029, Dec 2019. doi: 10.1088/1475-7516/2019/12/029. 93

[297] Katherine J. Mack and Daniel H. Wesley. Primordial black holes in the Dark Ages: Observational prospects for future 21cm surveys. *arXiv e-prints*, art. arXiv:0805.1531, May 2008. 93, 94, 109

[298] Hiroyuki Tashiro and Naoshi Sugiyama. The effect of primordial black holes on 21 cm fluctuations. *Mon. Not. Roy. Astron. Soc.*, 435:3001, 2013. doi: 10.1093/mnras/stt1493. 102

[299] Olga Mena, Sergio Palomares-Ruiz, Pablo Villanueva-Domingo, and Samuel J. Witte. Constraining the primordial black hole abundance with 21-cm cosmology. *Physical Review D*, 100(4):043540, Aug 2019. doi: 10.1103/PhysRevD.100.043540. 93, 94, 102, 109

[300] N. Afshordi, P. McDonald, and D. N. Spergel. Primordial black holes as dark matter: The Power spectrum and evaporation of early structures. *Astrophys. J.*, 594:L71–L74, 2003. doi: 10.1086/378763. 93, 95

[301] Riccardo Murgia, Giulio Scelfo, Matteo Viel, and Alvise Raccanelli. Lyman- forest constraints on primordial black holes as dark matter. *Physical Review Letters*, 123 (7), Aug 2019. ISSN 1079-7114. doi: 10.1103/physrevlett.123.071102. URL http://dx.doi.org/10.1103/PhysRevLett.123.071102. 95

[302] Jinn-Ouk Gong and Naoya Kitajima. Distribution of primordial black holes and

21cm signature. *JCAP*, 1811(11):041, 2018. doi: 10.1088/1475-7516/2018/11/041. 95, 102, 109

[303] Garrelt Mellema et al. Reionization and the Cosmic Dawn with the Square Kilometre Array. *Experimental Astronomy*, 36(1-2):235–318, August 2013. doi: 10.1007/s10686-013-9334-5. 99

[304] Steven R. Furlanetto, S. Peng Oh, and Elena Pierpaoli. The Effects of Dark Matter Decay and Annihilation on the High-Redshift 21 cm Background. *Phys. Rev.*, D74:103502, 2006. doi: 10.1103/PhysRevD.74.103502. 99

[305] Aviad Cohen, Anastasia Fialkov, Rennan Barkana, and Matan Lotem. Charting the Parameter Space of the Global 21-cm Signal. *Mon. Not. Roy. Astron. Soc.*, 472 (2):1915–1931, 2017. doi: 10.1093/mnras/stx2065. 102

[306] Adrian Liu and Aaron R. Parsons. Constraining cosmology and ionization history with combined 21cm power spectrum and global signal measurements. *Monthly Notices of the Royal Astronomical Society*, 457(2):1864â1877, Feb 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw071. URL http://dx.doi.org/10.1093/mnras/stw071. 103

[307] Julian B. Muñoz, Yacine Ali-Haïmoud, and Marc Kamionkowski. Primordial non-gaussianity from the bispectrum of 21-cm fluctuations in the dark ages. *Phys. Rev.*, D92(8):083508, 2015. doi: 10.1103/PhysRevD.92.083508. 103

[308] M Punturo et al. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity*, 27(8):084007, apr 2010. doi: 10.1088/0264-9381/27/8/084007. URL https://doi.org/10.1088%2F0264-9381%2F27%2F8%2F084007. 110

[309] Shaul Hanany et al. PICO: Probe of Inflation and Cosmic Origins. *arXiv e-prints*, art. arXiv:1902.10541, February 2019. 111

[310] L. F. Abbott and Jennie H. Traschen. Causality constraints on cosmological perturbations. *Astrophys. J.*, 302:39–42, 1986. doi: 10.1086/163970. 155

[311] Ruth Durrer and Chiara Caprini. Primordial magnetic fields and causality. *JCAP*, 0311:010, 2003. doi: 10.1088/1475-7516/2003/11/010. 155

[312] Kenji Kadota, Scott Dodelson, Wayne Hu, and Ewan D. Stewart. Precision of inflaton potential reconstruction from CMB using the general slow-roll

approximation. *Phys. Rev.*, D72:023510, 2005. doi: 10.1103/PhysRevD.72.023510. 156

[313] James M. Cline and Loison Hoi. Inflationary potential reconstruction for a wmap running power spectrum. *JCAP*, 0606:007, 2006. doi: 10.1088/1475-7516/2006/06/007.

[314] Rachel Bean, Daniel J. H. Chung, and Ghazal Geshnizjani. Reconstructing a general inflationary action. *Phys. Rev.*, D78:023517, 2008. doi: 10.1103/PhysRevD.78.023517. 156

Appendix A

The primordial power spectrum from matching

It is possible to arrive at an analytic understanding of various features of the shape of the primordial power spectrum generated by transiting into and out of a phase of ultra slow-roll (USR) inflation by approximating the evolution of η as a series of phases of constant η, and matching between these phases. In this appendix, we'll perform a series of matchings, culminating in a four-stage matching from $\eta \equiv 0 \to -2 \to -6 \to 2 \to 0$. We begin by matching from $\eta = 0$ to $\eta = -6$ (USR) and back to $\eta = 0$. For this we need the mode functions for inflation for each phase, which are obtained from the equation of motion for the Mukhanov-Sasaki variable, $v_k = z\mathcal{R}$, where $z^2 = 2a^2 M_{\rm pl}^2 \epsilon$ and $\mathcal{R}$ is the comoving curvature perturbation:

$$v_k'' + (k^2 - \frac{z''}{z})v_k = 0. \tag{A.1}$$

By directly differentiating z with respect to conformal time, it can be shown that

$$\frac{z''}{z} = (aH)^2 \left(2 - \epsilon + \frac{3}{2}\eta + \frac{1}{4}\eta^2 - \frac{1}{2}\epsilon\eta + \frac{1}{2}\frac{\dot{\eta}}{H}\right). \tag{A.2}$$

Equation (A.2) is exact to all orders. Assuming that $\epsilon \ll 1$, we can rewrite equation (A.1) as

$$v_k'' + (k^2 - \frac{\nu^2 - \frac{1}{4}}{\tau^2})v_k = 0, \tag{A.3}$$

with a new parameter ν defined as

$$\nu^2 = \frac{9}{4} + \frac{3}{2}\eta + \frac{1}{4}\eta^2 + \frac{\dot{\eta}}{2H}. \tag{A.4}$$

The solutions for the canonically normalised mode function are of the form

$$v_k = \frac{\sqrt{\pi}}{2} e^{i(\nu+\frac{1}{2})\frac{\pi}{2}} \sqrt{-\tau} H_\nu^{(1)}(-k\tau) \tag{A.5}$$

in linear combination with its complex conjugate, where $H_\nu^{(1)}(-k\tau)$ is the Hankel function of the first kind.

