A University of Sussex PhD thesis

Available online via Sussex Research Online:

http://sro.sussex.ac.uk/

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

A history of the universe in primordial black holes

Andrew David Gow

Submitted for the degree of Doctor of Philosophy
University of Sussex
August 2021

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.
Elements of this thesis have contributed (in part or in their entirety) to the following publications:

A. D. Gow, C. T. Byrnes, P. S. Cole, S. Young, *The power spectrum on small scales: Robust constraints and comparing PBH methodologies*, **JCAP 2021(02) (2021) 002**, arXiv:2008.03289

The idea for this paper was conceived by me and Christian Byrnes, before being modified in discussions with Sam Young and Pippa Cole. The calculations were carried out by me with suggestions from the other authors. The non-PBH constraints were calculated by Pippa Cole. The paper was written predominantly by me, with comments and additions from the other authors. Sections 2.4.2 and appendix B were written by Pippa Cole, and appendix C was written by Sam Young.

A. D. Gow, C. T. Byrnes, A. Hall, *An accurate model for the primordial black hole mass distribution from a peak in the power spectrum*, arXiv:2009.03204

The idea for this paper was suggested by me. The calculations were carried out by me with suggestions from the other authors. The paper was written by me, with comments and additions from the other authors.

A. D. Gow, C. T. Byrnes, A. Hall, J. A. Peacock, *Primordial black hole merger rates: distributions for multiple LIGO observables*, **JCAP 2020(01) (2020) 031**, arXiv:1911.12685

The idea for this paper was conceived by Christian Byrnes and John Peacock. The calculations were carried out by me with suggestions from the other authors. The paper was written by me with comments and additions from the other authors.

A. Hall, **A. D. Gow**, C. T. Byrnes, *Bayesian analysis of LIGO-Virgo mergers: Primordial vs. astrophysical black hole populations*, **PRD 102(12) (2020) 123524**, arXiv:2008.13704

Although not included in this thesis, this paper was contributed to as part of this PhD and builds on Gow et al. (arXiv:1911.12685), providing additional insight into the discussion of the PBH merger rate and the comparison to astrophysical black holes. The calculations and paper were completed primarily by Alex Hall, with comments, suggestions, and additions from me and Christian Byrnes.

Signature:

Andrew D. Gow

UNIVERSITY OF SUSSEX

Andrew D. Gow, Doctor of Philosophy

A history of the universe in primordial black holes

Primordial black holes (PBHs) are a unique type of astrophysical object. Formed in the early universe but persisting to the present day, they offer an incredible probe for studying the physics of both the early and late universe. Additionally, they provide a natural dark matter candidate without requiring physics beyond the standard model, although there are many constraints on the fraction they can contribute across a broad range of PBH masses.

In this thesis, PBHs are studied at both early and late times. Their formation from large overdensities is examined and choices in the calculation of their abundance are considered. These choices are shown to have a limited effect on the PBH abundance and mass distribution, although they will become important in the future. Additionally, robust constraints on the primordial power spectrum are calculated for present and future detections.

The detailed shape of the PBH mass distribution is important for constraining the population. Fitting late universe observables with the mass distribution calculated from the power spectrum would be computationally expensive, so it is necessary to use simple parametrisations that capture the underlying shape. A number of these parametrisations are tested against the numerical calculation, and it is found that two of these consistently outperform the often-assumed lognormal, at the cost of an extra fitting parameter.

Finally, the possibility that PBHs could explain the LIGO merger events is studied by applying a detailed model of the PBH merger rate and the detection probability of the LIGO instrument. Distributions of the merger rate are produced for a number of observables, indicating that the mass ratio could be an important quantity for distinguishing between astrophysical and primordial black hole mergers. A simple statistical test is carried out to demonstrate that PBHs can explain the totality of the LIGO events, with an appropriate abundance and mass distribution.

Acknowledgements

My greatest thanks must go to my supervisor Chris. You have always treated me as an equal, or even deferred to my expertise in certain areas. You have given me plenty of great advice, and were always happy to extend our quick update meetings (sometimes considerably) to work through a problem. You have made my PhD experience entirely enjoyable, and if I choose to stay in academia, you will have played a major role.

To my collaborators Alex, John, Pippa, Sam, and Subodh: thank you for useful advice and friendly chats, and for making the writing of papers a much less painful task than it doubtless would have been otherwise.

To my office mates Jussi, Kareem, and Luke: thanks for providing banter and jokes, as well as useful research advice. To Jussi: hopefully you will not find too many egregious acronyms within these pages. To everyone in the astronomy group, past and present: you make the group so vibrant and fun, and I look forward to returning sometime soon for a Friday bender.

To Alex, Ewan, Sam, and Poppy; forever the freshers, despite having all graduated: thank you for giving me somewhere to go when work was stressful, and for helping me understand that things I find trivial may not be so to an undergrad (an essential for teaching). To Poppy: I eagerly await your PhD thesis, which I assume will contain only TikZ figures!

To Ryan: thank you for being a part of my entire time at Sussex, and for your many puns over the years. Hopefully I never give you cause to become an angry coil. To Alex, Adam, and Sam: it's been amazing getting back in touch over the last few months.

To Mitch and Eleanor: thanks for games, laughs, and many madcap schemes (we'll solve Navier–Stokes one day, I promise!), always rounded off with a McFlurry run.

To everyone at Sussex Harmony: thank you for introducing me to an amazing tradition, and for being such a wonderful (and strange!) collection of people; however, I fear modern Christmas carols will be ruined for me forever.

To the Star Trek "society": thanks for all the entertainment over many years, and I promise that one day we'll be able to discuss the end of DS9. Until then: Live long and experience bIj.

To the members of the Doctor Who society across the years: thank you for accepting me as member, Lord President, and friend during my time at Sussex. It's the end, but the moment has been prepared for.

To my family: thanks for supporting me in every aspect of my life, and trusting that I've made the right decision (even if I only tell you after making it). Special thanks must go to Mum and Grandma, who willingly subjected themselves to a presentation about my research, despite having little interest in physics.

Finally, to Annie: you have provided me with love, support, stability, encouragement, and so much more for so long, and I can't imagine my life without you. Diolch am bopeth, fy nghath bach.

You came, and look'd and loved the view
Long-known and loved by me,
Green Sussex fading into blue
With one gray glimpse of sea

Alfred, Lord Tennyson

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis concerns the study of primordial black holes (PBHs) as a dark matter candidate and a probe of the early universe on small scales, as well as a potential explanation for the gravitational waves seen by the LIGO–Virgo collaboration. Chapter 1 presents a broad overview of the relevant background for this thesis. Section 1.1 considers modern cosmology, the need for inflation, and the method of characterising early universe physics with the primordial power spectrum. A discussion of PBH formation and properties is given in section 1.2, contrasted with equivalent information about astrophysical black holes in section 1.3. Finally, see section 1.4 for an introduction to gravitational wave sources and their detection.

Chapters 2–4 contain the research carried out for this thesis. In chapter 2, we consider the formation of PBHs from inflation, including a detailed calculation of the PBH mass distribution and robust constraints on the power spectrum. In chapter 3, we extend our study of the PBH mass distribution, considering its shape and demonstrating the need for accurate models for use in Markov Chain Monte Carlo (MCMC) analyses. Finally, chapter 4 considers the possibility of PBH mergers providing the signals detected by the LIGO–Virgo collaboration, demonstrating that PBHs can explain all 10 events in the O1O2 catalogue, and that the mass ratio may be a significant observable in distinguishing between primordial and astrophysical mergers. Throughout this thesis we set $c = 1$.

1.1 Cosmology, inflation, and the power spectrum

1.1.1 Hot big bang cosmology

Any work in modern cosmology must begin with Einstein's general theory of relativity [1]. The calculation of gravitational interactions in this theory ultimately begins with the spacetime metric $g_{\mu\nu}$. The application of this theory to the entire universe relies on the cosmological principle: the universe is homogeneous and isotropic on large scales. Homogeneity means that the universe is the same everywhere, and isotropy means that the universe looks the same in every direction. These postulates have been shown to hold by simulation and observation [2,3], although there is some uncertainty about the scale below which homogeneity breaks down [4–6]. In terms of the metric, the cosmological principle forces the spatial shape to be one of three choices: flat, spherical, or hyperbolic. Additionally, it ensures that the temporal part of the metric can have no coordinate dependence, and that the spatial part can only grow or shrink with time. This leads to the Friedmann-Lemaître-Robertson-Walker (FLRW) metric, whose line element $\mathrm{d}s^2$ is

$$\mathrm{d}s^2 = \mathrm{d}t^2 - a^2(t)\left[\frac{1}{1-Kr^2}\mathrm{d}r^2 + r^2(\mathrm{d}\theta^2 + \sin^2\mathrm{d}\phi^2)\right], \tag{1.1}$$

where t, r, θ, and ϕ are the temporal and spatial coordinates, K describes the spatial curvature and $a(t)$ is the scale factor. Application of the Einstein equations using this metric leads to the conclusion that the universe is expanding. A key result of this analysis is that the early universe was very hot, before cooling as it expanded. This led to the prediction of the cosmic microwave background (CMB), whose discovery cemented the "big bang" theory as the standard cosmological model.

1.1.2 The cosmic microwave background

The cosmic microwave background (CMB) is one of the oldest and most studied of all cosmological signals. Predicted in 1948 by Gamow, Alpher, and Herman [7, 8], its discovery in 1965 by Penzias and Wilson[1] [10] was a turning point in the path to modern cosmology, and it provides information on the content and history of our universe to the present day. The CMB comes from the epoch of recombination, the point at which the universe cooled enough to allow neutral hydrogen to form. Before this point, continual scattering with charged particles made the universe opaque to photons. After the formation of hydrogen atoms, the universe became transparent and the photons could decouple from the other particles and travel freely through space. These photons constitute the CMB, and represent the furthest back in time the universe can be studied using electromagnetic radiation.

Before recombination, the photons were in thermal equilibrium with the other particles, which gives the CMB a blackbody spectrum. Since its emission in the early universe, the CMB photons have been redshifted, which corresponds to a reduction of the blackbody temperature. The CMB temperature today is 2.735 ± 0.06 K [11], which provides information on when recombination happened. However, there could be deviations to the blackbody spectrum caused by processes that affect the thermal equilibrium before recombination. These are known as spectral distortions, and are relevant when considering the formation of very heavy primordial black holes (see chapter 2). Additionally, the CMB shows very small anisotropies that provide information about the early universe, as well as the components of the universe such as dark matter and dark energy [12].

While the prediction and discovery of the CMB provided clear evidence for the big bang theory, it also presented a problem in the uniformity of the blackbody across the sky. This suggests that the whole universe was in thermal equilibrium at the epoch of recombination. However, application of the traditional big bang theory

[1] The CMB had previously been detected in 1941 by Andrew McKellar who measured the coldest temperature in interstellar space as 2.3 K using CN doublet excitation lines, although this was not identified as the CMB temperature until much later [9].

implies that at this time only small patches of the universe would be in thermal equilibrium. This was termed the horizon problem, and demonstrated that the big bang model must be modified to create a period when the whole universe was in thermal equilibrium.

1.1.3 Cosmic inflation

In the 15 years following the discovery of the CMB, two additional problems were identified with the big bang theory: the flatness problem and the monopole problem. The flatness problem arises from the observation that the density of the universe is very close to the critical density required to make the universe spatially flat [13]. According to big bang cosmology, the expansion of the universe causes an increase in spatial curvature, implying that to achieve the flatness observed today, the universe must have begun in a highly fine-tuned state of flatness. The monopole problem is based on the fact that Grand Unified Theories (GUTs) are expected to produce a significant density of magnetic monopoles. However, no such monopoles have been observed, necessitating a method to reduce their density throughout the universe.

To solve all of these problems, a period of exponential expansion in the very early universe was posited by a number of people in the early 1980s, termed inflation by Alan Guth [14–18]. Problems in this theory were alleviated independently by a number of groups, in what became known as new inflation [19–21]. This exponential period pushes different parts of space out of causal contact, drives the spatial curvature to flat, and reduces the density of monopoles by rapidly increasing the volume of space they inhabit. Therefore, the inflationary theory elegantly solves the horizon, flatness, and monopole problems.

The most commonly used inflationary scenario introduces a scalar field ϕ, known as the inflaton, minimally coupled to gravity through the Einstein-Hilbert action,

$$S = \int \mathrm{d}^4x \, \sqrt{-g} \left[\frac{M_P^2}{2} R + \mathscr{L}_\phi \right], \tag{1.2}$$

where g is the determinant of the metric, M_P is the reduced Planck mass, R is the Ricci scalar and $\mathscr{L}_\phi$ is the inflaton Lagrangian, given by

$$\mathscr{L}_\phi = \frac{1}{2}\partial_\mu\phi\partial^\mu\phi - V(\phi), \tag{1.3}$$

where $\partial_\mu = \frac{\partial}{\partial x^\mu}$ and $V(\phi)$ is the inflaton potential. In eqs. (1.2) and (1.3) we use the mostly negative metric signature corresponding to the line element given in eq. (1.1). During inflation, the inflaton dominates the energy density of the universe, so other fields in the action have been ignored. Extremisation of this action leads to the usual Friedmann equations,

$$H^2 = \frac{1}{3M_P^2}\rho, \tag{1.4}$$

$$\dot{H} = -\frac{1}{2M_P^2}(\rho + P), \tag{1.5}$$

where $H = \dot{a}/a$ is the Hubble parameter, ρ is the background energy density, P is the background pressure and a dot indicates a derivative with respect to time. As stated above, the energy density and pressure are dominated by the inflaton field during inflation, and can be written in terms of the inflaton field and its potential as

$$\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi), \tag{1.6}$$

$$P = \frac{1}{2}\dot{\phi}^2 - V(\phi). \tag{1.7}$$

The action also gives the equation for the dynamics of the inflaton field,

$$\ddot{\phi} + 3H\dot{\phi} + \partial_\phi V = 0, \tag{1.8}$$

typically referred to as the (coupled) Klein-Gordon equation for the inflaton.

The inflaton potential is the part of this model that determines how inflation works. The main goal of inflation is to produce a period of exponential expansion long enough to explain the cosmological problems discussed above. Recalling the definition of the Hubble parameter

$$H = \frac{\dot{a}}{a} = \frac{\mathrm{d}\ln(a)}{\mathrm{d}t}, \tag{1.9}$$

where a is the scale factor, it is clear that exponential expansion can be achieved if H is constant. Comparing to eq. (1.5), this corresponds to $P = -\rho$, which can be realised if the kinetic energy can be neglected compared to the potential energy in eqs. (1.6) and (1.7). This leads to the slow-roll approximation, where the $\ddot{\phi}$ term in eq. (1.8) can be neglected, corresponding to an overdamped system.

Of great importance during inflation are the causal (or particle) horizon and the Hubble horizon. The causal horizon describes the maximum distance from which a particle travelling at the speed of light could have reached the observer at the current time, while the Hubble horizon is the boundary between objects moving away from the observer slower or faster than the speed of light. If two objects are separated by a causal horizon larger than the Hubble horizon, they are said to be out of causal contact, and cannot interact with one another. During inflation, the Hubble horizon is constant, while the expansion of the universe grows the causal horizon, leading to objects that were once in causal contact becoming causally separated from one another. This is how inflation solves the horizon problem.

The inflationary dynamics are commonly studied through the first two quantities in the Hubble hierarchy,

$$\epsilon_H = -\frac{\dot{H}}{H^2}, \tag{1.10}$$

$$\eta_H = \frac{\dot{\epsilon_H}}{\epsilon_H H}. \tag{1.11}$$

In the slow-roll approximation, these reduce to simpler forms derived from the inflaton potential,

$$\epsilon_H = \epsilon_V, \tag{1.12}$$

$$\eta_H = -2\eta_V + 4\epsilon_V, \tag{1.13}$$

with

$$\epsilon_V = \frac{M_P^2}{2}\left(\frac{\partial_\phi V}{V}\right)^2, \tag{1.14}$$

$$\eta_V = M_P^2 \frac{\partial_\phi^2 V}{V}. \tag{1.15}$$

These parameters control the expansion rate and duration of inflation respectively. To be in the slow-roll regime, $\epsilon_V << 1$ and $|\eta_V| << 1$. Care must be taken to use the general definitions ϵ_H and η_H in cases where inflation leaves the slow-roll regime, as is required for the production of primordial black holes.

In addition to solving the horizon, flatness, and monopole problems, inflation provides a means to explain the overdensities that seed large scale structure, through the medium of quantum perturbations on top of the background evolution described above. The magnitude and scale of these overdensities are typically considered using the primordial power spectrum.

1.1.4 The primordial power spectrum

The quantum fluctuations that seed large scale structure can be treated as perturbations of the inflaton field,

$$\phi(\boldsymbol{x}, t) = \phi(t) + \delta\phi(\boldsymbol{x}, t). \tag{1.16}$$

Adding this perturbation to the action and carrying out the extremisation at second order gives the equation of motion for the field perturbations,

$$\ddot{\delta\phi} + 3H\dot{\delta\phi} - \frac{1}{a^2}\boldsymbol{\nabla}^2\delta\phi + \partial_\phi^2 V\delta\phi = 0. \tag{1.17}$$

It is common to express these equations in Fourier space, allowing different k modes to be decoupled. This gives

$$\ddot{\delta\phi} + 3H\dot{\delta\phi} + \frac{k^2}{a^2}\delta\phi + \partial_\phi^2 V\delta\phi = 0. \tag{1.18}$$

Since these perturbations are random, their effect must be considered in a statistical manner. This is done using the correlation of the perturbations at two different spatial positions, but at equal times. In Fourier space, this is identified with the power spectrum $\mathcal{P}_{\delta\phi}(k,t)$, defined by

$$\langle\delta\phi(\boldsymbol{k},t)\delta\phi(\boldsymbol{k}',t)\rangle = \frac{2\pi^2}{k^3}\mathcal{P}_{\delta\phi}(k,t)\delta(\boldsymbol{k}+\boldsymbol{k}'), \tag{1.19}$$

where $\delta(\boldsymbol{k}+\boldsymbol{k}')$ is a Dirac delta function.

While $\mathcal{P}_{\delta\phi}$ is a relevant quantity during inflation, it becomes less convenient afterwards, when the inflaton has decayed into other fields. Therefore, it is necessary to construct a parameter that describes these overdensities after the end of inflation. This is done by considering each perturbation to be a separate FLRW universe with a curvature K, more commonly parametrised in terms of the curvature perturbation ζ as

$$\mathrm{d}s^2 = \mathrm{d}t^2 - a^2(t)e^{2\zeta}\mathrm{d}\boldsymbol{x}^2. \tag{1.20}$$

The curvature perturbation is an improvement over the inflaton perturbation, but it is still problematic in that it is a gauge variant quantity. To overcome this, a gauge invariant version can be determined. This is known as the comoving curvature

perturbation, and is related to the curvature perturbation by

$$\mathcal{R} = \zeta + \frac{\delta\rho}{3(\rho + P)}. \tag{1.21}$$

In the comoving gauge, $\delta\rho = 0$ and hence $\mathcal{R} = \zeta$. Additionally, in the super-horizon regime, where the causal horizon at scale k is larger than the Hubble horizon, $\delta\rho$ can be neglected and the two definitions are again equivalent. Therefore, it is common to see $\mathcal{R}$ and ζ used interchangeably, as is done throughout this thesis.

Converting eq. (1.18) to this variable gives the equation of motion for $\mathcal{R}$,

$$\ddot{\mathcal{R}} + (3 + \eta_H)H\dot{\mathcal{R}} + \frac{k^2}{a^2}\mathcal{R} = 0, \tag{1.22}$$

known as the Mukhanov-Sasaki equation [22, 23]. The comoving curvature perturbation power spectrum is then defined as for the inflaton perturbation,

$$\langle \mathcal{R}(\boldsymbol{k}, t)\mathcal{R}(\boldsymbol{k}', t)\rangle = \frac{2\pi^2}{k^3}\mathcal{P}_\mathcal{R}(k, t)\delta(\boldsymbol{k} + \boldsymbol{k}'). \tag{1.23}$$

The time dependence of this power spectrum is not important, since we only care about the impact these perturbations have on the evolution of the universe long after inflation. In slow-roll, this late-time effect is fixed by the phenomenon of freeze-out, where the evolution stops once the causal horizon at a scale k overtakes the horizon corresponding to the scale of the Hubble parameter H. Therefore, the late-time perturbations are described by the power spectrum evaluated at freeze-out,

$$\mathcal{P}_\mathcal{R}(k) = \mathcal{P}_\mathcal{R}(k, t)|_{k=aH}. \tag{1.24}$$

When not in slow-roll, such as in the ultra-slow-roll phase described in section 1.2.1, freeze-out does not necessarily apply and the power spectrum must be evaluated after the end of the non-slow-roll phase, i.e.

$$\mathcal{P}_\mathcal{R}(k) = \mathcal{P}_\mathcal{R}(k, t \geqslant t_{\text{end}}). \tag{1.25}$$

Observation of the CMB anisotropies provides a measurement of this power spectrum on large scales of $\mathcal{P}_{\mathcal{R}}(k) \approx 2 \times 10^{-9}$ with a very mild k dependence. However, no such measurements exist on smaller scales, allowing freedom for the power spectrum to grow to much larger values, potentially generating interesting observables, or objects such as primordial black holes.

1.2 Primordial black holes

Primordial black holes (PBHs) [24–26] inhabit a unique position in both cosmology and astrophysics. Their formation in the early universe (see section 1.2.1) allows them to act as a probe of the physics at these times. However, the no-hair theorem states that information about the formation of a black hole is lost, and the resulting black hole can be described in terms of only three parameters: its mass, spin, and electric charge [27–29]. This means that PBHs in the late universe are only distinguishable from black holes formed from stars by these three properties (see section 1.2.2). PBHs also have the potential to explain some or all of the dark matter content of the universe, characterised through f_{PBH}, the fraction of dark matter in the form of PBHs. Constraints on this fraction cover a broad range of masses, and are discussed in section 1.2.3. The detection of binary black hole merger events by the LIGO–Virgo collaboration (LVC) over the last 6 years has renewed interest in PBHs, due to the observation of some properties not trivially explained by astrophysics [30–35].

1.2.1 Formation mechanisms

Collapse of inflationary overdensities

The most commonly studied mechanism for forming PBHs is the collapse of large overdensities generated by the inflaton perturbations. These overdensities are characterised by the density contrast $\delta = \frac{\delta\rho}{\rho}$, where $\delta\rho$ is the local change in pressure

relative to ρ, the background density[2]. If the density contrast exceeds a critical value δ_c, then a PBH will be formed. This critical value is very large, so it is quite improbable to achieve a density contrast large enough to form a PBH. To connect this back to the inflationary physics discussed in sections 1.1.3 and 1.1.4 the variable δ must be related to the comoving curvature perturbation $\mathcal{R}$, which is done using eq. (C.1). In standard slow-roll inflation, the power spectrum for the comoving curvature perturbation $\mathcal{R}$ is approximately constant at $\mathcal{P}_\mathcal{R} \sim 10^{-9}$, and no PBHs form. To generate even a single PBH in a Hubble volume, the power spectrum must grow to $\mathcal{P}_\mathcal{R} > 10^{-3}$, as can be seen in fig. 2.6. As discussed in section 1.1.4, there are no measurements on the small scales relevant to PBHs, although there are some constraints, as can be seen in fig. 2.5. For a given power spectrum, there is a well-defined procedure for obtaining the PBH abundance and mass distribution, described in section 2.2, where there is also a discussion of the different choices and how they affect the final result.

To grow the power spectrum to the amplitudes required for PBHs to form requires going beyond slow-roll inflation. An enhancement of the power spectrum at a scale k corresponds to the inflaton spending more time at that scale than at other scales. This means the inflaton must slow down, and so this period is referred to as ultra-slow-roll (USR). Typically, USR inflation is realised by modifying the inflaton potential $V(\phi)$ to become much shallower at the relevant scales, by including a feature such as an inflection point or even a short period where the potential increases [36–41]. This type of potential is shown in fig. 1.1. It has recently been shown that USR can also be realised by increasing the inflaton velocity with a steeper potential before slowing it back down to the original velocity, since the quick transition is what is required to obtain USR [42].

[2]Often, the compaction function C is used, which acts as a smoothed version of δ.

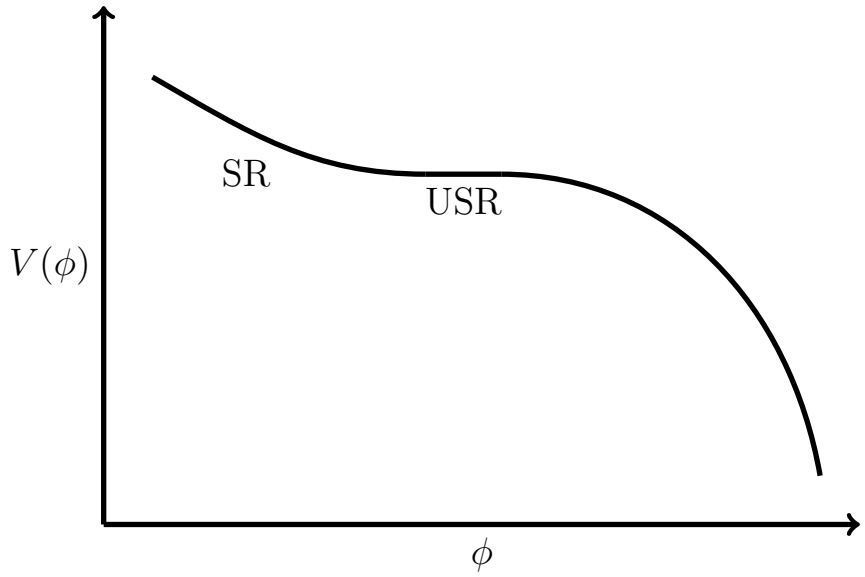

Figure 1.1: Inflaton potential $V(\phi)$ with an ultra-slow-roll (USR) period. The inflaton begins in a slow-roll (SR) phase, generating the power spectrum on CMB scales. It then passes through a plateau, causing a phase of USR and power spectrum growth, before speeding up again to end the inflationary period.

If there is primordial non-Gaussianity, the distribution of δ may be skewed towards higher values, making PBH formation more probable. Since the PBH abundance is exponentially sensitive to the overdensities, even a small amount of non-Gaussianity can significantly increase the population of PBHs.

While inflaton perturbations is the most studied method to produce PBHs and is the one considered throughout this thesis, there are a number of alternative formation scenarios.

Alternative formation scenarios

There are a number of mechanisms to form PBHs that do not involve the inflaton perturbations, and hence would evade the requirement of a large power spectrum amplitude. These include, but are not limited to, multi-field inflation, collapse of cosmic strings, bubble wall collisions, scalar condensate collapse, and collapse of domain walls. See [43] for a review of these methods.

Equation of state effects

A change in the equation of state can affect the critical density δ_c, and hence the abundance of PBHs. Any reduction in pressure will enhance the PBH abundance on the relevant scales. Such a reduction would be provided by a phase transition, either a non-standard model transition that creates a period of early matter domination [44], or standard model cases such as the QCD phase transition, which enhances the production of PBHs with masses of order the solar mass [45]. Alternatively, a period of early matter domination can be realised as part of the reheating phase after inflation [46].

1.2.2 Physical properties

The no-hair theorem ensures that in the late universe long after their formation, PBHs are mathematically identical to any other black hole, i.e. their behaviour depends only on their mass and spin[3]. Even so, PBHs can behave very differently to astrophysical black holes, in terms of the mass and spin values they can have.

Mass

The mass of a PBH can be significantly different to that of an astrophysical black hole in a number of ways. As described in section 1.3.1, astrophysical black hole formation is complex, and there are suggestions of two gaps in their allowed masses. On the contrary, PBH formation is much simpler, and there are no disallowed regions in their mass distribution. Additionally, the degeneracy pressure that prevents the collapse of stars into black holes is not present when PBHs form, meaning that while an astrophysical black hole could never have a mass smaller than $M_\odot$, a PBH could form with a mass as little as 10^{-18} $M_\odot$ and still persist to the present day. These are important distinctions, as significant evidence of compact objects in the astrophysical mass gaps or strong evidence of a single compact object with $m < M_\odot$ could be

[3] In principle the electric charge can also be important, but electromagnetic repulsion prevents black holes from forming with significant charge, and any charge attained by accretion of charged particles would quickly neutralise by attraction of particles of the opposite charge.

"smoking gun" evidence for PBHs. A few mass gap objects have been observed by the LVC [32–34], although these could be attributed to second generation mergers. No sub-solar mass objects have been observed so far [47–49], although these are computationally challenging to discover and there are ways to obtain mergers with non-detectable sub-solar mass components [50].

The PBH mass distribution ψ is a central theme of this thesis. A good description of ψ is essential for determining the behaviour of PBHs and how they relate to late-time observables and constraints, and also for tracking back those constraints to their formation and hence to constraints on the power spectrum. The precise method of calculating the mass distribution is a matter of some discussion, with a number of choices affecting the result. These choices are considered in chapter 2, where we demonstrate that the differences are currently unimportant, although as more accurate constraining mechanisms become available, they will have to be taken into account.

One thing that is universally agreed upon is that the initial mass m of a PBH follows a critical scaling in the density contrast,

$$m = K M_H (\delta - \delta_c)^\gamma, \tag{1.26}$$

where K and γ are constant scaling factors, δ_c is the critical overdensity required to form a PBH and M_H is the mass contained within the Hubble horizon at the formation time. This critical collapse ensures that, while there may be a maximum PBH mass determined by the position and shape of the peak in the power spectrum, the mass can extend to arbitrarily small values, creating an enhancement of the low-mass tail of the mass distribution. This is examined in chapter 3, where we demonstrate that the commonly used lognormal model for the PBH mass distribution fails to describe the shape when critical collapse dominates, and propose some alternative models that may match the numerical mass distribution better for narrow as well as broad power spectrum widths.

There is an assumption inherent in the above reasoning, which is that the PBH mass distribution does not change between the formation time and the late universe. One obvious case where this will not be true is if accretion is significant for PBHs. This has been considered for a number of years [51–53], and has the effect of skewing the mass distribution to heavier masses, although there are still a number of uncertainties in the result.

Spin

Much like the mass, the PBH spin is expected to be quite different to that of astrophysical black holes. While the collapse of stars is expected to produce a black hole with a very large spin, the formation of PBHs instead predicts a very small spin [54–57]. This is an important result, and could act as a discriminator between astrophysical and primordial black holes for measurements such as the LVC merger events [58,59]. The LIGO instrument is sensitive to a combination of the spins called the effective spin parameter $\chi_{\rm eff}$, and the LVC has found that the vast majority of merger events have $\chi_{\rm eff}$ posteriors consistent with zero [30, 35]. As with the mass, the potential effect of accretion will cause the PBH spins to increase significantly away from zero [60].

Clustering

Constraints on PBHs in the late universe typically assume that they are not clustered, i.e. that they follow a Poissonian spatial distribution. Whether any clustering would alleviate or tighten constraints has been a matter of debate, but it is now universally agreed that PBHs formed from inflaton perturbations should be unclustered if the initial perturbations are Gaussian in nature [61–66]. In the case of non-Gaussianity, or clustering after formation, there are significant implications for late-time observables such as the merger rate [67–74]

1.2.3 Constraints

The abundance of PBHs, or equivalently the fraction of dark matter they constitute, $f_{\rm PBH}$, can be constrained by a number of techniques that span the entire range of masses they could have. The constraints fall into two categories: those for PBHs that evaporated before the present day, and constraints for PBHs that have not yet evaporated. While only the second category of PBHs is of interest as an explanation of dark matter, evaporated PBHs can still provide useful constraints on the early universe physics when they formed, particularly the amplitude of the primordial power spectrum. Constraints on evaporated PBHs are focused on examining their effect on early universe events, such as Big Bang nucleosynthesis (BBN) and the CMB, and their contribution to energetic backgrounds, such as those for γ-rays and cosmic rays.

For non-evaporated PBHs, there are many more methods to apply constraints that vary depending on the mass being observed. Firstly, low mass PBHs will still be evaporating, and will have constraints from the γ-ray and cosmic ray backgrounds. This provides a steeply sloped constraint below 10^{-16} $M_\odot$. Between 10^{-10} $M_\odot$ and $M_\odot$, the dominant constraints are from lensing searches, looking for compact objects passing in front of background sources. These provide a constraint of $f_{\rm PBH} \lesssim 10^{-2}$. Between 10^2 $M_\odot$ and 10^5 $M_\odot$, the lack of CMB anisotropies caused by PBH accretion provides an extremely tight constraint of $f_{\rm PBH} \lesssim 10^{-8}$. From 10^5 $M_\odot$ to 10^{10} $M_\odot$, the dominant constraints mostly come from dynamical effects, based on encounters between PBHs and other astrophysical objects. These encounters subdivide into collisions, pass-bys that disrupt orbits or cause tidal disruption depending on the mass of the PBH, and PBHs clustering together due to dynamical friction. Most constraints in this mass range give $f_{\rm PBH} \lesssim 10^{-3}$, although dynamical friction reaches down to 10^{-4}. At the largest masses, there are constraints from ensuring that cosmic structures do not form too early due to the presence of PBHs. This gives a constraint of $f_{\rm PBH} \lesssim 10^{-3}$ at 10^{11} $M_\odot$. Finally, there is a constraint from binary black hole merger events observed by the LVC, which gives $f_{\rm PBH} \lesssim 10^{-2}$ at $\sim 10^2$ $M_\odot$ and from

the non-detection of sub-solar mass mergers, giving $f_{\rm PBH} \lesssim 10^{-1}$ at $\sim 10^{-1}$–1 $M_\odot$. Potential constraints in the 1–10^2 $M_\odot$ range are discussed in [43] but are not shown in fig. 1.2. With future detectors and upgrades to the existing detectors, gravitational waves will provide a significantly tighter constraint compared to the other methods in this mass range. It should be noted that many of these constraints are subject to astrophysical uncertainties or assumptions that may be incorrect (see [43] for further discussion of the assumptions and uncertainties inherent in each constraint), and some have appeared and disappeared over time. Recently, a window has opened up between 10^{-16} $M_\odot$ and 10^{-10} $M_\odot$ in which PBHs could constitute the entirety of the dark matter [75]. The tightest constraints from each method are shown in fig. 1.2.

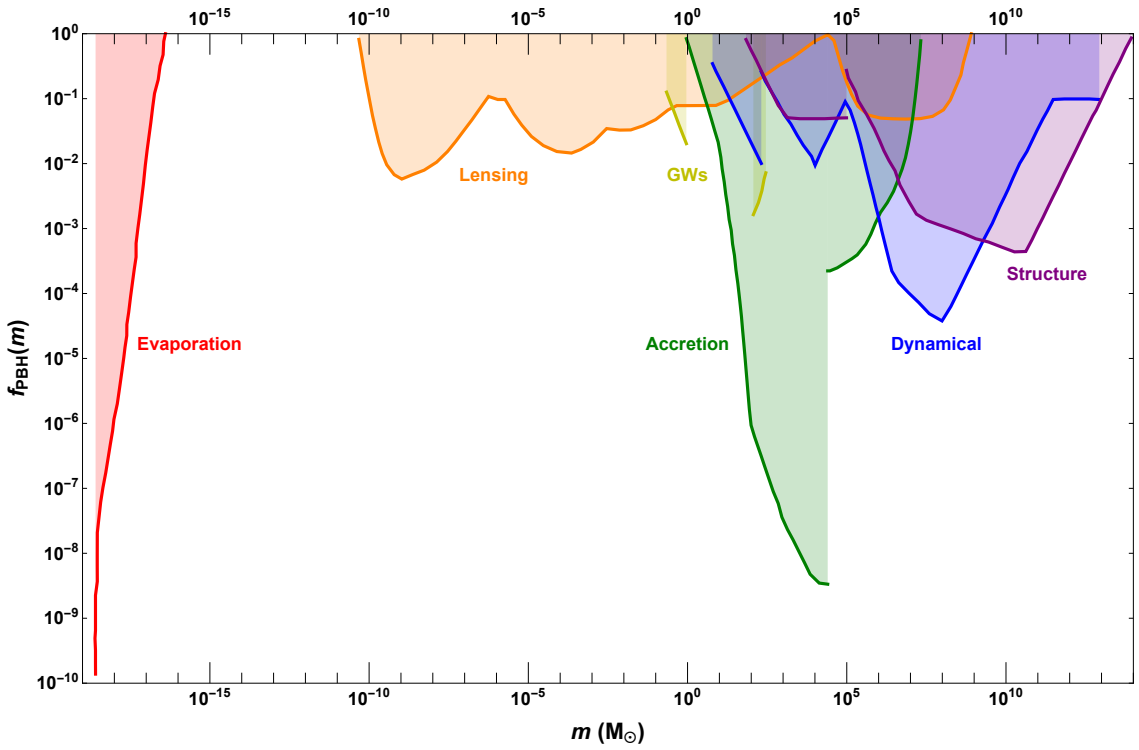

Figure 1.2: Constraints on $f_{\rm PBH}$ for PBHs that have not evaporated by the present day. Evaporation constraints are shown in red, lensing constraints in orange, gravitational wave constraints in yellow, constraints from accretion in green, constraints from dynamical effects in blue and constraints from cosmic structure in purple. Data from [43] and references therein.

This is not quite the whole story, as there are also constraints that can be applied to PBHs through other effects caused by a large power spectrum amplitude. These fall into two categories: spectral distortions in the CMB, and a contribution to the stochastic gravitational-wave background from scalar-induced gravitational waves (SIGWs). These effects are discussed in chapter 2, and show that PBHs cannot form with masses greater than $\sim 10^4$ $M_\odot$. Future constraints will be even more prescriptive, and will rule out the possibility of PBHs across the entire late-time mass range, apart from two small windows at $\sim 10^{-16}$ $M_\odot$ and $\sim 10^{-4}$ $M_\odot$.

Constraints on $f_{\rm PBH}$ are typically presented for a monochromatic PBH mass distribution, i.e. all PBHs have the same mass. This is not possible to generate in practice, due to the effects of critical collapse ensuring an extended mass distribution even for a delta function peak in the power spectrum (which is itself unphysical). Therefore, the effect of moving to an extended mass distribution has been studied [76, 77], showing that the constraints can change, typically broadening them, and hence closing windows that are open when considering the monochromatic case. For an extensive review of all the constraining mechanisms for $f_{\rm PBH}$, as well as other discussion about the constraints, see [43].

1.3 Astrophysical black holes

While primordial black holes are still hypothetical, there is clear evidence that some form of black hole exists in the universe. The only emission directly from a black hole is Hawking radiation, which is extremely difficult to observe. Therefore, black holes must be detected through mediating effects [78]. These can be broadly classified as gravitational effects (modification of the orbits of stars around the black hole [79], gravitational waves [80]) and signals from matter surrounding the black hole (X-ray emission from the accretion disk [81], Event Horizon Telescope image [82]). These detected black holes are typically assumed to have formed astrophysically, from the collapse of heavy stars at the end of their life cycles.

1.3.1 Formation and mass gaps

There are three types of black holes predicted or known to exist. At the heavy end are supermassive black holes (SMBHs), with $m \gtrsim 10^5$ $M_\odot$. These are believed to inhabit the centre of all galaxies, and form by accretion onto seeds from the earliest stars, although there are some problems growing these seeds quickly enough to explain very distant SMBHs [83, 84]. From $\sim 10^2$–10^5 $M_\odot$ are the intermediate mass black holes (IMBHs), candidates of which have been detected in our galaxy and nearby galaxies [33, 85–87]. Finally, and of most interest to this thesis, are stellar-mass black holes, with masses $\sim$ 1–100 $M_\odot$.

Astrophysical black holes are formed at the end of the stellar life cycle. When a star runs out of light elements to fuse, the radiation pressure falls off, and the stellar matter is no longer supported against the gravitational effects of its mass. The star then collapses in on itself, creating a region of very high density. For heavy stars, the outer layers of stellar matter are typically ejected in a supernova event, and a stellar remnant is left behind. The mass of this remnant will be significantly smaller than the mass of the original star before collapse. Below a certain mass, known as the Tolman-Oppenheimer-Volkoff (TOV) limit [88], the degeneracy pressure between free neutrons is enough to prevent further gravitational collapse. However, if the remnant exceeds the TOV limit, then nothing can prevent the gravitational collapse and a black hole will be formed. There are many uncertainties regarding the behaviour of these stellar remnants, so the TOV limit is not precisely known. Recent calculations place it at ~ 2.3 $M_\odot$ [89, 90], but there are claimed detections of heavier neutron stars [91, 92].

Through observational and theoretical means, it has been suggested that the stellar-mass black hole distribution may have two gaps: a lower one between the heaviest neutron star and the lightest black hole ($\sim$ 2.5–5 $M_\odot$), and an upper one at $\sim$ 50–150 $M_\odot$. As stated above, there are many uncertainties about stellar remnants close to the TOV limit, but observational evidence from X-ray binaries has implied the existence of the lower mass gap [93–95], although it has been suggested that

systematic errors could remove the evidence [96], or that the objects in this gap may be unobservable with the techniques applied so far [97]. The upper mass gap is motivated by more theoretical arguments. As the star collapses, pair production (the creation of a particle and its corresponding antiparticle) happens very efficiently, destabilising the star [98]. This ends in one of two ways, depending on the mass of the helium core in the star. If the helium core is $\sim$ 30–64 $M_\odot$, then the stellar matter pulses, shedding mass until stability can be reached again. This is known as a pulsational pair-instability supernova [99–104], and leads to a reduced stellar remnant mass, and hence a reduced black hole mass. If the helium core is $\sim$ 64–135 $M_\odot$, then the destabilisation incurred by pair-production is so large that it causes the entire star to be ripped apart. This is known as a pair-instability supernova [105–107], and leaves no stellar remnant behind, and hence no black hole. The combination of these two effects leads to the upper gap in the black hole mass distribution. The LVC has detected objects that fall into each of these mass gaps [33, 34], although it is possible that these objects are the result of previous mergers, such as the remnant object produced from the binary neutron star event GW170817 [108].