For a constant η phase, the last term in (A.4) vanishes, and we find that for $\eta = -6$ and $\eta = 0$, $\nu^2 = 9/4$, with $\nu = -3/2$ corresponding to USR and $\nu = 3/2$ corresponding to SR. The curvature perturbation is obtained via $\mathcal{R}_k = v/z$ with $z^2 = 2a^2 M_{\rm pl}^2 \epsilon$, and we now find the mode equations for each phase. For $\nu = 3/2$ (phase 1, $\eta = 0$), the curvature perturbation is given by

$$\mathcal{R}_k^{(1)} = i\frac{H}{M_{\rm pl}}\frac{1}{\sqrt{4\epsilon_1 k^3}}\left[c(1+ik\tau)e^{-ik\tau} - s(1-ik\tau)e^{ik\tau}\right], \tag{A.6}$$

where $\epsilon(\tau) = \epsilon_1$ is treated as constant and c, s are constant coefficients to be found via the matching, and they should satisfy the Wronskian condition. This needs to be matched to the Bunch-Davies vacuum in the limit $\tau \to -\infty$, so the mode equation during the first phase of $\eta = 0$ reduces to

$$\mathcal{R}_k^{(1)} = i\frac{H}{M_{\rm pl}}\frac{e^{-ik\tau}}{\sqrt{4\epsilon_1 k^3}}(1+ik\tau) \tag{A.7}$$

i.e. $c = 1$ and $s = 0$. Using the relation $H_{-3/2}^{(1)} = -iH_{3/2}^{(1)}$, and writing $\epsilon(\tau)$ during USR as

$$\epsilon(\tau) = \epsilon_1\left(\frac{a(\tau_1)}{a(\tau)}\right)^6 = \epsilon_1\left(\frac{\tau}{\tau_1}\right)^6 \tag{A.8}$$

where τ_1 is the time of transition between $\eta = 0$ and $\eta = -6$ and the second equality comes from $aH = -1/\tau$, we find the canonically normalised mode functions during the phase of $\eta = -6$ (phase 2, USR) to be

$$\mathcal{R}_k^{(2)} = i\frac{H}{M_{\rm pl}}\frac{(\tau_1/\tau)^3}{\sqrt{4\epsilon_1 k^3}}\left[c_1(1+ik\tau)e^{-ik\tau} - s_1(1-ik\tau)e^{ik\tau}\right], \tag{A.9}$$

where τ_1 and ϵ_1 are fixed, and the coefficients c_1 and s_1 will be determined by the matching from the $\eta = 0$ phase, and hence will be in terms of k and τ_1.

The matching conditions between the two phases are given by the Israel junction conditions [220, 221]

$$[\mathcal{R}_k]_\pm = 0, \ \left[z^2\mathcal{R}_k'\right]_\pm = 0. \tag{A.10}$$

The first of these is determined by requiring that the metric be continuous across the transition. The second follows from the equation for the mode function $\mathcal{R}_k$:

$$\mathcal{R}_k'' + \frac{(z^2)'}{z^2}\mathcal{R}_k' + k^2\mathcal{R}_k = 0, \tag{A.11}$$

which implies that

$$\left(z^2\mathcal{R}_k'\right)' = -z^2k^2\mathcal{R}_k. \tag{A.12}$$

Integrating the above over an infinitesimal interval around the transition and recalling the continuity of $\mathcal{R}_k$ and z^2 results in the second condition in (A.10). So, continuity between $\mathcal{R}_k^{(1)}$ and $\mathcal{R}_k^{(2)}$ at τ_1 results in the equation

$$(1+ik\tau_1)e^{-ik\tau_1} = c_1(1+ik\tau_1)e^{-ik\tau_1} - s_1(1-ik\tau_1)e^{ik\tau_1}, \tag{A.13}$$

and continuity of the time derivatives $\mathcal{R}_k'^{(1)}$ and $\mathcal{R}_k'^{(2)}$ at τ_1 requires

$$k^2\tau_1 e^{-ik\tau_1} = c_1 e^{-ik\tau_1}\left(k^2\tau_1 - \frac{3}{\tau_1}(1+ik\tau_1)\right) - s_1 e^{ik\tau_1}\left(k^2\tau_1 - \frac{3}{\tau_1}(1-ik\tau_1)\right), \tag{A.14}$$

which together imply that

$$s_1 = \frac{3\,i\,e^{-2ik\tau_1}}{2(k\tau_1)^3}(1+ik\tau_1)^2 \tag{A.15}$$

and

$$c_1 = 1 + \frac{3i(1+k^2\tau_1^2)}{2(k\tau_1)^3}. \tag{A.16}$$

We see that in the limit $\tau_1 k \to -\infty$, that is for modes that are deep within the Hubble radius at τ_1, $\theta_k \to 0$, which means that the corresponding modes are still in the BD vacuum during USR.

In order to meaningfully talk about a late time power spectrum, we need to end USR, otherwise the modes grow unboundedly. To model this, we consider a transition from USR back to a phase of $\eta = 0$. In the final phase, the mode function corresponds to the usual case (A.6) but with constant ϵ given by $\epsilon_2 = \epsilon_1(a_1/a_2)^6 = \epsilon_1(\tau_2/\tau_1)^6$, where $\log(a_2/a_1)$ is the total number of e-foldings of USR:

$$\mathcal{R}_k^{(3)} = i\frac{H}{M_{\rm pl}}\frac{(\tau_1/\tau_2)^3}{\sqrt{4\epsilon_1 k^3}}\left[c_2(1+ik\tau)e^{-ik\tau} - s_2(1-ik\tau)e^{ik\tau}\right]. \tag{A.17}$$

We therefore need to compute another matching between the mode functions in (A.9) and (A.17) at time τ_2, which is when USR ends, in order to determine c_2 and s_2 which will be functions of k and τ_2. From requiring continuity of $\mathcal{R}_k$ at τ_2, we find that

$$c_2 - s_2\frac{1-ik\tau_2}{1+ik\tau_2}e^{2ik\tau_2} = c_1 - s_1\frac{1-ik\tau_2}{1+ik\tau_2}e^{2ik\tau_2} \tag{A.18}$$

whereas continuity of $\mathcal{R}_k'$ at τ_2 implies

$$c_2k^2\tau_2 - s_2k^2\tau_2\,e^{2ik\tau_2} = c_1\left[k^2\tau_2 - \frac{3}{\tau_2}(1+ik\tau_2)\right] - s_1e^{2ik\tau_2}\left[k^2\tau_2 - \frac{3}{\tau_2}(1-ik\tau_2)\right] \tag{A.19}$$

which gives

$$\begin{aligned} c_2 = -\frac{1}{4k^6\tau_1^3\tau_2^3}\big\{&9e^{2ik(\tau_2-\tau_1)}\,(k\tau_1-i)^2\,(k\tau_2+i)^2 \\ &-\left(k^2\tau_1^2\,(2k\tau_1+3i)+3i\right)\left(k^2\tau_2^2\,(2k\tau_2-3i)-3i\right)\big\} \end{aligned} \tag{A.20}$$

and

$$s_2 = \frac{e^{-2ik(\tau_1+\tau_2)}}{4k^6\tau_1^3\tau_2^3}\Big\{3e^{2ik\tau_2}\left(3+k^2\tau_2^2\left(3-2ik\tau_2\right)\right)\left(k\tau_1-i\right)^2 + 3ie^{2ik\tau_1}\left(k^2\tau_1^2\left(2k\tau_1+3i\right)+3i\right)\left(k\tau_2-i\right)^2\Big\}. \tag{A.21}$$

The power spectrum for the curvature perturbation at late times (during slow roll again) is

$$\mathcal{P}_\mathcal{R} = \lim_{\tau\to 0^-}\frac{k^3}{2\pi^2}|\mathcal{R}_k^{(3)}|^2 = \frac{H^2}{8\pi^2 M_{\rm pl}^2\epsilon_3}\left[c_2^*c_2 + s_2^*s_2 - s_2^*c_2 - s_2c_2^*\right], \tag{A.22}$$

where

$$\epsilon_3 := \epsilon_1\, e^{-6N_{\rm USR}} \tag{A.23}$$

is the fixed, final value of ϵ during the second phase of $\eta = 0$, determined by ϵ_1 during the initial phase of $\eta = 0$, and $N_{\rm USR}$ which is the total number of e-folds of USR. The resulting late time spectrum is the black line plotted in Fig. 3.2 with $N_{USR} = 2.3$. We generalise this matching to go from $\eta = 0$ to arbitrary constant $\eta < 0$ and back to $\eta = 0$ in order to plot the other lines in Fig. 3.2. We do this in the same way as just described for $\eta = 0$ to $\eta = -6$ and back to $\eta = 0$, but replace the mode equation $\mathcal{R}_k^{(2)}$ with the appropriate solution from equation (A.5) for each value of ν using

$$\nu^2 = \frac{9}{4} + \frac{3\eta}{2} + \frac{\eta^2}{4} = \left(\frac{3+\eta}{2}\right)^2 \tag{A.24}$$

for constant η. We also note that

$$\epsilon = \epsilon_1\left(\frac{a}{a_1}\right)^\eta = \epsilon_1\left(\frac{\tau_1}{\tau}\right)^\eta \tag{A.25}$$

during the constant η phase. We then do the matching using exactly the same method as before, and find that the late time power spectra plotted in Fig. 3.2 are finally given by:

$$\mathcal{P}_\mathcal{R} = \lim_{\tau\to 0^-}\frac{k^3}{2\pi^2}|\mathcal{R}_k^{(3)}|^2 = \frac{H^2}{8\pi^2 M_{\rm pl}^2\epsilon_3}\left[c_2^*c_2 + s_2^*s_2 - s_2^*c_2 - s_2c_2^*\right], \tag{A.26}$$

with $\epsilon_3 = \epsilon_1 e^{-\eta N_{\eta=const}}$, where $N_{\eta=const}$ is set by the duration of the constant η phase and the coefficients c_2 and s_2 are given in general by:

$$c_2 = \frac{i\pi e^{-ik(\tau_1-\tau_2)}}{8k^{\frac{13}{2}}\sqrt{-\tau_1}\sqrt{-\tau_2}}\left(\begin{array}{l}\sqrt{k^5}(k\tau_2+i)\left(\begin{array}{l}\left(k^4\tau_1 H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_1)+k^3(1+ik\tau_1)H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)\right)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2)\\ +\left(-k^4\tau_1 H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_1)+k^3(-1-ik\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_1)\right)H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_2)\end{array}\right)\\ +k^{\frac{7}{2}}\tau_2\left(\begin{array}{l}\left(k^3(k\tau_1-i)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_1)-ik^4\tau_1 H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_1)\right)H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_2)\\ +\left(k^3(-k\tau_1+i)H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)+ik^4\tau_1 H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_1)\right)H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_2)\end{array}\right)\end{array}\right)$$

(A.27)

$$s_2 = -\frac{\pi e^{-ik(\tau_1+\tau_2)}}{8k^{\frac{7}{2}}\sqrt{-\tau_1}\sqrt{-\tau_2}} \begin{pmatrix} ik^{\frac{7}{2}}\tau_1 H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) + ik^{\frac{7}{2}}\tau_2 H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) \\ + \sqrt{k^5}H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) + ik^{\frac{9}{2}}\tau_1\tau_2 H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) \\ + ik^{\frac{9}{2}}\tau_1\tau_2 H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_2) - \sqrt{k^5}k^2\tau_1\tau_2 H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) \\ + k^{\frac{9}{2}}\tau_1\tau_2 H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_2) + \sqrt{k^3}k^2\tau_1 H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_2) \\ + k^{\frac{7}{2}}\tau_2 H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_1)H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_2) \\ + k^2\tau_2\left(-\sqrt{k^5}\tau_1 H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_1) + \sqrt{k^3}(-1-ik\tau_1)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_1)\right)H^{(1)}_{\frac{\eta+3}{2}}(-k\tau_2) \\ + H^{(1)}_{\frac{\eta+1}{2}}(-k\tau_2)\left(\sqrt{k^5}(k\tau_1-i)(k\tau_2-i)H^{(2)}_{\frac{\eta+1}{2}}(-k\tau_1) - k^{\frac{7}{2}}\tau_1(1+ik\tau_2)H^{(2)}_{\frac{\eta+3}{2}}(-k\tau_1)\right) \end{pmatrix}. \tag{A.28}$$

We now move on to a more realistic matching, wherein one transitions in and out of USR with intermediate phases that interpolate between USR and SR.

A.1 SR $\rightarrow \eta \equiv -2 \rightarrow$ USR $\rightarrow \eta \equiv 2 \rightarrow$ SR matching

By now, we see that the matching calculations involve nothing but sequentially solving a series of linear equations. We now attempt to model two additional intermediate phases to transition into USR, via a phase of $\eta = -2$, and out of USR, via a phase of $\eta = 2$. We match the mode functions given by equation (A.5) for each of the five phases of constant η, at four successive transition times, and the final power spectrum is given by:

$$\mathcal{P}_{\mathcal{R}} = \lim_{\tau\to 0^-}\frac{k^3}{2\pi^2}|\mathcal{R}^{(3)}_k|^2 = \frac{H^2}{8\pi^2 M_{\rm pl}^2\epsilon_5}\left[c_4^*c_4 + s_4^*s_4 - s_4^*c_4 - s_4c_4^*\right], \tag{A.29}$$

with

$$\epsilon_5 = \epsilon_1[(a_3/a_2)^3(a_3/a_4)(a_2/a_1)]^{-2} = \epsilon_1[(\tau_2/\tau_3)^3(\tau_4/\tau_3)(\tau_1/\tau_2)]^{-2}, \tag{A.30}$$

and the coefficients c_4 and s_4 given by:

$$c_4 = -\frac{ie^{ik\tau_4}}{16k^9\tau_1\tau_2^2\tau_3^4\tau_4^2} \begin{pmatrix} (2k\tau_1-i)(3+2ik\tau_2)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4+2i)e^{-ik(2\tau_2-\tau_4)} \\ + (2k\tau_2(2+ik\tau_2)-3i)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4+2i)e^{-ik(2\tau_1-\tau_4)} \\ + (2k\tau_2+3i)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4(3+2ik\tau_4)-2i)e^{-ik(2\tau_1-2\tau_2+\tau_4)} \\ + (1+2ik\tau_1)(-3+2k\tau_2(k\tau_2+2i))\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)e^{-ik\tau_4}(-2+k\tau_4(2k\tau_4-3i)) \\ + (2k\tau_2+3i)(9+2k\tau_3(k\tau_3(-7-2ik\tau_3)+9i))(k\tau_4+2i)e^{-ik(2(\tau_1-\tau_2+\tau_3)-\tau_4)} \\ + (2k\tau_1-i)(-3+2k\tau_2(k\tau_2+2i))(9+2k\tau_3(k\tau_3(-7-2ik\tau_3)+9i))(k\tau_4+2i)e^{-ik(2\tau_3-\tau_4)} \\ + (2k\tau_1-i)(2k\tau_2-3i)(3-2ik\tau_3)(-3+2k\tau_3(k\tau_3+2i))(-2+k\tau_4(2k\tau_4-3i))e^{-ik(2\tau_2-2\tau_3+\tau_4)} \\ + (-3+2k\tau_2(k\tau_2-2i))(3-2ik\tau_3)(-3+2k\tau_3(k\tau_3+2i))(-2+k\tau_4(2k\tau_4-3i))e^{-ik(2\tau_1-2\tau_3+\tau_4)} \end{pmatrix} \tag{A.31}$$