1.3.2 Binaries

Astrophysical binary black holes (BBHs) can form in two ways. In highly dense environments, such as globular clusters, binaries can form from the encounter of two independently formed black holes [109, 110]. This is referred to as dynamical formation, and is aided by dynamical friction causing black holes to sink to the centre of the clusters. These binaries are also subject to scattering with other black holes or binaries due to the high density of these objects, the importance of which is discussed below.

The other way of producing a BBH is through the common envelope scenario, where each star in a binary star system becomes a black hole in turn [111, 112]. This process is more complicated than dynamical formation, as it is tied to the stellar

evolution as well as the black hole physics. The dynamics are described by the Roche lobe of the two stars, and mass transfer between them by Roche lobe overflow. The Roche lobe describes the area that contains all material gravitationally bound to a star in a binary system. At the end of a heavy star's life, it will expand and its size will exceed its Roche lobe, at which point mass will start to transfer to the companion in the binary. If the companion cannot accrete matter fast enough, then the two Roche lobes will merge and form a common envelope of matter around both stars. Each of the stars will then accrete matter in turn until they collapse into black holes.

One property of both these methods is that the resulting BBHs will have similar masses [113]. In the common envelope scenario, this is because the stars that eventually form the black holes are sharing the same surrounding matter, leading to an equilibrium in the masses. For the dynamical formation mechanism, one might imagine that there is no preferred mass ratio, since the binaries are formed by chance encounters. However, since the black holes collect due to dynamical friction, the most massive ones fall to the centre first, followed by smaller and smaller ones in turn. This means that the black holes that form binaries tend to be of similar masses. Additionally, the highly dense environment means that scattering is common, either with another binary or a single black hole. These scattering events tend to exchange black holes between binaries, favouring the formation of binaries with similar mass black holes.

1.4 Gravitational waves

Gravitational waves are propagating solutions of the Einstein equations. First proposed by Heaviside in 1893 by analogy to the electromagnetic waves that arise from Maxwell's equations [114], they were formalised and named as "gravitational

waves" by Poincaré in 1905 based on the Lorentz transformations[4] [115], and finally derived explicitly from Einstein's general theory of relativity in 1916 [119]. Doubts about the physical reality of these solutions persisted for decades, with even Einstein switching sides multiple times. By the late 1950s, it was generally accepted that the solutions were physical, and work began on detecting them, beginning with Weber bars [120, 121] and culminating in the modern gravitational wave detectors and the first detection by LIGO in 2015 [80]. See [122] for a comprehensive discussion of the history of gravitational waves.

1.4.1 Gravitational wave sources

Gravitational waves can be produced by a number of astrophysical mechanisms, and typically fall into three categories: transient, continuous, and stochastic. Transient signals appear for a short time and then disappear again, continuous signals are approximately constant in frequency and persist over a very long time period, and the stochastic background is the result of the superposition of many high redshift sources that are impossible to resolve individually. Of interest to this thesis are transient signals from compact binary coalescences and the stochastic background.

Compact binary mergers

Transient signals from compact binary mergers are the most discussed gravitational wave signals at the moment, thanks to the LVC detections [30, 35]. These signals are produced from the final stages of the orbit and subsequent coalescence of two compact objects. The signal is split into three phases: inspiral, merger, and ring-down. The inspiral phase occurs as the components orbit each other, moving closer together as they lose energy by the emission of the gravitational waves. This leads to a distinctive signal known as a "chirp", characterised by an increase in frequency and amplitude as the components near each other and speed up. There is a mass

[4] Poincaré actually wrote two articles with the same title: the first in June 1905 [115], followed by a significantly extended version written in July 1905 and published in January 1906 [116]. Both include the term "gravitational wave", and downplay the significance of their results to the developing field of relativity as minor additions to Lorentz's 1904 paper [117, 118].

scale associated with this phase known as the chirp mass, which can be written in terms of the component masses m_1 and m_2, as

$$\mathcal{M} = \frac{(m_1 m_2)^{\frac{3}{5}}}{(m_1 + m_2)^{\frac{1}{5}}}, \tag{1.27}$$

and can also be approximately related to the gravitational wave frequency f and its time derivative $\dot{f}$ as

$$\mathcal{M} \approx \left(\frac{5}{96}\pi^{-\frac{8}{3}} f^{-\frac{11}{3}} \dot{f}\right)^{\frac{3}{5}}. \tag{1.28}$$

Being directly related to the gravitational wave frequency, this is often the best-constrained mass scale for a merger event. When the surfaces of the compact objects touch, the merger phase begins, as the two components coalesce into one final remnant object. This happens very quickly and produces the strongest gravitational waves. Afterwards, there is a ringdown phase, during which the remnant object releases excess energy and settles into its final state. These phases provide information on the properties of the final remnant object, and can be used to test the general theory of relativity by comparing the remnant properties predicted from the inspiral phase with those determined from the merger and ringdown phases. Figure 1.3 shows the waveform for the LIGO merger event GW150914, with the chirp effect clearly visible.

The compact objects involved in binary merger events can be black holes, neutron stars, or white dwarfs. Mergers involving white dwarfs are outside the frequency band of current detectors, so compact binary mergers can only fall into three categories: binary black holes (BBHs), binary neutron stars (BNSs), and a mixed neutron star–black hole merger (NSBH). Distinguishing between the categories is typically done using a combination of the chirp mass and the tidal deformability. A chirp mass $\mathcal{M} \lesssim 2$ $M_\odot$ is assumed to be a BNS, $\mathcal{M} \gtrsim 4.5$ $M_\odot$ is assumed to be a BBH, and any other events are considered as NSBHs [30]. Tidal deformability measures how much the shape of the compact object changes due to the gravitational

forces imparted by its companion. Black holes should have zero tidal deformability, whereas neutron stars will deform shortly before the merger phase [108]. Detection of an electromagnetic counterpart also indicates that at least one of the component objects must be a neutron star, as in the case of GW170817 [123].

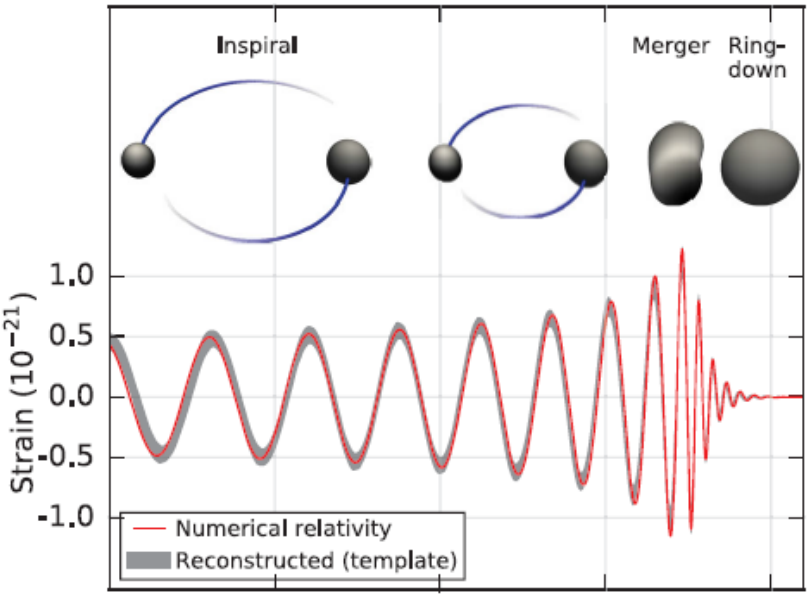

Figure 1.3: Waveform of the gravitational wave emitted by the inspiral, merger, and ringdown of the binary black hole coalescence GW150914 detected by the LVC. Image adapted from [80].

Stochastic background

The stochastic gravitational wave background (SGWB) is produced by all the signals that are unable to be individually detected. There are many possible contributions to the SGWB, both astrophysical and cosmological. The astrophysical background includes binary mergers, non-spherical gravitational collapse, and unstable rapidly rotating objects such as neutron stars. The cosmological background contributions are more hypothetical, and typically represent a departure from standard cosmology. These include details of the inflationary and reheating phases, phase transitions, and cosmic strings. The primary component of interest to this thesis are scalar-induced gravitational waves (SIGWs) arising from large amplitudes of the primordial power spectrum [124–126], which are also required for the production of PBHs.

A non-detection of the SGWB at a certain amplitude is an indication that any potential contributions producing such an amplitude cannot exist. For SIGWs, this allows constraints to be placed on the primordial power spectrum amplitude. However, the opposite statement is not quite so simple. If a signal is detected, it may arise from a number of possible components, either astrophysical or cosmological, and distinguishing between components can be complicated. An excess of gravitational waves was recently observed by the NANOGrav collaboration using the method of pulsar timing arrays [127] (see section 1.4.2 for details of how this method works). If this were an SGWB sourced by SIGWs, it would imply a power spectrum amplitude large enough that there would be an associated PBH production [128–130].

1.4.2 Gravitational wave detectors

The detection of gravitational waves is often categorised as either direct or indirect. A direct detector searches for the effect of the local passing of a gravitational wave, while an indirect detection is found by examining the properties of astrophysical objects that emit gravitational waves. The first claimed detections of gravitational waves were made by Weber in 1968 and 1969 using a Weber bar, an example of a resonant antenna [120, 121]. However, issues regarding the rate of energy loss from the galactic centre were raised, and future attempts by other groups failed to reproduce the results [122]. A more robust detection was made in 1979 using an indirect method. In 1974, Hulse and Taylor discovered a binary system consisting of a neutron star and a pulsar. Doppler shifting of the radio signals emitted by the pulsar allowed the orbital properties of the binary to be determined, showing a decay of the orbit consistent with the energy lost by the emission of gravitational waves [131]. Most modern gravitational wave detectors utilise interferometry or pulsar timing arrays. These methods are described in the following sections.

Interferometry and the LIGO–Virgo collaboration

Interferometry is a technique that has been used in physics for over a century. This method involves splitting light with a well-defined wavelength into two beams that travel over different paths, then measuring the phase shift caused by the different optical path lengths (OPLs) once the beams have been recombined. The first interferometer was constructed by Michelson in 1881 as a means to detect the luminiferous æther [132], and later used in the famous Michelson–Morley experiment of 1887 [133]. This type of interferometer consists of two arms that are perpendicular to each other, with a beam splitter at the intersection of the two arms and mirrors at the other ends. Initially designed to observe the changing speed of light moving either perpendicular or parallel to the æther, this mechanism is also sensitive to a change in the length of the arms, such as that caused by a passing gravitational wave.

The application of interferometry to gravitational wave detection was first suggested by Gertenshteĭn and Putsovoĭt in 1962 [134], and again by Braginskiĭ in 1966 [135], although neither of these suggestions were acted upon. A number of groups independently had the same idea, and began building prototypes in the 1970s, notably Forward and Weiss [122]. This method then grew in popularity as larger and larger interferometers were funded and constructed, leading to the worldwide network of detectors in use today. The method is identical to the Michelson interferometer, but with the change in OPL sourced by a different phenomenon. As a gravitational wave passes, it stretches space in one direction and shrinks it in the perpendicular direction. Shortly afterwards, the stretched and shrunk directions swap as a different part of the gravitational wave moves past. For the interferometer, this results in one arm being longer than the other, and hence a phase shift when the light is recombined. To develop an observable signal, the OPL must be approximately a quarter of the wavelength of the gravitational wave, which requires

extremely long arms (∼750 km for binary mergers of stellar-mass compact objects). This is obviously infeasible to construct, so the light is cycled up and down the arms many times to produce the required OPL.

The most famous interferometer detector is of course the Laser Interferometer Gravitational-wave Observatory (LIGO) in the USA, funded in 1988, with construction beginning in 1994. This consists of two detectors, at Hanford, Washington and Livingston, Louisiana, each of which is a Michelson interferometer with 4 km arms. These detectors are ∼3000 km apart, such that there is a significant time difference between the arrival of a gravitational wave at each detector. This helps to ensure that the signal is astrophysical rather than terrestrial, and also allows the source to be localised, although not with great precision. This localisation can be improved with the addition of more detectors across a larger area. The first of these additional detectors is the Virgo detector, located at Santo Stefano a Macerata in Italy, and consisting of 2 km arms. The significant improvement in localisation provided by this additional detector can be seen by comparing the GW170814 event to previous events detected only by the LIGO detectors [136]. These two collaborations are jointly referred to as the LIGO–Virgo collaboration (LVC), and have released details of 50 gravitational wave candidates, comprising BBH and BNS mergers, as well as some unclassified events [30, 35]. The detections from the first two observing runs and notable events from the O3a run are shown in fig. 1.4.

More detectors will allow better localisation and classification of signals, as well as the detection of weaker sources by combining the signals from multiple detectors. The KAGRA detector is located in the Kamioka mine in Japan, and has 3 km arms. It began operation in February 2020, and combined with LIGO and Virgo has detected two mixed NSBH mergers [138]. A further detector known as LIGO-India is planned to be set up in Aundha Nagnath, India, moved from the LIGO Hanford site where there are currently two identical detectors. This will greatly

improve source localisation, improving the chances of successfully carrying out electromagnetic follow-up observations, which will open up multimessenger studies of astrophysical phenomena [139, 140].

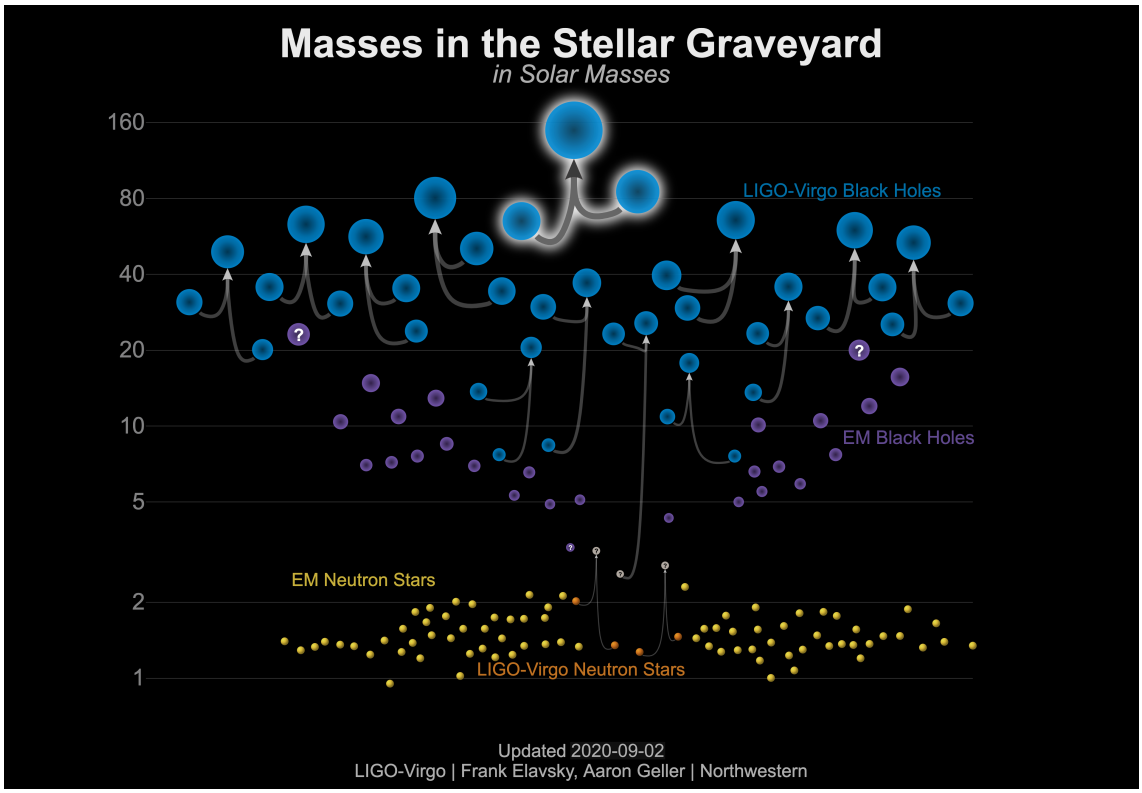

Figure 1.4: Binary merger events detected by the LIGO–Virgo Collaboration. Previous detections of black holes and neutron stars by electromagnetic means are shown in purple and yellow. The events shown here are from the first two observing runs (O1 and O2), as well as notable events from the O3a run, including three that have constituents in or close to one of the astrophysical mass gaps. Image from [137].

Pulsar timing arrays

An alternative method to detect gravitational waves comes from the timing of pulsars. Pulsars are rapidly rotating compact objects (mostly neutron stars) with strong magnetic fields. The rotating magnetic field leads to the emission of electromagnetic beams from the poles of the pulsar. If the earth is in the path of these beams, pulsars are observable as a series of pulses of electromagnetic radiation at fixed time intervals. As discussed above, pulsars will lose energy by the emission of gravitational waves and their rotation will slow. However, on shorter timescales, they can be treated as having a constant period and can therefore be used as precise

clocks. If a gravitational wave from another source passes between the pulsar and the earth, the time between subsequent pulses will be modified slightly, analogous to the changing length of the arms in an interferometric detector. Once the period of a pulsar is known, this tiny deviation can be detected and used to infer the existence of the passing gravitational wave. However, there are many astrophysical uncertainties that make this process complicated. To overcome this, observers can look for a signal that appears in many pulsars by correlating the individual signals together. Combining many pulsars in this way is described as a pulsar timing array (PTA) [141–143].

There are three main collaborations currently applying the PTA technique to search for gravitational waves. These are the Parkes Pulsar Timing Array (PPTA), the European Pulsar Timing Array (EPTA), and the North American Nanohertz Observatory for Gravitational waves (NANOGrav) collaboration. The three groups have also collaborated under the name International Pulsar Timing Array (IPTA), which also includes the planned Indian Pulsar Timing Array (InPTA). Additional projects are also in the works [144]. It must be made clear that the gravitational waves detected using this method are not those produced by the pulsars themselves, but unrelated ones passing between the pulsars and the earth. The requirement of correlating many pulsars means that the gravitational wave signals must last long enough to be seen multiple times. There are a number of sources potentially detectable by this method, but in this thesis the stochastic background is of primary concern. As discussed in section 1.4.1, constraints can be placed on the primordial power spectrum via the mechanism of SIGWs. Constraints from the NANOGrav collaboration are shown in fig. 2.5.

Future detectors

If the current state of gravitational wave detectors is exciting, then the future is even more bright, with a plethora of ground- and space-based detectors utilising interferometric and PTA techniques proposed or planned for the next few decades

ready to significantly improve the amount of gravitational wave signals and the science based on them. Although many are yet to receive funding, the number of proposals indicates the interest in the field and the likelihood of future projects receiving funding.

Two proposed ground-based interferometer detectors are the Einstein Telescope (ET) and the LIGO Cosmic Explorer (CE), both of which hope to begin observations in the 2030s at frequencies similar to the current LIGO band. The ET will have 10 km arms in a triangular setup and the CE will have 40 km arms in the traditional L-shape. These will be sensitive to the same signals as the current interferometer detectors such as LIGO, but with greatly enhanced sensitivity [145, 146]. A further ground-based detector is the Square Kilometer Array (SKA). This is a large radio telescope project with many science goals, including the detection of gravitational waves using the PTA method [147]. Construction is scheduled for the end of 2021.

Moving to space offers a number of benefits for gravitational wave detection, particularly the removal of many noise sources produced terrestrially. Many detectors have been proposed (all utilising interferometry), most notably the Laser Interferometer Space Antenna (LISA), which is funded and is planned to launch in 2034 [148]. This consists of three arms arranged in an equilateral triangle, with an approximate arm length of 2.5×10^6 km. This will primarily probe the mergers of compact binaries at the massive and supermassive scales, as well as the stochastic gravitational background [149]. Other proposed interferometric space-based detectors include the Japanese DECIGO [150] and the Chinese TianQin [151] and Taiji [152].

These future detectors will have the ability to confirm or deny the existence of PBHs across virtually the entire un-evaporated mass range, as described in chapter 2 (see fig. 2.6).

Chapter 2

The power spectrum on small scales: robust constraints and comparing PBH methodologies

Andrew D. Gow[1], Christian T. Byrnes[1], Philippa S. Cole[1], and Sam Young[2]

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

2) Max Planck Institute for Astrophysics, Karl-Schwarzschild-Strasse 1, 85748 Garching bei Muenchen, Germany

We compare primordial black hole (PBH) constraints on the power spectrum and mass distributions using the traditional Press–Schechter formalism, peaks theory, and a recently developed version of peaks theory relevant to PBHs. We show that, provided the PBH formation criteria and the power spectrum smoothing are treated consistently, the constraints only vary by $\sim 10\%$ between methods (a difference that will become increasingly important with better data). Our robust constraints from PBHs take into account the effects of critical collapse, the non-linear relation between ζ and δ, and the shift from the PBH mass to the power spectrum peak scale. We show that these constraints are remarkably similar to the pulsar timing array (PTA) constraints impacting the black hole masses detected by LIGO and Virgo, but that

the μ-distortion constraints rule out supermassive black hole (SMBH) formation and potentially even the much lighter mass range of $\sim$ (1–100) $M_\odot$ that LIGO–Virgo probes.

2.1 Introduction

Primordial black holes (PBHs) could have formed in the early universe from the collapse of density perturbations [24–26]. Although there are no confirmed detections of PBHs, there are tentative hints for their existence and in particular a lot of recent interest has focused on whether the Laser Interferometer Gravitational-wave Observatory (LIGO) has detected PBHs [153, 154]. Assuming that PBHs formed from the collapse of large amplitude perturbations shortly after horizon entry during radiation domination, there is an approximate one-to-one relation between the scale at which the primordial power spectrum has a large amplitude peak and the mass of PBHs that form. See [43, 155, 156] for reviews.

In order for PBHs to form, the amplitude of the power spectrum must become orders of magnitude larger than the value of 2×10^{-9} measured on large scales, e.g. via observations of the cosmic microwave background (CMB) [157]. Precisely how much larger it must become is a matter of active research, with significantly differing values being quoted in the literature, typically varying between $\mathcal{O}(10^{-3})$ and $\mathcal{O}(10^{-2})$ with values at the lower end quoted in e.g. [158, 159]. $\mathcal{O}(10^{-1})$ values have also been considered in e.g. [160]. Since the power spectrum amplitude is only logarithmically sensitive to the allowed energy density fraction of PBHs, this variation has little to do with the different PBH masses or constraints being considered and instead is primarily due to differences in the theoretical techniques being used to relate the power spectrum amplitude to the abundance of PBHs. Primordial non-Gaussianity also has an important impact on the required power spectrum amplitude, see e.g. [161–168], but we will not consider that issue further in this paper.

However, we do include an accurate approximation for the significant correction arising due to the non-linear relation between the density contrast and curvature perturbation, the importance of which has only recently been quantified [169–173].

In this paper we make the first detailed study of how the PBH mass distribution differs when using Press–Schechter or peaks theory as well as a recently developed treatment of peaks theory [174], which solves a problem for PBHs related to the cloud-in-cloud problem. When a PBH forms, the final mass depends on both the amplitude and scale of the perturbation from which it forms [175], and the new treatment of peaks theory ensures that the amplitude of peaks are evaluated at the correct scale, giving the correct mass distribution and abundance. We also consider the sensitivity to the choice of the window function. We show that, provided that all quantities are calculated in a self-consistent way—for example, the choice of window function must be reflected in the collapse threshold δ_c—all techniques and window functions lead to quite consistent results whereby the uncertainty in the power spectrum amplitude is only of order 10%. This is a much smaller variation than [176] found even due to just the choice of the window function alone, consistent with the corrections accounted for in [177]. We also note that, throughout this paper, we assume a fixed value for the collapse threshold of primordial perturbations. In reality, the exact value of the collapse threshold depends on the specific shape of each individual perturbation, and neglecting this gives an additional uncertainty of order a few percent [177–180].

The uncertainty in the initial conditions necessary to generate a required number of PBHs has important implications for relating observations of PBHs to observations of the associated enhanced amplitude of the primordial perturbations. This can be done, for example, via the observation of a stochastic background of gravitational waves measurable by pulsar timing arrays (PTAs) which measure frequencies corresponding to a horizon scale which could have formed the black holes observed

by LIGO and Virgo. In general, understanding how to map from a PBH abundance to a power spectrum constraint is important for our understanding of the initial conditions of the universe and the constraints on models of inflation.

In the next section we introduce the calculation of the PBH mass distribution. In section 2.3 we discuss how the result depends on the calculation technique and window function and we use these results to calculate robust constraints on the primordial power spectrum in section 2.4, in particular showing that the pulsar timing array constraints are not inconsistent with the formation of LIGO mass PBHs. We conclude in section 2.5, and some technical details of the observational constraints and the non-linear mapping from the curvature perturbation to the density contrast are contained in appendices B and C.

2.2 Obtaining the PBH mass distribution

The procedure for obtaining the mass distribution from the power spectrum is similar for all three methods considered, and is based on connecting the PBH abundance $\Omega_{\rm PBH}$ to the mass fraction $\beta = \frac{\rho_{\rm PBH}(t_i)}{\rho(t_i)}$, where $\rho_{\rm PBH}$ is the mean energy density in PBHs, ρ is the total background energy density, and t_i is the time at which the PBHs form. This mass fraction is then related to the power spectrum. In every case, the PBH abundance is calculated from the mass fraction using

$$\Omega_{\rm PBH} = \int \mathrm{d}(\ln R)\, \frac{R_{\rm eq}}{R} \beta(R), \tag{2.1}$$

where R is the horizon scale at the time the PBH is forming, $R_{\rm eq}$ is the horizon scale at matter-radiation equality and the ratio takes into account the relative growth of the PBH fraction during radiation domination. The form of $\beta(R)$ is different for each method, see eqs. (2.6), (2.8), and (2.10). The abundance is then related to the PBH mass function $f(m)$ through

$$f(m) = \frac{1}{\Omega_{\rm CDM}} \frac{\mathrm{d}\Omega_{\rm PBH}}{\mathrm{d}(\ln m)}, \tag{2.2}$$

which satisfies the normalisation condition

$$\int \mathrm{d}(\ln m)\ f(m) = f_{\mathrm{PBH}} = \frac{\Omega_{\mathrm{PBH}}}{\Omega_{\mathrm{CDM}}}. \tag{2.3}$$

This can then be related to the mass distribution $\psi(m)$ through

$$\psi(m) = \frac{1}{f_{\mathrm{PBH}}} \frac{f(m)}{m}, \tag{2.4}$$

which is a PDF and hence satisfies the normalisation condition

$$\int \mathrm{d}m\ \psi(m) = 1. \tag{2.5}$$

The relation between $\beta(R)$ and the power spectrum then depends on the method used. In this paper, three methods are considered: a Press–Schechter-like calculation (PS), the traditional peaks theory method (TP) described in the classic BBKS paper [181], and a modified peaks theory derived by Young and Musso (YM) [174].

Recently, other variations of peaks theory have also been developed and applied to PBHs. Ref. [172] proposed a method relating peaks in the curvature perturbation to peaks in the density field, with the caveat that the power spectrum is sufficiently narrow such that peaks of only one scale exist. Since we will here consider peaks in the power spectrum with a non-negligible width, we will not further consider the calculations presented in [172][1]. Ref. [183] proposed a similar method to [174], with two major differences. The first is that a top-hat window function is used instead of a Gaussian window function. The second difference is that, as well as extending peaks theory itself, [183] simultaneously attempted to account for the non-linear relation for the density contrast[2]. However, as discussed further in appendix A, the top-hat window function has significant drawbacks making it unsuitable for use

[1] A new paper, released at a similar date to this work, claims to have solved this issue [182], and applies peaks theory to $\Delta\zeta$, which is proportional to the (linear component of the) density contrast—meaning that it is similar to the peaks theory calculation considered here.

[2] The authors made use of an analytic relationship between the linear and non-linear fields. The expression is valid only at the centre of spherically symmetric peaks when a top-hat window function is used, and it is not clear this is a valid equation to use to represent the entire field.

in this paper without an additional cut-off. We will therefore focus on comparing previous "traditional" calculations with the YM calculation, which is expected to be an improvement on traditional peaks theory by correctly accounting for the initial scale and amplitude of perturbations in calculating the final PBH mass.

Ref. [173] discusses many complex points related to calculating the PBH abundance from the primordial power spectrum in detail. However, the calculation of the PBH abundance in [173] makes numerous simplifying assumptions, using peaks theory in a method similarly to that presented in [184]. The calculations used in this paper improve upon this by accounting for the non-linearity of the density contrast, a non-zero width of the power spectrum, and the dependence of PBH mass upon both the scale and amplitude of the perturbations from which it formed.

In the Press–Schechter formalism, the mass fraction is related to a probability distribution in the compaction function C by

$$\beta(R) = 2\int_{C_c}^{\infty} \mathrm{d}C \ \frac{m}{M_H} P(C), \tag{2.6}$$

where the compaction is a smoothed version of the density contrast δ (see eq. (C.4)). The probability density function is given by

$$P(C) = \frac{1}{\sqrt{2\pi}\sigma_0(R)} \exp\left(-\frac{C^2}{2\sigma_0(R)^2}\right), \tag{2.7}$$

and the mass ratio m/M_H takes into account the effect of critical collapse. In traditional peaks theory, the mass fraction is related to the number density of peaks, n, through

$$\beta(R) = (2\pi)^{\frac{3}{2}} R^3 \int_{C_c}^{\infty} \mathrm{d}C \ \frac{m}{M_H} \ n\left(\frac{C}{\sigma_0(R)}\right), \tag{2.8}$$

where the number density is a function of $\nu = C/\sigma_0$, given by [181]

$$n(\nu) = \frac{1}{3^{3/2}(2\pi)^2} \left(\frac{\sigma_1}{\sigma_0}\right)^3 \nu^3 \exp\left(-\frac{1}{2}\nu^2\right). \tag{2.9}$$

The modified peaks theory developed in [174] also has β related to n in a similar way to eq. (2.8), but with a factor of R^4 rather than R^3, i.e.

$$\beta(R) = (2\pi)^{\frac{3}{2}} R^4 \int_{C_c}^{\infty} \mathrm{d}C \; \frac{m}{M_H} \; n\left(\frac{C}{\sigma_0(R)}\right). \tag{2.10}$$

This is required to counteract an extra inverse spatial dimension in the number density, given by

$$n(\nu) = \frac{16\sqrt{2}}{3^{3/2}\pi^{5/2}} \frac{\sigma_{RR}}{\sigma_2\sqrt{1-\gamma_{0,2}^2}R^7} \left(\frac{\sigma_0}{\sigma_1}\right)^3 \alpha\nu^4 \exp\left(-\frac{1+\frac{16\sigma_0^2}{R^4\sigma_2^2}-\frac{8\sigma_0\gamma_{0,2}}{R^2\sigma_2}}{1-\gamma_{0,2}^2}\frac{\nu^2}{2}\right), \tag{2.11}$$

where σ_{RR}, $\gamma_{0,2}$, and α are related to the width parameters $\sigma_n(R)$ (see [174] for more details). These width parameters relate the probability density (in Press–Schechter) or number density (in the peaks theories) to the power spectrum through the relation

$$\sigma_n^2(R) = \int_0^{\infty} \frac{\mathrm{d}k}{k} \; k^{2n} \, \mathcal{P}_{\delta_R}(k), \tag{2.12}$$

where $\mathcal{P}_{\delta_R}(k)$ is the compaction power spectrum, related to the power spectrum for ζ through

$$\mathcal{P}_{\delta_R}(k) = \frac{16}{81}(kR)^4 W^2(k,R)\mathcal{P}_\zeta(k). \tag{2.13}$$

$W(k,R)$ is a window function applied to the power spectrum. In this paper, two window functions are considered: a real-space top-hat[3], given in Fourier-space by

$$W_{\mathrm{TH}}(k,R) = 3\frac{\sin(kR) - kR\cos(kR)}{(kR)^3}, \tag{2.14}$$

[3]It should be noted that we have modified the top-hat window function to remove a ringing effect at large-R (see appendix A for details).

and a Gaussian window function modified by a factor of 2 in the exponent as suggested in [177],

$$W_{\rm G}(k,R) = \exp\left(-\frac{(kR)^2}{4}\right). \tag{2.15}$$

It should be noted that, in the case of the modified Gaussian window function, the compaction referred to by C above is not technically the compaction, but is rather a "compaction-like" function. The compaction (or compaction-like function) is related to the PBH mass through the critical collapse equation,

$$m = KM_H(C - C_c)^\gamma, \tag{2.16}$$

where K, C_c, and γ are numerical factors that depend on the window function used to smooth the power spectrum, as well as the shape of the density perturbation [173, 178, 180]. The values $K \approx 3.3$, $C_c \approx 0.45$, and $\gamma \approx 0.36$ (commonly referred to as the Musco criteria) were derived for the top-hat window function [175, 185, 186], but are regularly used for other window functions. This has been highlighted in recent work, where different window functions cause a large deviation in the amplitude of power spectrum constraints, but this difference is not so significant if these numerical values are handled consistently for each window function [177]. We will take the values stated in [174]: $K = 4$ and $C_c = 0.55$ for the top-hat window function, and $K = 10$ and $C_c = 0.25$ for the modified Gaussian window function. For both window functions we take $\gamma = 0.36$.

In this paper we will frequently consider a power spectrum with a lognormal peak, as a simple parametrisation of a peaked power spectrum with a position and width that can be easily tuned. The form is

$$\mathcal{P}_\zeta = A\frac{1}{\sqrt{2\pi}\Delta}\exp\left(-\frac{\ln^2(k/k_p)}{2\Delta^2}\right) \tag{2.17}$$

with amplitude A, peak position k_p, and width Δ. This has been appropriately normalised such that the constraint on A becomes independent of Δ in the limit of a narrow peak, and it matches the delta function power spectrum $A\,\delta(\ln(k/k_p))$ in this limit. We show this later in table 2.2. The integral of this power spectrum over $\ln k$ is A, independent of the value of Δ. The width Δ is a free parameter, and we will normally choose two representative values for the width, $\Delta = 0.3$ as a narrow peak which results in a PBH mass distribution not very different from that due to a delta-function power spectrum, and $\Delta = 1$ as a broad peak which is roughly what one would expect if the inflaton field dynamics change over a time-scale of 1 e-folding during inflation. We note that such a peak should not be extrapolated to values of k very different in magnitude from k_p (and of course the power spectrum needs to match the quasi scale-invariant spectrum observed on CMB scales), but in practice we have checked that both the power spectrum constraints and the PBH mass distribution do not depend on the shape of the peak when sufficiently far from the peak position (where the power spectrum amplitude is significantly smaller than the peak value). We are therefore not concerned (for the values of Δ we focus on) that a lognormal peak exhibits a growth steeper than k^4 on scales far from the peak, even though this is the approximate maximum growth rate of the power spectrum in canonical single-field inflation [187–189]. A steeper growth can be achieved in e.g. multifield inflation [190, 191].

It is convenient to state the peak scale k_p in terms of the horizon mass it corresponds to, using the relation derived by comparing the temperature of the radiation within the horizon mass with the temperature at matter-radiation equality in [192],

$$M_H = \frac{1}{\sqrt{2}} M_{\rm eq} \left(\frac{g_{\rm eq}}{g}\right)^{1/6} \left(\frac{k_{\rm eq}}{k}\right)^2 \approx 17 \left(\frac{g_*}{10.75}\right)^{-\frac{1}{6}} \left(\frac{k}{10^6\ {\rm Mpc}^{-1}}\right)^{-2} {\rm M}_\odot, \qquad (2.18)$$

where g_* is the number of relativistic degrees of freedom. We define the horizon mass at the peak of the power spectrum as

$$M_{H,\mathcal{P}} = M_H(k_p). \qquad (2.19)$$

2.3 Variability of the mass distribution

2.3.1 Effect of the calculation method and window function

Constraints on the PBH abundance can be used to place constraints on the amplitude of the primordial power spectrum. If the black holes in the LIGO merger events are considered to be primordial in origin, a fit of the masses and number of events can be used to constrain the PBH mass distribution, and hence the power spectrum. Recent studies have shown that, in this case, $f_{\rm PBH}$ would have to lie between 10^{-2} and 10^{-3}, and would be closer to the lower of these two values [69, 193–195].

See, however, recent papers [72, 74, 196, 197] discussing the effect of interactions between binary and single PBHs, which suggests that a much larger value for $f_{\rm PBH}$ is possible provided that PBH binaries are sufficiently disrupted by other PBHs. Ref. [72] studied such 3-body interactions within extremely dense PBH clusters thought to form at high redshift when $f_{\rm PBH} \approx 1$, finding that the large majority of binaries in such clusters are expected to be disrupted, therefore not contributing to the merger rate observable today, implying that PBHs could make up the entirety of dark matter. Ref. [196] studied similar interactions within *Milky Way*-type haloes, finding that the coalescence times can change significantly due to the interactions, especially when the PBH abundance is low. In addition, the effect of initial clustering of PBHs (due to primordial non-Gaussianity) on the merger rate was studied in [71], showing that this results in large uncertainties in the merger rate. Combined, these papers cast significant doubt on constraints on the PBH abundance coming from the observed merger rate.

However, in order to proceed with the comparison presented here, we will assume that the constraints are valid. Therefore, for each method and window function described above, we determine the power spectrum amplitude required to generate an $f_{\rm PBH}$ in the range $10^{-2} < f_{\rm PBH} < 10^{-3}$, chosen as $f_{\rm PBH} = 2 \times 10^{-3}$. The resulting amplitudes are shown in table 2.1. It should be noted that these amplitudes are

defined for power spectrum peaks centered on the LIGO mass range, and would be significantly different on different scales. The full procedure for obtaining constraints on the power spectrum across all scales is described in section 2.4.

Table 2.1: Power spectrum amplitudes required to generate $f_{\rm PBH} = 2 \times 10^{-3}$, with masses in the LIGO range. The two window functions are Gaussian (G) and top-hat (TH), and the three methods are Press–Schechter (PS), traditional peaks theory (TP), and the modification to peaks theory calculated in [174] (YM). The modified peaks theory cannot be applied in the case of a delta function peak, or with the top-hat window function, so these combinations are not shown.

	Window Function, Method				
$\boldsymbol{\mathcal{P}}$ **peak**	G, PS	G, TP	G, YM	TH, PS	TH, TP
δ-function	3.21×10^{-3}	2.93×10^{-3}	N/A	3.47×10^{-3}	2.94×10^{-3}
$\Delta = 0.3$	4.14×10^{-3}	3.78×10^{-3}	3.55×10^{-3}	4.84×10^{-3}	4.13×10^{-3}
$\Delta = 1.0$	8.92×10^{-3}	8.14×10^{-3}	7.70×10^{-3}	1.11×10^{-2}	9.56×10^{-3}

It can be seen that, when being careful with the combination of the window function and the corresponding critical collapse values, all the amplitudes are of the same order. When changing either the method or the window function while keeping the other fixed, the difference in the required amplitude is $\lesssim 20\%$. The biggest difference when taking both the window function and the calculation method into account is $\sim 32\%$. We note that the maximum value of the power spectrum does not vary nearly as much when Δ changes as suggested by table 2.1 due to our parametrisation of the power spectrum definition (2.17). Choosing a different normalisation by leaving out the division by Δ would instead lead to a divergent value of the power spectrum amplitude in the limit $\Delta \to 0$, instead of a value which matches the delta function power spectrum.

We can also examine the amount of variability in the shapes of the mass distribution generated with different methods/window functions. The effect of changing the method is shown in fig. 2.1, for the Gaussian window function. The results for the top-hat case are similar. The mass distribution generated by a delta peak is shown in red, and the distribution for a lognormal peak with $\Delta = 1$ in blue, with both peaks centered on $M_{H,\mathcal{P}} = 4\ \mathrm{M}_\odot$ because this generates PBHs in the LIGO mass range. All the distributions are normalised to one, and correspond to

$f_{\rm PBH} = 2 \times 10^{-3}$. We find that the Press–Schechter (PS, solid) and peaks theory (TP, dashed) methods yield very similar results, while the modified peaks theory (YM, dotted) yields a marginally taller and narrower mass distribution.

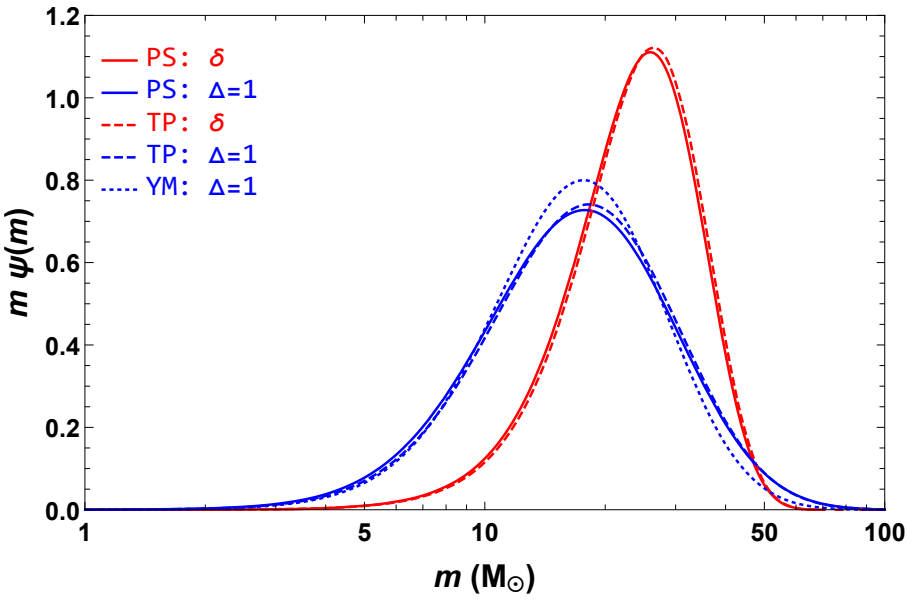

Figure 2.1: Difference between PBH mass distributions calculated using different methods, while keeping the window function fixed. The Gaussian window function is used in every case. The red curves are for the delta function peak in the power spectrum, and the blue curves are for the lognormal peak with $\Delta = 1$. The Press–Schechter (PS), traditional peaks (TP), and modified peaks (YM) methods are shown with solid, dashed, and dotted lines respectively. All lines have $f_{\rm PBH} = 2 \times 10^{-3}$.