$$s_4 = \frac{e^{-2ik(4\tau_1+\tau_2+\tau_3+\tau_4)}}{16k^9\tau_1\tau_2^2\tau_3^4\tau_4^2}\begin{pmatrix}(3-2ik\tau_2)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4-2i)e^{2ik(3\tau_1+2\tau_2+\tau_3)}\\+(2k\tau_1-i)(2k\tau_2(2-ik\tau_2)+3i)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4-2i)e^{2ik(4\tau_1+\tau_2+\tau_3)}\\+(2k\tau_1-i)(2k\tau_2-3i)\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4(3-2ik\tau_4)+2i)e^{2ik(4\tau_1+\tau_3+\tau_4)}\\+(-3+2k\tau_2(k\tau_2-2i))\left(2k^4\tau_3^4+4k^2\tau_3^2+9\right)(k\tau_4(3-2ik\tau_4)+2i)e^{2ik(3\tau_1+\tau_2+\tau_3+\tau_4)}\\+(-3+2k\tau_2(k\tau_2-2i))(2k\tau_3+3i)(-3+2k\tau_3(k\tau_3+2i))(2+ik\tau_4)e^{2ik(3\tau_1+\tau_2+2\tau_3)}\\+(2k\tau_1-i)(3+2ik\tau_2)(2k\tau_3+3i)(-3+2k\tau_3(k\tau_3+2i))(k\tau_4-2i)e^{4ik(2\tau_1+\tau_3)}\\+(2k\tau_2+3i)(2k\tau_3-3i)(2k\tau_3(2+ik\tau_3)-3i)(-2+k\tau_4(2k\tau_4+3i))e^{2ik(3\tau_1+2\tau_2+\tau_4)}\\-(2k\tau_1-i)(3-2k\tau_2(k\tau_2+2i))(3+2ik\tau_3)(3-2k\tau_3(k\tau_3-2i))(2-k\tau_4(2k\tau_4+3i))e^{2ik(4\tau_1+\tau_2+\tau_4)}\end{pmatrix} \tag{A.32}$$

The late-time power spectrum is shown in Fig. A.1. The phase of $\eta = +2$ causes a decrease in power for large k, which we have chosen to return to the small-k amplitude of 2×10^{-9} for the red, yellow and green lines in Fig. A.1, rather than the scale-invariant spectrum produced by matching straight back to $\eta = 0$ as in previous sections, and shown by the blue line in Fig. A.1. Also notice that the effect of the $\eta = -2$ phase is only visible if it lasts considerably longer than the phase of $\eta = -6$, otherwise the k^4 growth is dominant on the scales that the $\eta = -2$ phase affects.

A.2 Peak amplitude sensitivity to late times

The amplitude of the peak of the power spectrum depends on how ultra slow roll finishes. How η transitions back to 0 from a phase of $\eta = -6$ can shave off power from the peak. For example, if we set $\tau_1 = \tau_2$ in the matching calculation from section A.1, then we can plot the power spectrum for constant phases of η from 0 → -6 → 2 → 0 so as to focus on the transition out of USR. In Fig. A.2 the power spectrum is plotted for 6 different durations of $\eta = 2$ – all other parameters are kept the same – with the different spectra being normalised such that the large scale amplitude is 2×10^{-9}. There is almost a factor of 2 difference in the peak amplitude between no $\eta = 2$ phase and 1 e-folding of $\eta = 2$ following ultra slow roll. However, the amplitude of the power spectrum is unaffected any further by increasing the duration of the $\eta = 2$ phase beyond 1 e-folding. While this is unlikely to have significance in terms of avoiding power spectrum constraints, it may have a large effect on the predicted number of PBHs produced, since the mass fraction is exponentially sensitive to the amplitude of the power spectrum. Note that this is for a sharp transition in η, and the effect may not be present for a smooth transition. This was investigated for the bispectrum in [238], where it was found that local non-Gaussianity is erased during a smooth exit from ultra slow roll, but that it can survive a sharp transition.

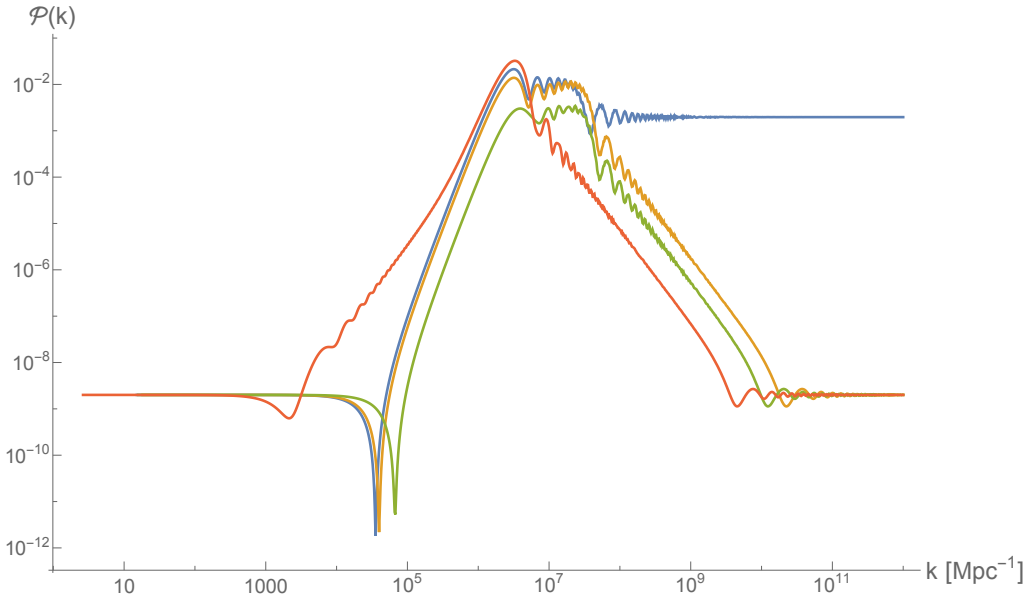

Figure A.1: Four power spectra involving different matchings between constant η phases. The blue line is the same as the blue line plotted in Fig. 3.2, matching from $\eta = 0$ to $\eta = -6$ for 2.3 e-folds and back to $\eta = 0$. The yellow line is a matching from $\eta = 0$ to $\eta = -6$, then to $\eta = 2$ and back to $\eta = 0$. Notice that the peak amplitude decreases slightly when the positive η phase is included - we comment on this further in App. A.2. The green line is a 5-phase matching from $\eta = 0$ to $\eta = -2$, then $\eta = -6$, then $\eta = 2$ and back to $\eta = 0$. The $\eta = -2$ phase does not decrease the slope of the power spectrum because the phase of $\eta = -6$ affects the scales that exit before the onset of the $\eta = -2$ phase, however it does cause the dip to occur at a larger value of k, and for the peak amplitude to be reduced. The red line is the same set-up as for the green line, but with a longer duration of $\eta = -2$ and shorter duration of $\eta = -6$ so that the k^2 growth is visible before the onset of the k^4 spectrum due to USR.

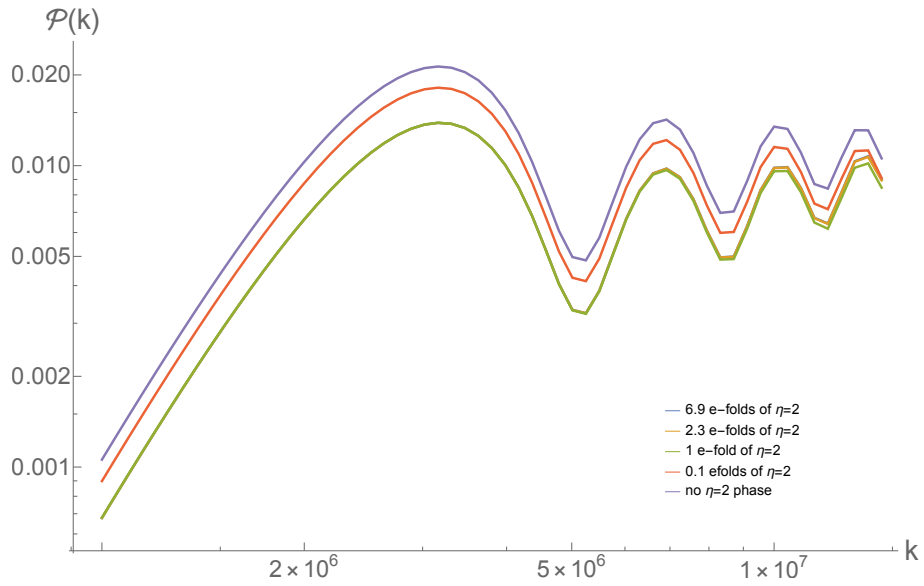

Figure A.2: Analytical power spectra with 4 phases of constant η: 0, -6, 2, 0. The only difference between the lines is the duration of the $\eta = 2$ phase. The longer the phase of $\eta = 2$, the less power at the very peak of the power spectrum, showing that how ultra slow roll ends has an effect on the amplitude of the peak. Notice that the spectra quickly converge to the amplitude for longer than 1 e-fold of $\eta = 2$, the blue and yellow lines are hidden beneath the green line.

Appendix B

The steepest constant η spectrum

If we consider the toy situation of an epoch of inflation defined by a constant, but non-zero $\eta < 0$ without matching to another epoch of inflation, we can arrive at simple bounds on how fast the power spectrum can grow given a constant ν and $\epsilon \ll 1$. In this case, the general solution to the mode function corresponding to the Bunch-Davies vacuum is given by (A.5). If the late-time limit is taken directly without matching to any other phases, then

$$\mathcal{P}_{\mathcal{R}} \propto k^{3-2\nu}, \tag{B.1}$$

and the spectral index is given by

$$n_s - 1 = 3 - |3 + \eta|, \tag{B.2}$$

which gives a scale-invariant spectrum for $\eta = 0, -6$ and the strongest possible positive scaling is $n_s - 1 = 3$ for $\eta = -3$.

The steepest possible growth follows from setting $\eta = -3 \Leftrightarrow \nu = 0$ but because both modes are important in this case, the approximation (B.2) overestimates the actual slope, with the complete late time solution being

$$\mathcal{P}_{\mathcal{R}} \propto k^3 \left(1 + \frac{4}{\pi^2}\left(\gamma + \ln\left(\frac{k}{2k_e}\right)\right)^2\right), \tag{B.3}$$

$$n_s - 1 = 3 + \frac{8\left(\log\left(\frac{k}{2k_e}\right) + \gamma\right)}{\pi^2\left(\frac{4\left(\log\left(\frac{k}{2k_e}\right)+\gamma\right)^2}{\pi^2} + 1\right)}, \tag{B.4}$$

where $\gamma = 0.5772$ and $k_e = (aH)_e$ is the value of k when this period of inflation at the boundary of USR ends (meaning that the decaying and growing modes are both important), assuming that the curvature perturbation freezes out afterwards. The

correction to the k^3 scaling in (B.3) comes from the "decaying" mode which scales as

$$\mathcal{R}_{\rm decaying} \sim \int \frac{dN}{a^3\epsilon} \propto \int dN \sim N = \log(k/k_e). \tag{B.5}$$

The solution (B.3) agrees with [40] in the limit $\ln(k/k_e) \gg 1$. The potential giving rise to this growth rate of the power spectrum in the limit $\epsilon \to 0$, is

$$V = M^4 e^{\frac{3}{8}\frac{\phi^2}{M_{Pl}^2}}. \tag{B.6}$$

which can be calculated by using $\phi' \propto \sqrt{\epsilon} \propto e^{-3N}$ and the equation of motion (3.4). However, (B.3) implies a weaker bound than the k^4 steepest growth index for single-field inflation shown via a more realistic matching calculation in App. A. A complementary perspective is obtained by reconsidering the what a power spectrum with a constant growth index implies in position space. In order to do this, we consider the following form for the power spectrum

$$\mathcal{P}_{\mathcal{R}} \propto k^n e^{-\alpha k}, \tag{B.7}$$

which needs to be regulated for certain values of the index by a non-zero α which we take to zero at the end of the calculation. We recall that the position space two-point function of the curvature perturbation at late times is given by Fourier transform

$$\int \frac{d^3k}{(2\pi)^3} e^{i\vec{k}\cdot\vec{x}} |\mathcal{R}_k(0)|^2 = \langle \mathcal{R}(\vec{x},0)\mathcal{R}(0,0)\rangle \tag{B.8}$$

where

$$\mathcal{P}_{\mathcal{R}} = \lim_{\tau\to 0^-} \frac{k^3}{2\pi^2} |\mathcal{R}_k(\tau)|^2, \tag{B.9}$$

and therefore

$$\langle \mathcal{R}(\vec{x},0)\mathcal{R}(0,0)\rangle \propto \lim_{\alpha\to 0} \int \frac{d^3k}{4\pi} e^{i\vec{k}\cdot\vec{x}}\, k^{n-3} e^{-\alpha k}. \tag{B.10}$$

For $n = 0$ we recover the usual logarithmically divergent position space correlation function (an artefact of us working in the strict dS limit). For $n > 3$ we find

$$\langle \mathcal{R}(\vec{x},0)\mathcal{R}(0,0)\rangle \propto \lim_{\alpha\to 0} \int \frac{d^3k}{4\pi} e^{i\vec{k}\cdot\vec{x}}\, k^{n-3} e^{-\alpha k} \propto \frac{1}{|x|^n}. \tag{B.11}$$

Therefore, asking why a power spectrum with a constant index n can't have an index greater than $n = 4$ is the same as asking why the position space two-point function for the curvature perturbation can't diverge in the coincident limit faster than the fourth power of the distance between the two operators. The reason for this boils down to

dimensional analysis. In a mass dependent regularisation scheme (i.e. regulating divergences with a hard cut-off Λ), the more divergent a correlation function is in position space, the greater the power of divergence in momentum space. Two-point functions that diverge as the inverse square of the distance require counterterms proportional to Λ^2. Since our theory has no other UV mass scale, one cannot have a dependence on the r.h.s. of (B.11) where n is greater than 4, since this would require a counterterm that goes as $\Lambda^{n>4}$, which is not possible in four dimensions. However, this does not completely account for the steepest growth shown in App. A, since for any finite α, the spectrum cuts off and the corresponding divergence is automatically regulated, invalidating the above argument. Although causality and analyticity arguments have been invoked in different contexts to argue for a particular bounds on the growth index of various cosmological perturbations[1], none of these appear to apply to our present context. The physical origin of the steepest growth index over a *finite* range of modes that we've uncovered is still something we're investigating.