Figure 2.2 shows the effect of changing the window function, again for the delta function (red) and $\Delta = 1$ lognormal (blue) cases, both with $M_{H,\mathcal{P}} = 4\ \mathrm{M}_\odot$. All the distributions have been calculated using traditional peaks theory. The distributions calculated using the Gaussian and top-hat window functions are shown as solid and dashed lines respectively. The distributions from the two window functions are similar, but with a small shift in the peak position. Additionally, it can be seen from figs. 2.1 and 2.2 that there is a shift in the peak mass between the delta function power spectrum, and the $\Delta = 1$ case. In the next section, we examine this shift in more detail for a range of power spectrum widths.

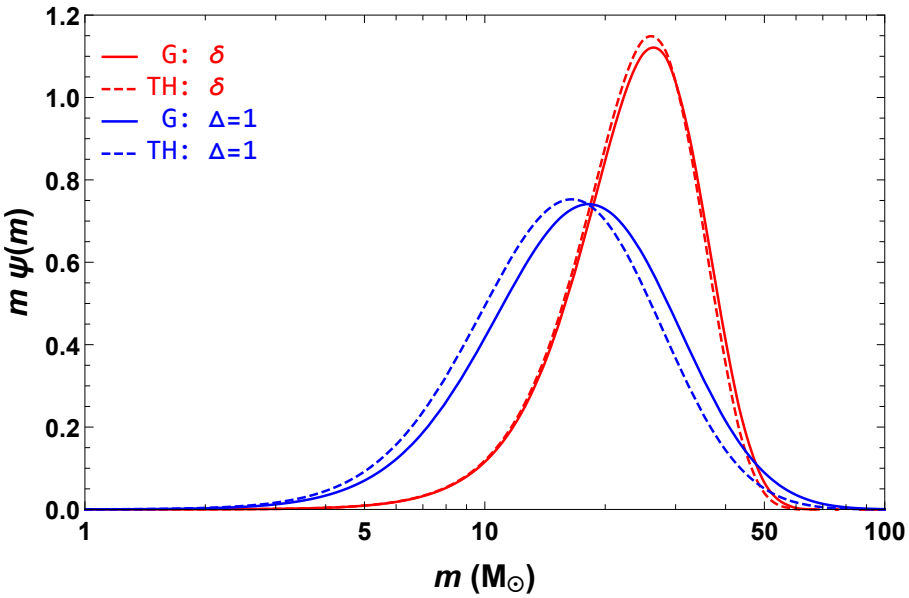

Figure 2.2: Difference between PBH mass distributions calculated using different window functions, using the traditional peaks theory (TP) method. The red and blue curves correspond to a delta function power spectrum and a lognormal with $\Delta = 1$ respectively. The solid and dashed lines are calculated using the Gaussian and top-hat window functions respectively. All lines have $f_{\rm PBH} = 2 \times 10^{-3}$.

We have shown that the different calculation methods result in an $\mathcal{O}(10\%)$ shift in the required power spectrum amplitude, and a small difference in the shape and position of the mass distribution. We expect the BBKS peaks method (TP) to provide a more accurate result than the Press–Schechter (PS) case, since it can be viewed as a generalisation and collapses to the PS case under certain assumptions [198], and that the modified version (YM) be better than TP, since it is a direct extension. Although the differences are small, they will become important in the future as experiments that can probe the PBH mass distribution become more accurate. For the remainder of this work, we will use the modified Gaussian window function in eq. (2.15) and the traditional peaks theory (TP) method. This allows comparisons between other works that use the TP method and the results in this paper, which can then be compared between the different methods based on the differences highlighted here.

2.3.2 Effect of the peak width Δ

As shown in section 2.3.1, the calculated mass distributions have a shift in the peak position which depends on the width of the power spectrum peak used. Additionally, we expect the width of the mass distribution to increase. We can demonstrate these effects by calculating the mass distributions for a range of values of Δ between zero (i.e. a delta function peak) and two. The result of these calculations is shown in fig. 2.3.

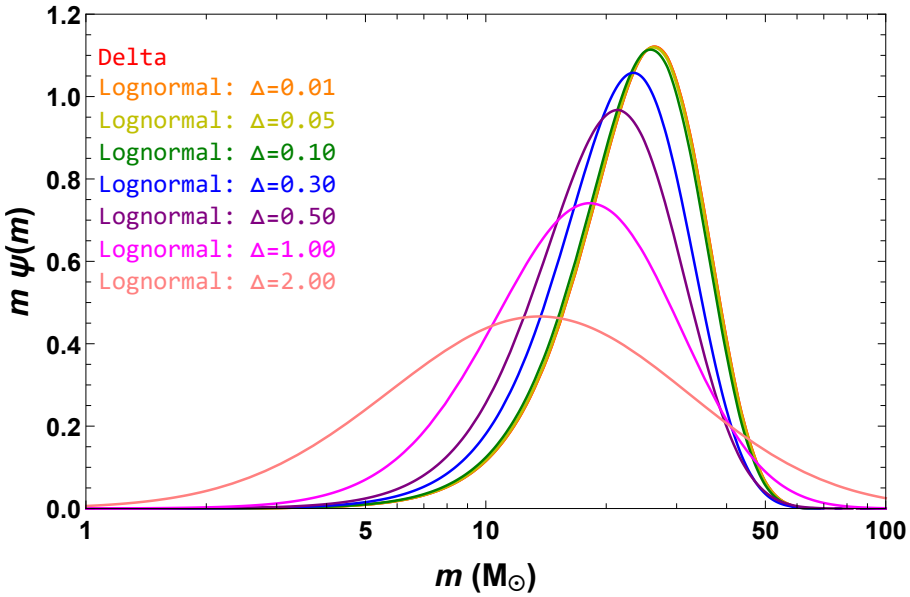

Figure 2.3: PBH mass distribution for different power spectrum peak widths Δ. The peak position k_p of the power spectrum is the same in every case, and corresponds to $M_{H,\mathcal{P}} = 4\ \mathrm{M}_\odot$. All lines have $f_{\mathrm{PBH}} = 2 \times 10^{-3}$. As Δ increases, the peak in the mass distribution shifts to smaller masses and spans a broader range of mass scales.

It is immediately apparent that, even for the unphysical choice of a delta function peak in the power spectrum, there is a minimum width in the mass distribution, associated with the critical collapse effect described in section 2.2. It can also be seen that for very narrow peaks in the primordial power spectrum, the resulting mass distribution hardly varies until $\Delta \gtrsim 0.1$. Beyond that point, the shift of the peak and the increased width become apparent. This means that whilst a monochromatic mass spectrum is unrealistic, studying a mass distribution with the minimum width due to critical collapse and a delta function power spectrum may be a good approximation to a physically realisable PBH mass distribution. The

increasing width is also obvious, and can be quantified by fitting a lognormal mass distribution (the shape expected for PBHs arising from a smooth, symmetric peak) to data generated from the curves, and comparing the widths of these lognormals. The lognormal mass distribution is given by

$$\psi(m) = \frac{1}{\sqrt{2\pi}\sigma_\psi m} \exp\left(-\frac{\ln^2(m/m_c)}{2\sigma_\psi^2}\right), \tag{2.20}$$

where m_c is the mean of the distribution and σ_ψ is the width (note the subscript to avoid confusion with the $\sigma_n(R)$ parameters appearing in section 2.2). The resulting lognormal parameters are shown in table 2.2, and show that, as expected, the width of the calculated mass distribution increases with the peak width, as well as the amplitude required to keep $f_{\rm PBH}$ fixed. This minimum width appears to be much larger than is required in order for PBH decay to result in a sufficiently rapid transition from an early matter dominated era (caused by low mass PBHs) to radiation domination to generate an observable stochastic background of gravitational waves [199].

Table 2.2: Comparison of the amplitude A required to generate $f_{\rm PBH} = 2 \times 10^{-3}$, the ratio of the mean PBH mass m_c to the power spectrum peak mass $M_{H,\mathcal{P}}$, and the mass distribution width σ_ψ for different power spectrum peak widths Δ.

$\boldsymbol{\Delta}$	$\boldsymbol{A}$	$\boldsymbol{m_c/M_{H,\mathcal{P}}}$	$\boldsymbol{\sigma_\psi}$
0 (Delta)	2.93×10^{-3}	6.21	0.374
0.01	2.94×10^{-3}	6.21	0.374
0.05	2.96×10^{-3}	6.17	0.375
0.10	3.04×10^{-3}	6.09	0.377
0.30	3.78×10^{-3}	5.52	0.395
0.50	4.89×10^{-3}	5.07	0.430
1.00	8.14×10^{-3}	4.39	0.553
2.00	1.51×10^{-2}	3.35	0.864

A noteworthy point here is that the typical mass of a PBH is actually significantly larger than the horizon mass corresponding to the scale at which the power spectrum peaks, $m_c/M_{H,\mathcal{P}} > 1$. At first glance, this statement may seem to be in disagreement with previous works where the expected PBH mass has been shown to be smaller

than the horizon mass at re-entry. Physically, this apparent discrepancy is due to the fact that, if there is a narrow peak in the ζ power spectrum at a scale k_p, the resultant perturbations will, on average, have a significantly larger characteristic scale r_m. In the calculation presented here, this manifests itself in the fact that the variance $\sigma_0^2(R)$ peaks at a larger value of R than that corresponding to the scale k_p (as calculated in [180] for example). Thus, the final mass of PBHs is smaller than the horizon mass corresponding to r_m, but larger than the horizon mass corresponding to k_p. The important conclusion drawn from this is that constraints on the PBH abundance for a given mass of PBH correspond to constraints on the primordial power spectrum at a larger value of k than have previously been calculated.

Now we have a clear picture of how the different method and window function choices affect the mass distribution ψ and the amplitude required to generate a fixed $f_{\rm PBH}$, we can calculate the constraints on the power spectrum from PBHs, being careful about the consistency of our window function and critical collapse choices. We show the procedure for obtaining these constraints, and the final constraint plots, in the next section.

2.4 The constraints on the power spectrum

2.4.1 Relevant constraints and how they are calculated

Whilst calculating the PBH abundance with different methods has a huge effect on the abundance and mass distribution, we have shown that the resultant uncertainty in constraints on the power spectrum is relatively small. We will now consider how observational limits on the PBH abundance, as well as a swathe of other observational probes, constrain the amplitude of the primordial power spectrum. The key additional constraints on small scales come from cosmic μ-distortions [200] and a stochastic background of gravitational waves, which could be generated with a large amplitude due to the non-linear coupling between the scalar and tensor perturbations around the time of horizon entry [201, 202]. The calculation of many of these

constraints follows closely the procedure presented in [187], and we therefore relegate the details to appendix B. However, we describe the constraints from PBHs in detail here, and we also highlight that constraints from PTAs have been updated to use the improved analysis of the NANOGrav 11 year data set [203]. There are additional small-scale constraints on the power spectrum, including for example those from y-distortions [204, 205], 21 cm observations [206–210] and the non-detection of ultra-compact minihaloes [160, 211–214]. We do not display the former because the combination of CMB constraints and μ-distortion constraints are more competitive on commensurate scales, and we do not display either of the latter because they depend on the dark matter model. Big Bang nucleosynthesis constraints are discussed in e.g. [215–217].

2.4.2 Constraints due to the gravitational wave background

Large amplitude scalar perturbations re-entering the horizon after inflation induce gravitational waves as a second-order effect. These contribute to the stochastic gravitational background, which pulsar timing arrays (PTAs) are trying to detect and/or constrain by looking for global changes in the time of arrival of pulses from a population of millisecond pulsars over a period of $\mathcal{O}(10)$ years. Details of the calculation of the GW power spectrum are contained in appendix B.2.

Translating this power spectrum to $\Omega_{\rm GW}h^2$ with eq. (B.4), we can then compare the predicted signal with PTA constraints from the NANOGrav 11 year data set.[4] We choose this data set because the new analysis takes errors in the modelling of the solar system ephemeris into account. This can have a large effect on the constraints which will need to be factored into the previous NANOGrav 9 year constraints [218], as well as those from other arrays such as the European Pulsar Timing Array (EPTA) [219] which have previously been used to constrain the primordial power spectrum with induced gravitational waves. Those constraints should now be revised upwards,

[4]During the refereeing process, NANOGrav released their 12.5 year dataset [127] which showed *possible* evidence for a signal due to a stochastic gravitational wave background. This is unconfirmed, but understanding the origin of this signal could have significant implications for the induced gravitational wave constraints discussed in this work.

but the analysis would need to be redone in each case to quantify by exactly how much. Based on the current analyses, the constraint on the characteristic strain h_c improves by a factor of a third at the frequency of the tightest constraint, and improves by up to a factor of 5 at the highest frequencies between the 9 year and 11 year datasets. The resulting improvement on the primordial power spectrum constraint is shown in figure B.1 of appendix B.3. Since the NANOGrav data set has pulsar timing data for 11 years of observations, it does not extend to quite as large scales as does the EPTA data, which is from 18 years of observations. This means that our constraints do not span as wide a range of scales (and hence PBH masses) as previous constraints in the literature show, but the constraints we do show are more robust to errors in solar system ephemeris modelling. We also avoid confusion over different analyses from different data sets, and are able to use the free spectrum constraints on $\Omega_{\rm GW}h^2$ consistently throughout.

These constraints (taken from the bottom panel of fig. 3 in [203]) are the 2-σ constraints derived as a function of frequency so as to represent the sensitivity to monochromatic signals. This means that we will construct our constraints based on finding the limiting amplitude of the lognormal power spectrum to which the NANOGrav constraints would be sensitive. One could do a more sophisticated analysis, taking into account the fact that confidence in a detection would become even stronger if there are also weaker detections of a given signal on larger or smaller frequencies than where the strongest detection would come. We choose to show just the 2-σ constraints for clarity. We convert from frequency to scale with $k = 2\pi c/f$ and then find the minimum value of A for which $\Omega_{\rm GW,NG}h^2 = \Omega_{\rm GW,signal}h^2$, i.e.

$$A_{\rm constraint} = \text{Min}\left(\sqrt{\frac{\Omega_{\rm GW,NG}(k)h^2}{\Omega_{\rm GW,signal}(k,k_p)h^2}}\right) \tag{2.21}$$

for each k_p. The minimum value of A for each k_p is found by scanning over all values of k for which NANOGrav has sensitivity. We plot the results in figs. 2.5 and 2.6 for $\Delta = 0.3$ and $\Delta = 1$, where again to be clear, the constraint on $\mathcal{P}_\mathcal{R}$ at a given

k represents the maximum amplitude A for a lognormal power spectrum centered at $k = k_p$ such that the induced second-order gravitational waves would not be in conflict with the PTA constraints from the NANOGrav 11 year data set.

2.4.3 Constraints from PBHs

Constraints on primordial black holes are normally presented in terms of either $f_{\rm PBH}$ or the mass fraction β, so a method is required to relate these to the power spectrum amplitude. A relation between $f_{\rm PBH}$ (or equivalently $\Omega_{\rm PBH}$) and A is complicated by the fact that the redshifting factor in eq. (2.1) means that the required amplitude to generate a fixed $f_{\rm PBH}$ varies with the peak positions (as demonstrated in section 2.3.1). In general, the best way to overcome this would be to produce a relation for A as a function of both $f_{\rm PBH}$ and the relevant mass scale. However, this is computationally expensive, and so a simplified approach is necessary. We can find an approximation by relating the power spectrum amplitude to a parameter that does not vary with the peak position, which we achieve by modifying eq. (2.1), adjusting the redshift factor by introducing a new scale R_*, such that

$$\Omega_{\rm PBH*} = \int \mathrm{d}(\ln R)\, \frac{R_*}{R} \beta(R). \tag{2.22}$$

If R_* is chosen to be close enough to the peak scale in the power spectrum, then the relation between this quantity and the power spectrum amplitude will be independent of the peak position. This quantity cannot be treated exactly as the abundance, because the abundance is calculated in the super-horizon regime before PBHs form, whereas this is at some later time, corresponding to when the horizon scale is R_*. This quantity can be related to the constraints for PBHs using

$$\Omega_{\rm PBH} = \frac{R_{\rm eq}}{R_*} \Omega_{\rm PBH*}. \tag{2.23}$$

The relation between the power spectrum amplitude and $\Omega_{\mathrm{PBH}*}$ for all three methods is shown in fig. 2.4 for the $\Delta = 1$ (left) and $\Delta = 0.3$ (right) cases. The modified Gaussian window function is used in every case. It can be seen that there is a shift in the amplitude required between the methods, as was observed earlier. However, comparing the scale of changes to the power spectrum amplitude between the CMB value of 10^{-9} and these values, the differences are unimportant. For the constraint plots shown in figs. 2.5 and 2.6, the traditional peaks theory method (TP) is used.

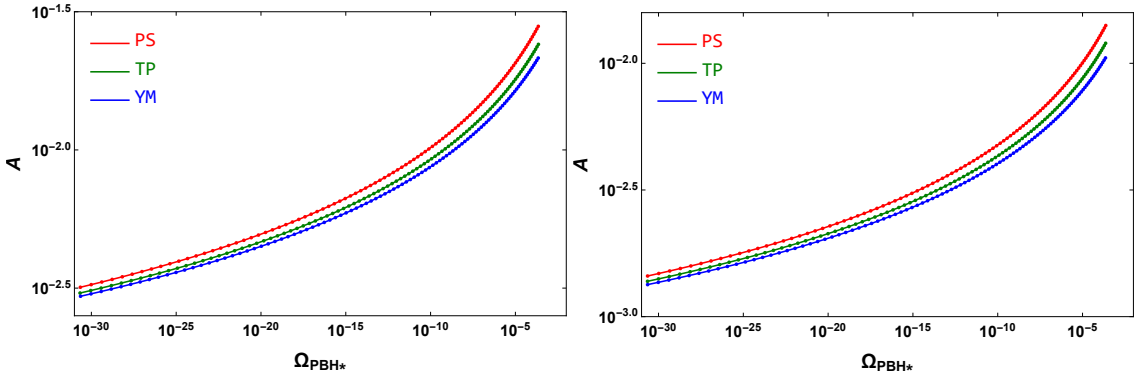

Figure 2.4: Relation between power spectrum amplitude A and $\Omega_{\mathrm{PBH}*}$ for the three methods. The power spectrum peak widths are $\Delta = 1$ (left) and $\Delta = 0.3$ (right). All lines use the Gaussian window function.

To obtain constraints on the power spectrum, $\Omega_{\mathrm{PBH}*}$ must be related to constraints on either β or f_{PBH}. We will use the PBH constraints stated in [43] for β', which is a version of the mass fraction β with common parameters normalised out. These constraints are calculated assuming that all the PBHs form at the same time (or equivalently, the same scale R), but it is possible to relate the constraints to $\Omega_{\mathrm{PBH}*}$, and hence determine the constraints on the amplitude for the calculation used throughout this paper, where PBHs form over a range of different scales. We obtain this relation from eqs. (6) and (8) from [43] (reproduced here for clarity):

$$\beta(m_c) = 7.06 \times 10^{-18} \gamma^{-1/2} \left(\frac{h}{0.67}\right)^2 \left(\frac{g_{*,i}}{106.75}\right)^{1/2} \left(\frac{m_c}{10^{15}\ \mathrm{g}}\right)^{1/2} \Omega_{\mathrm{PBH}}(m_c), \tag{2.24}$$

$$\beta'(m_c) = \gamma^{1/2} \left(\frac{h}{0.67}\right)^{-2} \left(\frac{g_{*,i}}{106.75}\right)^{-1/2} \beta(m_c), \tag{2.25}$$

where the monochromatic PBH mass M in [43] has been substituted for the mean lognormal mass m_c (the constraints do not change significantly when considering a reasonably narrow PBH mass distribution [76, 77]). It can immediately be seen that, combining eqs. (2.24) and (2.25),

$$\beta'(m_c) = 7.06 \times 10^{-18} \left(\frac{m_c}{10^{15}\ \mathrm{g}}\right)^{1/2} \Omega_{\mathrm{PBH}}(m_c). \tag{2.26}$$

Since solar mass PBHs are of special interest, it is sensible to rescale the mass fraction to be in terms of solar masses, giving

$$\beta'(m_c) = 7.06 \times 10^{-18} \left(2 \times 10^{18} \frac{m_c}{\mathrm{M}_\odot}\right)^{1/2} \Omega_{\mathrm{PBH}}(m_c) \tag{2.27}$$

$$= 10^{-8} \left(\frac{m_c}{\mathrm{M}_\odot}\right)^{1/2} \Omega_{\mathrm{PBH}}(m_c). \tag{2.28}$$

Inverting this relation gives Ω_{PBH} as a function of m_c in solar masses,

$$\Omega_{\mathrm{PBH}}(m_c) = 10^8 \left(\frac{m_c}{\mathrm{M}_\odot}\right)^{-1/2} \beta'(m_c). \tag{2.29}$$

We can then be relate this to the quantity $\Omega_{\mathrm{PBH}*}$ using eq. (2.23) to give

$$\Omega_{\mathrm{PBH}*}(m_c) = 10^8 \frac{R_*}{R_{\mathrm{eq}}} \left(\frac{m_c}{\mathrm{M}_\odot}\right)^{-1/2} \beta'(m_c). \tag{2.30}$$

For convenience, we have chosen R_* such that the corresponding mass scale M_* is approximately m_c. Therefore,

$$\Omega_{\mathrm{PBH}*}(m_c) = 10^8 \left(\frac{m_c}{M_{\mathrm{eq}}}\right)^{1/2} \left(\frac{m_c}{\mathrm{M}_\odot}\right)^{-1/2} \beta'(m_c) \tag{2.31}$$

$$= 10^8 \left(\frac{M_{\mathrm{eq}}}{\mathrm{M}_\odot}\right)^{-1/2} \beta'(m_c). \tag{2.32}$$

Substituting in the value of the horizon mass at matter-radiation equality, $M_{\rm eq} = 2.8 \times 10^{17}\ {\rm M}_\odot$, the relation becomes

$$\Omega_{\rm PBH*}(m_c) \approx 0.2\ \beta'(m_c). \tag{2.33}$$

Recent papers [169–173] have discussed the effect of the non-linear relation between the curvature perturbation ζ and the density contrast δ on the PBH abundance. The point is that, even if the level of primordial non-Gaussianity of ζ is taken to be zero, δ will not have a Gaussian distribution, and subsequently nor will the compaction. The non-linearity is difficult to account for, especially if window functions other than a top-hat are considered. This is discussed in some detail in appendix C, with the conclusion that constraints on the power spectrum will be approximately 1.98 times weaker once the non-linearity is included in the calculation. We include this factor in the PBH lines in figs. 2.5 and 2.6.

By applying the method described in this section, we are taking into account the effects of critical collapse (making sure it is treated consistently with the choice of window function), the shift between the PBH mass and the peak scale k_p, and the non-linear relation between ζ and δ. This is the first time that all of these effects have been captured simultaneously.

2.4.4 Summarising all the constraints

In fig. 2.5 we put together the key observational constraints to show the principal current constraints on the primordial power spectrum. The power spectrum has been accurately measured on large scales whilst PBHs constrain—albeit weakly—a far larger range of scales. We do not show PBH constraints on masses close to matter-radiation equality because we always assume PBHs form during radiation domination, and the smallest scale constrained corresponds to a PBH with $m_c \sim 10^{-24}\ {\rm M}_\odot$, which evaporates around the time of Big Bang nucleosynthesis.

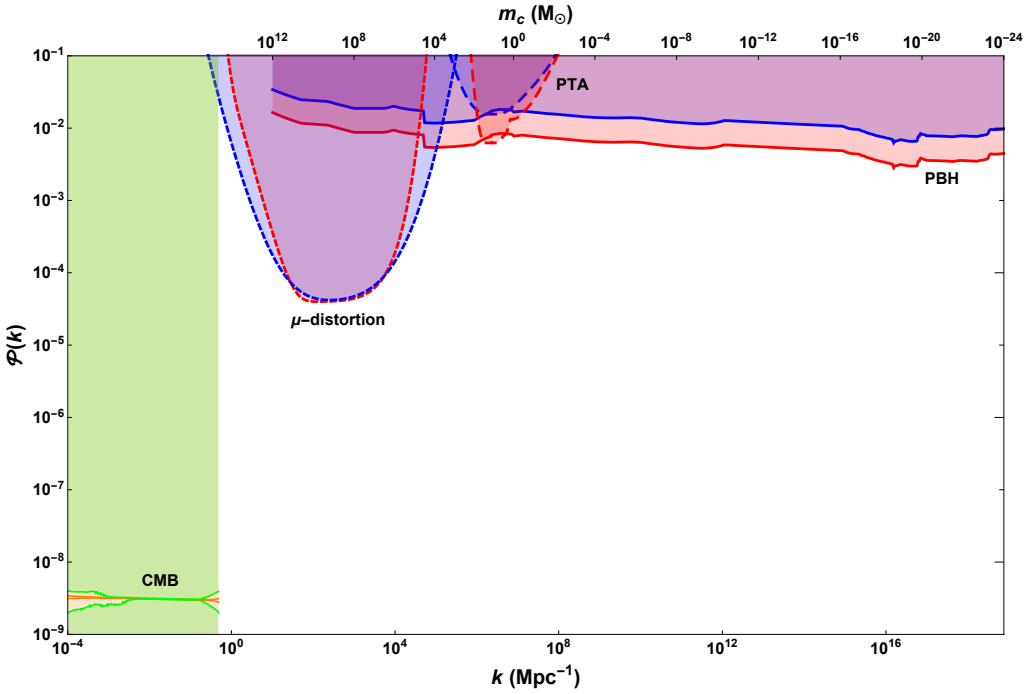

Figure 2.5: Current constraints on the power spectrum amplitude from PBH, PTA, and μ-distortion sources, as well as the measured one- and three-sigma constraints from the CMB. The constraints for $\Delta = 0.3$ (which are tighter for the PBH constraints, and narrower for the other constraints) are shown in red, and the constraints for $\Delta = 1$ are shown in blue. The PBH, PTA, and μ-distortion constraints are shown with solid, long-dashed, and short-dashed lines respectively.

By coincidence the PTA measurements constrain the power spectrum amplitude to almost the same amplitude as the non-detection of PBHs, meaning that there is a potential tension between the PTA bounds and any claim that LIGO detected PBHs (see fig. 2.5). This has been studied by various groups [187, 220–228], with no consensus reached on how severe the tension is. For example, [224] claim that $f_{\rm PBH} < 10^{-6}$ over a significant range of PBH masses and the power spectrum constraint plots in [187] appear to show a significant tension. The impact of the PBH density profile was studied in depth in [173] but the PTA constraint was not varied to reflect changes in the shape of the primordial power spectrum. By making a careful study of the power spectrum amplitude required to generate PBHs, including the important reduction in the PBH constraining power due to the non-linear relation

between ζ and δ, and using improved NANOGrav constraints, we have shown that there is no significant tension between generating LIGO mass PBHs and the PTA constraints.

We note that the slight overlap between the PBH and PTA constraint lines is not significant given the remaining $\mathcal{O}(10\%)$ uncertainty in the amplitude of the PBH constraint, and that there should also be about an $\mathcal{O}(10\%)$ reduction in the PBH line at about the $M_\odot$ scale caused by the reduction in the equation-of-state parameter during the QCD transition. See [45] for further discussion, and [229] for extensions to other masses where there is a smaller reduction in pressure within standard model physics. A study of non-standard expansion histories (such as an early matter dominated epoch) are beyond the scope of this paper [230]. Nonetheless, because the PBH amplitude only depends very weakly on the value of $f_{\rm PBH}$ it is clear that the PTA collaborations should be very close to detecting a stochastic gravitational wave background even if only one of the compact objects which LIGO has detected was a PBH, for example the secondary mass object in the recently detected event which falls into the mass gap between neutron stars and astrophysical black holes [34]. It seems plausible that the associated stochastic background could be detectable with current PTA data if a dedicated search was made by using specific GW templates generated by power spectra that cause LIGO mass PBHs to form.

The cosmic μ-distortion places an upper limit on the maximum PBH mass which can be generated by the collapse of large amplitude perturbations shortly after horizon re-entry. The maximum mass decreases as the power spectrum width Δ increases, but even for a narrow peak with $\Delta = 0.3$ the initial PBH mass cannot be much greater than 10^4 $M_\odot$, which is much smaller than the supermassive BHs seen in the centre of most galaxies even at high redshift, with masses 10^6–10^9 $M_\odot$, whose origin remains a mystery. However, such large PBHs could still act as a seed to the SMBHs [207], and the constraints can be evaded if the initial perturbations are extremely non-Gaussian [231] although one then needs to evade the strong Planck constraints on dark matter isocurvature modes [232, 233]. For even broader power

spectra the μ-distortion constraints rule out an ever greater range of PBH masses, and for $\Delta = 2$ they extend as far as the peak PTA constraint and thereby even rule out LIGO mass PBHs. Since such a wide peak in the primordial power spectrum provides the preferred PBH mass distribution width when fitting to LIGO data, it appears that the μ-distortions may surprisingly provide a stronger constraint on models in which all LIGO black holes are PBHs than the PTA constraints. Of course this conclusion may also depend on the assumed shape of the power spectrum peak.

Future constraints from μ-distortions and the gravitational wave background will significantly affect the PBH landscape. To examine the maximum extent of these future constraints, we calculate the PBH lines in the case that zero PBHs form in the observable universe. This is done using the method described in [234], particularly eq. (7) of that paper, but with β replaced with the $\Omega_{\mathrm{PBH}*}$ parameter used in this paper. For reasons summarised in [234], these extreme constraints might actually apply to the case of evaporated PBHs. Extremely tight constraints on f_{PBH} for $M_{\mathrm{PBH}} \gtrsim 10^{-6}$ $\mathrm{M}_\odot$ are also possible if the majority of dark matter consists of "standard" WIMPs [153, 235–239]. We show the tightest possible PBH constraints in fig. 2.6, as well as future μ-distortion constraints from a detector like the Primordial Inflation Explorer (PIXIE) [240], and future gravitational wave background constraints from the Square Kilometre Array (SKA), the Laser Interferometer Space Antenna (LISA), and the Einstein Telescope (ET)[5]. The SKA constraints are derived from the sensitivity curve calculated in [241], the LISA constraints are derived from the most optimistic sensitivity curve in fig. 1 of [242], and the ET constraints are derived from fig. 13 of [145].

It can be seen that the SKA constraints are so tight that a non-detection will indicate that no PBHs can exist in the LIGO range of masses, and hence that the LIGO merger events cannot possibly be explained with a primordial origin. Additionally, the combined effect of the μ-distortion, SKA, LISA, and ET constraints

[5]Note that free spectrum sensitivity curves, as were used to calculate the PTA constraints, are not available for the future detectors SKA, LISA, and ET, so instead we have used the sensitivity curves that are derived assuming a power-law for the gravitational wave frequency spectrum.

removes the possibility of any PBHs existing over an extremely broad range of masses in the case of a non-detection, leaving only the space below $\sim 10^{-22}$ $M_\odot$, and two small pockets at $\sim 10^{-17}$–10^{-14} $M_\odot$ and $\sim 10^{-6}$–10^{-3} $M_\odot$.

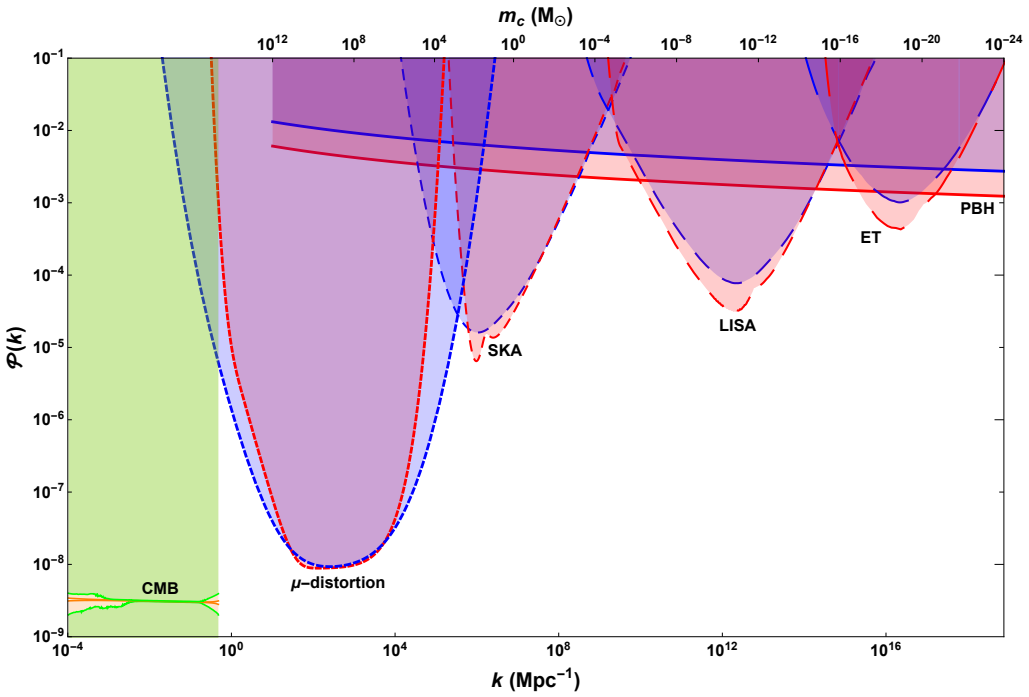

Figure 2.6: Future constraints on the power spectrum amplitude from PBH, gravitational wave background, and μ-distortion sources, as well as the measured one- and three-sigma values from the CMB. The PBH curves indicate the amplitude required to generate only a single PBH in the observable universe. The constraints for $\Delta = 0.3$ (which are tighter for the PBH constraints, and narrower for the other constraints) are shown in red, and the constraints for $\Delta = 1$ are shown in blue. The PBH constraints are shown with a solid line, and the ET, LISA, SKA, and μ-distortion constraints are shown with longest to shortest dashes respectively.

2.5 Conclusions

We have made the first detailed analysis of how the PBH mass distribution shape and amplitude varies between three different techniques to calculate the primordial mass distribution: Press–Schechter, traditional peaks theory and a newly developed peaks theory variation. We also consider two choices of the window function, a real-space top-hat and a modified Gaussian. We show that the amplitude of the primordial power spectrum only varies by $\mathcal{O}(10\%)$ for different choices, far smaller than may have been expected based on the large range of values of the power spectrum

amplitude considered in the literature. A substantial variation remains depending on the shape of the peak in the primordial power spectrum, but this reflects a change in the physical theory rather than a change in methodology. The results are summarised in table 2.1 while fig. 2.1 shows that the mass distribution shape hardly changes depending on the calculation technique. These differences, while not significant now, will be important for future data that probe the PBH mass distribution accurately, at which point an improvement of the TP method, such as the Young–Musso technique, should be used. We also show that the PBH mass distribution becomes broader as the power spectrum peak becomes broader, as highlighted in fig. 2.3. In the limit of a narrow lognormal peak ($\Delta \lesssim 0.3$) the mass distribution tends to a constant width which is set by critical collapse, making a peak of this width a well-motivated choice.

We have also calculated robust constraints on the primordial power spectrum from PBHs, taking into account the effects of critical collapse and the non-linear relation between ζ and δ, as well as the choice of window function and the relation between the PBH mass scale and the peak power spectrum scale. This leads to tighter constraints that are shifted to different values of k compared to those presented in [43]. We show a summary of all of the key bounds on the amplitude of the primordial power spectrum in fig. 2.5. We stress that all the constraints must be recalculated when the shape of the primordial power spectrum peak is varied, and in the figure we choose $\Delta = 0.3$ as a representative narrow peak and $\Delta = 1$ as a broader peak. In both cases the PTA constraints (we use a recently improved data set from the NANOGrav collaboration) are almost identical to those from PBHs in the mass range that LIGO also probes. This interesting coincidence means that it is premature to rule out the possibility that LIGO detected PBHs that formed from large amplitude density perturbations during radiation domination, but if that is the case then there is a realistic hope that the PTA measurements will detect a stochastic background of gravitational waves in the near future and a dedicated analysis should be made. We note that the non-linear relation between ζ and δ

weakens the PBH constraints by about a factor of 2, and had we not taken this into account (and normally it is not taken into account) we would have erroneously concluded that the PTA constraints do not come close to ruling out the formation of LIGO mass PBHs. However, we caution that if all BH binaries detected by LIGO were due to PBHs then the PBH mass distribution should be so broad ($\sigma_\psi \approx 0.8$ corresponding to $\Delta = 2$) that the cosmic μ-distortion constraints spread to relatively small masses and alternative shapes of the primordial power spectrum which are more "top-hat"-like than the lognormal power spectrum studied here should be considered.

In fig. 2.6 we show constraints on the primordial power spectrum that could be achieved in the foreseeable future (assuming there is no detection) from a PIXIE-like experiment measuring μ-distortions and searches for a stochastic background of gravitational waves. The gravitational wave constraints show SKA constraints on pulsar timings, plus LISA and ET constraints. The PBH constraints show the amplitude required to generate a single PBH within the observable universe, provided that they form from Gaussian-distributed perturbations entering the horizon during radiation domination. This shows that apart from two narrow mass ranges around 10^{-4} $M_\odot$ and 10^{-16} $M_\odot$, there will be no remaining window for un-evaporated PBHs to exist today.

Acknowledgements

CB thanks Qing-Guo Huang for correspondence, and we thank Eiichiro Komatsu for useful comments on a draft of this paper. AG is funded by a Royal Society Studentship by means of a Royal Society Enhancement Award. CB acknowledges support from the Science and Technology Facilities Council [grant number ST/T000473/1]. PC acknowledges support from the Science and Technology Facilities Council [grant number ST/N504452/1]. SY is supported by a Humboldt Research Fellowship for Postdoctoral researchers.

Chapter 3

An accurate model for the primordial black hole mass distribution from a peak in the power spectrum

Andrew D. Gow[1], Christian T. Byrnes[1], and Alex Hall[2]

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

2) Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, United Kingdom

We examine the shape of the primordial black hole mass distribution arising from a peak in the primordial power spectrum. In light of improvements to the modelling, we revisit the claim that the effects of critical collapse produce a distribution that is not described by the commonly assumed lognormal, showing that this conclusion remains valid. We propose some alternative models that may better describe the shape, both for the case of narrow power spectrum peaks where critical collapse determines a minimum width of the mass distribution, and for much broader peaks where the peak shape is significant. We highlight the skew-lognormal and a gener-

alised model motivated by the physics of critical collapse as the best of these possible alternatives. These models can be used as an accurate and fast approximation to the numerically calculated mass distribution, allowing for efficient implementation in an MCMC analysis. We advocate the use of one of these two models instead of the lognormal with sufficiently accurate data, such as future LIGO–Virgo observations, or when considering strongly mass dependent constraints on the PBH abundance.

3.1 Introduction

Since the idea of primordial black holes (PBHs) was first postulated half a century ago [24–26], a lot of progress has been made in studying constraints on their abundance as well as possible signs that they have been detected. Until quite recently, most constraints on the PBH abundance assumed a monochromatic mass distribution which has the advantage of simplicity, since this is the unique case where a constraint at any given mass can be made without considering the constraints on other, similar masses. See e.g. [43, 155, 156, 243, 244] for reviews. However, the phenomenon of critical collapse means that a range of PBH masses are generated from large amplitude perturbations re-entering the horizon even if the perturbation spectrum has power at only one wavenumber [175, 185, 245, 246], due to the spread in amplitudes of modes at that scale. Therefore, as one would intuitively expect, a monochromatic mass distribution is not physically realistic, no matter how narrowly peaked the primordial power spectrum might be[1].

While the community was focused on making order-of-magnitude constraints to the PBH abundance and simple "yes/no" answers to whether PBHs of a given mass could constitute the entirety of dark matter, the approximation of a monochromatic mass distribution was adequate. However, in recent times there has been a vigorous debate about exactly what fraction of the dark matter could be contained in PBHs with a mass of order the solar mass, for example to fit to lensing surveys or the

[1] In practise there is also a limit to how narrow the primordial power spectrum can be, with the limits depending on the model of inflation, see e.g. [187–191, 247, 248].

LIGO–Virgo detection of gravitational waves. Many, but not all, constraints allowed an order one fraction of PBHs to be in dark matter. See e.g. [153,154] for very recent reviews. These constraints come from a wide range of methods as well as probing a wide range of redshifts, and there is the possibility that accretion makes the constraints time dependent in a mass dependent manner [53,195,249]. Finally there are some hints that LIGO–Virgo may have detected PBHs, for example due to the low spin of most of the detected events [58,60] as well as some objects which fall into or close to the lower and upper mass gaps commonly considered for astrophysical formation channels [32–34, 197, 250], although these can be explained with specific astrophysical models [251, 252].