[1] In the context of density perturbations produced from a causal collapse process in a non-inflationary context, Traschen and Abbott have derived a *minimum* growth index of k^4 [310]. For primordial magnetic fields, Durrer and Caprini have shown that the two-point function must scale at least as k^2 at large scales [311].

Appendix C

On the background potential

In the first part of this appendix, we show how one can reconstruct a potential having specified an arbitrary time-dependence for ϵ. Note that this is a much simpler problem than reconstructing the inflaton potential (more generally, action) from CMB data, a process that is necessarily hamstrung by a variety of degeneracies [312–314]. Our goal is simply to show that one can in principle design a potential assuming a minimally coupled scalar field with a canonical kinetic term to reproduce an arbitrary time-dependent profile for ϵ. In the second part of this appendix, we show how one cannot engineer an arbitrarily abrupt end to inflation in terms of e-folds without introducing additional hierarchies that will be radiatively unstable.

C.1 Reconstructing V from ϵ

We begin with the equation of motion for a minimally coupled scalar ϕ, switching to e-folding number $\mathcal{N}$ as the time variable

$$H^2 \frac{d^2\phi}{d\mathcal{N}^2} + \left(3H^2 + H\frac{dH}{d\mathcal{N}}\right)\frac{d\phi}{d\mathcal{N}} + \frac{\partial V}{\partial \phi} = 0. \tag{C.1}$$

Given that $H\frac{dH}{d\mathcal{N}} = \dot{H}$, one can use the Friedmann equations $3H^2 = \rho$, $\dot{H} = -(\rho + p)/2M_{\rm pl}^2$ to obtain

$$3H^2 + H\frac{dH}{d\mathcal{N}} = 3H^2 + \dot{H} = \frac{\rho - p}{2M_{\rm pl}^2} = \frac{V}{M_{\rm pl}^2}. \tag{C.2}$$

Furthermore the Einstein constraint equation becomes

$$H^2\,(3 - \epsilon) = \frac{V}{M_{\rm pl}^2}. \tag{C.3}$$

Inserting these relations into (C.1) results in the equation of motion

$$\frac{d^2\phi}{d\mathcal{N}^2} + \left[\frac{d\phi}{d\mathcal{N}} + \frac{M_{\rm pl}^2}{V}\frac{\partial V}{\partial \phi}\right](3 - \epsilon) = 0 \tag{C.4}$$

or

$$\frac{d\epsilon}{d\mathcal{N}} = -(3-\epsilon)\left[2\epsilon + \frac{d\phi}{d\mathcal{N}}\frac{\partial_\phi V}{V}\right] \tag{C.5}$$

where we have used $\epsilon = \frac{(d\phi/d\mathcal{N})^2}{2M_{\rm pl}^2}$. So far, the above relations are exact. We now presume that $\epsilon \ll 3$ so that the above can be approximated as[1]

$$\frac{d\epsilon}{d\mathcal{N}} = -6\epsilon + \frac{d\log V^{-3}}{d\mathcal{N}}. \tag{C.6}$$

Using the definition of ϵ and (C.6) we find

$$\phi(\mathcal{N}) = \phi_* \pm M_{\rm pl}\int_{\mathcal{N}_*}^{\mathcal{N}} d\mathcal{N}'\sqrt{2\epsilon(\mathcal{N}')}, \tag{C.7}$$

and

$$V(\mathcal{N}) = V(\mathcal{N}_*)\exp\left[-\frac{1}{3}\int_{\mathcal{N}_*}^{\mathcal{N}} d\mathcal{N}'\left(\frac{d\epsilon}{d\mathcal{N}'} + 6\epsilon\right)\right], \tag{C.8}$$

giving us ϕ and V as functions of $\mathcal{N}$ determined entirely by the evolution of ϵ that we take as an input. It remains to figure out what V is as a function of ϕ. To do this, we observe that if

$$V(\mathcal{N}) = \sum_{n=0} c_n f_n[\phi(\mathcal{N})] \tag{C.9}$$

where the f_n are some complete basis of functions[2], and if we know $V(\mathcal{N}_i)$ and $\phi(\mathcal{N}_i)$ for $0 \leq i \leq m$ discrete values, then if we demand than $V(\phi)$ truncate at some finite order m, we have a system of $m+1$ linear equations in $m+1$ unknowns which will be possible to invert given the presumption of monotonicity of ϕ and linear independence of the basis functions, allowing us to calculate the coefficients c_i for $0 \leq i \leq m$, thus reconstructing an approximation to the potential to order m. For a limited enough field excursion it suffices to truncate to some small finite order e.g. at $m = 6$ for a monomial basis: the typical order to which we need to know the potential in order to have a handle on the η problem (see for instance, the treatment in [208]).

However, for simple enough time dependence for ϵ one can explicitly perform a direct reconstruction. We first match a phase of constant ϵ slow roll to a phase of USR. First, note that using the definition $d\epsilon/d\mathcal{N} = \epsilon\eta$, (C.8) can be recast as

$$V(\mathcal{N}) = V_*\exp\left[-\frac{1}{3}\int_{\mathcal{N}_*}^{\mathcal{N}} d\mathcal{N}'\epsilon\,(\eta+6)\right]. \tag{C.10}$$

[1] Note that one can straightforwardly generalise the above derivation to the case of multi-field inflation, where the final equation (C.6) would also result.

[2] e.g. $f_n = \phi^n$ or $f_n = e^{n\lambda\phi}$ for some fixed λ etc. In general, the convergence of the reconstructed potential to the true potential will depend greatly on choice of basis functions adopted, and the range in field space one wants the approximation to be valid.