For all of the above reasons, it has now become commonplace to consider extended mass distributions. By far the most commonly considered case is the lognormal mass distribution, and constraints for this distribution were made by e.g. [76,77] (see [253] for the first related reference to this mass distribution in the PBH context). Broad mass distributions, such as a power law, or one with a spike at around one solar mass motivated by the QCD transition have also been considered [45,254], but in this paper we will focus on the more commonly studied case of a mass distribution generated by a single symmetric peak in the primordial power spectrum.

The lognormal mass distribution is frequently applied irrespective of its width, either in the form of priors allowing narrow widths (see e.g. [255–257]), or in the case of explicitly considering a very narrow case (see e.g. [199]). However, it has been known for almost 25 years that for sufficiently narrow peaks in the power spectrum, the effects of critical collapse dominate. This creates a minimum width for the mass distribution, as discussed in [258]; table II in that work gives a minimum lognormal width of 0.37 based on a simple least squares fit to the numerical mass distribution calculated for a delta function peak in the power spectrum[2]. Additionally, it is known that critical collapse causes the mass distribution shape to be significantly

[2] A different value of 0.26 was stated in [76], although this did not involve a full calculation of the mass distribution from a power spectrum peak, instead applying a method of moments approach to compare the lognormal with the critical collapse motivated shape in [245].

non-lognormal [185, 245]. A large amount of work has been carried out on the mass distribution calculation since this deviation was first demonstrated, including the integration over all formation epochs mentioned but not pursued in [185] (see [45], and e.g. [174, 177, 183, 258–261] for further discussion of the calculation). This leads to two questions: does the conclusion of non-lognormality for narrow power spectrum peaks still hold and if so, is there a model for the PBH mass distribution that can accurately describe its behaviour for a broad range of power spectrum peak widths?

While the most rigorous choice is carrying out the full calculation of the mass distribution from the power spectrum, this can be computationally expensive, making it unsuitable for e.g. Bayesian model selection calculations. Therefore, it is necessary to use models which allow an approximate capturing of the numerical mass distribution. Different constraints require the mass distribution to be narrow or broad, so it is essential to use a model that describes the numerical mass distribution for all these cases. In the following, we show that the lognormal assumption does indeed still break down for the narrowest widths, and propose some alternative models that can achieve a better fit over a large range of widths.

3.2 Modelling the PBH mass distribution

3.2.1 The numerical mass distribution

In order to test the validity of the lognormal mass distribution, we need a robust method of determining the PBH mass distribution corresponding to a particular peak in the primordial power spectrum. For this purpose, we use an accurate model for PBH formation described in Gow et al. [258]. The procedure is to first relate the power spectrum peak to the PBH abundance $\Omega_{\rm PBH}(m)$, and then determine the mass distribution, given by

$$\psi(m) = \frac{1}{\Omega_{\rm PBH}} \frac{{\rm d}\Omega_{\rm PBH}}{{\rm d}m}. \tag{3.1}$$

This is a probability distribution, and hence satisfies the condition $\int \mathrm{d}m\ \psi(m) = 1$, as will all the models we consider later.

The procedure to obtain the mass distribution is described in detail in section 2 of [258]. It incorporates the effects of critical collapse, and is robust to modelling choices at the 10% level. In this paper, we will use the traditional peaks theory method with the modified Gaussian window function stated in eq. (15) of [258, arXiv version]. We additionally choose the same lognormal form for the primordial power spectrum peak,

$$\mathcal{P}_\zeta = A\frac{1}{\sqrt{2\pi}\Delta}\exp\left(-\frac{\ln^2(k/k_p)}{2\Delta^2}\right), \tag{3.2}$$

which has a peak at k_p and a width Δ. The normalisation is chosen such that $\int \frac{\mathrm{d}k}{k}\ \mathcal{P}_\zeta(k) = A$, and means that this peak matches the case of a (Dirac) delta function $A\delta(\ln(k/k_p))$ in the limit $\Delta \to 0$. The peak position is chosen such that the mass distribution peaks at $\sim 35\ \mathrm{M}_\odot$. As noted in [258], a broader power spectrum peak not only results in a broader mass distribution, but also a shift of the peak to lower masses. To ensure that the calculated mass distributions all peak at approximately the same mass, the position of the power spectrum peak is shifted accordingly. For the delta function case, $k_p = 1.6 \times 10^6\ \mathrm{Mpc}^{-1}$, corresponding to a horizon mass of $M_H = 7\ \mathrm{M}_\odot$. It should also be noted that in the LIGO mass range, there is an enhancement caused by the softening of the equation of state during the QCD phase transition [45, 229, 262, 263]. We have neglected this effect so that the results regarding the optimal models are reliable at other mass scales. When considering a given mass range, any relevant thermal effects should be taken into account, which may mean that the models presented here need to be modified, but the relative ability of these models to fit the underlying distribution are not expected to change.

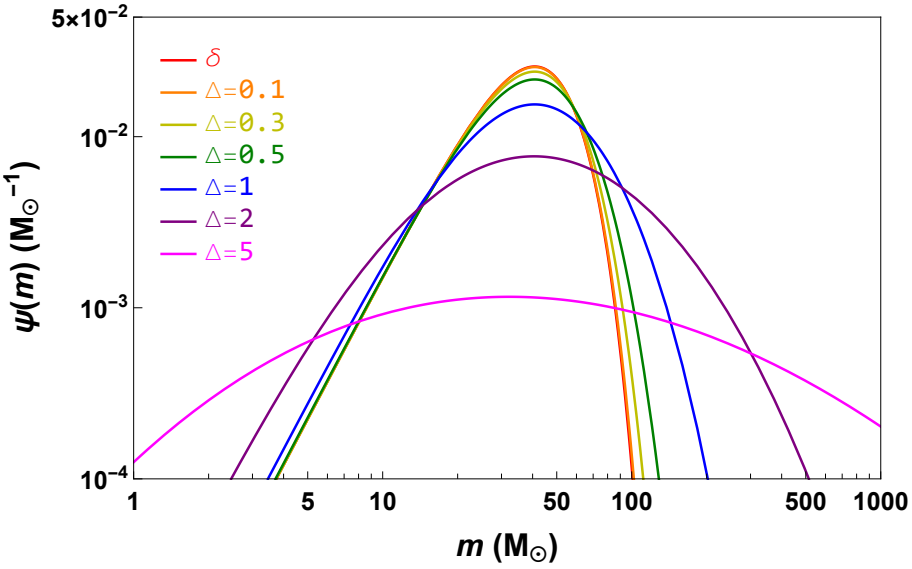

Figure 3.1: The numerical mass distribution calculated for a range of power spectrum peak widths. The peak positions are chosen for each width such that the resulting distribution peaks at ~ 35 $M_\odot$.

We can see from fig. 3.1 that, taking into account the changes to the mass distribution calculation over the last 25 years, the non-lognormality seen in [185, 245] is still valid, with the narrowest peaks showing significant deviation from the symmetric shape expected for a lognormal mass distribution. We must now consider the question of whether an alternative model can capture the detailed shape of the mass distribution over a large range of widths significantly better than the lognormal case.

3.2.2 Fitting the mass distribution

To find the best fitting mass distribution model, we use a χ^2 statistic,

$$\chi^2 = \frac{1}{\psi^2_{\rm peak}} \sum_i [\psi_{\rm num}(m_i) - \psi_{\rm model}(m_i)]^2 w_i, \tag{3.3}$$

with the weightings given by

$$w_i = \left(\frac{\psi_{\rm num}(m_i)}{\psi_{\rm peak}} \right)^2. \tag{3.4}$$

We make this choice such that the peak receives more weight than the tails, since the majority of observational techniques are most sensitive to the masses around the peak. The squared power is motivated by considering a fit to the LIGO merger rate data, where the merger rate is (roughly) proportional to ψ^2. The overall normalisation by $\psi^2_{\rm peak}$ is similar to fitting $\psi/\psi_{\rm peak}$, and ensures that the χ^2 values can be compared not only between models, but also for the same model with different widths.

The data are drawn from the numerical mass distribution, and consist of 100 values spaced equally in log mass. We choose a log-spacing because the constraints on the mass distribution stretch over many orders of magnitude, and the low-mass tail must be fitted with comparable weight to the high-mass tail in order to ensure that relevant constraints are not missed. The lower mass tail is especially relevant for e.g. microlensing constraints if we want a peak in the LIGO range, or evaporation constraints for a peak in the asteroid mass band. The limits are set arbitrarily to encompass the top four orders of magnitude of the distribution. The weighting applied to the χ^2 statistic should mean that any part of the mass distribution outside of these limits will contribute negligibly to the best fit.

3.2.3 Models

In this section, we present various parametrisations considered for the PBH mass distribution.

Lognormal

The de-facto standard mass distribution considered for PBHs generated from a reasonably narrow, smooth, symmetric peak in the power spectrum is the lognormal, given by

$$\psi_{\rm L}(m) = \frac{1}{\sqrt{2\pi}\sigma m} \exp\left(-\frac{\ln^2(m/m_c)}{2\sigma^2}\right), \tag{3.5}$$

where m_c is the mean of $m\psi(m)$ and σ is the width. There are a number of alternative distributions to the lognormal that may fit the numerical mass distribution better over the whole range of power spectrum peak widths. The ones chosen for testing in this work are described in the following sections.

Gaussian

This is simply a standard Gaussian distribution, given by

$$\psi_{\rm G}(m) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(m-m_c)^2}{2\sigma^2}\right), \tag{3.6}$$

with m_c the mean and σ the width. It should be noted that this distribution allows for negative masses, which are clearly unphysical. However, if the fit is good, the fraction of negative masses should be negligible.

Skew-normal

The skew-normal is a modification to the Gaussian distribution which introduces skewness by multiplying the Gaussian PDF with a Gaussian CDF modified with a parameter α. The definition is

$$\psi_{\rm SN}(m) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(m-m_c)^2}{2\sigma^2}\right) \left[1 + \mathrm{erf}\left(\alpha\frac{m-m_c}{\sqrt{2}\sigma}\right)\right]. \tag{3.7}$$

As for the Gaussian, this distribution can produce negative masses, although the fraction is expected to be small for a good fit.

Skew-lognormal

The skew-lognormal is virtually identical to the skew-normal, but with the mass terms switched for log-mass terms, and an additional factor of $1/m$ to preserve the normalisation over mass, i.e.,

$$\psi_{\rm SL}(m) = \frac{1}{\sqrt{2\pi}\sigma m} \exp\left(-\frac{\ln^2(m/m_c)}{2\sigma^2}\right) \left[1 + \mathrm{erf}\left(\alpha\frac{\ln(m/m_c)}{\sqrt{2}\sigma}\right)\right]. \tag{3.8}$$

It can be seen that, excluding the last bracket, this is simply the lognormal mass distribution, hence the name skew-lognormal. Since this is defined in log-space, it is superior to the skew-normal in that it avoids producing negative masses.

Critical collapse models

Motivated by the mass distribution dominated by critical collapse effects calculated in [185, 245], and later models based upon this form [76, 129], we introduce a critical collapse model class, given in general by

$$\psi_{\rm CC}(m) = \frac{\beta}{m_f}\left[\Gamma\left(\frac{\alpha+1}{\beta}\right)\right]^{-1}\left(\frac{m}{m_f}\right)^{\alpha}\exp\left[-\left(\frac{m}{m_f}\right)^{\beta}\right], \tag{3.9}$$

where the PBH mass is given by the critical collapse equation

$$m = K M_H(\delta - \delta_c)^{\gamma}, \tag{3.10}$$

where M_H is the horizon mass at formation, K is a dimensionless constant, $\gamma \approx 0.36$ is a universal scaling exponent which is independent of the initial shape of the density fluctuations and δ_c is the minimum overdensity required for PBH formation [175]. In this class, we consider three models, defined as follows:

- CC1: $\alpha = \beta = \gamma^{-1}$, $\gamma = 0.36$

  This is the most simple model, motivated entirely by the critical collapse calculations. It is identical to the form stated in [245], and has been numerically checked for small $\delta - \delta_c$ in [175]. It has just one parameter, to fit the location of the distribution.

- CC2: $\alpha = \beta = \gamma^{-1}$, γ variable

  The shape of this model is identical to the above case, in that both tails are described by γ. However, in this case, we allow γ to float to find the best fit. This is motivated by the demonstration in [264] that the value $\gamma = 0.36$ does

not hold for larger values of $\delta - \delta_c$, and a modification to the critical collapse parameters K and γ could yield a better fit across the whole range. This model has two parameters, to fit the location and shape of the distribution.

- CC3: α, β variable

  This is a generalisation of the critical collapse model, disconnecting the behaviour of the two tails. It has three parameters, to fit the location and the shape of each tail.

Location parameter

The location parameters stated in the parametrisations above can be extremely sensitive to the width of the mass distribution, causing problems in the fitting procedure. To overcome this, we reparametrise most of the models in terms of their peak mass m_p, which we have held approximately fixed for all of the numerical mass distributions. The transformations between the peak mass and the location parameters defined above are given below.

Table 3.1: Transformation of location parameter to peak mass m_p for each model.

Model	**Transformation**
Lognormal	$m_c = m_p e^{\sigma^2}$
Gaussian	$m_c = m_p$
Skew-normal	N/A
Skew-lognormal	N/A
CC1&2	$m_f = m_p$
CC3	$m_f = m_p \left(\frac{\beta}{\alpha}\right)^{1/\beta}$

For the skew-normal and skew-lognormal there is no analytical form for the peak mass. There is an approximate transformation derived from numerical fits [265], but this does not hold for the broadest cases. Therefore, for the skew-normal and skew-lognormal we retain the location parameters m_c and $\ln(m_c)$ defined above. For these two models these location parameters are stable enough to avoid numerical errors during fitting.

3.2.4 Fit results

We obtain fits to the numerical mass distribution calculated from a peak in the power spectrum by minimisation of eq. (3.3). We consider a large range of power spectrum peak widths, from the limiting case of a delta function up to a very broad case of $\Delta = 5$. The optimised model fits are shown in fig. 3.2 for three representative cases: a delta function, $\Delta = 1$, and $\Delta = 5$. It is immediately apparent that the lognormal is outperformed by the vast majority of the models for the narrowest case. However, it can be seen that many of these models begin to fail as the width increases, and are completely wrong for the broadest case.

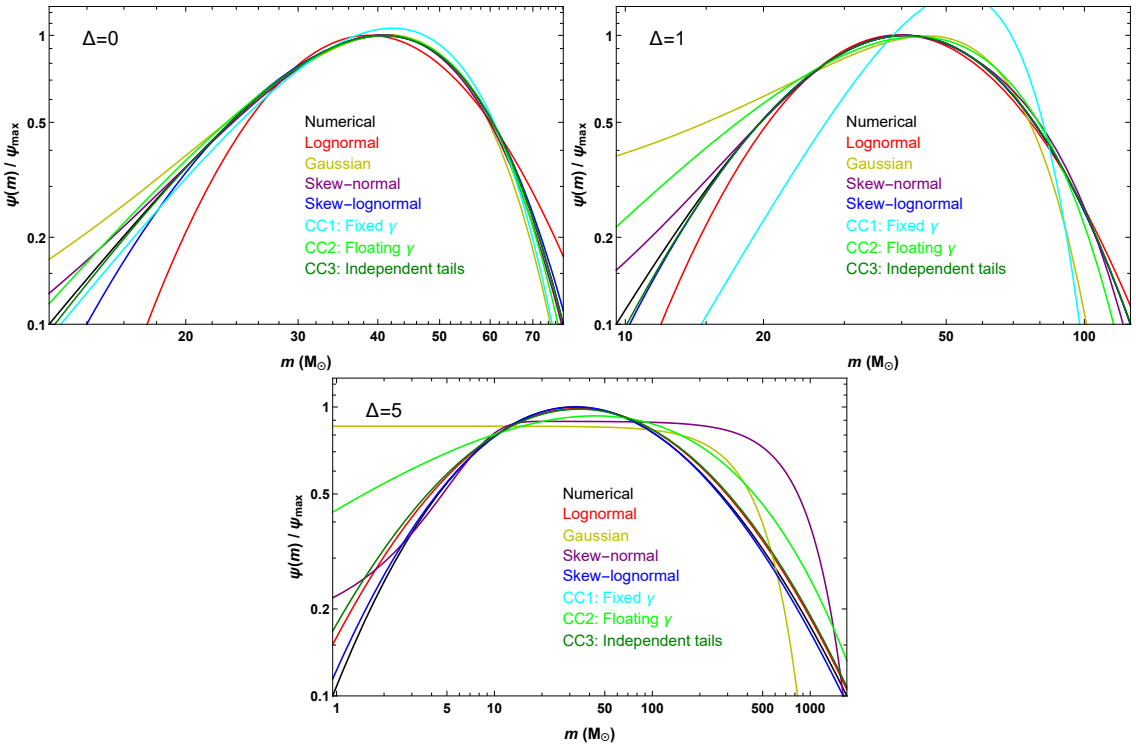

Figure 3.2: Optimal model fits to the numerical mass distribution for three representative power spectrum widths $\Delta = 0$ (delta function), $\Delta = 1$, and $\Delta = 5$. The mass limits are chosen to contain the top 10% of the numerical mass distribution, to highlight the deviation of the models near the peak.

We can compare the models more carefully by examining their reduced χ^2 values. Figure 3.3 shows the χ^2_ν values for all the models and widths considered. Here we can see again that, while there are many models that outperform the lognormal for the narrowest cases, a large number of them fail as the width increases, where they

cannot generate the appropriate skewness. However, it can be seen that two models, the skew-lognormal and the generalised critical collapse model, consistently provide a more accurate fit than the lognormal. These models also have the benefit of not producing negative masses, although the models which do allow this are deemed irrelevant by their failure to fit the broadest cases anyway. The reduced χ^2 values are provided in table D.1.

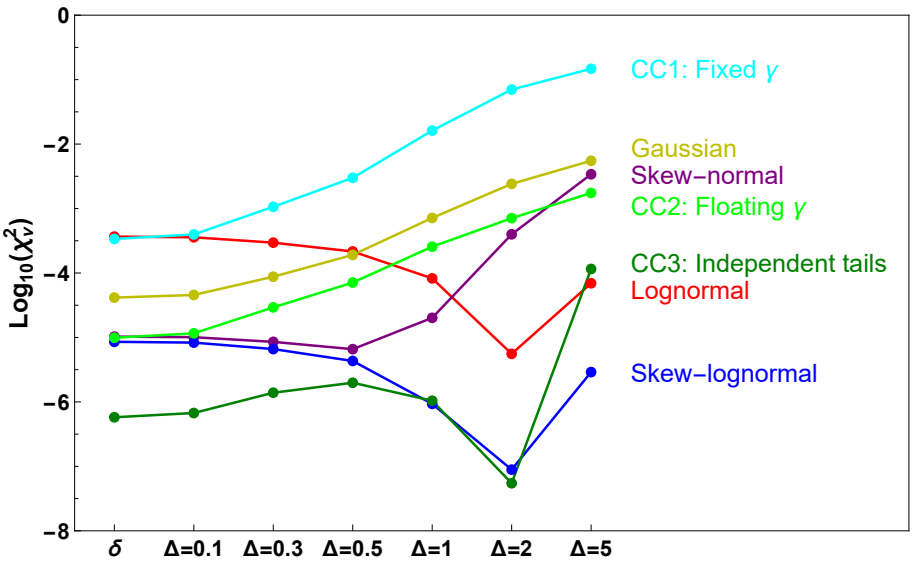

Figure 3.3: Reduced χ^2 values for the models and widths considered. Lower (more negative) values indicate a better fit.

The comparison between the lognormal and the two models that consistently outperform it can be seen graphically in fig. 3.4, where the best fit lognormal is shown with a long-dashed red line, the skew-lognormal with a mid-dashed blue line, and the generalised critical collapse model with a short-dashed green line. The numerical distribution calculated from the power spectrum is shown with a solid black line.

It is evident from these plots that modelling the PBH mass distribution across a broad range of widths is a challenging task, as it requires negative skewness in log-space for the narrowest cases, before a change to symmetrical and then positively skewed distributions. The model best suited for the job is the generalised critical collapse model, which can produce the negative skewness exceptionally well, but be-

gins to fail when positive skewness is required. The skew-lognormal acts oppositely, with a failure to produce enough negative skewness, but an improvement for the positive skewness regime.

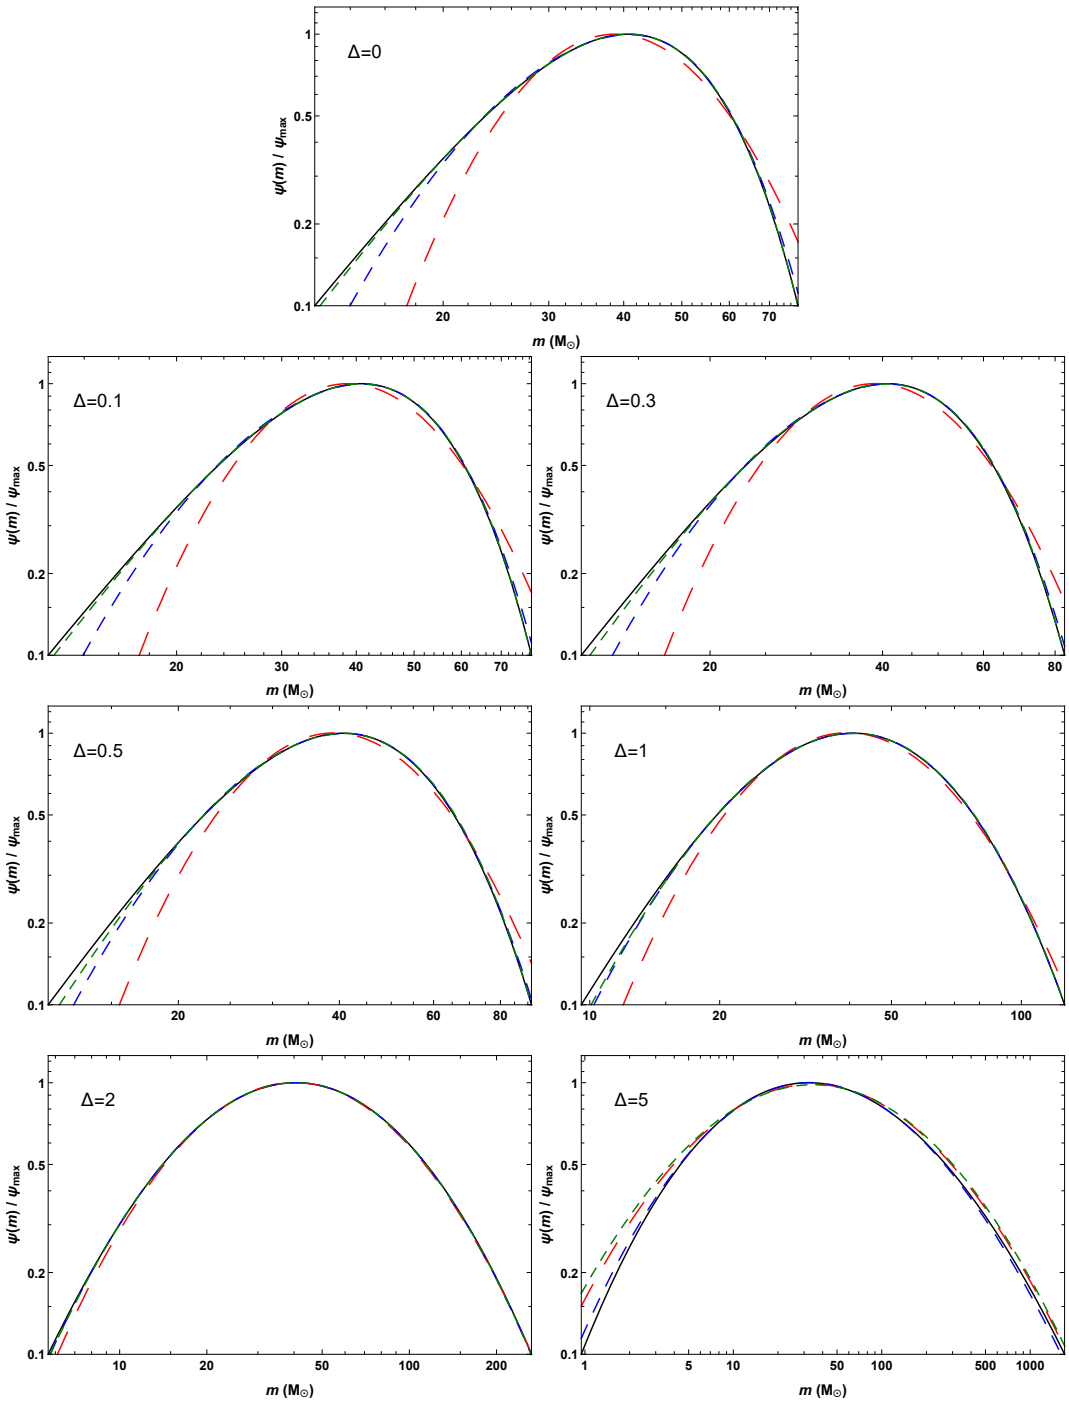

Figure 3.4: Plots of the lognormal (red, long-dashed), skew-lognormal (blue, mid-dashed) and generalised critical collapse (green, short-dashed) fits to the numerical mass distribution generated by a lognormal peak in the power spectrum (black, solid). The mass limits are chosen to contain the top 10% of the numerical mass distribution, to highlight the deviation of the models near the peak.

Of course, there is a price to pay for achieving this matching, and that is the introduction of an additional parameter. Both of the best-fitting models have three parameters, as opposed to the two required for the lognormal. However, the importance of accurately describing the shape of the tails of the distribution cannot be overstated. Failure to capture this shape can result in incorrect conclusions about the acceptability of a particular model. For example, if we are looking for PBHs in the LIGO mass range, we must ensure that the tails of the distribution do not conflict with the microlensing constraints on the low-mass side and the CMB anisotropy constraints on the high-mass end. Similarly, for PBHs in the recently reopened asteroid mass window, a significant underfitting of the low-mass tail, such as that displayed by the lognormal model for the narrower widths, could suggest that PBHs can evade all the constraints, whereas a more accurate model would show that they are in tension with the evaporation limits.

Table 3.2: Fitted parameter values for the skew-lognormal and generalised critical collapse distributions with different power spectrum widths. For the skew-lognormal model, we also provide the peak mass m_p determined by numerical maximisation. It should be noted that $\ln(m_c)$ and m_p are not independent and only $\ln(m_c)$ is determined by the fit for the skew-lognormal model. The peak mass is included only for comparison to the fitted parameter in the critical collapse model.

	Parameters						
	SL				**CC3**		
Width Δ	$\ln(m_c)$	σ	α	m_p	m_p	α	β
δ	4.13	0.55	−2.27	40.9	40.8	3.06	2.12
0.1	4.13	0.55	−2.24	40.9	40.8	3.09	2.08
0.3	4.15	0.57	−2.07	40.9	40.7	3.34	1.72
0.5	4.21	0.60	−1.82	40.8	40.7	3.82	1.27
1.0	4.40	0.71	−1.31	40.8	40.8	5.76	0.51
2.0	4.88	0.97	−0.66	40.6	40.6	18.9	0.0669
5.0	5.41	2.77	1.39	32.9	35.1	13.9	0.0206

In table 3.2, we provide the optimal model parameters for the best two models, the skew-lognormal and the generalised critical collapse model, for the widths considered. This allows fits of these analytical approximations to be crudely compared to the power spectrum details without the necessity of recalculating the full mass distribution. However, it should be noted that although these more accurate mod-

els provide a significant improvement over the lognormal, even they fail to capture the detailed shape deep into the tails, and the only truly rigorous way to determine whether PBHs are not excluded in a given range is to calculate the mass distribution from the power spectrum peak.

3.3 Conclusions

We have carried out a thorough examination of the PBH mass distribution arising from a peak in the primordial power spectrum, re-evaluating the validity of the lognormal approximation to the mass distribution. We confirm that the modifications to the PBH mass distribution calculation over the last 25 years do not change the conclusion that the lognormal model is still unable to accurately capture the shape of the distribution generated from sufficiently narrow peaks, with $\Delta < 1$. We compare a set of alternative models using a weighted χ^2 statistic, and show that over a large range of peak widths, the lognormal is outperformed by the skew-lognormal and a generalised form motivated by the effects of critical collapse.

This deviation between the lognormal assumption and the PBH mass distribution calculated for a specific power spectrum peak will have important consequences for physical inferences made from accurate data, such as the LIGO–Virgo observations. In a related previous paper [255], we considered the skew-lognormal as part of a detailed Bayesian analysis of the LIGO–Virgo O1O2 dataset. The limited sample size means that the difference in the mass distribution does not significantly affect the results, but the difference will become increasingly important with the accurate data in the O3 run and future runs.

An accurate model of the PBH mass distribution will also be relevant in other areas, such as making accurate constraints on the PBH abundance. These constraints are typically presented for a monochromatic mass distribution, but extended mass distributions have been considered, see e.g. [76, 77]. The constraints for extended mass distributions are typically similar to the monochromatic case, but the differences become important when determining the validity of specific extended mass

distributions, particularly in the case of $f_{\rm PBH} \sim 1$, where the tails of the distributions may be in tension with constraints. In these cases, an accurate model of the mass distribution is essential, to avoid drawing an incorrect conclusion about the validity of the distribution. This is especially important in areas where there are extremely tight constraints, such as those from CMB anisotropies and evaporation, of particular interest for the LIGO and asteroid mass windows respectively. For cases involving fitting to accurate data or tight constraints, we advocate the use of the skew-lognormal or generalised critical collapse model, to ensure that the conclusions drawn are valid.

If the shape of the power spectrum peak deviates from that considered here, either by considering other symmetric peaks or non-symmetric peaks, the shape of the numerical mass distribution will naturally alter as well. It is expected that for the case of $\Delta \lesssim 1$, this would not change the conclusions, since critical collapse dominates the mass distribution shape in this regime. However, for much broader peaks, the detailed shape of the peak will be important, and may affect the results stated here. Nonetheless, we believe that in general, a three-parameter model will be required to capture the full shape of the mass distribution across a broad range of widths.

Acknowledgements

AG is funded by a Royal Society Studentship by means of a Royal Society Enhancement Award. CB acknowledges support from the Science and Technology Facilities Council [grant number ST/T000473/1]. AH is supported by a Science and Technology Facilities Council Consolidated Grant.

Chapter 4

Primordial black hole merger rates: distributions for multiple LIGO observables

Andrew D. Gow[1], Christian T. Byrnes[1], Alex Hall[2], and John A. Peacock[2]

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

2) Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, United Kingdom

We have calculated the detectable merger rate of primordial black holes (PBHs), as a function of the redshift, as well as the binary's mass ratio, total mass and chirp mass (observables that have not previously been explored in great detail for PBHs). We consider both the current and design sensitivity of LIGO and five different PBH mass functions, as well as showing a comparison to a predicted astrophysical black hole merger rate. We show that the empirical preference for nearly equal-mass binaries in current LIGO–Virgo data can be consistent with a PBH hypothesis once observational selection effects are taken into account. However, current data do exclude some PBH mass distributions, and future data may be able to rule out the possibility that all observed BH mergers had a primordial origin.

4.1 Introduction

Primordial black holes (PBHs) were first considered by Zel'dovich and Novikov [24], and were heavily studied by Hawking and Carr [25, 26]. Since then, the field has generated an extensive literature; for a review of the current state of research, see [155] and [156]. As they interact only via gravity, PBHs are a natural dark matter (DM) candidate without requiring physics beyond the standard model. The fraction of DM that can be composed of PBHs, $f_{\rm PBH}$, has been well constrained by a number of methods [156]. One of these methods is provided by the detection of gravitational wave signals by the Laser Interferometer Gravitational-wave Observatory (LIGO). During the O1 and O2 runs, LIGO detected 10 binary black hole (BBH) mergers [30]. The detector has recently finished its O3a run (the first half of the O3 sensitivity run) lasting from 1 April to 1 October 2019, and has detected 21 mergers with $> 90\%$ probability of being BBHs. Some unexpected properties of the detected mergers in the O1O2 dataset, such as the high mass and low effective spin, led to the suggestion that the mergers may be primordial in origin [67, 266, 267], and that they could explain an excess of power in the cosmic infrared background, although this requires $f_{\rm PBH} \sim 1$ [268]. The aim of our current paper is to investigate this possibility in more detail, attempting to model properties such as the BBH mass ratios, allowing for observational selection effects that bias the rates with which different binaries are detected.

A major goal of any PBH analysis is to place constraints on $f_{\rm PBH}$. There are a number of methods for placing limits on this parameter, over a large range of mass scales. In the LIGO range of $\sim$1–100 $M_\odot$, the relevant constraining techniques are microlensing events [269–272] and CMB accretion effects [273,274]. A summary plot of these limits (and other limits on different mass scales) can be seen in fig. 10 of [156, published version]. In this range, the tightest constraint on $f_{\rm PBH}$ comes from either CMB accretion limits or from the LIGO events themselves (fig. 17 of [156, published version]), depending on the sound speed of baryonic matter compared to the relative baryon to dark matter velocity [274]. This means that, even with $f_{\rm PBH} < 1$, all of

the LIGO events could be of primordial origin. However, the constraints discussed above have all been determined for a monochromatic PBH mass distribution. This case is unrealistic based on the typical formation mechanisms, and so some effort has been made into determining equivalent limits for extended mass functions [76, 77].

If PBHs exist, it is of extreme importance that their mass distribution be accurately characterised. We have made the first study of several parameters related to the masses in PBH scenarios, such as the total mass M, the chirp mass $\mathcal{M}$, and the mass ratio $q = m_2/m_1$. The LIGO convention is that m_2 is the smaller mass, and hence $0 < q \leqslant 1$. The PBH mass distribution is already observationally constrained, and with future data it will be possible to determine the distribution and its parameters to a high degree of accuracy, or perhaps even rule out any possible PBH mass distribution as the origin of all the detected LIGO events.

A particular motivation for considering the mass ratio q is that the LIGO data have central q values that are all statistically consistent with equal mass mergers. One may wonder whether such a strong correlation in mass is plausible for PBHs with a broad mass function, and much of our paper is devoted to considering this question. Of course, the same issue of principle arises if the BHs are of astrophysical origin, but it seems that q could naturally be close to unity in this case [109,110,112,113], while PBH binaries, having a very different formation mechanism, would not necessarily have such a strong tendency. It is interesting to consider whether the current LIGO data favour a particular q value, but the LIGO selection effects discussed below (which have a preference for equal mass mergers) must be taken into account. A recent analysis of the LIGO data showed that, for a mass distribution based on a power-law form, $q > 0.6$ is favoured [275].

Additional observables could also be used to distinguish between mergers of primordial and astrophysical origin. A recent paper by Gerosa et al. [276] obtained merger rate distributions for astrophysical black holes against three observables:

the total mass M, redshift z, and the mass ratio q. It is desirable to have the same distributions for PBHs so a comparison between the primordial and astrophysical cases can be drawn.

The intrinsic merger rate for PBHs is obtained by considering the number density of PBHs, and their interactions. A binary is formed when the gravitational attraction between two adjacent PBHs dominates over the Hubble flow. The surrounding PBHs, as well as other forms of DM, then generate a tidal force that determines the angular momentum of the binary, which in turn determines how long the binary takes to merge. The intrinsic merger rate as a function of time can then be obtained. This calculation has been carried out in a number of ways by various groups, for monochromatic [67, 277–280] and extended PBH mass distributions [68, 193], most recently by Raidal et al. [69], whose method we use for the following calculations.

In section 4.2 the theoretical process for obtaining the intrinsic merger rate for PBHs is described, and is briefly shown in its numerical form. Section 4.3 explains how to determine the rate of detections expected by LIGO for a given intrinsic merger rate. The resulting distributions for the detectable merger rate are shown in section 4.4, and a comparison of different PBH mass distributions is considered in section 4.5. Finally, the merger rate expected for the LIGO O1O2 sensitivity is compared to the detected merger events in section 4.6.

4.2 Intrinsic merger rate from PBHs

To determine the intrinsic merger rate of PBHs, a number of factors must be considered. First, there is the number density of PBHs of a given mass, which is related to the mass distribution $\psi(m)$. Then, the fraction of these that form binaries must be determined, and the angular momentum distribution must be taken into account to determine the number of PBH binaries that will result in mergers at time t. This procedure will give the merger rate assuming that the binaries are not disrupted between their formation and merger. Even a small alteration in the angular momentum j of the binary will cause a significant change in the merger time τ, due

to the relation $\tau \propto j^7$ [281]. This assumption was considered by Ali-Haïmoud et al. [278], who estimated that little disruption occurs, but Raidal et al. [69] carried out simulations and argued that significant disruption may occur for $f_{\rm PBH} \gtrsim 10^{-1}$. A more recent work by Vaskonen and Veermäe determined the lower bound on the merger rate including the impact of disruptions for $f_{\rm PBH} \gtrsim 0.1$, and found that it remained too large compared with the LIGO observed merger rate [282].

An additional consideration in finding the merger rate is the clustering of PBHs. If this is an important factor, then it could considerably alter the merger rate at a given time. This has been a topic of some debate, but it is now generally agreed that, for Gaussian initial conditions, the spatial distribution of PBHs is Poissonian [61,62,65,283]. Primordial non-Gaussianity can strongly change the initial clustering of PBHs [64, 66, 70, 232, 233] and the subsequent merger rate [71] (see also [63]).

The merger rate calculation performed by Raidal et al. [69] yielded the following equations, reproduced here for convenience:

$$\mathrm{d}R = S \times \mathrm{d}R_0, \tag{4.1}$$

where $\mathrm{d}R$ is the differential merger rate, S is a suppression factor (given by eq. (2.37) in [69, published version]) that depends on the component masses m_1 and m_2, the fraction of dark matter in PBHs $f_{\rm PBH}$, and the rescaled deviation of matter density perturbations $\sigma_{\rm M}$, and

$$\mathrm{d}R_0 = \frac{1.6 \times 10^6}{\mathrm{Gpc}^3\ \mathrm{yr}} f_{\rm PBH}^{\frac{53}{37}} \eta^{-\frac{34}{37}} \left(\frac{M}{\mathrm{M}_\odot}\right)^{-\frac{32}{37}} \left(\frac{t}{t_0}\right)^{-\frac{34}{37}} \psi(m_1)\, \psi(m_2)\ \mathrm{d}m_1\, \mathrm{d}m_2 \tag{4.2}$$

is the unsuppressed differential merger rate, where $\eta = m_1 m_2 / M^2$ is the symmetric mass ratio, M is the total mass of the system, t is the proper time, t_0 is the age of the universe, and $\psi(m)$ is the mass distribution of PBHs, normalised to unity. The suppression factor S depends on the average number $\bar{N}(y)$ of PBHs in a spherical

shell of radius y. Raidal et al. determine a value of this to ensure minimal disruption of the binary for $f_{\rm PBH} < 10^{-1}$, given by eq. (3.5) in [69, published version]. We use this value for the following calculations.

The mass distribution used in [69] is a lognormal, given by the form

$$\psi(m) = \frac{1}{\sqrt{2\pi}\sigma m} \exp\left(-\frac{\ln^2(m/m_c)}{2\sigma^2}\right), \tag{4.3}$$

where m_c is the median of the distribution (also the mean of $m\psi(m)$) and σ describes the width. This is a common choice for the PBH mass distribution, as it well approximates the class of distributions for PBHs formed from peaks in the power spectrum [76]. Raidal et al. [69] carried out a fit to the LIGO data for this mass distribution, although they did not incorporate the full detectability procedure described in section 4.3, instead using a step function in the signal to noise ratio. Their best fit parameters are $m_c = 20\ \mathrm{M}_\odot$ and $\sigma = 0.6$, and we will begin by considering these values for the mass distribution. The mass distribution with these parameters is shown in fig. 4.1.

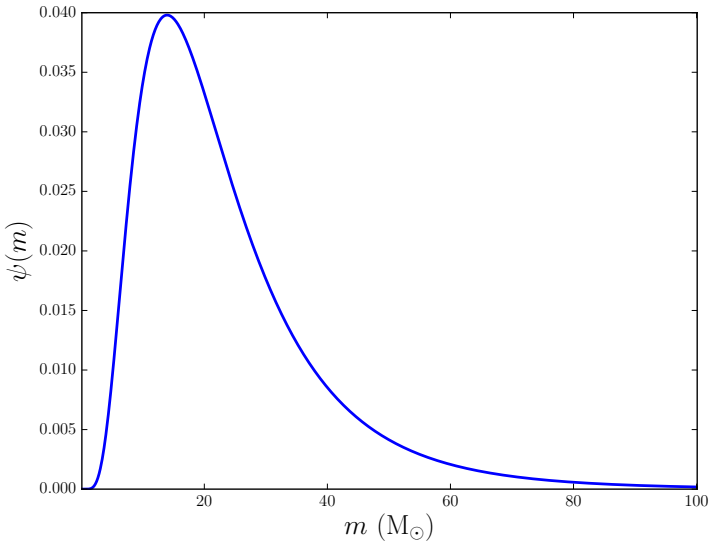

Figure 4.1: Lognormal mass distribution with $m_c = 20\ \mathrm{M}_\odot$ and $\sigma = 0.6$.

A common alternative form for the mass distribution is $f(m)$, which satisfies

$$\int \mathrm{d}\ln(m)\ f(m) = f_{\mathrm{PBH}}. \tag{4.4}$$

This is related to the above mass distribution by the relation

$$\psi(m) = \frac{1}{f_{\mathrm{PBH}}}\frac{f(m)}{m}. \tag{4.5}$$

The total intrinsic (source-frame) merger rate in $\mathrm{Gpc}^{-3}\ \mathrm{yr}^{-1}$ can be obtained from the differential merger rate by applying the integration

$$R = \int \mathrm{d}m_1 \int \mathrm{d}m_2\ \frac{\mathrm{d}R}{\mathrm{d}m_1 \mathrm{d}m_2} \tag{4.6}$$

for source-frame masses m_1 and m_2, and with a fixed value of proper time t. The result of this (or a similar) equation with $t = t_0$ is often compared with the LIGO estimate for the intrinsic merger rate [67, 193, 267, 278–280]. However, in obtaining their estimate of the intrinsic merger rate, the LIGO collaboration assumes a mass distribution, so this estimate could differ significantly compared to the PBH calculation if a very different mass distribution is used. For further details on this estimation method, see section VII of [30, published version]. To overcome this problem, the ground-based rate of detections in yr^{-1} can be found, which can then be directly compared to the LIGO measurements, rather than their intrinsic rate estimate. Distributions of this type in the component masses can be seen in [193] with f_{PBH} chosen to fix the intrinsic merger rate to $R = 100\ \mathrm{Gpc}^{-3}\ \mathrm{yr}^{-1}$, although these distributions do not take into account the dependence of the detectability on the component masses. The process for finding the ground-based detection rate is described in the following section.

4.3 Detectability and ground-based detection rate

To obtain the ground-based detection rate, we must weight the different parts of the intrinsic merger rate with the ability of the LIGO instrument to detect the resulting waveform. This detectable merger rate is calculated using

$$R_{\rm det} = \int {\rm d}z \int {\rm d}m_1 \int {\rm d}m_2 \, \frac{1}{1+z} \frac{{\rm d}V_c}{{\rm d}z} \, p_{\rm det}(z, m_1, m_2) \, \frac{{\rm d}R}{{\rm d}m_1 {\rm d}m_2}(z), \tag{4.7}$$

where V_c is the comoving volume and $p_{\rm det}$ is the detection probability [284]. This detection probability is obtained by simulating merger waveforms and passing them through the LIGO detection pipeline to find the signal to noise ratio (SNR) in a single detector for a certain set of parameters. The angular dependence of the detection probability may be well approximated by the function $p(\omega)$, where ω is the projection parameter defined in eq. (2) in [276, arXiv version], with the detection probability given by

$$p_{\rm det}(z, m_1, m_2) = \int_{\rho_{\rm thr}/\rho_{\rm opt}(z,m_1,m_2)}^{1} {\rm d}\omega \; p(\omega), \tag{4.8}$$

where $\rho_{\rm opt}$ is the SNR for a merger happening face-on to the detector located directly above the detector and $\rho_{\rm thr}$ is a threshold SNR above which it is assumed that the signal is detected, typically taken as $\rho_{\rm thr} = 8$ [276]. The noise curve used is the design sensitivity curve (`aLIGOZeroDetHighPower`). This process is carried out using the public code `gwdet` written by Davide Gerosa [285]. The resulting probability is plotted against the component masses in fig. 4.2, at $z = 0.2$ and $z = 0.5$.

The method described above does not take into account the spin of the component BHs. In principle the dependence of the SNR on these spins should be taken into account. Since PBH spins are expected to be very small at formation [54,57,286], we avoid this computationally expensive step by computing waveform approximants having zero spin. The difference in spin has also been considered as another observable that could be used for distinguishing between mergers of astrophysical and

primordial origin [58, 287]. There is the possibility that, although the PBH spin is small at formation, they could spin up between formation and merger. However, this is likely to be a small effect [56]. Assuming the detection probability varies little over the range around zero where PBH spins are expected to lie, then taking zero spin is a good approximation.

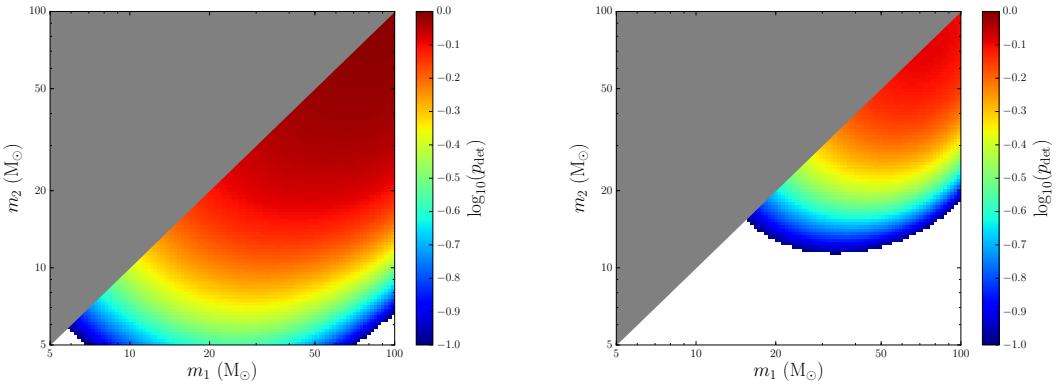

Figure 4.2: Detection probability $p_{\rm det}(m_1, m_2)$ at $z = 0.2$ (left) and $z = 0.5$ (right). Note that all three scales are in log-space. The white area indicates $p_{\rm det} < 0.1$, and the grey triangle indicates that the case $m_2 < m_1$ is chosen by LIGO for their analysis.

4.4 Detectable merger rate for LIGO observables

Distributions of the detectable merger rate against four observables have been generated by Monte Carlo integration of eq. (4.7) using the `AdaptiveQuasiMonteCarlo` integrator in Mathematica . These distributions are shown in fig. 4.3. The four observables are total mass M, redshift z, mass ratio q, and chirp mass $\mathcal{M}$. The first three of these observables are chosen for comparison with detectable merger rate distributions recently determined by Gerosa et al. [276] for astrophysical black holes, using the StarTrack code for stellar evolution and the Precession code to add spins, with the same detectability process described above [276]. The fourth observable, $\mathcal{M}$, is chosen because this is the observable best constrained by LIGO for lower mass mergers. There is no astrophysical curve publicly available at the time of writing for this observable, as it was not calculated in [276]. Three values

of $f_{\rm PBH}$ are shown, with the largest being 10^{-1}. Above this value, the merger rate calculation is unreliable due to the high probability of the binary being disrupted between formation and merger [69].

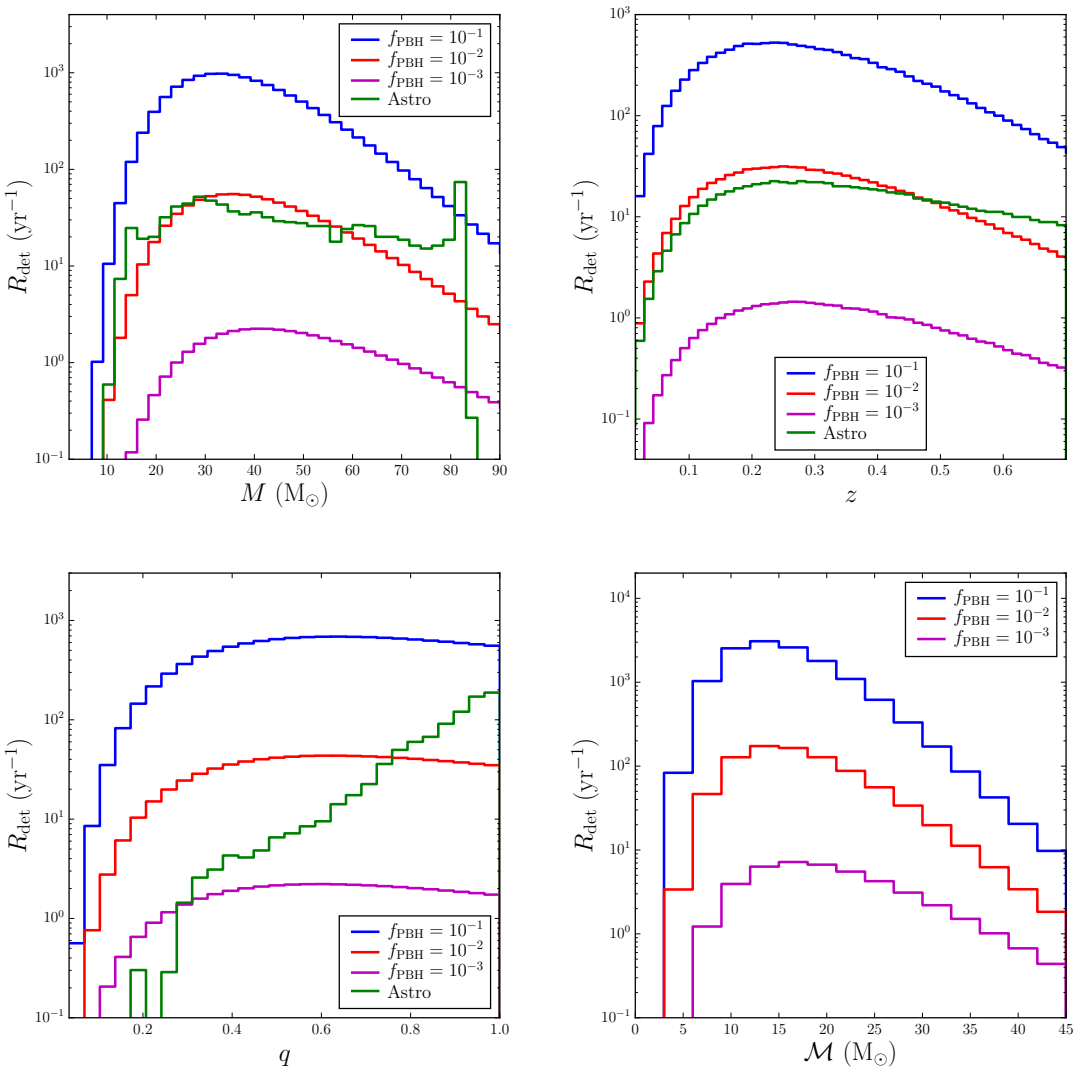

Figure 4.3: Merger rate distributions in total mass M, redshift z, mass ratio q, and chirp mass $\mathcal{M}$ for a lognormal mass distribution with $\sigma = 0.6$, at design sensitivity. The distributions for astrophysical black holes from [276] are shown in green for the first three plots.

The distributions for the total mass M and the chirp mass $\mathcal{M}$ follow the component mass distribution shown in fig. 4.1 closely, with the peaks lying where one would expect by taking the peak of the individual lognormal mass distribution ($m_c = 20$ M$_\odot$, $\sigma = 0.6$) and calculating the resulting values of M and $\mathcal{M}$. The distribution for the mass ratio q seems to favour $q \approx 0.6$, but is fairly flat from $q = 1$

down to $q \approx 0.4$. After this, there is a steep drop-off, due to a combination of the width and detectability factors. It can be seen that the major dependence on $f_{\rm PBH}$ is simply a global multiplier, scaling the curves up or down. However, there are other dependencies, such as the $f_{\rm PBH} = 10^{-3}$ curve being flatter at high observable values than the curves with higher $f_{\rm PBH}$ values for the total mass M and chirp mass $\mathcal{M}$.

As can be seen in fig. 4.3, the distributions for total mass M and redshift z seem to match the astrophysical distribution quite closely at low redshift, but the astrophysical rate drops for redshifts above $z \approx 1.5$ as it follows the stellar formation rate, while the primordial rate continues growing and becomes the dominant merger source, as can be seen in fig. 10 of [69, published version]. All the rates tend to zero as the redshift tends to zero, due to the volume factor in eq. (4.7). In contrast to the above two cases, the distribution for the mass ratio q shows a clear difference between the astrophysical distribution and any of the primordial curves. This could therefore be a useful observable for distinguishing between mergers of astrophysical and primordial origin. As can be seen, the astrophysical distribution tends to favour higher q values, which is to be expected considering the formation mechanisms [109, 110, 112, 113]. The current LIGO data also favour high mass ratios [275], and with future data, the allowed width and shape of the PBH mass distribution could be seriously constrained.

It is also interesting to know how the merger rate is distributed across multiple observables. 2D plots of the merger rate against four sets of parameters were created. These are the component masses (m_1, m_2), and the three combinations of the redshift z, the chirp mass $\mathcal{M}$, and the mass ratio q. These distributions are shown in fig. 4.4. The distribution in (m_1, m_2) exhibits the expected behaviour following the mass distribution, with a peak between 10 $\rm M_\odot$ and 20 $\rm M_\odot$, and then a drop off to higher masses. The distribution in (q, $\mathcal{M}$) shows that, away from the chirp mass peak, there is no detectable merger rate for low q values.

The distributions involving redshift show that nothing can be detected past $z \approx 1.6$, even at design sensitivity. The $(z, \mathcal{M})$ distribution peaks at the same chirp mass value as the other 1D and 2D distributions, but drops off very rapidly with redshift. Also, it can be seen that the larger chirp mass values can be detected out to much higher redshifts, due to the larger amplitude of the gravitational wave produced. The same is true of the higher q values in the (z, q) distribution.

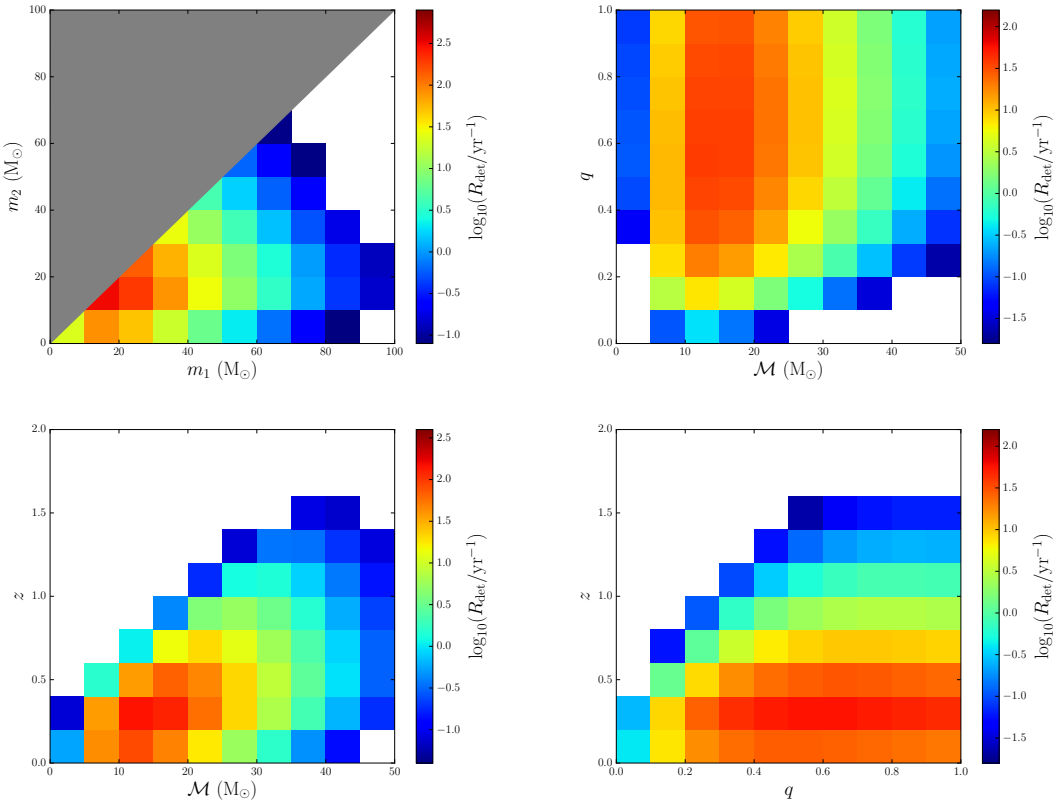

Figure 4.4: 2D merger rate distributions in individual masses, mass ratio q, chirp mass $\mathcal{M}$ and redshift z for a lognormal mass distribution with $\sigma = 0.6$, at design sensitivity. All plots have $f_{\rm PBH} = 10^{-2}$. The white area corresponds to no significant merger rate, and the grey triangle indicates the LIGO choice $m_2 < m_1$.

For the two distributions involving the chirp mass $\mathcal{M}$, we can take vertical slices and produce 1D distributions over a given chirp mass range, to better demonstrate the sensitivity of the merger rate to the chirp mass. These are shown in fig. 4.5. It can be seen for both of the observables that the low chirp mass curve begins above the medium chirp mass curve, but then drops below as the value on the respective

horizontal axes is increased. For the redshift distribution, this is because the low chirp mass binaries have a low detection probability, and so have a very limited z-range.

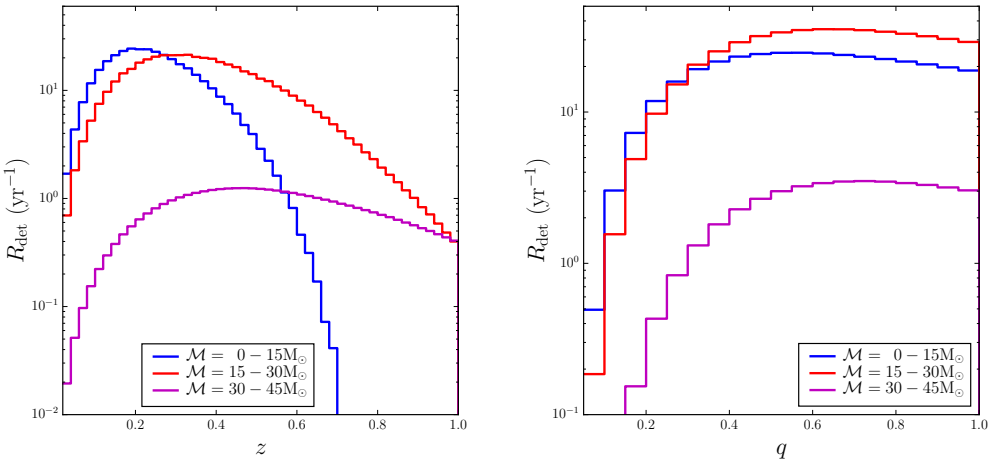

Figure 4.5: Merger rate distributions in redshift z and mass ratio q for a lognormal mass distribution with $\sigma = 0.6$ at design sensitivity, binned by chirp mass $\mathcal{M}$. Both plots have $f_{\rm PBH} = 10^{-2}$.

1D and 2D distributions of the detectable merger rate have been generated for the lognormal mass distribution in eq. (4.3), with the parameters $m_c = 20\ M_\odot$ and $\sigma = 0.6$ given in [69]. However, while this is a plausible form for the mass function and its parameters, it is not the only option. Therefore, it is interesting to consider how these merger rate distributions change for other PBH mass distributions.

4.5 Comparison of different PBH mass distributions

4.5.1 Lognormal width parameter σ

While the best fit values in [69] for the lognormal mass distribution are $m_c = 20\ M_\odot$ and $\sigma = 0.6$, the full angular dependence for the detectability was not incorporated, and so it is of interest to consider the merger rate distributions for other widths. The method above was carried out for a second lognormal distribution with the same m_c

but with $\sigma = 0.3$, and also for a monochromatic distribution $\delta(m - m_c)$ which is the limit of the lognormal distribution as $\sigma \to 0$. Both of these additional distributions had the same $m_c = 20\ \mathrm{M}_\odot$. For typical PBH formation scenarios, critical collapse imposes a minimum width on the mass distribution, and so a monochromatic distribution is not realistic. However, the monochromatic distribution is still useful for comparison and demonstration of the properties affecting the merger rate. A comparison of the distributions is shown in fig. 4.6.

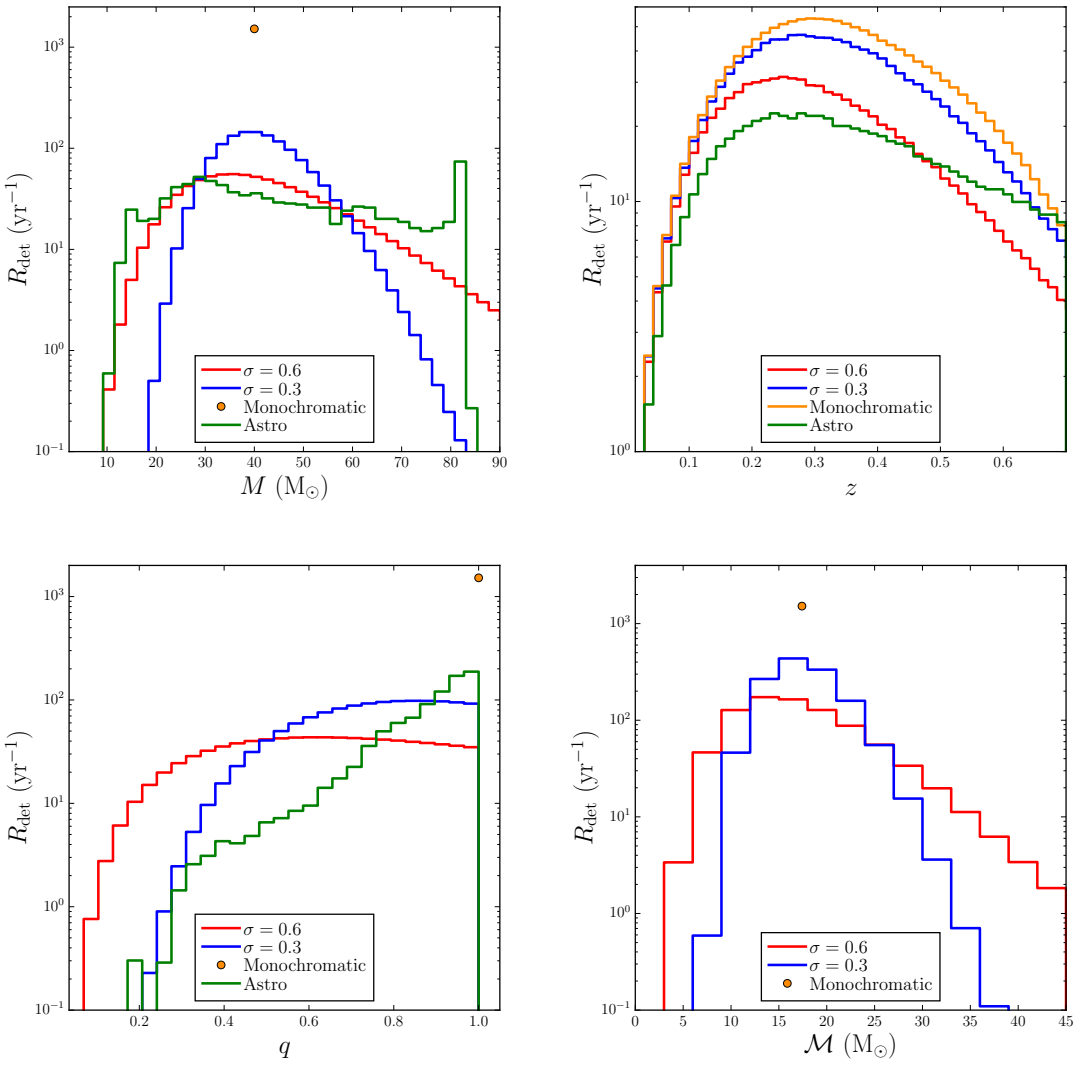

Figure 4.6: Comparison of merger rate distributions in total mass M, redshift z, mass ratio q and chirp mass $\mathcal{M}$ for different widths of mass distribution, at design sensitivity. All plots have $f_{\mathrm{PBH}} = 10^{-2}$.

For M, q, and $\mathcal{M}$, the monochromatic distribution is represented by a single point, since there is only one value of each observable it can take. In M and $\mathcal{M}$, it can be seen that reducing the width of the mass distribution leads to a reduction of the width of the merger rate distribution, as expected. It also leads to an enhancement of the peak. For q, the width reduction suppresses low q values, also to be expected, since a narrower distribution has a smaller difference between the highest and lowest probable masses. For z, the narrower mass distribution leads to an enhancement across the whole range because, for a given total mass, equal mass mergers are easier to detect. 2D distributions for the lognormal distribution with $\sigma = 0.3$ are shown in fig. E.2.

For the monochromatic distribution, there are some further parameters that can be considered due to the simplicity of the function. These are $f_{\rm PBH}$, the mass of the monochromatic distribution m_c, and the rescaled deviation of matter density perturbations at the time of binary formation $\sigma_{\rm M}$, given just after eq. (2.24) in [69, published version]. The value of $\sigma_{\rm M}$ is usually taken as 0.006 on scales relating to black hole masses of order 1–10^3 $\mathrm{M}_\odot$, corresponding to the deviation of density perturbations $\sigma_{\rm eq} = 0.005$ in [278] and [279]. However, this value is found by extrapolating the power spectrum amplitude and spectral index measured from the CMB. Since PBH formation typically requires some type of peak in the power spectrum on relevant scales, this could quite dramatically change the value of $\sigma_{\rm M}$, or even give it a strong dependence on the black hole mass, and so this value is very uncertain [69, 278]. A detailed study on the angular momentum sources excluding PBHs outside the binary, including the variability of $\sigma_{\rm M}$, was carried out by Garriga et al. [288].

Figure 4.7 shows the dependence of the detectable merger rate for a monochromatic distribution on the three parameters m_c, $f_{\rm PBH}$, and $\sigma_{\rm M}$. The dependence on $f_{\rm PBH}$ looks very similar to the plot of the intrinsic merger rate on this parameter

in fig. 3 of [69, published version]. This is expected, since the detectability and comoving volume factors do not introduce any additional dependence on $f_{\rm PBH}$, so the only difference is the monochromatic vs. lognormal mass distribution.

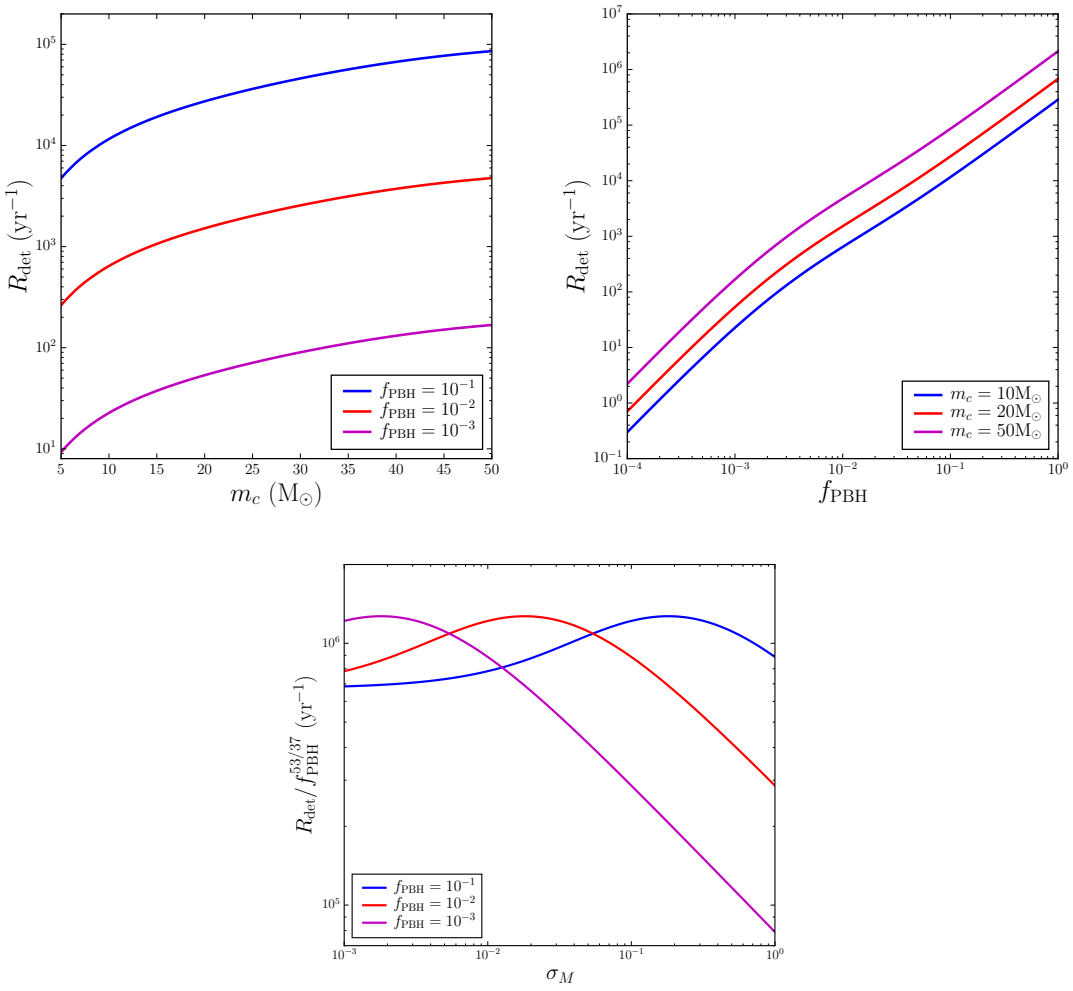

Figure 4.7: Merger rate dependence on the PBH mass m_c, the fraction of dark matter in PBHs $f_{\rm PBH}$, and the rescaled variance of matter density perturbations $\sigma_{\rm M}$ for a monochromatic mass distribution. The top two plots have $\sigma_{\rm M} = 0.006$ and the bottom plot has $m_c = 20$ $M_\odot$. The plots are generated using the method in [69], and would change with the additional effects considered in [288].

In the bottom panel of fig. 4.7, the detectable merger rate has been normalised by removing the global dependence on $f_{\rm PBH}$, leaving only the $f_{\rm PBH}$ dependence in the suppression factor S. This is to better highlight the relationship between the merger rate and the two parameters $\sigma_{\rm M}$ and $f_{\rm PBH}$, which have a degeneracy in S. The curves follow the same shape as each other, but with the peak in a different place that is determined by the relationship between $f_{\rm PBH}$ and $\sigma_{\rm M}$. To the right of

the peak, the torque on the binary is dominated by the matter perturbations, leading to binaries that have not yet merged. To the left of the peak, there is insufficient torquing from matter perturbations, and the binaries merged at redshifts too high to be detected by LIGO. However, on the left of the peak, the merger rate tends to a constant value, determined by the torquing generated by other PBHs, which is fixed by the value of $f_{\rm PBH}$. It can also be seen that the normalised merger rate varies by up to an order of magnitude for different values of $\sigma_{\rm M}$. This is enough to shift the LIGO constraints yielding different optimised mass distribution parameters, and so further study of the degeneracy of observables with the currently unknown value of $\sigma_{\rm M}$ is required.

4.5.2 Power-law distribution

Another commonly considered distribution for the masses of primordial (and astrophysical) black holes is a power-law ($\propto m^{-\alpha}$). The parameters for this model are the power to which the mass is raised and the lower/upper mass cutoffs if applicable. A scale-invariant primordial power spectrum generates $\alpha = 3/2$, due to the enhancement of the PBH energy density relative to the background radiation energy density after they have formed [76, 244]. With both a lower and upper mass cutoff[1], the normalised mass distribution is

$$\psi(m) = (\alpha - 1)\left[m_{\rm min}^{-(\alpha-1)} - m_{\rm max}^{-(\alpha-1)}\right]^{-1} m^{-\alpha}. \tag{4.9}$$

However, if the minimum mass is chosen to be too small, the suppression factor S in the merger rate calculation described above heavily suppresses the result. This is probably because a power-law mass distribution heavily favours the lighter end of the mass spectrum, meaning there are far more of these than there are heavier black holes. Physically, if there is a large population of lighter black holes and a smaller population of heavier black holes, then it will still be the heavier black holes that will merge as the lighter ones will not contribute significantly to the gravitational

[1] In practice, unless the mass function is close to scale invariant, only one cutoff is important.

force that causes a binary to form. However, the calculation described in [69], while being very thorough, does not capture this effect because it assumes that a PBH will form a binary with its nearest neighbour, rather than the neighbour contributing the largest gravitational force. In the equations, this manifests itself as a strong dependence on the average mass $\bar{m}$, and a suppression of the resulting merger rate.

To determine if this effect is important, and if so how problematic it is, the merger rates of three power-law mass distributions were calculated. Each mass distribution had the same values of $\alpha = 3/2$ and $m_{\max} = 100\ \mathrm{M}_\odot$, but had different $m_{\min}$ values ($5\ \mathrm{M}_\odot$, $1\ \mathrm{M}_\odot$, and $0.1\ \mathrm{M}_\odot$). For each distribution, the value of f_{PBH} was chosen such that the number density of PBHs with masses in the range 1–100 $\mathrm{M}_\odot$ was the same. Physical intuition would suggest that the total number of merger events with masses in this range would be similar for all three distributions. However, the intrinsic merger rate calculated using eq. (4.6) varied by two orders of magnitude between the distributions with minimum masses of $5\ \mathrm{M}_\odot$ and $1\ \mathrm{M}_\odot$, and by 30 orders of magnitude between $5\ \mathrm{M}_\odot$ and $0.1\ \mathrm{M}_\odot$. This is clearly a very significant problem that prevents the study of very broad mass distributions.

While it remains unclear how broad the mass distribution can be before this effect starts to become important, it is still desirable to compare the lognormal distribution with a power-law distribution. Therefore, the analysis was rerun with two power-law distributions. Both had $\alpha = 3/2$ and $m_{\max} = 100\ \mathrm{M}_\odot$. The minimum mass for the two distributions was $m_{\min} = 5\ \mathrm{M}_\odot$ and $10\ \mathrm{M}_\odot$ respectively. However, due to the problem described above, it is not clear if these results are reliable. The 1D merger rate distributions for these mass distribution can be seen in fig. E.1, and the 2D distributions are shown in figs. E.3 and E.4.

Another mass distribution of interest is that of PBHs generated at the QCD phase transition. For a scale invariant power spectrum, this formation gives a mass distribution like a power-law with $\alpha = 3/2$, but with an excess at around $1\ \mathrm{M}_\odot$, caused by the reduction in pressure at the QCD transition [45]. However, the relevant range of this mass distribution is extremely large (four orders of magnitude),

and so it is affected by the broadness problem discussed above. This mass distribution remains of interest due to its physical motivation, and should be studied once a reliable calculation of the merger rate for broad mass distributions is available.

4.6 Current LIGO data and constraints

In the above sections, the detection probability $p_{\rm det}$ was calculated using the default power spectral density (PSD) for the LIGO noise. This is the design sensitivity noise curve (`aLIGOZeroDetHighPower`). To compare to current LIGO data, a different PSD must be used. Therefore, the process above was carried out again with a detection probability generated using the `aLIGOEarlyHighSensitivityP1200087` PSD, which is a good approximation of the O1 and O2 sensitivities.

Figure 4.8 shows the same plots as in fig. 4.3, but for the O1O2 detectability. The posterior probability distributions for the 10 LIGO binary black hole (BBH) events are shown in the top panel of each plot [30]. For all four observables, we can see that none of the posteriors has a distribution that drastically disagrees with the shapes of the merger rate curves. The expected number of events in a given observable range can be found by summing the merger rate curves over this range and multiplying by the total observing time of the O1O2 dataset, which is 0.46 yr. We can also consider the other mass distributions. Figure 4.9 shows the O1O2 merger rate distributions for the lognormal distribution with the widths $\sigma = 0.6$ and 0.3, and the power-law distribution with $m_{\rm min} = 5\ {\rm M}_\odot$ and $m_{\rm min} = 10\ {\rm M}_\odot$.

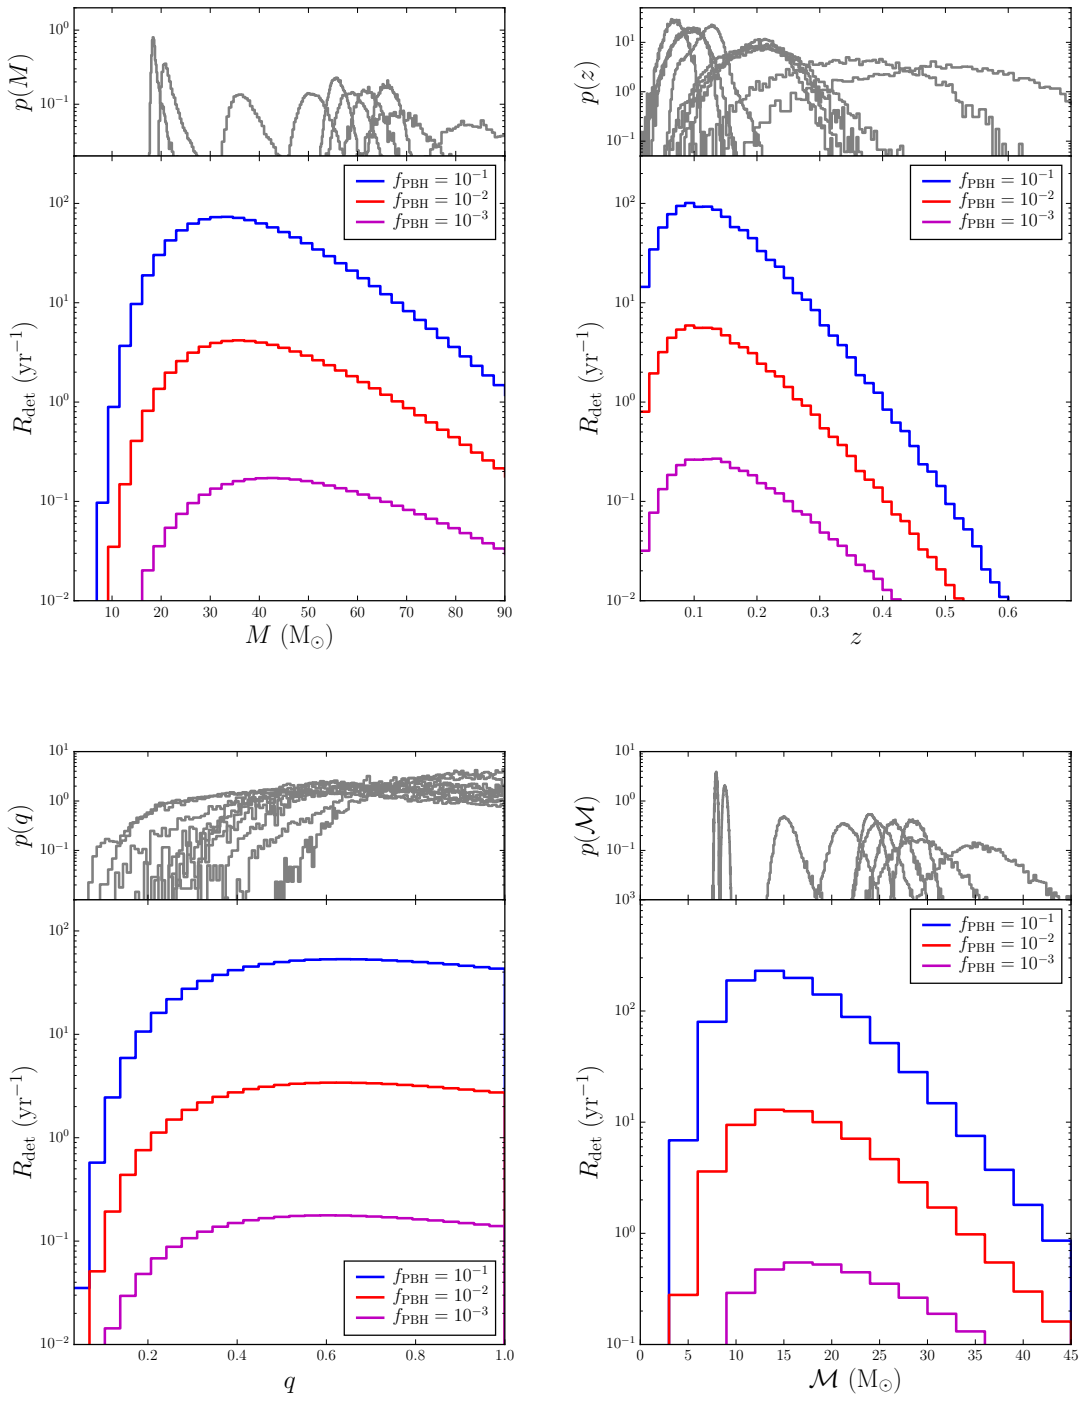

Figure 4.8: Merger rate distributions in total mass M, redshift z, mass ratio q, and chirp mass $\mathcal{M}$ for the lognormal mass distribution with $\sigma = 0.6$, at O1O2 sensitivity. The top panel in each plot shows the LIGO posteriors for the 10 BBH events.