Therefore, it is clear that during USR, $V(\mathcal{N})$ remains constant as inflation progresses. Furthermore, during constant ϵ slow roll, $\epsilon(\mathcal{N}) \equiv \epsilon_0$ and $\eta = 0$, so that during this phase

$$V(\mathcal{N}) = V_* e^{-2\epsilon_0(\mathcal{N}-\mathcal{N}_*)}. \tag{C.11}$$

Picking the + branch of the solution (C.7)

$$\phi(\mathcal{N}) - \phi_* = M_{\rm pl}\sqrt{2\epsilon_0}(\mathcal{N} - \mathcal{N}_*) \tag{C.12}$$

we find that we can straightforwardly invert ϕ for $\mathcal{N}$, resulting in the potential

$$V(\phi) = V_* e^{-\frac{\sqrt{2\epsilon_0}}{M_{\rm pl}}(\phi-\phi_*)}, \tag{C.13}$$

which is consistent with the fact that the only constant ϵ attractors are given by exponential potentials. Next, we note that during USR, the argument of (C.10) vanishes identically, so that the potential during this phase has a constant value set by the value at the end of the constant SR epoch –

$$V(\mathcal{N}) = V_* e^{-2\epsilon_0(\mathcal{N}_1-\mathcal{N}_*)} = \text{const.} \quad \mathcal{N} > \mathcal{N}_1. \tag{C.14}$$

Similarly, given that $\epsilon(\mathcal{N}) = \epsilon_0 e^{-6(\mathcal{N}-\mathcal{N}_1)}$ during USR, we find from (C.7) that

$$\frac{3(\phi-\phi_1)}{\sqrt{2\epsilon_0}M_{\rm pl}} = 1 - e^{-3(\mathcal{N}-\mathcal{N}_1)} \tag{C.15}$$

with ϕ_1 given by (C.12) evaluated at $\mathcal{N}_1$. The only way some polynomial function of the above can result in a constant is if it were itself a constant. Hence the reconstructed potential that transitions from slow roll to USR is a piecewise potential that glues an exponential potential to a constant. This is not particularly physical, so we can try to suitably smooth the transition from SR to USR.

We now reproduce a potential that can mimic the matching calculation done in the previous appendix. Namely, from $\eta = 0$ slow roll to $\eta = -2 \to \eta = -6 \to \eta = +2$ back to $\eta = 0$ slow roll. When $\eta = 0 \to \eta = -2$ at $\mathcal{N} = \mathcal{N}_1$, we can repeat the steps above for $\mathcal{N} > \mathcal{N}_1$ to find

$$V(\mathcal{N}) = V(\mathcal{N}_1) e^{\frac{2\epsilon_0}{3}\left[e^{-2(\mathcal{N}-\mathcal{N}_1)}-1\right]} \tag{C.16}$$

and similarly for the field profile,

$$\frac{\phi-\phi_1}{M_{\rm pl}\sqrt{2\epsilon_0}} = 1 - e^{-(\mathcal{N}-\mathcal{N}_1)}. \tag{C.17}$$

Substituting the above into the exponent of (C.16) results in

$$V(\phi) = V(\phi_1) e^{\frac{(\phi-\phi_1)^2}{3M_{\rm pl}^2} - \sqrt{2\epsilon_0}\frac{2(\phi-\phi_1)}{3M_{\rm pl}}}. \tag{C.18}$$

Note that from (C.17) the field can only asymptote to $\phi - \phi_1 = \sqrt{2\epsilon_0} M_{\rm pl}$, in which case the potential goes to zero smoothly. From the previous discussion, we see that to match to $\eta = -6$ is to splice this potential to a constant piece at $\mathcal{N} = \mathcal{N}_2$. To subsequently match from this phase to $\eta = +2$ at $\mathcal{N}_3$ results in (for $\mathcal{N} > \mathcal{N}_3$)

$$V(\mathcal{N}) = V(\mathcal{N}_3) e^{-\frac{4\epsilon_2}{3}\left[e^{2(\mathcal{N}-\mathcal{N}_3)}-1\right]} \tag{C.19}$$

during which time the field evolves as

$$\phi - \phi_3 = M_{\rm pl}\sqrt{2\epsilon_2}\left[e^{\mathcal{N}-\mathcal{N}_3} - 1\right] \tag{C.20}$$

so that the potential this corresponds to is given by

$$V(\phi) = V(\phi_3) e^{\frac{-2(\phi-\phi_3)^2}{3M_{\rm pl}^2} - \sqrt{2\epsilon_2}\frac{4(\phi-\phi_3)}{3M_{\rm pl}}} \tag{C.21}$$

with

$$\epsilon_2 = \epsilon_0 e^{-2(\mathcal{N}_2-\mathcal{N}_1)} e^{-6(\mathcal{N}_3-\mathcal{N}_2)}. \tag{C.22}$$

Finally, one would like to match to a slow roll phase again, where

$$V(\phi) = V(\phi_4) e^{-\frac{\sqrt{2\epsilon_3}}{M_{\rm pl}}(\phi-\phi_4)} \tag{C.23}$$

with

$$\epsilon_3 = \epsilon_0 e^{-2(\mathcal{N}_2-\mathcal{N}_1)} e^{-6(\mathcal{N}_3-\mathcal{N}_2)} e^{2(\mathcal{N}_4-\mathcal{N}_3)} \tag{C.24}$$

Therefore we summarize that the piecewise continuous potential that reproduces the matching $\eta = 0 \to \eta = -2 \to \eta = -6 \to \eta = +2 \to \eta = 0$ is given by

$$\begin{aligned}
V_1(\phi) &= V_* e^{-\frac{\sqrt{2\epsilon_0}}{M_{\rm pl}}(\phi-\phi_*)} && \phi < \phi_1,\ \eta = 0 \\
V_2(\phi) &= V_1(\phi_1) e^{\frac{(\phi-\phi_1)^2}{3M_{\rm pl}^2} - \sqrt{2\epsilon_0}\frac{2(\phi-\phi_1)}{3M_{\rm pl}}} && \phi_1 < \phi < \phi_2,\ \eta = -2 \\
V_3(\phi) &= V_2(\phi_2) = \text{constant} && \phi_2 < \phi < \phi_3,\ \eta = -6 \\
V_4(\phi) &= V_3 e^{\frac{-2(\phi-\phi_3)^2}{3M_{\rm pl}^2} - \sqrt{2\epsilon_2}\frac{4(\phi-\phi_3)}{3M_{\rm pl}}} && \phi_3 < \phi < \phi_4,\ \eta = +2 \\
V_5(\phi) &= V_4(\phi_4) e^{-\frac{\sqrt{2\epsilon_3}}{M_{\rm pl}}(\phi-\phi_4)} && \phi_4 < \phi,\ \eta = 0
\end{aligned} \tag{C.25}$$

with ϵ_2 and ϵ_3 given by (C.22) and (C.24), and where the fixed field intervals in terms of the number of e-folds of the different phases as

$$\begin{aligned}
\phi_2 - \phi_1 &= M_{\rm pl}\sqrt{2\epsilon_0}\left[1 - e^{-(\mathcal{N}_2-\mathcal{N}_1)}\right] \\
\phi_3 - \phi_2 &= \frac{M_{\rm pl}}{3}\sqrt{2\epsilon_0} e^{-(\mathcal{N}_2-\mathcal{N}_1)}\left[1 - e^{-3(\mathcal{N}_3-\mathcal{N}_2)}\right] \\
\phi_4 - \phi_3 &= M_{\rm pl}\sqrt{2\epsilon_0} e^{-(\mathcal{N}_2-\mathcal{N}_1)} e^{-3(\mathcal{N}_3-\mathcal{N}_2)}\left[e^{(\mathcal{N}_4-\mathcal{N}_3)} - 1\right].
\end{aligned} \tag{C.26}$$

We plot the reconstructed potential below for specific values of the $\mathcal{N}_i$:

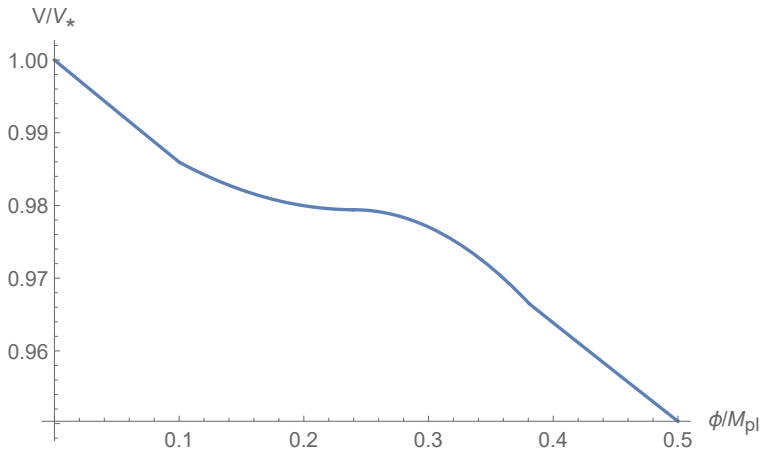

Figure C.1: The reconstructed potential (C.25) for $\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3, \mathcal{N}_4 = 10, 14, 18, 34$ respectively, with $\phi_* = 0$ and $\epsilon_0 = 0.01$. Note that the field range over which USR occurs ($\phi_3 - \phi_2 \simeq 0.0009 M_{\rm pl}$) is parametrically much smaller than the phases where $\eta = \pm 2$, so as to effectively appear as an inflection point in the above plot.

C.2 The quickest possible end to inflation

Having demonstrated a steepest possible growth for the primordial power spectrum, one might wonder about the complementary question – how quickly can it fall off? An accurate estimate for this can be inferred from rephrasing the question as how quickly inflation can end, or transition to another phase of inflation. To understand this, we first recall (C.4), but now generalise to multi-field inflation

$$\frac{d^2\phi^a}{d\mathcal{N}^2} + \left[\frac{d\phi^a}{d\mathcal{N}} + \frac{M_{\rm pl}^2}{V}\frac{\partial V}{\partial \phi_a}\right](3-\epsilon) = 0 \tag{C.27}$$

where ϕ^a denotes coordinates in some general field space. For simplicity, we assume a flat field space metric (and so accord no significance to raised or lowered indices) although this can be straightforwardly generalised. The slow-roll parameter ϵ is now defined as

$$\epsilon = \frac{1}{2M_{\rm pl}^2}\frac{d\phi^a}{d\mathcal{N}}\frac{d\phi_a}{d\mathcal{N}}. \tag{C.28}$$

Multiplying (C.27) by $d\phi_a/d\mathcal{N}$ results in the analog of (C.5)

$$\frac{d\epsilon}{d\mathcal{N}} = -(3-\epsilon)\left[2\epsilon + \frac{d\phi^a}{d\mathcal{N}}\frac{\partial_a V}{V}\right]. \tag{C.29}$$

We now consider the situation where over some interval, ϵ increases monotonically from some initial $\epsilon_0 \ll 1$ to $\epsilon_f = 1$ over an interval $\Delta\mathcal{N}_{\rm end}$. Application of the mean value

theorem of calculus[3] then implies that

$$\left|\frac{d\epsilon}{d\mathcal{N}}\right|_{\rm int} \gtrsim \frac{1-\epsilon_0}{\Delta\mathcal{N}_{\rm end}} \sim \frac{1}{\Delta\mathcal{N}_{\rm end}} \tag{C.30}$$

at some intermediate $\mathcal{N}_{\rm int}$. Inserting the rhs of (C.29) into the above, assuming $\Delta\mathcal{N}_{\rm end} \ll 1$ and applying the triangle inequality results, after some manipulation, in the lower bound

$$M_{\rm pl}\left|\frac{\nabla_T V}{V}\right|_{\rm max} \gtrsim \frac{1}{3\sqrt{2}\Delta\mathcal{N}_{\rm end}} \tag{C.31}$$

where $\nabla_T V$ is the tangential derivative of the potential with respect to the trajectory, defined as $\nabla_T V := T^a \partial_a V$ and $T^a := \frac{d\phi^a}{d\mathcal{N}} \Big/ \left(\frac{d\phi^b}{d\mathcal{N}}\frac{d\phi_b}{d\mathcal{N}}\right)^{1/2}$. In the single-field case, it reduces to the more familiar expression

$$M_{\rm pl}\left|\frac{V_{,\phi}}{V}\right|_{\rm max} \gtrsim \frac{1}{3\sqrt{2}\Delta\mathcal{N}_{\rm end}}. \tag{C.32}$$

Therefore, if we would like inflation to end in $\Delta\mathcal{N}_{\rm end} \ll 1$ e-folds or less, we necessarily require the gradient of the potential along the trajectory as inflation ends to be bounded from below according to (C.31). Although classically we are entitled to make the transition out of inflation as sharp as we desire, one cannot make it arbitrarily sharp without introducing additional hierarchies that will be unstable under quantum corrections, since these corrections spoil the flatness of the potential away from the transition, in effect ending inflation earlier and restoring the smoothness of the transition. Nevertheless, from (C.31) we see that a transition that lasts an order unity fraction of an e-fold can easily be accommodated without introducing additional hierarchies, and for the purposes of our discussion, justifies any approximation that cuts off the primordial power spectrum at some fixed comoving scale.

For completeness, we illustrate the considerations above with a concrete example. Consider the following prototype potential for a rapid exit from inflation

$$V(\phi) = \frac{V_*}{2} e^{-\gamma\phi/M_{\rm pl}} \left(1 - \tanh\left[\mu(\phi-\phi_*)/M_{\rm pl}\right]\right). \tag{C.33}$$

When $\gamma \ll 1$, one has power law inflation in the region $\phi \ll \phi_*$. At $\phi = \phi_*$, there is a transition (that can be made arbitrarily abrupt as the dimensionless parameter $\mu \to \infty$).

[3] Recalling that if f and f' are continuous functions on the interval $[a,b]$, then there exists some point $c \in [a,b]$ such that $f'(c) = \frac{f(b)-f(a)}{b-a}$. Since f' is also continuous, $f'(c)$ must lie between the minimum and maximum of f' in the interval $[a,b]$. That is

$$\min_{a\le x\le b} f'(x) \le \frac{f(b)-f(a)}{b-a} = f'(c) \le \max_{a\le x\le b} f'(x).$$

Requiring the transition to last less than 1/100th of an e-fold requires for example μ to be at least of order 10^2 through (C.31), which would imply that the hyperbolic tangent is an operator expansion in odd powers of effective operators with very large Wilson coefficients:

$$\mathcal{L} \supset \frac{\mu^n \phi^n}{M_{\rm pl}^n}; \quad \mu \sim 10^2, \ \Delta\mathcal{N}_{\rm end} \sim 10^{-2}. \tag{C.34}$$

Calculating loop corrections to the potential (C.33) expanded around ϕ_* for $\mu \sim 10^2$ would result in a deformation of the inflationary part of the potential. If one were to try and approximate it close enough to ϕ_* as an exponential, one would find an effectively renormalised γ that is no longer $\ll 1$. On the other hand, requiring $\Delta\mathcal{N}_{\rm int} \sim 10^{-1}$ is possible for values of $\mu \sim 1$, resulting in a renormalisation group improved potential where the hierarchy $\gamma \ll 1$ is preserved. We stress however that the bound (C.31) is completely general and can be applied to multi-field inflation as well.