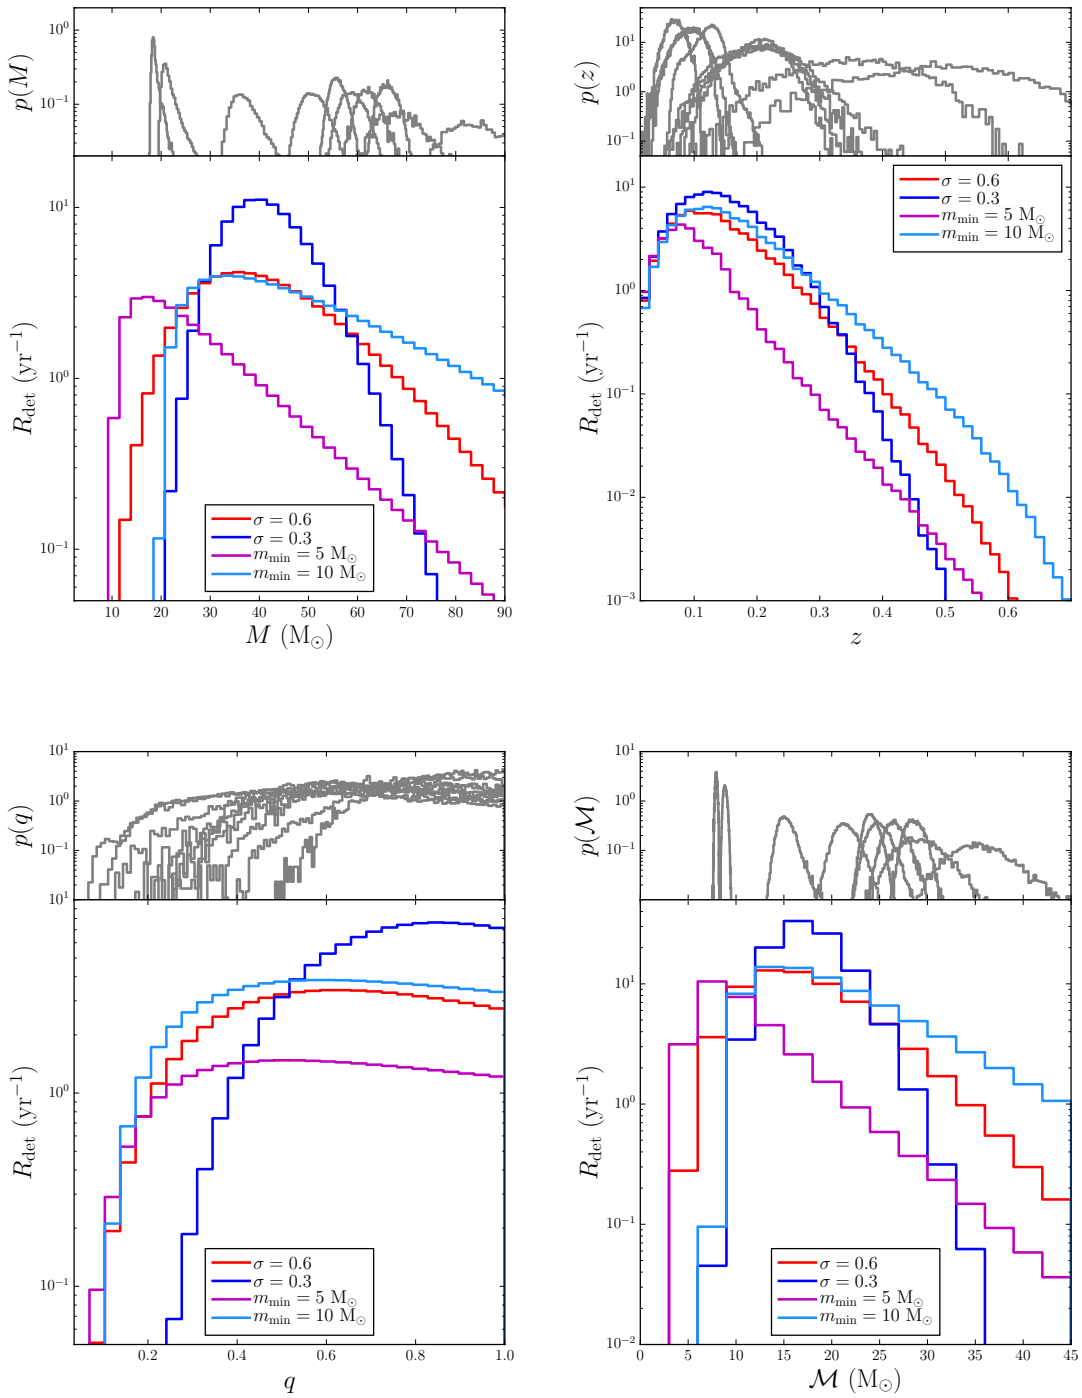

Figure 4.9: Merger rate distributions in total mass M, redshift z, mass ratio q, and chirp mass $\mathcal{M}$ for a lognormal mass distribution with $\sigma = 0.3$ and $\sigma = 0.6$, and a power-law mass distribution with $m_{\rm min} = 5$ M$_\odot$ and $m_{\rm min} = 10$ M$_\odot$, at O1O2 sensitivity. The top panel in each plot shows the LIGO posteriors for the 10 BBH events. All plots have $f_{\rm PBH} = 10^{-2}$.

2D distributions were also produced for this sensitivity. The first of these, $R_{\rm det}$ vs. (m_1, m_2) is shown in fig. 4.10 for the four mass distributions considered (the two lognormal distributions with different widths, and the two power-law distribu-

tions with different minimum masses). The black points show the LIGO values from the ten events and their 90% confidence ranges (note that these are 1D marginalised error bars, and the full contour would not just follow the shape of the errors). All the data points lie in the lower triangle due to the LIGO analysis imposing $m_2 < m_1$.

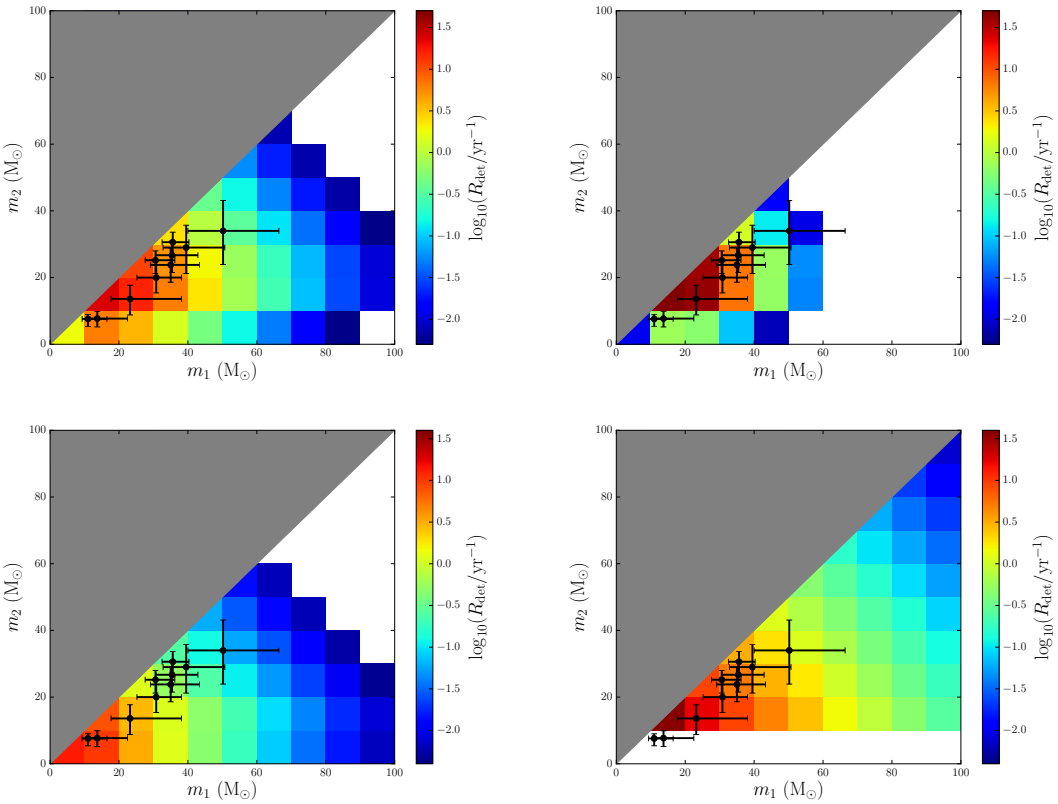

Figure 4.10: 2D merger rate distributions in individual masses for the four mass distributions considered, at O1O2 sensitivity. The top row has the lognormal mass distribution with widths $\sigma = 0.6$ (left) and 0.3 (right), and the bottom row has the power-law mass distribution with minimum mass 5 $M_\odot$ (left) and 10 $M_\odot$ (right). All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate and the grey triangle indicates the choice $m_2 < m_1$. The LIGO values and their 90% confidence limits are shown in black.

It can be seen by eye that some of these PBH models fail quite badly to match the observed locations of the LIGO events, whereas others look more acceptable. A statistical procedure is required to compare the observed and predicted distributions on the m_1–m_2 plane so that we can quantify the acceptability of the different models. This is relatively straightforward, but is made harder by the complex shape of the LIGO event posteriors in the m_1–m_2 plane that depends on the true parameters of the event (since these dictate the SNR of the detection). We are currently pursuing

a Bayesian inference of the PBH scenario using the LIGO events, but we present here a simplified attempt to quantify how well the PBH model fits the data. We attempt to capture the posteriors by assuming that all events have measurement error distributions that are independently lognormal in m_1 and m_2, adopting a typical rms of 0.2 in $\ln m$ (the final p-values are not highly sensitive to this choice). With this assumption, we can smooth our distribution on the m_1–m_2 plane and hence convert it to a 2D function from which the observed data can be treated as random error-free samples (note that the smoothing is performed in the symmetric m_1–m_2 plane before folding to impose the convention $m_2 < m_1$, and preserves the total merger rate over the m_1–m_2 plane).

In order to carry out the statistical comparison, there are two relevant questions: (1) is the observed number of events consistent with expectation? and (2) is the distribution over the m_1–m_2 plane correct? For the former, we have 10 events, and the likelihood of this number is to be computed using a Poisson distribution based on the predicted number from integration over the m_1–m_2 plane and multiplying by the observing time. As the Poisson distribution is exponentially sensitive to the expected number μ, the probability can drop off rapidly as μ moves away from the observed number of events. The expected value for each case is shown in table 4.1, with the corresponding probabilities adjacent. For the 2D distribution, we normalise the rate distribution over the m_1–m_2 plane to obtain a 2D probability density. This could then be compared with the data using the 2D Kolmogorov-Smirnov (KS) test [289]. The only drawback with this approach is that the KS test is rather insensitive to whether a few points lie in an area of the plane with zero density, as does seem to be the case here. We therefore prefer a simpler statistic, which is just the overall likelihood of the data (product of the 2D density at the location of the 10 events). The expected distribution of this statistic can be readily obtained by drawing 10 points independently and at random from the 2D distribution multiple times. In this way, we can identify models whose likelihood is sufficiently low that they can be ruled out at an interesting level of significance. Thus we obtain two

two-tailed frequentist p-values based on the absolute number of events and on their distribution. For the present purpose, it is probably the second of these that is of more interest, since it addresses directly the initial question of whether the PBH model can consistently generate nearly equal mass binaries.

Table 4.1: Expected number μ and probability p of the LIGO results based on the number of observed mergers $N = 10$ with observing time $T = 0.46$ yr, and their distribution in the m_1–m_2 plane for four PBH mass distributions: A = Lognormal ($m_c = 20$ M$_\odot$, $\sigma = 0.6$), B = Lognormal ($m_c = 20$ M$_\odot$, $\sigma = 0.3$), C = Power-law ($m_{\min} = 5$ M$_\odot$, $\alpha = 3/2$), D = Power-law ($m_{\min} = 10$ M$_\odot$, $\alpha = 3/2$).

Test	N				m_1–m_2	
$f_{\rm PBH}$	10^{-2}		10^{-3}		10^{-2}	10^{-3}
Model	μ	p	μ	p	p	p
A	31	1.2×10^{-5}	1.6	1.3×10^{-6}	0.15	0.37
B	47	6.4×10^{-11}	1.9	4.1×10^{-6}	$\lesssim 10^{-4}$	5×10^{-4}
C	15	0.12	0.96	6.3×10^{-9}	2×10^{-4}	0.30
D	37	1.5×10^{-7}	2.0	9.3×10^{-6}	0.15	0.45

The results of this exercise are collected in table 4.1 for the four extended mass distributions discussed above, labelled A–D as in the caption to table 4.1, and for two values of $f_{\rm PBH}$ each (10^{-2} and 10^{-3}). It can be seen that only model C is compatible at the 5% level with the LIGO number of observed mergers for either of the values of $f_{\rm PBH}$. For the other models, $f_{\rm PBH} = 10^{-2}$ significantly overproduces merger events, while for all the models $f_{\rm PBH} = 10^{-3}$ does not produce enough, assuming that all the LIGO events are of primordial origin. For any of these models, it will be possible to choose a value of $f_{\rm PBH}$ between 10^{-2} and 10^{-3} that will match the observed number of events. The test of the m_1–m_2 plane shows better agreement, with only model B and model C at $f_{\rm PBH} = 10^{-2}$ disfavoured at the 5% level. These three probabilities are limited by shot noise due to the number of samples. While the global factor of $f_{\rm PBH}$ has been normalised out, there is still a non-trivial degeneracy between this parameter and the shape of the merger rate distribution in the m_1–m_2 plane. The p-values for model C vary greatly between the 10^{-2} and 10^{-3} case, due to the rapid suppression of the merger rate for higher values of m_1 and m_2 with $f_{\rm PBH} = 10^{-2}$. This is likely to be an effect of the broadness problem discussed in section 4.5.2,

and so the small p-value for model C with $f_{\rm PBH} = 10^{-2}$ should be taken with some hesitation. It is clear however, that a PBH scenario can explain the shape of the merger rate distribution in the m_1–m_2 plane and, with an appropriate $f_{\rm PBH}$ for normalisation, the total number of observed mergers.

4.7 Conclusions

Since the LIGO–Virgo collaborations have begun the era of direct gravitational wave detection, there has been great interest in the origin of the black holes whose mergers they have detected. We have focused on observational methods to discriminate between primordial and astrophysical black holes. We have explored a larger range of observables as a probe of the PBH scenario than have been greatly explored previously, with a particular focus on the mass ratio of the BHs that merged. Astrophysical BHs which form from a common envelope may dynamically equalise their masses and hence predict $q \approx 1$, while still forming a large range of masses between different binary pairs [109, 110, 112, 113]. In contrast, PBHs form before they become part of a binary system, suggesting that $q \sim 1$ is only likely to occur if the PBH mass distribution is narrow. However, a narrow mass distribution may be in tension with the broad range of chirp (or total) masses observed in the 10 binary black hole merger events detected by LIGO–Virgo to date.

In order to analyse this problem, we have made merger rate distributions incorporating the LIGO detectability for the O1O2 LIGO–Virgo sensitivity curves, and compared to the LIGO data from the O1 and O2 runs. A rough analysis shows that the LIGO data have begun to apply constraints on the form and parameters of the PBH mass distribution, which is only possible using the detectable merger rate. Three types of PBH mass distribution were considered: lognormal, power-law, and monochromatic, although the monochromatic distribution is already ruled out by the variation in masses detected by LIGO. Table 4.1 shows the results of this analysis, indicating that the narrow lognormal (model B) is disfavoured at the 5% level.

We have also calculated the expected detectable merger rate distributions at LIGO design sensitivity for the mass ratio q, redshift z, total mass M and the chirp mass $\mathcal{M}$, and compared the results to distributions generated for astrophysical BHs. These distributions, which take into account the detection probability of any given merger, overcome the problem of comparing the intrinsic PBH merger rate to the intrinsic merger rate estimated by LIGO, each of which assumes a different mass distribution for the component black holes. The LIGO estimation method is described in section VII of [30].

With the many new events expected to be detected by LIGO in the future, the PBH mass distribution can be probed in great detail, and following on from the methods developed in this paper, the best fits for any form of the PBH mass distribution can be found. A complication in these fits would be the uncertainty of the source of any given merger. In the future, a fit simultaneously incorporating the two potential BH populations (primordial and astrophysical) should be made, and it may be possible to rule out all of the BHs being primordial in origin. One potential discriminant is the spin of the BHs, which is expected to be negligibly small for PBHs formed during radiation domination [58]. Even if the spin cannot discriminate, it must be taken into account for astrophysical BHs, which may have significant spins. We are currently pursuing a fully Bayesian inference of the PBH merger scenario with current (and future) LIGO data.

Uncertainties in the PBH merger rate calculation remain an open issue. The method applied in this paper from [69] builds a strong framework for the merger rate calculation, but there are further considerations, such as the torquing effects from matter and radiation perturbations [288], and the uncertainty of how frequently binaries are disrupted between formation and merger [69, 278]. We have also found that the current calculation cannot be used in the case of a very broad mass distribution. Detailed simulations and further analytic developments of PBH binary formation, disruption, and merger events are essential to ensure that the fits to the current and future LIGO data are accurate.

Acknowledgements

AG thanks Ville Vaskonen, Davide Gerosa, and Zu-Cheng Chen for discussions and clarification. AH thanks Derek Inman for helpful discussions. AG is funded by a Royal Society Studentship by means of a Royal Society Enhancement Award. CB is funded by a Royal Society University Research Fellowship. AH is supported by an STFC Consolidated Grant.

Chapter 5

Conclusions and future work

Many questions about our universe remain unanswered. What are the natures of dark matter and dark energy? How did the universe begin and evolve? What provided the seeds for the structure we see today? Primordial black holes may provide the answers to some of these challenges. They can provide an explanation for some or all of the dark matter in the universe today. Additionally, their formation in the early universe means that they are a unique probe of the physics at those times, such as inflation. Crucially, they provide insight on the small-scale behaviour, far away from the knowledge we have from the CMB. In this thesis, we have studied PBHs at different stages in the history of the universe, from their formation at very early times to binary mergers detectable at present with instruments such as LIGO.

Beginning in the past, we examined the formation of PBHs from large inflaton overdensities in chapter 2. While there are well-defined procedures to calculate the PBH abundance and mass distribution from a known peak in the primordial power spectrum, there are a number of effects that must be taken into account carefully, such as critical collapse and the non-linear relation between ζ and δ. Additionally, there are different methods, e.g. Press–Schechter vs. peaks theory, as well as the choice of window function applied to the power spectrum, that may modify the result of the calculation. We carried out a thorough test of these choices, finding that the power spectrum required to produce a fixed PBH abundance only differed by $\mathcal{O}(10\%)$ between the different choices, with a similar shape for the resulting mass

distribution. This shows that these choices are currently unimportant, although these differences will become important when considering accurate data from future experiments. Furthermore, we calculated robust constraints on the primordial power spectrum from PBH sources, and compared with those from pulsar timing arrays and μ-distortions, showing that PBHs with masses less than $\sim 10^4$ $M_\odot$ are not in tension with other constraints. However, future constraints from μ-distortions and the stochastic gravitational wave background will change this result, and a non-detection in a PIXIE-like experiment, as well as the SKA, LISA, and the Einstein telescope will rule out the possibility of PBHs existing across the vast majority of masses which have not evaporated by the present day.

Although the procedure to obtain the PBH mass distribution is well known and currently resistant to alternative choices in the method and window function, it is computationally expensive to calculate, making it infeasible for use in optimisation procedures such as an MCMC analysis. Therefore, models that capture the shape of the distribution in a simple function are required. The most commonly used model is the lognormal, although work from almost 25 years ago suggested that this will not capture the underlying shape when the effects of critical collapse dominate the calculation [185, 245]. In chapter 3, we confirmed that this 25 year old conclusion remains valid despite modifications to the PBH mass distribution calculation such as the integration over all formation times, and tested a number of alternative models that have the potential to capture the underlying shape more accurately. We found that two three-parameter models (the skew-lognormal and a generalised model motivated by the physics of critical collapse) provide a significant improvement over the lognormal across a broad range of power spectrum peak widths, and suggest that one of these models be used in the case of accurate datasets or strongly mass-dependent constraints, such as those from PBH evaporation.

Moving towards the present, we considered the mergers of PBHs and the possibility that they can explain the LIGO gravitational wave detections in chapter 4. Utilising a detailed model of PBH binary formation and merger, we produced distri-

butions of the detectable PBH merger rate across a number of key observables, such as the chirp mass and mass ratio, for the LIGO O1O2 and design sensitivities. We compared the design sensitivity distributions to publicly available analogues for the mergers of astrophysical black holes, highlighting the mass ratio as the observable most likely to be able to distinguish between primordial and astrophysical merger events. We also carried out a statistical test comparing the O1O2 merger rates to the 10 detected events in LIGO's GWTC-1 [30], showing that PBHs could satisfactorily explain these signals, although the PBH abundance and mass distribution have to be chosen quite carefully. In a related work not included in this thesis [255], we carried out a Bayesian model comparison for the 10 O1O2 events between the primordial scenario and two simple parametrisations used by LIGO with some astrophysical motivation, We tested a range of PBH mass distributions based on the lognormal, including the skew-lognormal discussed in chapter 3. We found that the PBH scenario is heavily disfavoured compared to the astrophysical parametrisations, although a mixed model may still be required to explain these and future LIGO signals.

If the past and present have been exciting for PBHs, the future looks even brighter. As discussed in chapter 2, future gravitational wave detectors have the sensitivity required to rule out large power spectrum values, and hence PBHs formed from inflaton perturbations across a very large range of masses. This would force a shift towards considering other formation methods, such as those briefly discussed in section 1.2.1. Alternatively, this can be framed in a more optimistic way for the field, as the statement that if even a single PBH formed from inflation exists within these sensitive ranges, the associated gravitational waves will be detected by one of these experiments. Such a detection would fundamentally change our entire view of inflation and early universe cosmology, requiring a new paradigm of models that can provide these large power spectrum amplitudes. If the NANOGrav excess of gravitational waves is demonstrated to be a stochastic background caused by scalar-induced gravitational waves, it is very likely there will be PBHs in this mass

range. The late-time effects of PBHs will also become more and more probeable, with the increasing network of gravitational wave detectors leading to a large growth in the number of binary black hole merger events in the 1–100 $M_\odot$ range. As these events become more numerous, it will become apparent whether astrophysical black holes from various channels can explain them, or whether a primordial component is required as well. The PBH mass distribution will be crucial in making such a statement, and more work must be done to ensure that its calculation is valid. The choice of method and window function considered in chapter 2 will become important, and the physically correct options will need to be determined, probably through the use of simulations. Other factors will also have to be taken into account, such as non-Gaussianities and inflationary effects such as quantum diffusion. For a field that is over 50 years old, there are no signs that PBHs are done yet.

Bibliography

[1] A. Einstein, *Die Grundlage der allgemeinen Relativitätstheorie*, *Annalen der Physik* **354(7)** (1916) 769. Cited on 2.

[2] J. K. Yadav, J. S. Bagla and N. Khandai, *Fractal dimension as a measure of the scale of homogeneity*, *Monthly Notices of the Royal Astronomical Society* **405(3)** (2010) 2009 [1001.0617]. Cited on 2.

[3] D. Saadeh et al., *How Isotropic is the Universe?*, *Physical Review Letters* **117(13)** (2016) 131302 [1605.07178]. Cited on 2.

[4] F. S. Labini, D. Tekhanovich and Yu. V. Baryshev, *Spatial density fluctuations and selection effects in galaxy redshift surveys*, *Journal of Cosmology and Astroparticle Physics* **2014(07)** (2014) 035 [1406.5899]. Cited on 2.

[5] D. Alonso et al., *Homogeneity and isotropy in the Two Micron All Sky Survey Photometric Redshift catalogue*, *Monthly Notices of the Royal Astronomical Society* **449(1)** (2015) 670 [1412.5151]. Cited on 2.

[6] G. De Marzo, F. Sylos Labini and L. Pietronero, *Zipf's law for cosmic structures: How large are the greatest structures in the universe?*, *Astronomy & Astrophysics* **651** (2021) A114 [2105.06110]. Cited on 2.

[7] G. Gamow, *The Evolution of the Universe*, *Nature* **162(4122)** (1948) 680. Cited on 3.

[8] R. A. Alpher and R. Herman, *Evolution of the Universe*, *Nature* **162** (1948) 774. Cited on 3.

[9] A. McKellar, *Molecular Lines from the Lowest States of Diatomic Molecules Composed of Atoms Probably Present in Interstellar Space*, *Publications of the Dominion Astrophysical Observatory* **7(15)** (1941) 251. Cited on 3.

[10] A. A. Penzias and R. W. Wilson, *A Measurement of Excess Antenna Temperature at 4080 Mc/s*, *The Astrophysical Journal* **142** (1965) 419. Cited on 3.

[11] J. C. Mather et al., *A Preliminary Measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (COBE) Satellite*, *The Astrophysical Journal Letters* **354** (1990) L37. Cited on 3.

[12] W. Hu and S. Dodelson, *Cosmic Microwave Background Anisotropies*, *Annual Review of Astronomy and Astrophysics* **40** (2002) 171 [astro-ph/0110414]. Cited on 3.

[13] R. H. Dicke, *Gravitation and the Universe*, American Philosophical Society (1970). Cited on 4.

[14] A. A. Starobinsky, *A new type of isotropic cosmological models without singularity*, *Physics Letters B* **91(1)** (1980) 99. Cited on 4.

[15] D. Kazanas, *Dynamics of the universe and spontaneous symmetry breaking*, *The Astrophysical Journal Letters* **241** (1980) L59. Cited on 4.

[16] A. H. Guth, *Inflationary universe: A possible solution to the horizon and flatness problems*, *Physical Review D* **23(2)** (1981) 347. Cited on 4.

[17] K. Sato, *Cosmological baryon-number domain structure and the first order phase transition of a vacuum*, *Physics Letters B* **99(1)** (1981) 66. Cited on 4.

[18] M. B. Einhorn and K. Sato, *Monopole production in the very early universe in a first-order phase transition*, *Nuclear Physics B* **180(3)** (1981) 385. Cited on 4.

[19] A. D. Linde, *A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems*, *Physics Letters B* **108(6)** (1982) 389. Cited on 4.

[20] S. W. Hawking and I. L. Moss, *Supercooled phase transitions in the very early universe*, *Physics Letters B* **110(1)** (1982) 35. Cited on 4.

[21] A. Albrecht and P. J. Steinhardt, *Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking*, *Physical Review Letters* **48(17)** (1982) 1220. Cited on 4.

[22] M. Sasaki, *Large Scale Quantum Fluctuations in the Inflationary Universe*, *Progress of Theoretical Physics* **76(5)** (1986) 1036. Cited on 9.

[23] V. F. Mukhanov, *Quantum theory of gauge-invariant cosmological perturbations*, *Journal of Experimental and Theoretical Physics* **67(7)** (1988) 1297. Cited on 9.

[24] Ya. B. Zel'dovich and I. D. Novikov, *The Hypothesis of Cores Retarded during Expansion and the Hot Cosmological Model*, *Soviet Astronomy* **10(4)** (1967) 602. Cited on 10, 32, 60, 76.

[25] S. Hawking, *Gravitationally Collapsed Objects of Very Low Mass*, *Monthly Notices of the Royal Astronomical Society* **152(1)** (1971) 75. Cited on 10, 32, 60, 76.

[26] B. J. Carr and S. W. Hawking, *Black Holes in the Early Universe*, *Monthly Notices of the Royal Astronomical Society* **168(2)** (1974) 399. Cited on 10, 32, 60, 76.

[27] W. Israel, *Event Horizons in Static Vacuum Space-Times*, *Physical Review* **164(5)** (1967) 1776. Cited on 10.

[28] W. Israel, *Event horizons in static electrovac space-times*, *Communications in Mathematical Physics* **8** (1968) 245. Cited on 10.

[29] B. Carter, *Axisymmetric Black Hole Has Only Two Degrees of Freedom*, *Physical Review Letters* **26(6)** (1971) 331. Cited on 10.

[30] LIGO Scientific Collaboration and Virgo Collaboration, *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)** (2019) 031040 [1811.12907]. Cited on 10, 15, 22, 23, 27, 76, 81, 93, 100, 104.

[31] LIGO Scientific Collaboration and Virgo Collaboration, *GW190412: Observation of a binary-black-hole coalescence with asymmetric masses*, *Physical Review D* **102(4)** (2020) 043015 [2004.08342]. Cited on 10.

[32] LIGO Scientific Collaboration and Virgo Collaboration, *GW190425: Observation of a Compact Binary Coalescence with Total Mass* $\sim 3.4 M_{\odot}$, *The Astrophysical Journal Letters* **892(1)** (2020) L3 [2001.01761]. Cited on 10, 14, 61.

[33] LIGO Scientific Collaboration and Virgo Collaboration, *GW190521: A Binary Black Hole Merger with a Total Mass of* 150 $M_{\odot}$, *Physical Review Letters* **125(10)** (2020) 101102 [2009.01075]. Cited on 10, 14, 19, 20, 61.

[34] LIGO Scientific Collaboration and Virgo Collaboration, *GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object*, *The Astrophysical Journal Letters* **896(2)** (2020) L44 [2006.12611]. Cited on 10, 14, 20, 54, 61.

[35] LIGO Scientific Collaboration and Virgo Collaboration, *GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run*, *Physical Review X* **11(2)** (2021) 021053 [2010.14527]. Cited on 10, 15, 22, 27.

[36] S. M. Leach, M. Sasaki, D. Wands and A. R. Liddle, *Enhancement of superhorizon scale inflationary curvature perturbations*, *Physical Review D* **64(2)** (2001) 023512 [astro-ph/0101406]. Cited on 11.

[37] C. Germani and T. Prokopec, *On primordial black holes from an inflection point*, *Physics of the Dark Universe* **18** (2017) 6 [1706.04226]. Cited on 11.

[38] K. Dimopoulos, *Ultra slow-roll inflation demystified*, *Physics Letters B* **775** (2017) 262 [1707.05644]. Cited on 11.

[39] M. P. Hertzberg and M. Yamada, *Primordial black holes from polynomial potentials in single field inflation*, *Physical Review D* **97(8)** (2018) 083509 [1712.09750]. Cited on 11.

[40] S. S. Mishra and V. Sahni, *Primordial black holes from a tiny bump/dip in the inflaton potential*, *Journal of Cosmology and Astroparticle Physics* **2020(04)** (2020) 007 [1911.00057]. Cited on 11.

[41] K.-W. Ng and Y.-P. Wu, *Constant-rate inflation: primordial black holes from conformal weight transitions*, (2021) [2102.05620]. Cited on 11.

[42] K. Inomata, E. McDonough and W. Hu, *Primordial Black Holes Arise When The Inflaton Falls*, (2021) [2104.03972]. Cited on 11.

[43] B. Carr, K. Kohri, Y. Sendouda and J. Yokoyama, *Constraints on Primordial Black Holes*, (2020) [2002.12778]. Cited on 12, 17, 18, 32, 50, 51, 57, 60.

[44] M. Yu. Khlopov and A. G. Polnarev, *Primordial black holes as a cosmological test of grand unification*, *Physics Letters B* **97(3–4)** (1980) 383. Cited on 13.

[45] C. T. Byrnes, M. Hindmarsh, S. Young and M. R. S. Hawkins, *Primordial black holes with an accurate QCD equation of state*, *Journal of Cosmology and Astroparticle Physics* **2018(08)** (2018) 041 [1801.06138]. Cited on 13, 54, 61, 62, 63, 92.

[46] B. J. Carr, J. H. Gilbert and J. E. Lidsey, *Black hole relics and inflation: Limits on blue perturbation spectra*, *Physical Review D* **50(8)** (1994) 4853 [astro-ph/9405027]. Cited on 13.

[47] LIGO Scientific Collaboration and Virgo Collaboration, *Search for Subsolar Mass Ultracompact Binaries in Advanced LIGO's Second Observing Run*, *Physical Review Letters* **123(16)** (2019) 161102 [1904.08976]. Cited on 14.

[48] A. H. Nitz and Y.-F. Wang, *Search for Gravitational Waves from the Coalescence of Subsolar Mass and Eccentric Compact Binaries*, *The Astrophysical Journal* **915(1)** (2021) 54 [2102.00868]. Cited on 14.

[49] A. H. Nitz and Y.-F. Wang, *Search for Gravitational Waves from the Coalescence of Subsolar-Mass Binaries in the First Half of Advanced LIGO and Virgo's Third Observing Run*, *Physical Review Letters* **127(15)** (2021) 151101 [2106.08979]. Cited on 14.

[50] K. S. Phukon et al., *The hunt for sub-solar primordial black holes in low mass ratio binaries is open*, (2021) [2105.11449]. Cited on 14.

[51] M. Ricotti, *Bondi Accretion in the Early Universe*, *The Astrophysical Journal* **662(1)** (2007) 53 [0706.0864]. Cited on 15.

[52] M. Ricotti, J. P. Ostriker and K. J. Mack, *Effect of Primordial Black Holes on the Cosmic Microwave Background and Cosmological Parameter Estimates*, *The Astrophysical Journal* **680(2)** (2008) 829 [0709.0524]. Cited on 15.

[53] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *Constraints on primordial black holes: The importance of accretion*, *Physical Review D* **102(4)** (2020) 043505 [2003.12589]. Cited on 15, 61.

[54] T. Chiba and S. Yokoyama, *Spin distribution of primordial black holes*, *Progress of Theoretical and Experimental Physics* **2017(8)** (2017) 083E01 [1704.06573]. Cited on 15, 82.

[55] V. De Luca et al., *The initial spin probability distribution of primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2019(05)** (2019) 018 [1903.01179]. Cited on 15.

[56] K. A. Postnov and N. A. Mitichkin, *Spins of primordial binary black holes before coalescence*, *Journal of Cosmology and Astroparticle Physics* **2019(06)** (2019) 044 [1904.00570]. Cited on 15, 83.

[57] M. Mirbabayi, A. Gruzinov and J. Noreña, *Spin of primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2020(03)** (2020) 017 [1901.05963]. Cited on 15, 82.

[58] N. Fernandez and S. Profumo, *Unraveling the origin of black holes from effective spin measurements with LIGO-Virgo*, *Journal of Cosmology and Astroparticle Physics* **2019(08)** (2019) 022 [1905.13019]. Cited on 15, 61, 83, 100.

[59] J. García-Bellido, J. F. Nuño Siles and E. Ruiz Morales, *Bayesian analysis of the spin distribution of LIGO/Virgo black holes*, *Physics of the Dark Universe* **31** (2021) 100791 [2010.13811]. Cited on 15.

[60] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *The evolution of primordial black holes and their final observable spins*, *Journal of Cosmology and Astroparticle Physics* **2020(04)** (2020) 052 [2003.02778]. Cited on 15, 61.

[61] Y. Ali-Haïmoud, *Correlation Function of High-Threshold Regions and Application to the Initial Small-Scale Clustering of Primordial Black Holes*, *Physical Review Letters* **121(8)** (2018) 081304 [1805.05912]. Cited on 15, 79.

[62] V. Desjacques and A. Riotto, *Spatial clustering of primordial black holes*, *Physical Review D* **98(12)** (2018) 123533 [1806.10414]. Cited on 15, 79.

[63] T. Bringmann, P. F. Depta, V. Domcke and K. Schmidt-Hoberg, *Towards closing the window of primordial black holes as dark matter: The case of large clustering*, *Physical Review D* **99(6)** (2019) 063532 [1808.05910]. Cited on 15, 79.

[64] T. Suyama and S. Yokoyama, *Clustering of primordial black holes with non-Gaussian initial fluctuations*, *Progress of Theoretical and Experimental Physics* **2019(10)** (2019) [1906.04958]. Cited on 15, 79.

[65] A. Moradinezhad Dizgah, G. Franciolini and A. Riotto, *Primordial black holes from broad spectra: abundance and clustering*, *Journal of Cosmology and Astroparticle Physics* **2019(11)** (2019) 001 [1906.08978]. Cited on 15, 79.

[66] T. Matsubara, T. Terada, K. Kohri and S. Yokoyama, *Clustering of primordial black holes formed in a matter-dominated epoch*, *Physical Review D* **100(12)** (2019) 123544 [1909.04053]. Cited on 15, 79.

[67] S. Clesse and J. García-Bellido, *The clustering of massive Primordial Black Holes as Dark Matter: Measuring their mass distribution with advanced LIGO*, *Physics of the Dark Universe* **15** (2017) 142 [1603.05234]. Cited on 15, 76, 78, 81.

[68] M. Raidal, V. Vaskonen and H. Veermäe, *Gravitational waves from primordial black hole mergers*, *Journal of Cosmology and Astroparticle Physics* **2017(09)** (2017) 037 [1707.01480]. Cited on 15, 78.

[69] M. Raidal, C. Spethmann, V. Vaskonen and H. Veermäe, *Formation and evolution of primordial black hole binaries in the early universe*, *Journal of Cosmology and Astroparticle Physics* **2019(02)** (2019) 018 [1812.01930]. Cited on 15, 40, 78, 79, 80, 84, 85, 87, 89, 90, 92, 100.

[70] Q. Ding, T. Nakama, J. Silk and Y. Wang, *Detectability of gravitational waves from the coalescence of massive primordial black holes with initial clustering*, *Physical Review D* **100(10)** (2019) 103003 [1903.07337]. Cited on 15, 79.

[71] S. Young and C. T. Byrnes, *Initial clustering and the primordial black hole merger rate*, *Journal of Cosmology and Astroparticle Physics* **2020(03)** (2020) 004 [1910.06077]. Cited on 15, 40, 79.

[72] K. Jedamzik, *Primordial black hole dark matter and the LIGO/Virgo observations*, *Journal of Cosmology and Astroparticle Physics* **2020(09)** (2020) 022 [2006.11172]. Cited on 15, 40.

[73] M. Trashorras, J. García-Bellido and S. Nesseris, *The Clustering Dynamics of Primordial Black Boles in N-Body Simulations*, *Universe* **7(1)** (2021) 1 [2006.15018]. Cited on 15.

[74] K. Jedamzik, *Consistency of Primordial Black Hole Dark Matter with LIGO/Virgo Merger Rates*, *Physical Review Letters* **126(5)** (2021) 051302 [2007.03565]. Cited on 15, 40.

[75] P. Montero-Camacho et al., *Revisiting constraints on asteroid-mass primordial black holes as dark matter candidates*, *Journal of Cosmology and Astroparticle Physics* **2019(08)** (2019) 031 [1906.05950]. Cited on 17.

[76] B. Carr et al., *Primordial black hole constraints for extended mass functions*, *Physical Review D* **96(2)** (2017) 023514 [1705.05567]. Cited on 18, 51, 61, 67, 73, 77, 80, 91.

[77] N. Bellomo, J. L. Bernal, A. Raccanelli and L. Verde, *Primordial black holes as dark matter: converting constraints from monochromatic to extended mass distributions*, *Journal of Cosmology and Astroparticle Physics* **2018(01)** (2018) 004 [1709.07467]. Cited on 18, 51, 61, 73, 77.

[78] A. Celotti, J. C. Miller and D. W. Sciama, *Astrophysical evidence for the existence of black holes*, *Classical and Quantum Gravity* **16(12A)** (1999) A3 [astro-ph/9912186]. Cited on 18.

[79] A. M. Ghez, B. L. Klein, M. Morris and E. E. Becklin, *High Proper-Motion Stars in the Vicinity of Sagittarius A*: Evidence for a Supermassive Black Hole at the Center of Our Galaxy*, *The Astrophysical Journal* **509(2)** (1998) 678 [astro-ph/9807210]. Cited on 18.

[80] LIGO Scientific Collaboration and Virgo Collaboration, *Observation of Gravitational Waves from a Binary Black Hole Merger*, *Physical Review Letters* **116(6)** (2016) 061102 [1602.03837]. Cited on 18, 22, 24.

[81] J. E. McClintock and R. A. Remillard, *Black hole binaries*, in *Compact Stellar X-ray Sources*, W. Lewin and M. van der Klis, eds., p. 157, Cambridge University Press, (2006) [astro-ph/0306213]. Cited on 18.

[82] Event Horizon Telescope Collaboration, *First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole*, *The Astrophysical Journal Letters* **875(1)** (2019) L4 [1906.11241]. Cited on 18.

[83] E. Bañados et al., *An 800-million-solar-mass black hole in a significantly neutral Universe at a redshift of 7.5*, *Nature* **553** (2017) 473 [1712.01860]. Cited on 19.

[84] F. Wang et al., *A Luminous Quasar at Redshift 7.642*, *The Astrophysical Journal Letters* **907(1)** (2021) L1 [2101.03179]. Cited on 19.

[85] A. Patruno, S. Portegies Zwart, J. Dewi and C. Hopman, *The ultraluminous X-ray source in M82: an intermediate-mass black hole with a giant companion*, *Monthly Notices of the Royal Astronomical Society: Letters* **370(1)** (2006) L6 [astro-ph/0602230]. Cited on 19.

[86] T. J. Maccarone, A. Kundu, S. E. Zepf and K. L. Rhode, *A black hole in a globular cluster*, *Nature* **445** (2007) 183 [astro-ph/0701310]. Cited on 19.

[87] I. V. Chilingarian et al., *A Population of Bona Fide Intermediate-mass Black Holes Identified as Low-luminosity Active Galactic Nuclei*, *The Astrophysical Journal* **863(1)** (2018) 1 [1805.01467]. Cited on 19.

[88] J. R. Oppenheimer and G. M. Volkoff, *On Massive Neutron Cores*, *Physical Review* **55(4)** (1939) 374. Cited on 19.

[89] H. Gao, B. Zhang and H.-J. Lü, *Constraints on binary neutron star merger product from short GRB observations*, *Physical Review D* **93(4)** (2016) 044065 [1511.00753]. Cited on 19.

[90] M. Shibata, E. Zhou, K. Kiuchi and S. Fujibayashi, *Constraint on the maximum mass of neutron stars using GW170817 event*, *Physical Review D* **100(2)** (2019) 023015 [1905.03656]. Cited on 19.

[91] J. M. Lattimer, *Introduction to neutron stars*, *AIP Conference Proceedings* **1645(1)** (2015) 61. Cited on 19.

[92] R. W. Romani et al., *PSR J1311–3430: A HEAVYWEIGHT NEUTRON STAR WITH A FLYWEIGHT HELIUM COMPANION*, *The Astrophysical Journal Letters* **760(2)** (2012) L36 [1210.6884]. Cited on 19.

[93] C. D. Bailyn, R. K. Jain, P. Coppi and J. A. Orosz, *The Mass Distribution of Stellar Black Holes*, *The Astrophysical Journal* **499(1)** (1998) 367 [astro-ph/9708032]. Cited on 19.

[94] F. Özel, D. Psaltis, R. Narayan and J. E. McClintock, *THE BLACK HOLE MASS DISTRIBUTION IN THE GALAXY*, *The Astrophysical Journal* **725(2)** (2010) 1918 [1006.2834]. Cited on 19.

[95] W. M. Farr et al., *THE MASS DISTRIBUTION OF STELLAR-MASS BLACK HOLES*, *The Astrophysical Journal* **741(2)** (2011) 103 [1011.1459]. Cited on 19.

[96] L. Kreidberg, C. D. Bailyn, W. M. Farr and V. Kalogera, *MASS MEASUREMENTS OF BLACK HOLES IN X-RAY TRANSIENTS: IS THERE A MASS GAP?*, *The Astrophysical Journal* **757(1)** (2012) 36 [1205.1805]. Cited on 20.

[97] T. A. Thompson et al., *A noninteracting low-mass black hole–giant star binary system*, *Science* **366(6465)** (2019) 637 [1806.02751]. Cited on 20.

[98] LIGO Scientific Collaboration and Virgo Collaboration, *Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo*, *The Astrophysical Journal Letters* **882(2)** (2019) L24 [1811.12940]. Cited on 20.

[99] A. Heger and S. E. Woosley, *The Nucleosynthetic Signature of Population III*, *The Astrophysical Journal* **567(1)** (2002) 532 [astro-ph/0107037]. Cited on 20.

[100] S. E. Woosley, S. Blinnikov and A. Heger, *Pulsational pair instability as an explanation for the most luminous supernovae*, *Nature* **450** (2007) 390 [0710.3314]. Cited on 20.

[101] K. Belczynski et al., *The effect of pair-instability mass loss on black-hole mergers*, *Astronomy & Astrophysics* **594** (2016) A97 [1607.03116]. Cited on 20.

[102] S. E. Woosley, *Pulsational Pair-instability Supernovae*, *The Astrophysical Journal* **836(2)** (2017) 244 [1608.08939]. Cited on 20.

[103] M. Spera and M. Mapelli, *Very massive stars, pair-instability supernovae and intermediate-mass black holes with the SEVN code*, *Monthly Notices of the Royal Astronomical Society* **470(4)** (2017) 4739 [1706.06109]. Cited on 20.

[104] P. Marchant et al., *Pulsational Pair-instability Supernovae in Very Close Binaries*, *The Astrophysical Journal* **882(1)** (2019) 36 [1810.13412]. Cited on 20.

[105] W. A. Fowler and F. Hoyle, *Neutrino Processes and Pair Formation in Massive Stars and Supernovae*, *The Astrophysical Journal Supplement* **9** (1964) 201. Cited on 20.

[106] Z. Barkat, G. Rakavy and N. Sack, *Dynamics of Supernova Explosion Resulting from Pair Formation*, *Physical Review Letters* **18(10)** (1967) 379. Cited on 20.

[107] G. Rakavy and G. Shaviv, *Instabilities in Highly Evolved Stellar Models*, *The Astrophysical Journal* **148** (1967) 803. Cited on 20.

[108] LIGO Scientific Collaboration and Virgo Collaboration, *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, *Physical Review Letters* **119(16)** (2017) 161101 [1710.05832]. Cited on 20, 24.

[109] C. L. Rodriguez, S. Chatterjee and F. A. Rasio, *Binary black hole mergers from globular clusters: Masses, merger rates, and the impact of stellar evolution*, *Physical Review D* **93(8)** (2016) 084029 [1602.02444]. Cited on 20, 77, 85, 99.

[110] R. M. O'Leary, Y. Meiron and B. Kocsis, *DYNAMICAL FORMATION SIGNATURES OF BLACK HOLE BINARIES IN THE FIRST DETECTED MERGERS BY LIGO*, *The Astrophysical Journal Letters* **824(1)** (2016) L12 [1602.02809]. Cited on 20, 77, 85, 99.

[111] K. Belczynski, V. Kalogera and T. Bulik, *A Comprehensive Study of Binary Compact Objects as Gravitational Wave Sources: Evolutionary Channels, Rates, and Physical Properties*, *The Astrophysical Journal* **572(1)** (2002) 407 [astro-ph/0111452]. Cited on 20.

[112] P. Marchant et al., *A new route towards merging massive black holes*, *Astronomy & Astrophysics* **588** (2016) A50 [1601.03718]. Cited on 20, 77, 85, 99.

[113] E. D. Kovetz, I. Cholis, P. C. Breysse and M. Kamionkowski, *Black hole mass function from gravitational wave measurements*, *Physical Review D* **95(10)** (2017) 103010 [1611.01157]. Cited on 21, 77, 85, 99.

[114] O. Heaviside, *A GRAVITATIONAL AND ELECTROMAGNETIC ANALOGY*, *Electromagnetic Theory* **1** (1893) 455. Cited on 21.

[115] H. Poincaré, *Sur la dynamique de l'électron*, *Comptes Rendus de l'Académie des Sciences* **140** (1905) 1504. Cited on 22.

[116] H. Poincaré, *Sur la dynamique de l'électron*, *Rendiconti del Circolo matematico di Palermo* **21** (1906) 129. Cited on 22.

[117] A. I. Miller, *A Study of Henri Poincaré's "Sur la Dynamique de l'Électron"*, in *Frontiers of Physics: 1900–1911: Selected Essays*, p. 29, Birkhäuser Boston, (1986). Cited on 22.

[118] T. Damour, *Poincaré, the dynamics of the electron, and relativity*, *Comptes Rendus Physique* **18(9)** (2017) 551 [1710.00706]. Cited on 22.

[119] A. Einstein, *Näherungsweise Integration der Feldgleichungen der Gravitation*, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* **XXXII** (1916) 688. Cited on 22.

[120] J. Weber, *Gravitational-Wave-Detector Events*, *Physical Review Letters* **20(23)** (1968) 1307. Cited on 22, 25.

[121] J. Weber, *Evidence for Discovery of Gravitational Radiation*, *Physical Review Letters* **22(24)** (1969) 1320. Cited on 22, 25.

[122] J. L. Cervantes-Cota, S. Galindo-Uribarri and G. F. Smoot, *A Brief History of Gravitational Waves*, *Universe* **2(3)** (2016) 22 [1609.09400]. Cited on 22, 25, 26.

[123] LIGO Scientific Collaboration et al., *Multi-messenger Observations of a Binary Neutron Star Merger*, *The Astrophysical Journal Letters* **848(2)** (2017) L12 [1710.05833]. Cited on 24.

[124] K. Kohri and T. Terada, *Semianalytic calculation of gravitational wave spectrum nonlinearly induced from primordial curvature perturbations*, *Physical Review D* **97(12)** (2018) 123532 [1804.08577]. Cited on 24, 143.

[125] J. R. Espinosa, D. Racco and A. Riotto, *A cosmological signature of the SM Higgs instability: gravitational waves*, *Journal of Cosmology and Astroparticle Physics* **2018(09)** (2018) 012 [1804.07732]. Cited on 24, 143.

[126] S. Pi and M. Sasaki, *Gravitational waves induced by scalar perturbations with a lognormal peak*, *Journal of Cosmology and Astroparticle Physics* **2020(09)** (2020) 037 [2005.12306]. Cited on 24, 143.

[127] NANOGrav Collaboration, *The NANOGrav 12.5 yr Data Set: Search for an Isotropic Stochastic Gravitational-wave Background*, *The Astrophysical Journal Letters* **905(2)** (2020) L34 [2009.04496]. Cited on 25, 47.

[128] V. De Luca, G. Franciolini and A. Riotto, *NANOGrav Data Hints at Primordial Black Holes as Dark Matter*, *Physical Review Letters* **126(4)** (2021) 041303 [2009.08268]. Cited on 25.

[129] V. Vaskonen and H. Veermäe, *Did NANOGrav See a Signal from Primordial Black Hole Formation?*, *Physical Review Letters* **126(5)** (2021) 051303 [2009.07832]. Cited on 25, 67.

[130] K. Kohri and T. Terada, *Solar-mass primordial black holes explain NANOGrav hint of gravitational waves*, *Physics Letters B* **813** (2021) 136040 [2009.11853]. Cited on 25.

[131] J. H. Taylor, L. A. Fowler and P. M. McCulloch, *Measurements of general relativistic effects in the binary pulsar PSR1913+16*, *Nature* **277** (1979) 437. Cited on 25.

[132] A. A. Michelson, *The relative motion of the Earth and of the luminiferous ether*, *American Journal of Science* **s3-22(128)** (1881) 120. Cited on 26.

[133] A. A. Michelson and E. W. Morley, *On the relative motion of the Earth and the luminiferous ether*, *American Journal of Science* **s3-34(203)** (1887) 333. Cited on 26.

[134] M. E. Gertsenshteĭn and V. I. Pustovoĭt, *On the detection of low frequency gravitational waves*, *Journal of Experimental and Theoretical Physics* **16(2)** (1963) 433. Cited on 26.

[135] V. B. Braginskiĭ, *GRAVITATIONAL RADIATION AND THE PROSPECT OF ITS EXPERIMENTAL DISCOVERY*, *Physics-Uspekhi* **8(4)** (1966) 513. Cited on 26.

[136] LIGO Scientific Collaboration and Virgo Collaboration, *GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence*, *Physical Review Letters* **119(14)** (2017) 141101 [1709.09660]. Cited on 27.

[137] LIGO Scientific Collaboration, *GW190521*, (2020). Accessed on 2021-07-29. Cited on 28.

[138] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration, *Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences*, *The Astrophysical Journal Letters* **915(1)** (2021) L5 [2106.15163]. Cited on 27.

[139] B. F. Schutz, *Networks of gravitational wave detectors and three figures of merit*, *Classical and Quantum Gravity* **28(12)** (2011) 125023 [1102.5421]. Cited on 28.

[140] S. Fairhurst, *Improved source localization with LIGO-India*, *Journal of Physics: Conference Series* **484** (2014) 012007 [1205.6611]. Cited on 28.

[141] R. W. Hellings and G. S. Downs, *Upper limits on the isotropic gravitational radiation background from pulsar timing analysis*, *The Astrophysical Journal* **265** (1983) L39. Cited on 29.

[142] R. W. Romani, *Timing a Millisecond Pulsar Array*, in *Timing Neutron Stars*, H. Ögelman and E. P. J. van den Heuvel, eds., p. 113, Springer Netherlands, (1989). Cited on 29.

[143] R. S. Foster and D. C. Backer, *Constructing a Pulsar Timing Array*, *The Astrophysical Journal* **361** (1990) 300. Cited on 29.

[144] J. P. W. Verbiest, S. Osłowski and S. Burke-Spolaor, *Pulsar Timing Array Experiments*, in *Handbook of Gravitational Wave Astronomy*, C. Bambi, S. Katsanevas and K. D. Kokkotas, eds., p. 1, Springer Singapore, (2021) [2101.10081]. Cited on 29.

[145] M. Maggiore et al., *Science case for the Einstein telescope*, *Journal of Cosmology and Astroparticle Physics* **2020(03)** (2020) 050 [1912.02622]. Cited on 30, 55.

[146] D. Reitze et al., *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*, (2019) [1907.04833]. Cited on 30.

[147] F. Aharonian et al., *Pathway to the Square Kilometre Array - The German White Paper*, (2013) [1301.4124]. Cited on 30.

[148] A. Klein et al., *Science with the space-based interferometer eLISA: Supermassive black hole binaries*, *Physical Review D* **93(2)** (2016) 024003 [1511.05581]. Cited on 30.

[149] P. Amaro-Seoane et al., *Laser Interferometer Space Antenna*, (2017) [1702.00786]. Cited on 30.

[150] S. Kawamura et al., *The Japanese space gravitational wave antenna - DECIGO*, *Journal of Physics: Conference Series* **122** (2008) 012006. Cited on 30.

[151] J. Luo et al., *TianQin: a space-borne gravitational wave detector*, *Classical and Quantum Gravity* **33(3)** (2016) 035010 [1512.02076]. Cited on 30.

[152] W.-H. Ruan, Z.-K. Guo, R.-G. Cai and Y.-Z. Zhang, *Taiji program: Gravitational-wave sources*, *International Journal of Modern Physics A* **35(17)** (2020) 2050075 [1807.09495]. Cited on 30.

[153] B. Carr and F. Kühnel, *Primordial Black Holes as Dark Matter: Recent Developments*, *Annual Review of Nuclear and Particle Science* **70** (2020) [2006.02838]. Cited on 32, 55, 61.

[154] A. M. Green and B. J. Kavanagh, *Primordial black holes as a dark matter candidate*, *Journal of Physics G* **48(4)** (2021) 043001 [2007.10722]. Cited on 32, 61.

[155] A. M. Green, *Primordial Black Holes: Sirens of the Early Universe*, in *Quantum Aspects of Black Holes*, X. Calmet, ed., p. 129, Springer International, (2015) [1403.1198]. Cited on 32, 60, 76.

[156] M. Sasaki, T. Suyama, T. Tanaka and S. Yokoyama, *Primordial black holes—perspectives in gravitational wave astronomy*, *Classical and Quantum Gravity* **35(6)** (2018) 063001 [1801.05235]. Cited on 32, 60, 76.

[157] Planck Collaboration, *Planck 2018 results. X. Constraints on inflation*, *Astronomy & Astrophysics* **641** (2020) A10 [1807.06211]. Cited on 32.

[158] I. Dalianis, *Constraints on the curvature power spectrum from primordial black hole evaporation*, *Journal of Cosmology and Astroparticle Physics* **2019(08)** (2019) 032 [1812.09807]. Cited on 32.

[159] G. Sato-Polito, E. D. Kovetz and M. Kamionkowski, *Constraints on the primordial curvature power spectrum from primordial black holes*, *Physical Review D* **100(6)** (2019) 063521 [1904.10971]. Cited on 32.

[160] T. Bringmann, P. Scott and Y. Akrami, *Improved constraints on the primordial power spectrum at small scales from ultracompact minihalos*, *Physical Review D* **85(12)** (2012) 125027 [1110.2484]. Cited on 32, 47.

[161] P. P. Avelino, *Primordial black hole abundance in non-Gaussian inflation models*, *Physical Review D* **72(12)** (2005) 124004 [astro-ph/0510052]. Cited on 32.

[162] D. H. Lyth, *The hybrid inflation waterfall and the primordial curvature perturbation*, *Journal of Cosmology and Astroparticle Physics* **2012(05)** (2012) 022 [1201.4312]. Cited on 32.

[163] C. T. Byrnes, E. J. Copeland and A. M. Green, *Primordial black holes as a tool for constraining non-Gaussianity*, *Physical Review D* **86(4)** (2012) 043512 [1206.4188]. Cited on 32.

[164] S. Young and C. T. Byrnes, *Long-short wavelength mode coupling tightens primordial black hole constraints*, *Physical Review D* **91(8)** (2015) 083521 [1411.4620]. Cited on 32.

[165] S. Young, D. Regan and C. T. Byrnes, *Influence of large local and non-local bispectra on primordial black hole abundance*, *Journal of Cosmology and Astroparticle Physics* **2016(02)** (2016) 029 [1512.07224]. Cited on 32.

[166] G. Franciolini, A. Kehagias, S. Matarrese and A. Riotto, *Primordial black holes from inflation and non-Gaussianity*, *Journal of Cosmology and Astroparticle Physics* **2018(03)** (2018) 016 [1801.09415]. Cited on 32.

[167] C.-M. Yoo, J.-O. Gong and S. Yokoyama, *Abundance of primordial black holes with local non-Gaussianity in peak theory*, *Journal of Cosmology and Astroparticle Physics* **2019(09)** (2019) 033 [1906.06790]. Cited on 32.

[168] R.-G. Cai, S. Pi and M. Sasaki, *Gravitational Waves Induced by Non-Gaussian Scalar Perturbations*, *Physical Review Letters* **122(20)** (2019) 201101 [1810.11000]. Cited on 32.

[169] M. Kawasaki and H. Nakatsuka, *Effect of nonlinearity between density and curvature perturbations on the primordial black hole formation*, *Physical Review D* **99(12)** (2019) 123501 [1903.02994]. Cited on 33, 52, 145.

[170] S. Young, I. Musco and C. T. Byrnes, *Primordial black hole formation and abundance: contribution from the non-linear relation between the density and curvature perturbation*, *Journal of Cosmology and Astroparticle Physics* **2019(11)** (2019) 012 [1904.00984]. Cited on 33, 52, 145.

[171] V. De Luca et al., *The ineludible non-Gaussianity of the primordial black hole abundance*, *Journal of Cosmology and Astroparticle Physics* **2019(07)** (2019) 048 [1904.00970]. Cited on 33, 52, 145.

[172] C.-M. Yoo, T. Harada, J. Garriga and K. Kohri, *Primordial black hole abundance from random Gaussian curvature perturbations and a local density threshold*, *Progress of Theoretical and Experimental Physics* **2018(12)** (2018) 123E01 [1805.03946]. Cited on 33, 35, 52, 145.

[173] A. Kalaja et al., *From primordial black holes abundance to primordial curvature power spectrum (and back)*, *Journal of Cosmology and Astroparticle Physics* **2019(10)** (2019) 031 [1908.03596]. Cited on 33, 36, 38, 52, 53, 145.

[174] S. Young and M. Musso, *Application of peaks theory to the abundance of primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2020(11)** (2020) 022 [2001.06469]. Cited on 33, 35, 37, 38, 41, 62.

[175] I. Musco, J. C. Miller and A. G. Polnarev, *Primordial black hole formation in the radiative era: investigation of the critical nature of the collapse*, *Classical and Quantum Gravity* **26(23)** (2009) 235001 [0811.1452]. Cited on 33, 38, 60, 67.

[176] K. Ando, K. Inomata and M. Kawasaki, *Primordial black holes and uncertainties in the choice of the window function*, *Physical Review D* **97(10)** (2018) 103528 [1802.06393]. Cited on 33.

[177] S. Young, *The primordial black hole formation criterion re-examined: Parametrisation, timing and the choice of window function*, *International Journal of Modern Physics D* **29(02)** (2020) 2030002 [1905.01230]. Cited on 33, 38, 62, 140.

[178] I. Musco, *Threshold for primordial black holes: Dependence on the shape of the cosmological perturbations*, *Physical Review D* **100(12)** (2019) 123524 [1809.02127]. Cited on 33, 38.

[179] A. Escrivà, C. Germani and R. K. Sheth, *Analytical thresholds for black hole formation in general cosmological backgrounds*, *Journal of Cosmology and Astroparticle Physics* **2021(01)** (2021) 030 [2007.05564]. Cited on 33.

[180] C. Germani and I. Musco, *Abundance of Primordial Black Holes Depends on the Shape of the Inflationary Power Spectrum*, *Physical Review Letters* **122(14)** (2019) 141302 [1805.04087]. Cited on 33, 38, 46.

[181] J. M. Bardeen, J. R. Bond, N. Kaiser and A. S. Szalay, *The Statistics of Peaks of Gaussian Random Fields*, *The Astrophysical Journal* **304** (1986) 15. Cited on 35, 36, 146.

[182] C.-M. Yoo, T. Harada, S. Hirano and K. Kohri, *Abundance of Primordial Black Holes in Peak Theory for an Arbitrary Power Spectrum*, (2020) [2008.02425]. Cited on 35.

[183] C. Germani and R. K. Sheth, *Nonlinear statistics of primordial black holes from Gaussian curvature perturbations*, *Physical Review D* **101(6)** (2020) 063520 [1912.07072]. Cited on 35, 62.

[184] S. Young, C. T. Byrnes and M. Sasaki, *Calculating the mass fraction of primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2014(07)** (2014) 045 [1405.7023]. Cited on 36.

[185] J. C. Niemeyer and K. Jedamzik, *Near-Critical Gravitational Collapse and the Initial Mass Function of Primordial Black Holes*, *Physical Review Letters* **80(25)** (1998) 5481 [astro-ph/9709072]. Cited on 38, 60, 62, 64, 67, 103.

[186] I. Musco, J. C. Miller and L. Rezzolla, *Computations of primordial black-hole formation*, *Classical and Quantum Gravity* **22(7)** (2005) 1405 [gr-qc/0412063]. Cited on 38.

[187] C. T. Byrnes, P. S. Cole and S. P. Patil, *Steepest growth of the power spectrum and primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2019(06)** (2019) 028 [1811.11158]. Cited on 39, 47, 53, 60.

[188] P. Carrilho, K. A. Malik and D. J. Mulryne, *Dissecting the growth of the power spectrum for primordial black holes*, *Physical Review D* **100(10)** (2019) 103529 [1907.05237]. Cited on 39, 60.

[189] O. Özsoy and G. Tasinato, *On the slope of curvature power spectrum in non-attractor inflation*, *Journal of Cosmology and Astroparticle Physics* **2020(04)** (2020) 048 [1912.01061]. Cited on 39, 60.

[190] G. A. Palma, S. Sypsas and C. Zenteno, *Seeding Primordial Black Holes in Multifield Inflation*, *Physical Review Letters* **125(12)** (2020) 121301 [2004.06106]. Cited on 39, 60.

[191] J. Fumagalli, S. Renaux-Petel, J. W. Ronayne and L. T. Witkowski, *Turning in the landscape: a new mechanism for generating Primordial Black Holes*, (2020) [2004.08369]. Cited on 39, 60.

[192] T. Nakama, J. Silk and M. Kamionkowski, *Stochastic gravitational waves associated with the formation of primordial black holes*, *Physical Review D* **95(4)** (2017) 043511 [1612.06264]. Cited on 39.

[193] Z.-C. Chen and Q.-G. Huang, *Merger Rate Distribution of Primordial Black Hole Binaries*, *The Astrophysical Journal* **864(1)** (2018) 61 [1801.10327]. Cited on 40, 78, 81.

[194] A. D. Gow, C. T. Byrnes, A. Hall and J. A. Peacock, *Primordial black hole merger rates: distributions for multiple LIGO observables*, *Journal of Cosmology and Astroparticle Physics* **2020(01)** (2020) 031 [1911.12685]. Cited on 40.

[195] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *Primordial black holes confront LIGO/Virgo data: current situation*, *Journal of Cosmology and Astroparticle Physics* **2020(06)** (2020) 044 [2005.05641]. Cited on 40, 61.

[196] S. Young and A. S. Hamers, *The impact on distant fly-bys on the rate of binary primordial black hole mergers*, *Journal of Cosmology and Astroparticle Physics* **2020(10)** (2020) 036 [2006.15023]. Cited on 40.

[197] S. Clesse and J. Garcia-Bellido, *GW190425, GW190521 and GW190814: Three candidate mergers of primordial black holes from the QCD epoch*, (2021) [2007.06481]. Cited on 40, 61.

[198] Y.-P. Wu, *Peak statistics for the primordial black hole abundance*, *Physics of the Dark Universe* **30** (2020) 100654 [2005.00441]. Cited on 43.

[199] K. Inomata et al., *Gravitational wave production right after a primordial black hole evaporation*, *Physical Review D* **101(12)** (2020) 123533 [2003.10455]. Cited on 45, 61.

[200] J. Chluba, A. L. Erickcek and I. Ben-Dayan, *PROBING THE INFLATON: SMALL-SCALE POWER SPECTRUM CONSTRAINTS FROM MEASUREMENTS OF THE COSMIC MICROWAVE BACKGROUND ENERGY SPECTRUM*, *The Astrophysical Journal* **758(2)** (2012) 76 [1203.2681]. Cited on 46.

[201] K. Tomita, *Non-Linear Theory of Gravitational Instability in the Expanding Universe*, *Progress of Theoretical Physics* **37(5)** (1967) 831. Cited on 46.

[202] K. N. Ananda, C. Clarkson and D. Wands, *Cosmological gravitational wave background from primordial density perturbations*, *Physical Review D* **75(12)** (2007) 123518 [gr-qc/0612013]. Cited on 46.

[203] NANOGrav Collaboration, *The NANOGrav 11 Year Data Set: Pulsar-timing Constraints on the Stochastic Gravitational-wave Background*, *The Astrophysical Journal* **859(1)** (2018) 47 [1801.02617]. Cited on 47, 48, 144.

[204] J. Chluba, J. Hamann and S. P. Patil, *Features and new physical scales in primordial observables: Theory and observation*, *International Journal of Modern Physics D* **24(10)** (2015) 1530023 [1505.01834]. Cited on 47, 141.

[205] M. Lucca et al., *The synergy between CMB spectral distortions and anisotropies*, *Journal of Cosmology and Astroparticle Physics* **2020(02)** (2020) 026 [1910.04619]. Cited on 47.

[206] J.-O. Gong and N. Kitajima, *Small-scale structure and 21cm fluctuations by primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2017(08)** (2017) 017 [1704.04132]. Cited on 47.

[207] J. L. Bernal, A. Raccanelli, L. Verde and J. Silk, *Signatures of primordial black holes as seeds of supermassive black holes*, *Journal of Cosmology and Astroparticle Physics* **2018(05)** (2018) 017 [1712.01311]. Cited on 47, 54.

[208] O. Mena, S. Palomares-Ruiz, P. Villanueva-Domingo and S. J. Witte, *Constraining the primordial black hole abundance with 21-cm cosmology*, *Physical Review D* **100(4)** (2019) 043540 [1906.07735]. Cited on 47.

[209] J. B. Muñoz, C. Dvorkin and F.-Y. Cyr-Racine, *Probing the small-scale matter power spectrum with large-scale 21-cm data*, *Physical Review D* **101(6)** (2020) 063526 [1911.11144]. Cited on 47.

[210] P. S. Cole and J. Silk, *Small-scale primordial fluctuations in the 21 cm Dark Ages signal*, *Monthly Notices of the Royal Astronomical Society* **501(2)** (2020) 2627 [1912.02171]. Cited on 47.

[211] M. Gosenca, J. Adamek, C. T. Byrnes and S. Hotchkiss, *3D simulations with boosted primordial power spectra and ultracompact minihalos*, *Physical Review D* **96(12)** (2017) 123519 [1710.02055]. Cited on 47.

[212] M. Sten Delos, A. L. Erickcek, A. P. Bailey and M. A. Alvarez, *Are ultracompact minihalos really ultracompact?*, *Physical Review D* **97(4)** (2018) 041303(R) [1712.05421]. Cited on 47.

[213] M. Sten Delos, A. L. Erickcek, A. P. Bailey and M. A. Alvarez, *Density profiles of ultracompact minihalos: Implications for constraining the primordial power spectrum*, *Physical Review D* **98(6)** (2018) 063527 [1806.07389]. Cited on 47.

[214] K. Furugori et al., *The 21-cm signals from ultracompact minihalos as a probe of primordial small-scale fluctuations*, *Monthly Notices of the Royal Astronomical Society* **494(3)** (2020) 4334 [2002.04817]. Cited on 47.

[215] D. Jeong, J. Pradler, J. Chluba and M. Kamionkowski, *Silk Damping at a Redshift of a Billion: New Limit on Small-Scale Adiabatic Perturbations*, *Physical Review Letters* **113(6)** (2014) 061301 [1403.3697]. Cited on 47.

[216] T. Nakama, T. Suyama and J. Yokoyama, *Reheating the Universe Once More: The Dissipation of Acoustic Waves as a Novel Probe of Primordial Inhomogeneities on Even Smaller Scales*, *Physical Review Letters* **113(6)** (2014) 061302 [1403.5407]. Cited on 47.

[217] K. Inomata, M. Kawasaki and Y. Tada, *Revisiting constraints on small scale perturbations from big-bang nucleosynthesis*, *Physical Review D* **94(4)** (2016) 043527 [1605.04646]. Cited on 47.

[218] NANOGrav Collaboration, *THE NANOGRAV NINE-YEAR DATA SET: LIMITS ON THE ISOTROPIC STOCHASTIC GRAVITATIONAL WAVE BACKGROUND*, *The Astrophysical Journal* **821(1)** (2016) 13 [1508.03024]. Cited on 47, 144.

[219] L. Lentati et al., *European Pulsar Timing Array limits on an isotropic stochastic gravitational-wave background*, *Monthly Notices of the Royal Astronomical Society* **453(3)** (2015) 2576 [1504.03692]. Cited on 47, 144.

[220] K. Inomata et al., *Inflationary primordial black holes for the LIGO gravitational wave events and pulsar timing array experiments*, *Physical Review D* **95(12)** (2017) 123510 [1611.06130]. Cited on 53.

[221] K. Inomata and T. Nakama, *Gravitational waves induced by scalar perturbations as probes of the small-scale primordial spectrum*, *Physical Review D* **99(4)** (2019) 043511 [1812.00674]. Cited on 53.

[222] S. Clesse, J. García-Bellido and S. Orani, *Detecting the Stochastic Gravitational Wave Background from Primordial Black Hole Formation*, (2018) [1812.11011]. Cited on 53.

[223] Y. Lu, Y. Gong, Z. Yi and F. Zhang, *Constraints on primordial curvature perturbations from primordial black hole dark matter and secondary gravitational waves*, *Journal of Cosmology and Astroparticle Physics* **2019(12)** (2019) 031 [1907.11896]. Cited on 53.

[224] Z.-C. Chen, C. Yuan and Q.-G. Huang, *Pulsar Timing Array Constraints on Primordial Black Holes with NANOGrav 11-Year Dataset*, *Physical Review Letters* **124(25)** (2020) 251101 [1910.12239]. Cited on 53.

[225] S. Wang, T. Terada and K. Kohri, *Prospective constraints on the primordial black hole abundance from the stochastic gravitational-wave backgrounds produced by coalescing events and curvature perturbations*, *Physical Review D* **99(10)** (2019) 103531 [1903.05924]. Cited on 53.

[226] C. Yuan, Z.-C. Chen and Q.-G. Huang, *Probing primordial–black-hole dark matter with scalar induced gravitational waves*, *Physical Review D* **100(8)** (2019) 081301(R) [1906.11549]. Cited on 53.

[227] R.-G. Cai, S. Pi, S.-J. Wang and X.-Y. Yang, *Pulsar timing array constraints on the induced gravitational waves*, *Journal of Cosmology and Astroparticle Physics* **2019(10)** (2019) 059 [1907.06372]. Cited on 53.

[228] I. Dalianis and K. Kritos, *Exploring the spectral shape of gravitational waves induced by primordial scalar perturbations and connection with the primordial black hole scenarios*, *Physical Review D* **103(2)** (2021) 023505 [2007.07915]. Cited on 53.

[229] B. Carr, S. Clesse, J. García-Bellido and F. Kühnel, *Cosmic conundra explained by thermal history and primordial black holes*, *Physics of the Dark Universe* **31** (2021) 100755 [1906.08217]. Cited on 54, 63.

[230] R. Allahverdi et al., *The first three seconds: A Review of Possible Expansion Histories of the early Universe*, *The Open Journal of Astrophysics* **4** (2021) [2006.16182]. Cited on 54.

[231] T. Nakama, B. Carr and J. Silk, *Limits on primordial black holes from μ distortions in cosmic microwave background*, *Physical Review D* **97(4)** (2018) 043525 [1710.06945]. Cited on 54.

[232] Y. Tada and S. Yokoyama, *Primordial black holes as biased tracers*, *Physical Review D* **91(12)** (2015) 123534 [1502.01124]. Cited on 54, 79.

[233] S. Young and C. T. Byrnes, *Signatures of non-gaussianity in the isocurvature modes of primordial black hole dark matter*, *Journal of Cosmology and Astroparticle Physics* **2015(04)** (2015) 034 [1503.01505]. Cited on 54, 79.

[234] P. S. Cole and C. T. Byrnes, *Extreme scenarios: the tightest possible constraints on the power spectrum due to primordial black holes*, *Journal of Cosmology and Astroparticle Physics* **2018(02)** (2018) 019 [1706.10288]. Cited on 55.

[235] B. C. Lacki and J. F. Beacom, *PRIMORDIAL BLACK HOLES AS DARK MATTER: ALMOST ALL OR ALMOST NOTHING*, *The Astrophysical Journal Letters* **720(1)** (2010) L67 [1003.3466]. Cited on 55.

[236] Yu. N. Eroshenko, *Dark matter density spikes around primordial black holes*, *Astronomy Letters* **42(6)** (2016) 347 [1607.00612]. Cited on 55.

[237] S. M. Boucenna, F. Kühnel, T. Ohlsson and L. Visinelli, *Novel constraints on mixed dark-matter scenarios of primordial black holes and WIMPs*, *Journal of Cosmology and Astroparticle Physics* **2018(07)** (2018) 003 [1712.06383]. Cited on 55.

[238] J. Adamek, C. T. Byrnes, M. Gosenca and S. Hotchkiss, *WIMPs and stellar-mass primordial black holes are incompatible*, *Physical Review D* **100(2)** (2019) 023506 [1901.08528]. Cited on 55.

[239] G. Bertone et al., *Primordial black holes as silver bullets for new physics at the weak scale*, *Physical Review D* **100(12)** (2019) 123013 [1905.01238]. Cited on 55.

[240] 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)** (2012) 1129 [1202.0057]. Cited on 55.

[241] C. J. Moore, R. H. Cole and C. P. L. Berry, *Gravitational-wave sensitivity curves*, *Classical and Quantum Gravity* **32(1)** (2015) 015014 [1408.0740]. Cited on 55.

[242] N. Bartolo et al., *Science with the space-based interferometer LISA. IV: probing inflation with gravitational waves*, *Journal of Cosmology and Astroparticle Physics* **2016(12)** (2016) 026 [1610.06481]. Cited on 55.

[243] B. J. Carr, K. Kohri, Y. Sendouda and J. Yokoyama, *New cosmological constraints on primordial black holes*, *Physical Review D* **81(10)** (2010) 104019 [0912.5297]. Cited on 60.

[244] B. Carr, F. Kühnel and M. Sandstad, *Primordial black holes as dark matter*, *Physical Review D* **94(8)** (2016) 083504 [1607.06077]. Cited on 60, 91.

[245] J. Yokoyama, *Cosmological constraints on primordial black holes produced in the near-critical gravitational collapse*, *Physical Review D* **58(10)** (1998) 107502 [gr-qc/9804041]. Cited on 60, 61, 62, 64, 67, 103.

[246] F. Kühnel, C. Rampf and M. Sandstad, *Effects of critical collapse on primordial black-hole mass spectra*, *The European Physical Journal C* **76(2)** (2016) 93 [1512.00488]. Cited on 60.

[247] Y.-F. Cai, X. Tong, D.-G. Wang and S.-F. Yan, *Primordial Black Holes from Sound Speed Resonance during Inflation*, *Physical Review Letters* **121(8)** (2018) 081306 [1805.03639]. Cited on 60.

[248] A. Ashoorioon, A. Rostami and J. T. Firouzjaee, *EFT compatible PBHs: effective spawning of the seeds for primordial black holes during inflation*, *Journal of High Energy Physics* **2021(07)** (2021) 087 [1912.13326]. Cited on 60.

[249] V. Bosch-Ramon and N. Bellomo, *Mechanical feedback effects on primordial black hole accretion*, *Astronomy & Astrophysics* **638** (2020) A132 [2004.11224]. Cited on 61.

[250] K. Vattis, I. S. Goldstein and S. M. Koushiappas, *Could the* 2.6 $M_\odot$ *object in GW190814 be a primordial black hole?*, *Physical Review D* **102(6)** (2020) 061301(R) [2006.15675]. Cited on 61.

[251] K. Belczynski et al., *Evolutionary roads leading to low effective spins, high black hole masses, and O1/O2 rates for LIGO/Virgo binary black holes*, *Astronomy & Astrophysics* **636** (2020) A104 [1706.07053]. Cited on 61.

[252] K. Belczynski et al., *MISSING BLACK HOLES UNVEIL THE SUPERNOVA EXPLOSION MECHANISM*, *The Astrophysical Journal* **757(1)** (2012) 91 [1110.1635]. Cited on 61.

[253] A. Dolgov and J. Silk, *Baryon isocurvature fluctuations at small scales and baryonic dark matter*, *Physical Review D* **47(10)** (1993) 4244. Cited on 61.

[254] B. J. Carr, *The primordial black hole mass spectrum*, *The Astrophysical Journal* **201** (1975) 1. Cited on 61.

[255] A. Hall, A. D. Gow and C. T. Byrnes, *Bayesian analysis of LIGO-Virgo mergers: Primordial versus astrophysical black hole populations*, *Physical Review D* **102(12)** (2020) 123524 [2008.13704]. Cited on 61, 73, 104.

[256] G. Hütsi, M. Raidal, V. Vaskonen and H. Veermäe, *Two populations of LIGO-Virgo black holes*, *Journal of Cosmology and Astroparticle Physics* **2021(03)** (2021) 068 [2012.02786]. Cited on 61.

[257] G. Franciolini et al., *Quantifying the evidence for primordial black holes in LIGO/Virgo gravitational-wave data*, (2021) [2105.03349]. Cited on 61.

[258] A. D. Gow, C. T. Byrnes, P. S. Cole and S. Young, *The power spectrum on small scales: robust constraints and comparing PBH methodologies*, *Journal of Cosmology and Astroparticle Physics* **2021(02)** (2021) 002 [2008.03289]. Cited on 61, 62, 63.

[259] A. M. Green and K. A. Malik, *Primordial black hole production due to preheating*, *Physical Review D* **64(2)** (2001) 021301(R) [hep-ph/0008113]. Cited on 62.

[260] A. M. Green, A. R. Liddle, K. A. Malik and M. Sasaki, *New calculation of the mass fraction of primordial black holes*, *Physical Review D* **70(4)** (2004) 041502(R) [astro-ph/0403181]. Cited on 62.

[261] T. Suyama and S. Yokoyama, *A novel formulation of the PBH mass function*, (2019) [1912.04687]. Cited on 62.

[262] P. Widerin and C. Schmid, *Primordial Black Holes from the QCD Transition?*, (1998) [astro-ph/9808142]. Cited on 63.

[263] K. Jedamzik and J. C. Niemeyer, *Primordial black hole formation during first-order phase transitions*, *Physical Review D* **59(12)** (1999) 124014 [astro-ph/9901293]. Cited on 63.

[264] A. Escrivà, *Simulation of primordial black hole formation using pseudo-spectral methods*, *Physics of the Dark Universe* **27** (2020) 100466 [1907.13065]. Cited on 67.

[265] A. Azzalini, *The Skew-Normal and Related Families*, Cambridge University Press (2013). Cited on 68.

[266] S. Bird et al., *Did LIGO Detect Dark Matter?*, *Physical Review Letters* **116(20)** (2016) 201301 [1603.00464]. Cited on 76.

[267] M. Sasaki, T. Suyama, T. Tanaka and S. Yokoyama, *Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914*, *Physical Review Letters* **117(6)** (2016) 061101 [1603.08338]. Cited on 76, 81.

[268] A. Kashlinsky, *LIGO GRAVITATIONAL WAVE DETECTION, PRIMORDIAL BLACK HOLES, AND THE NEAR-IR COSMIC INFRARED BACKGROUND ANISOTROPIES*, *The Astrophysical Journal* **823(2)** (2016) L25 [1605.04023]. Cited on 76.

[269] EROS-2 Collaboration, *Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds*, *Astronomy & Astrophysics* **469(2)** (2007) 387 [astro-ph/0607207]. Cited on 76.

[270] Ł. Wyrzykowski et al., *The OGLE view of microlensing towards the Magellanic Clouds – II. OGLE-II Small Magellanic Cloud data*, *Monthly Notices of the Royal Astronomical Society* **407(1)** (2010) 189 [0905.2044]. Cited on 76.

[271] M. Oguri et al., *Understanding caustic crossings in giant arcs: Characteristic scales, event rates, and constraints on compact dark matter*, *Physical Review D* **97(2)** (2018) 023518 [1710.00148]. Cited on 76.

[272] M. Zumalacárregui and U. Seljak, *Limits on Stellar-Mass Compact Objects as Dark Matter from Gravitational Lensing of Type Ia Supernovae*, *Physical Review Letters* **121(14)** (2018) 141101 [1712.02240]. Cited on 76.

[273] Y. Ali-Haïmoud and M. Kamionkowski, *Cosmic microwave background limits on accreting primordial black holes*, *Physical Review D* **95(4)** (2017) 043534 [1612.05644]. Cited on 76.

[274] V. Poulin et al., *CMB bounds on disk-accreting massive primordial black holes*, *Physical Review D* **96(8)** (2017) 083524 [1707.04206]. Cited on 76.

[275] M. Fishbach and D. E. Holz, *Picky Partners: The Pairing of Component Masses in Binary Black Hole Mergers*, *The Astrophysical Journal* **891(1)** (2020) L27 [1905.12669]. Cited on 77, 85.

[276] D. Gerosa et al., *Multiband gravitational-wave event rates and stellar physics*, *Physical Review D* **99(10)** (2019) 103004 [1902.00021]. Cited on 77, 82, 83, 84, 149.

[277] T. Nakamura, M. Sasaki, T. Tanaka and K. S. Thorne, *Gravitational Waves from Coalescing Black Hole MACHO Binaries*, *The Astrophysical Journal* **487(2)** (1997) L139 [astro-ph/9708060]. Cited on 78.

[278] Y. Ali-Haïmoud, E. D. Kovetz and M. Kamionkowski, *Merger rate of primordial black-hole binaries*, *Physical Review D* **96(12)** (2017) 123523 [1709.06576]. Cited on 78, 79, 81, 89, 100.

[279] Yu. N. Eroshenko, *Gravitational waves from primordial black holes collisions in binary systems*, *Journal of Physics: Conference Series* **1051** (2018) 012010 [1604.04932]. Cited on 78, 81, 89.

[280] B. J. Kavanagh, D. Gaggero and G. Bertone, *Merger rate of a subdominant population of primordial black holes*, *Physical Review D* **98(2)** (2018) 023536 [1805.09034]. Cited on 78, 81.

[281] P. C. Peters, *Gravitational Radiation and the Motion of Two Point Masses*, *Physical Review* **136(4B)** (1964) B1224. Cited on 79.

[282] V. Vaskonen and H. Veermäe, *Lower bound on the primordial black hole merger rate*, *Physical Review D* **101(4)** (2020) 043015 [1908.09752]. Cited on 79.

[283] G. Ballesteros, P. D. Serpico and M. Taoso, *On the merger rate of primordial black holes: effects of nearest neighbours distribution and clustering*, *Journal of Cosmology and Astroparticle Physics* **2018(10)** (2018) 043 [1807.02084]. Cited on 79.

[284] M. Dominik et al., *DOUBLE COMPACT OBJECTS. III. GRAVITATIONAL-WAVE DETECTION RATES*, *The Astrophysical Journal* **806(2)** (2015) 263 [1405.7016]. Cited on 82.

[285] D. Gerosa, *dgerosa/gwdet: v0.1*, (2017). Cited on 82.

[286] M. He and T. Suyama, *Formation threshold of rotating primordial black holes*, *Physical Review D* **100(6)** (2019) 063520 [1906.10987]. Cited on 82.

[287] S. Clesse and J. García-Bellido, *Seven hints for primordial black hole dark matter*, *Physics of the Dark Universe* **22** (2018) 137 [1711.10458]. Cited on 83.

[288] J. Garriga and N. Triantafyllou, *Enhanced cosmological perturbations and the merger rate of PBH binaries*, *Journal of Cosmology and Astroparticle Physics* **2019(09)** (2019) 043 [1907.01455]. Cited on 89, 90, 100.

[289] J. A. Peacock, *Two-dimensional goodness-of-fit testing in astronomy*, *Monthly Notices of the Royal Astronomical Society* **202(3)** (1983) 615. Cited on 97.

[290] D. J. Fixsen et al., *The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set*, *The Astrophysical Journal* **473(2)** (1996) 576 [astro-ph/9605054]. Cited on 142.

[291] A. Kogut et al., *The Primordial Inflation Explorer (PIXIE)*, in *Space Telescopes and Instrumentation 2016: Optical, Infrared and Millimeter*, vol. 9904, p. 331, SPIE, (2016). Cited on 142.

[292] J. 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)** (2019) 184 [1903.04218]. Cited on 142.

Appendices

A Ringing in the top-hat window function

Here we explain our procedure to produce constraints when using a real-space top-hat window function, which corresponds to a rapidly oscillating window function in Fourier-space, with consequent convergence issues. The width parameter $\sigma_0(R)$ is shown in fig. A.1 for a delta function peak (left) and the lognormal widths $\Delta = 0.3$ (middle) and $\Delta = 1$ (right).

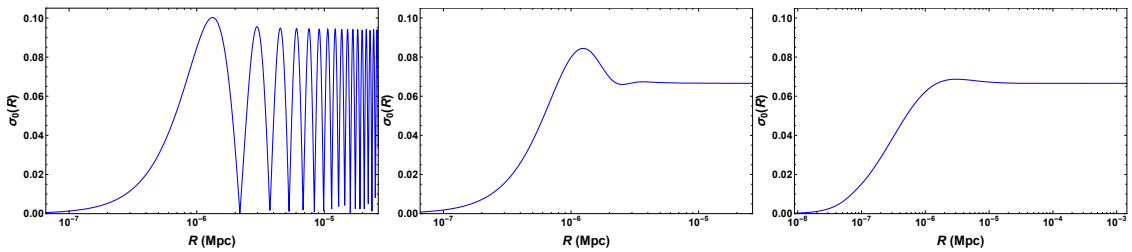

Figure A.1: Width parameter $\sigma_0(R)$ for a delta function power spectrum (left), and a lognormal peak with widths $\Delta = 0.3$ (middle) and $\Delta = 1$ (right). The ringing peaks visible in the delta case merge to a constant height as Δ increases.

It can be seen that the oscillatory nature of the top-hat window function leads to a ringing effect in the width parameter $\sigma_0(R)$. For broader peaks in the power spectrum, this ringing effect merges into a constant height for large values of R. This leads to a divergent integral when evaluating eq. (2.1), and so the mass distribution cannot be calculated with this window function without some form of adjustment. It is common to suppress the large-R constant effect using a transfer function, but this method is not compatible with other parts of our calculation (i.e. [177]). Therefore, we take an alternative approach, which is to adjust the calculation of $\sigma_n(R)$ in

eq. (2.12) with a large-k cutoff. This is placed at the point where the window function reaches its first trough, which is at $4.49/R$. This solves the divergence problem and removes the ringing/constant effect, but it must be noted that the window function is technically not a true top-hat any more.

B Observational constraints

B.1 Constraints due to spectral distortions of the CMB

Spectral distortions of the energy spectrum of the CMB are able to constrain the primordial power spectrum on small 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. 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 means 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. So-called y-distortions quantify late-time processes and place constraints on larger modes $k < 3$ Mpc^{-1}, whilst μ-distortions quantify earlier energy injection and removal and hence constrain the smaller scales, up to $k \approx 10^4$ Mpc^{-1} which will be most interesting for PBH production. The final μ-distortions induced by the scalar perturbations can be approximated by [204]

$$\mu \approx \int_{k_{\rm min}}^{\infty} \frac{dk}{k} \mathcal{P}_{\mathcal{R}}(k)\, W_{\mu}(k), \tag{B.1}$$

with k-space window functions of the form

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

where $\hat{k} = k/1\ \mathrm{Mpc}^{-1}$ and $k_{\mathrm{min}} \approx 1\ \mathrm{Mpc}^{-1}$. Given a particular form for the power spectrum, this can be used to compute the total induced μ- or y-distortion. Comparing this with observations then results in constraints on the primordial power spectrum.

The Far-InfraRed Absolute Spectrophotometer (FIRAS) instrument on board the COsmic Background Explorer (COBE) satellite measured spectral distortions to be smaller than $\Delta\rho_\gamma/\rho_\gamma < 6 \times 10^{-5}$ [290], and a proposed future detector such as the Primordial Inflation Explorer (PIXIE) [291], or a more recent proposal [292], aims for constraints of $\Delta\rho_\gamma/\rho_\gamma < 8 \times 10^{-9}$. To calculate the constraints on the amplitude of the power spectrum due to the COBE/FIRAS observations, we insert eq. (2.17) into eq. (B.1) and set $\mu = 9\times10^{-5}$ which is the 2-σ constraint. We can then rearrange for A and compute the integral over k, plotting the constraint on A for each k_p. Our results for lognormal power spectra of widths $\Delta = 0.3$ and $\Delta = 1$ are shown in fig. 2.5. For complete clarity, the constraint on $\mathcal{P}_\mathcal{R}$ at a given k represents the maximum amplitude A for a lognormal power spectrum centered at $k = k_p$ so as not to induce μ-distortions that would be in conflict with the COBE/FIRAS constraint of $\mu < 9 \times 10^{-5}$.

B.2 The stochastic gravitational wave background

Here we summarise how the GW background can be calculated given a primordial power spectrum, adding more details to section 2.4.2. The contribution to the tensor power spectrum from the square of the scalar power spectrum is given by [124, 125]

$$\mathcal{P}_h(\tau, k) = 4\int_0^\infty \mathrm{d}v \int_{|1-v|}^{1+v} \mathrm{d}u \left(\frac{4v^2 - (1+v^2-u^2)^2}{4vu}\right)^2 I^2(v, u, k\tau)\mathcal{P}_\mathcal{R}(kv)\mathcal{P}_\mathcal{R}(ku), \tag{B.3}$$

where $u = |\mathbf{k} - \tilde{\mathbf{k}}|/k$, $v = \tilde{k}/k$ and $\tilde{k}$ is the wavelength corresponding to the scalar source. $I(v, u, k\tau)$ is a highly oscillatory function which contains the source information. We solve this integral numerically but note that it can be solved analytically in some regimes [126]. The observational quantity related to this power spectrum is the energy density of gravitational waves given by

$$\Omega_{\mathrm{GW}}(\tau, k) = \frac{\rho_{\mathrm{GW}}(\tau, k)}{\rho_{\mathrm{tot}}(\tau)} = \frac{1}{24}\left(\frac{k}{aH}\right)^2 \mathcal{P}_h(\tau, k). \tag{B.4}$$

If we assume that the entire contribution to any stochastic background detection is from the tensor power spectrum in eq. (B.3), 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. To calculate the constraints on the primordial power spectrum, we first calculate $\Omega_{\mathrm{GW}}h^2$ today as a function of k by inserting the lognormal power spectrum in eq. (2.17) with given k_p and Δ into eq. (B.3), pulling out the amplitude A which is the quantity that we aim to constrain. We perform this integral numerically once for each value of Δ, and the results can be shifted post-integration for any value of k_p.

B.3 Updated NANOGrav dataset

The 11 year NANOGrav dataset [203] includes improved modelling of the solar system ephemeris which make the constraints on the stochastic gravitational wave background weaker than they would be with previous models of these effects. That makes the improvement on the primordial power spectrum constraints between the 11 year dataset and the 9 year dataset [218] not as large as one might hope based purely on the improved sensitivity. This solar system ephemeris modelling effect also applies to other pulsar timing array observations from, for example, EPTA [219]. Therefore all of the constraints from these datasets need to be revised upwards by taking into account the better model for the solar system ephemeris. For this reason, we choose to use the 11 year NANOGrav dataset alone, despite the fact that the EPTA dataset reaches lower frequencies, and as a guide to the improvement between datasets we show the constraint on the amplitude of the primordial power spectrum for a lognormal power spectrum with width $\Delta = 1$ for both the 9 year and 11 year datasets in figure B.1.

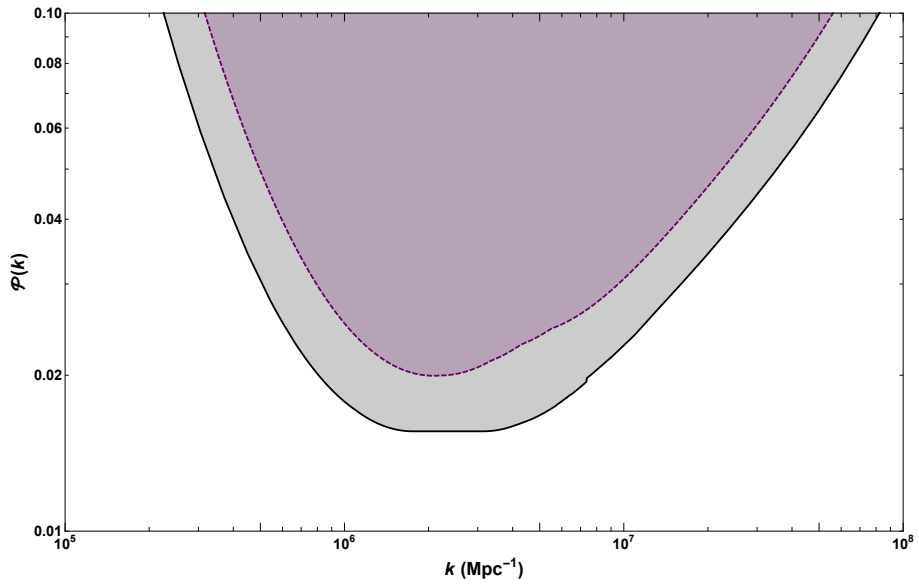

Figure B.1: Constraints on the amplitude of the primordial power spectrum due to NANOGrav pulsar timing array observations from the 9 year (purple, dashed) and 11 year (black, solid) datasets. For both datasets, constraints are for lognormal power spectra with width $\Delta = 1$.

C The non-linear relationship between ζ and δ

In recent years, there has been a large amount of literature discussing the fact that, even if the curvature perturbation ζ is Gaussian, the density contrast will not be [169–173], due to the non-linear relationship between the two parameters. In the super-horizon limit, the relationship between the two parameters can be calculated with a gradient-expansion approach. At first order in gradients, the full non-linear relationship, in polar coordinates and assuming spherical symmetry, is given by

$$\delta_{NL} = \frac{\delta\rho}{\rho_b}(r,t) = -\frac{4(1+\omega)}{5+3\omega}\left(\frac{1}{aH}\right)^2 e^{-5\zeta(r)/2}\nabla^2 e^{\zeta(r)/2}, \tag{C.1}$$

whilst the linear relation is

$$\delta_L = \frac{\delta\rho_l}{\rho_b} = -\frac{2(1+\omega)}{5+3\omega}\left(\frac{1}{aH}\right)^2 \nabla^2\zeta. \tag{C.2}$$

For simplicity, we will set the equation-of-state parameter $w = 1/3$ from here on.

We can define a time-independent component of the density contrast,

$$\delta_{\rm TI}(\mathbf{x}, R) = (R\ aH)^2\, \delta_{NL}, \tag{C.3}$$

where R is taken to be the scale of the perturbation. The compaction function $C(\mathbf{x}, R)$ is obtained by calculating the mass excess δM within a sphere of radius R, and dividing by R, which corresponds to smoothing the time-independent component of the density contrast with a top-hat smoothing function,

$$C(\mathbf{x}, R) = \frac{\delta M}{R} = \int \mathrm{d}^3 y\ \delta_{\rm TI}(x-y) W(y, R). \tag{C.4}$$

Performing this integral gives an expression for the compaction function at the centre of spherically symmetric peaks:

$$C(\mathbf{x}, R) = C_L - \frac{3}{8}C_L^2, \tag{C.5}$$

where C_L is the expression one would obtain using the linear relation above,

$$C_L = \frac{2}{3} R\zeta'(R), \tag{C.6}$$

where the prime denotes a derivative with respect to the smoothing scale R.

The rare, large-amplitude peaks from which PBHs form are well approximated by spherically-symmetric peaks [181], and so the above equation can be used to relate relevant peaks in C_L to peaks in the compaction C. We note that the compaction has a maximum value, $C_{\rm max} = 2/3$, corresponding to $C_L = 4/3$. For higher values of C_L, the compaction decreases—perturbations of this type correspond to a case for which PBH formation has not been simulated. For this reason, only perturbations with $C_L < 4/3$ are typically considered – although in practice this has little effect on the PBH abundance since such large values of C_L are exponentially suppressed.

If we then wish to calculate parameters related to the PBH abundance, we can simply replace the equation for the PBH mass, eq. (2.16), with a corresponding equation which relates the PBH mass to the linear, Gaussian component of the compaction instead

$$m = kM_H(C_L - \frac{3}{8}C_L^2 - C_c)^\gamma. \tag{C.7}$$

In order to make an analytic estimate for how constraints on the power spectrum are affected by this non-linearity, we can make a simple assumption that all peaks which form PBHs are close to the critical amplitude (since the abundance of significantly larger peaks is exponentially suppressed). In this simple case, and assuming $C_c = 0.55$ (the case for the top-hat window function, see eq. (2.14)), the critical amplitude for the linear component of the compaction is $C_{c,L} \approx 0.77$, i.e. we can assume that peaks in the linear field need to have an amplitude 1.41 times larger than if we assumed a linear relation between ζ and δ, as in eq. (C.2). Therefore, the power spectrum (which is proportional to the variance of perturbations) should be approximately $1.41^2 = 1.98$ times greater. We can test this approximation by com-

paring the full calculation of the amplitude required to generate a fixed abundance $f_{\rm PBH} = 2 \times 10^{-3}$ in the linear and non-linear cases. For the two lognormal power spectra considered in this paper, with widths $\Delta = 1$ and 0.3, the approximation holds to the precision of two decimal places stated above. Although this validity may vary with the position of the peak, we assume it holds globally for the results shown in figs. 2.5 and 2.6.

For the top-hat window function, there is a relatively simple analytic relationship relating the compaction function to the curvature perturbation (which we assume to be Gaussian). However, we note that if one instead uses a Gaussian window function, as we have considered in this paper, there is no analytic solution, and accounting for the non-linearity becomes complicated. When looking at individual perturbations, it is trivial to show that the amplitude of the compaction (or "compaction-like") function calculated with both a top-hat or Gaussian window function is proportional to the amplitude of the perturbation. Therefore, we expect the non-linearities described above to have a similar effect on constraints on the power spectrum, whether a top-hat or Gaussian function is used.

D Reduced χ^2 values

Table D.1 shows the reduced χ^2 values for each of the mass distribution models considered, across the full range of power spectrum peak widths Δ.

Table D.1: χ^2_ν values for different models and power spectrum widths.

	Model						
Width Δ	Lognormal	Gaussian	Skew-normal	Skew-lognormal	CC1	CC2	CC3
δ	3.67×10^{-4}	4.15×10^{-5}	1.03×10^{-5}	8.55×10^{-6}	3.37×10^{-4}	9.88×10^{-6}	5.78×10^{-7}
0.1	3.58×10^{-4}	4.56×10^{-5}	1.01×10^{-5}	8.32×10^{-6}	3.97×10^{-4}	1.16×10^{-5}	6.72×10^{-7}
0.3	2.96×10^{-4}	8.76×10^{-5}	8.54×10^{-6}	6.61×10^{-6}	1.06×10^{-3}	2.93×10^{-5}	1.39×10^{-6}
0.5	2.16×10^{-4}	1.91×10^{-4}	6.58×10^{-6}	4.32×10^{-6}	3.00×10^{-3}	7.11×10^{-5}	1.98×10^{-6}
1.0	8.25×10^{-5}	7.15×10^{-4}	2.02×10^{-5}	9.33×10^{-7}	1.62×10^{-2}	2.56×10^{-4}	1.04×10^{-6}
2.0	5.57×10^{-6}	2.41×10^{-3}	3.98×10^{-4}	8.90×10^{-8}	7.01×10^{-2}	7.09×10^{-4}	5.47×10^{-8}
5.0	6.93×10^{-5}	5.51×10^{-3}	3.40×10^{-3}	2.90×10^{-6}	1.47×10^{-1}	1.74×10^{-3}	1.11×10^{-4}

E Additional merger rate plots

E.1 Design sensitivity plots

Figure E.1 shows the design sensitivity merger rate distributions for the two lognormal mass distributions and the two power-law distributions. The astrophysical distributions are shown in green.

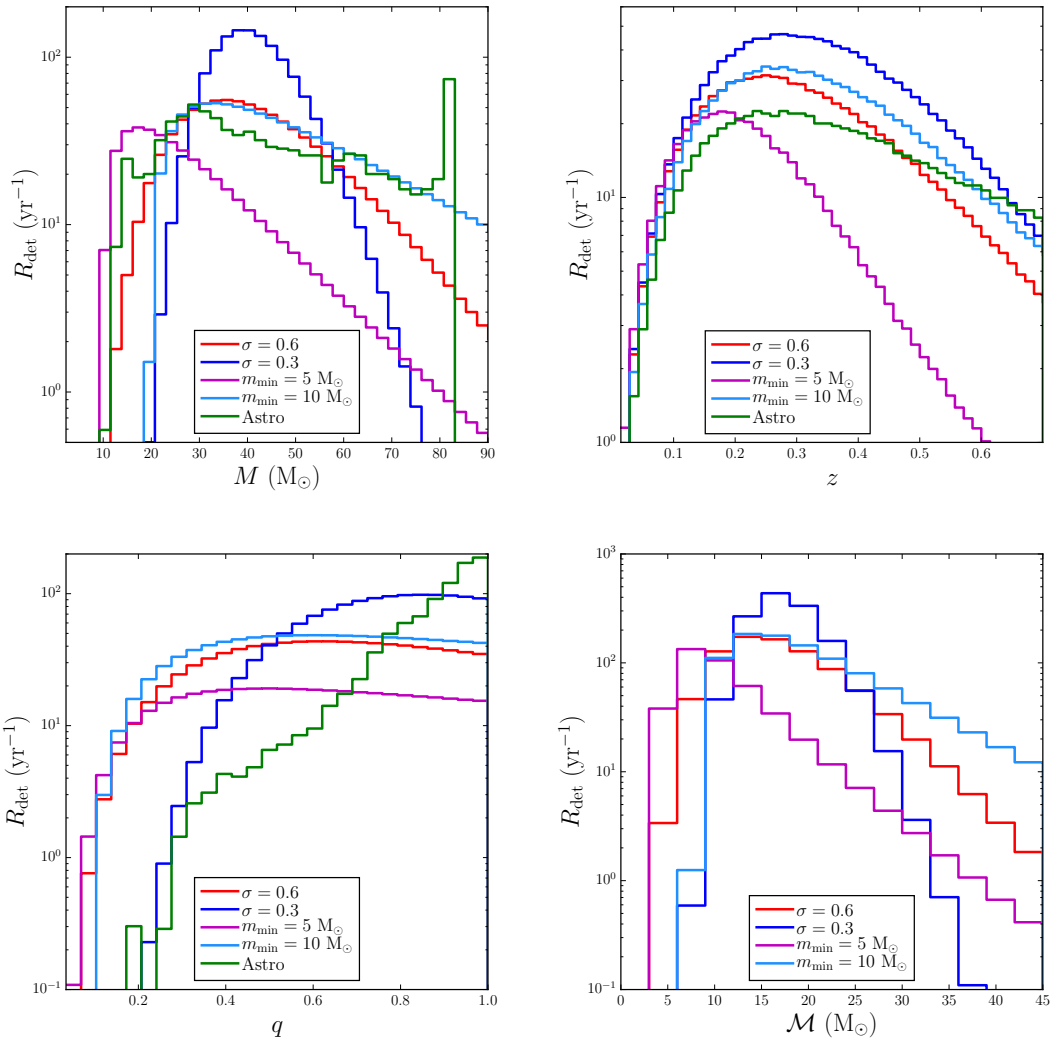

Figure E.1: Merger rate distributions in total mass M, redshift z, mass ratio q, and chirp mass $\mathcal{M}$ for a lognormal mass distribution with $\sigma = 0.3$ and $\sigma = 0.6$, and a power-law mass distribution with $m_{\rm min} = 5$ M$_\odot$ and $m_{\rm min} = 10$ M$_\odot$, at design sensitivity. The distributions for astrophysical black holes from [276] are shown in green for the first three plots. All plots have $f_{\rm PBH} = 10^{-2}$.

Figures E.2 to E.4 show the design sensitivity 2D distributions for a lognormal mass distribution with $m_c = 20\ \mathrm{M}_\odot$ and $\sigma = 0.3$, and two power-law mass distributions with $\alpha = 3/2$ and $m_{\mathrm{min}} = 5\ \mathrm{M}_\odot$ and $10\ \mathrm{M}_\odot$ respectively.

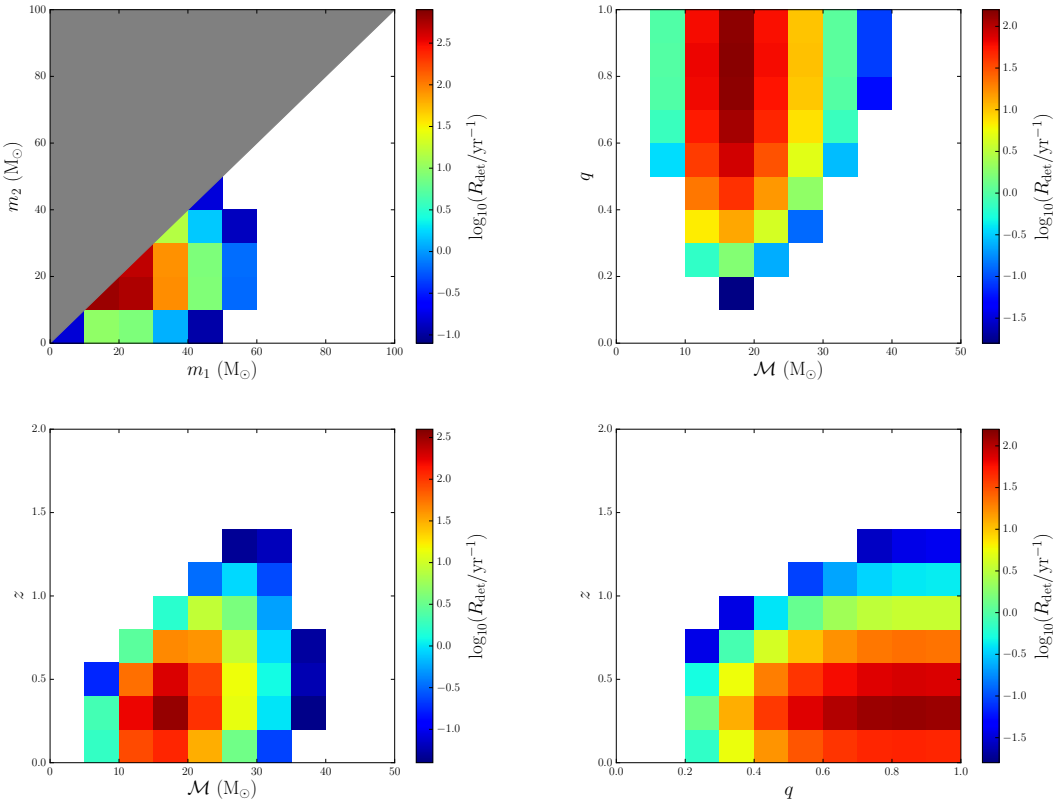

Figure E.2: 2D merger rate distributions in individual masses, mass ratio q, chirp mass $\mathcal{M}$ and redshift z for a lognormal mass distribution with $\sigma = 0.3$, at design sensitivity. All plots have $f_{\mathrm{PBH}} = 10^{-2}$. The white area indicates no significant merger rate and the grey triangle indicates the choice $m_2 < m_1$. The colorbar limits are the same as in fig. 4.4 for comparison.

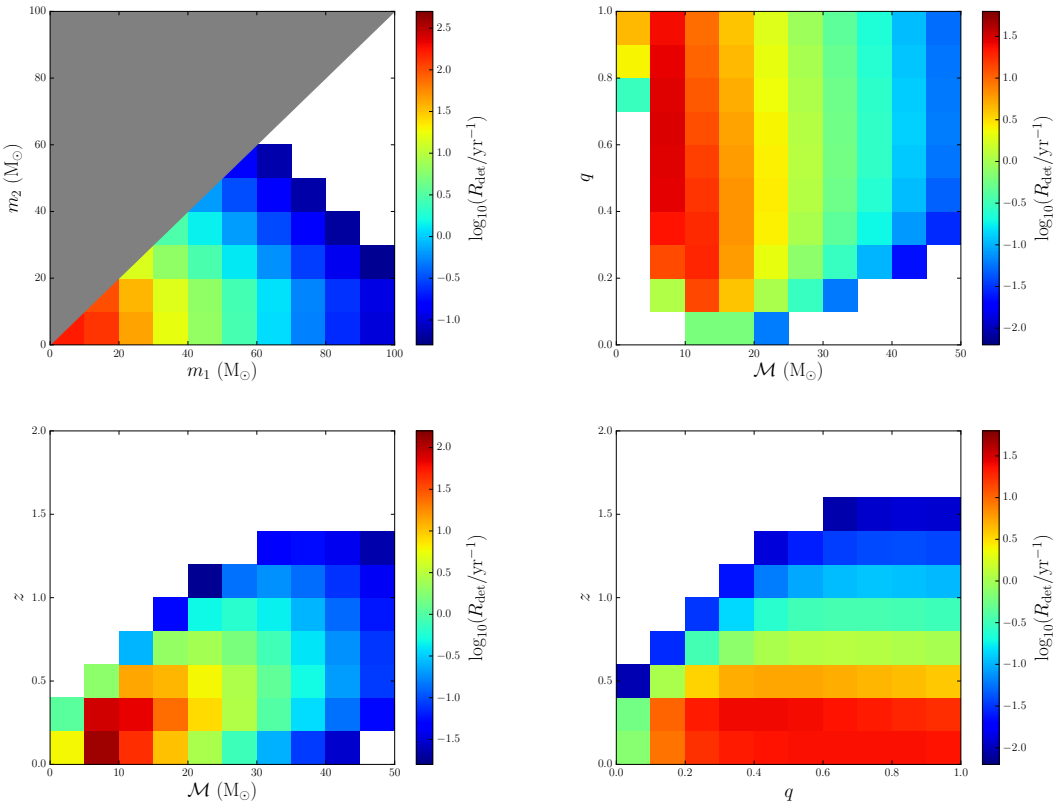

Figure E.3: 2D merger rate distributions in individual masses, mass ratio q, chirp mass $\mathcal{M}$ and redshift z for a power-law mass distribution with $m_{\rm min} = 5\ {\rm M}_\odot$, at design sensitivity. All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate and the grey triangle indicates the choice $m_2 < m_1$. The colorbar limits are the same as in fig. E.4 for comparison.

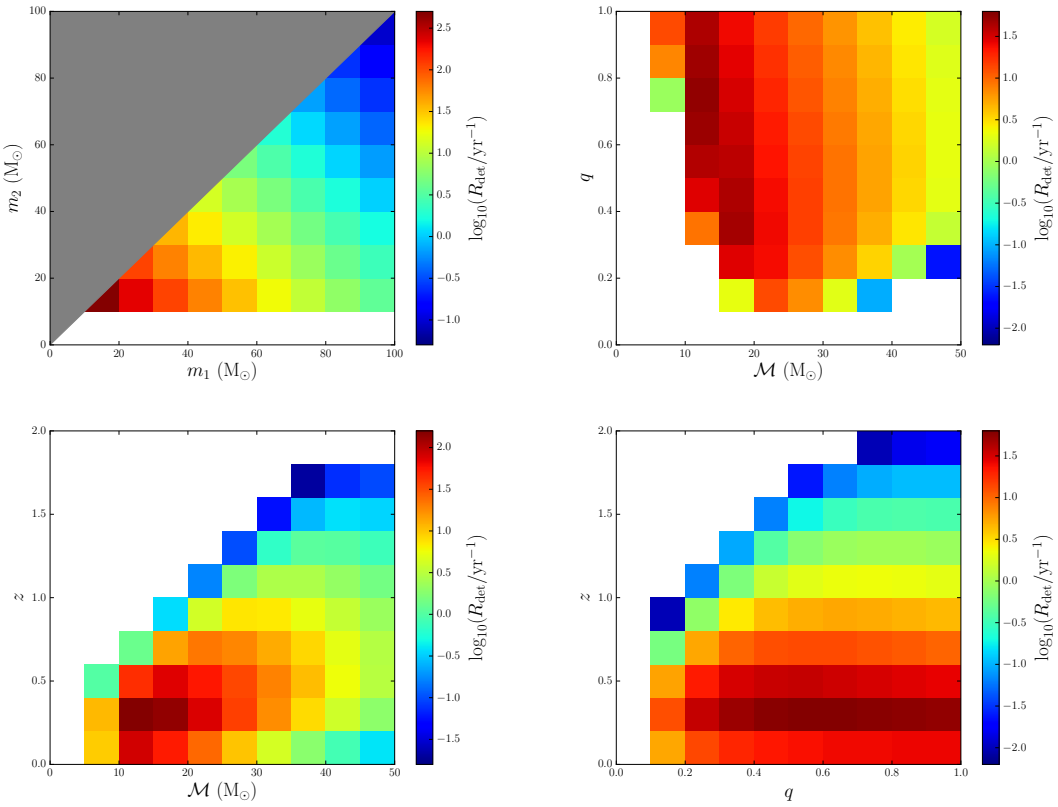

Figure E.4: 2D merger rate distributions in individual masses, mass ratio q, chirp mass $\mathcal{M}$ and redshift z for a power-law mass distribution with $m_{\rm min} = 10\ {\rm M}_\odot$, at design sensitivity. All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate and the grey triangle indicates the choice $m_2 < m_1$. The colorbar limits are the same as in fig. E.3 for comparison.

E.2 O1O2 sensitivity plots

Figures E.5 to E.7 show the O1O2 sensitivity 2D distributions for the four mass distributions in the observables ($\mathcal{M}$, q), ($\mathcal{M}$, z), and (q, z) respectively. The LIGO values and their 1D marginalised 90% confidence limits are shown by the black dots and their error bars.

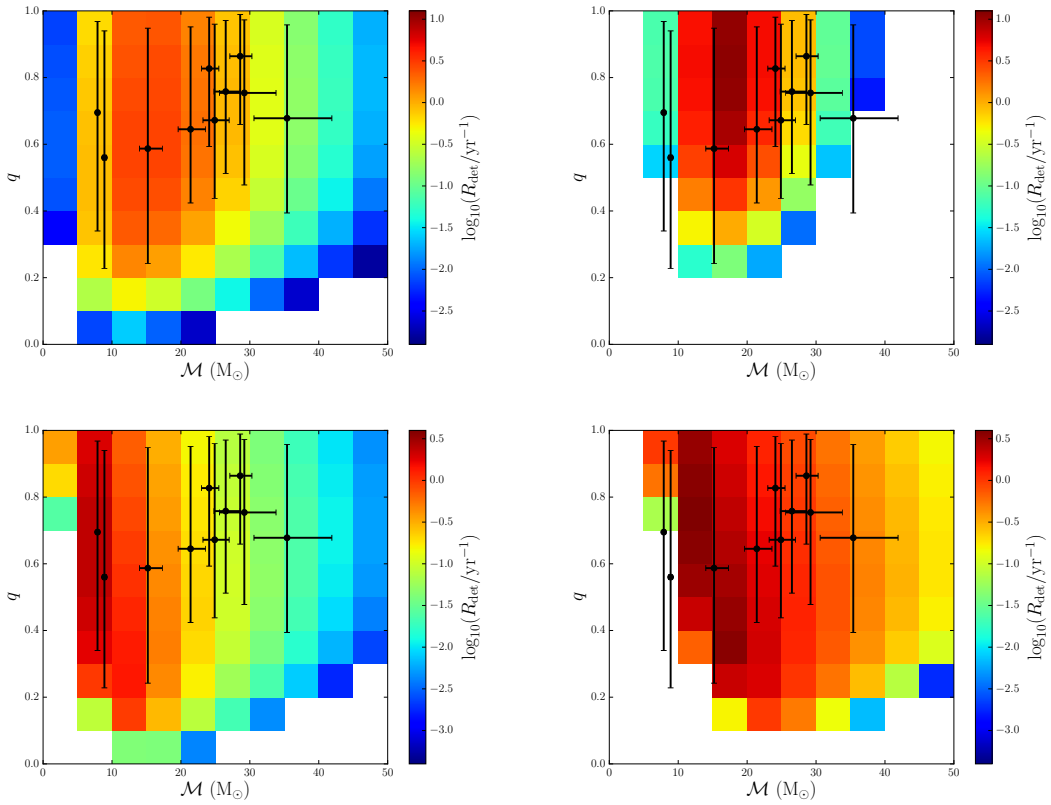

Figure E.5: 2D merger rate distributions in mass ratio q and chirp mass $\mathcal{M}$ for the four mass distributions considered, at O1O2 sensitivity. The top row has the lognormal mass distribution with widths $\sigma = 0.6$ (left) and 0.3 (right), and the bottom row has the power-law mass distribution with minimum mass 5 $M_\odot$ (left) and 10 $M_\odot$ (right). All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate. The colorbar limits are the same for the top two plots (lognormal distribution), and for the bottom two plots (power-law distribution). The LIGO values and their 90% confidence limits are shown in black.

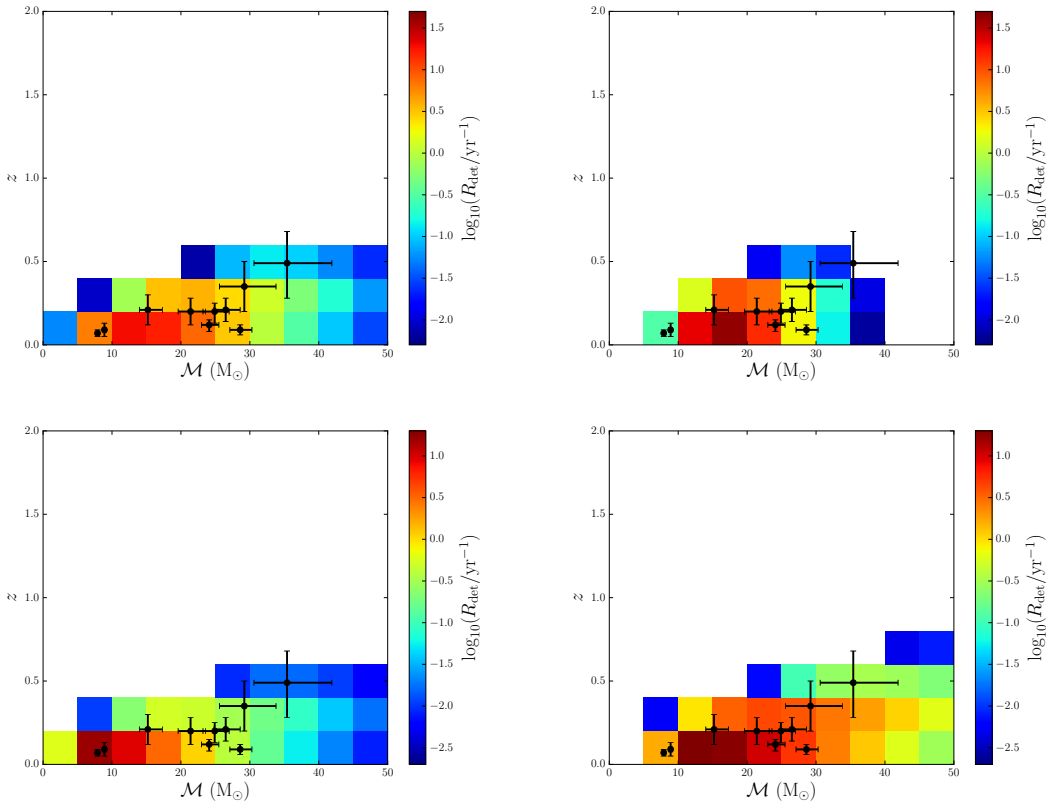

Figure E.6: 2D merger rate distributions in redshift z and chirp mass $\mathcal{M}$ for the four mass distributions considered, at O1O2 sensitivity. The top row has the lognormal mass distribution with widths $\sigma = 0.6$ (left) and 0.3 (right), and the bottom row has the power-law mass distribution with minimum mass 5 $M_\odot$ (left) and 10 $M_\odot$ (right). All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate. The colorbar limits are the same for the top two plots (lognormal distribution), and for the bottom two plots (power-law distribution). The LIGO values and their 90% confidence limits are shown in black.

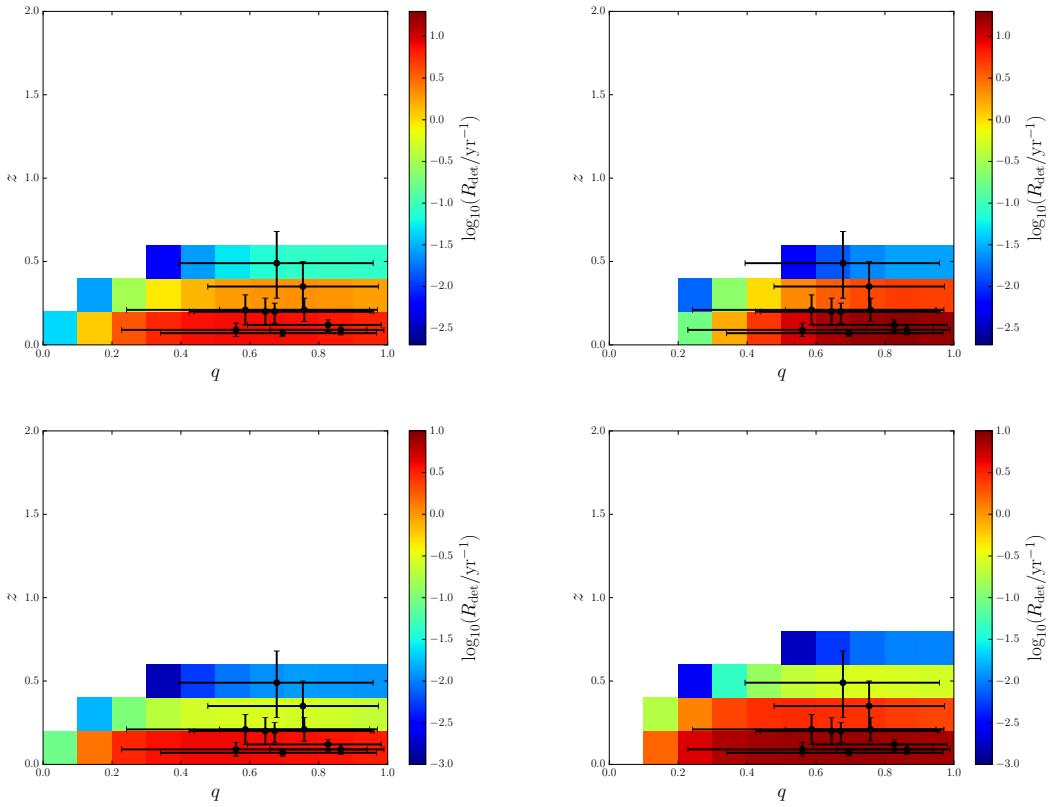

Figure E.7: 2D merger rate distributions in redshift z and mass ratio q for the four mass distributions considered, at O1O2 sensitivity. The top row has the lognormal mass distribution with widths $\sigma = 0.6$ (left) and 0.3 (right), and the bottom row has the power-law mass distribution with minimum mass 5 $M_\odot$ (left) and 10 $M_\odot$ (right). All plots have $f_{\rm PBH} = 10^{-2}$. The white area indicates no significant merger rate. The colorbar limits are the same for the top two plots (lognormal distribution), and for the bottom two plots (power-law distribution). The LIGO values and their 90% confidence limits are shown in black.