School of Physics
and Astronomy

You Spin Me Round:
Measuring Precession in the Black Hole Population

Charlie G. Hoy

Submitted for the degree of Doctor of Philosophy
School of Physics and Astronomy
Cardiff University

June 14, 2021

Summary of thesis

Black holes: cosmic bodies of extreme gravity that nothing, not even light, can escape. Einstein first told the world about the existence of these enigmatic objects in his groundbreaking theory of General Relativity. It was not until 2015 that black holes were probed directly through the first observation of gravitational waves. The treasure trove of knowledge gleaned from this observation revolutionised our understanding of the cosmos. Since this Nobel prize winning observation, the LIGO Scientific, Virgo and KAGRA collaborations have observed fifty gravitational waves.

The aim of this thesis has been to take advantage of the growing population of gravitational wave sources to answer the following fundamental question: do binary black holes undergo spin-induced orbital precession? This has significant implications on our understanding of how binary black holes form in nature.

To answer this question, I first introduced a brand new formalisation for modelling a gravitational wave that originated from a precessing system. I then introduced the "precession signal-to-noise ratio" which naturally followed from this unique description. This novel tool quantified, for the first time, the significance of precession in an observed gravitational wave. This elegant new formalisation is presented in Chapter 2 and verified in Chapter 3.

In the subsequent four chapters, I used the "precession signal-to-noise ratio" to answer the aforementioned question. In Chapter 4 I demonstrated that there is no evidence for precession in any of the binary black hole candidates from the first gravitational wave catalog. I then described how this lack of precession allows us to constrain the properties of black holes. In Chapter 5 I presented the properties of potentially the first neutron star-black hole binary observed – a system which is most likely to have measurable precession as a result of the asymmetric component masses. In Chapter 6 I calculated the "precession signal-to-noise ratio" for all gravitational wave candidates observed in the first half of the third gravitational wave observing run, and demonstrated that three observed gravitational waves could have originated from precessing systems. In Chapter 7, I used the gravitational wave data from the second gravitational wave catalog to determine the most likely spin distribution of black holes. By doing so, I was able to determine whether the population of binary black holes are likely to undergo spin-induced orbital precession.

Finally, in Chapter 8, I presented a new and innovative software package to analyse, display and combine posterior samples. This package has become one of the major workhorses of the LIGO Scientific, Virgo and KAGRA collaborations and is widely distributed through the gravitational wave data analysis computing environment.

Contents

List of Figures

List of Tables

List of Publications

1. **Charlie Hoy** and others. *Constraining black-hole spins with gravitational wave observations 2.* in preparation, Chapter 7 in this thesis.

2. **Charlie Hoy**, Cameron Mills, and others. *Precession and higher order multipoles in the first half of the third gravitational wave observing run.* in preparation, Chapter 6 in this thesis.

3. E. Fauchon-Jones and others inc. **Charlie Hoy**. *BAM catalogue of binary black hole simulations.* in preparation.

4. Rhys Green, **Charlie Hoy**, Stephen Fairhurst, Mark Hannam, Francesco Pannarale, and Cory Thomas. *Identifying when Precession can be Measured in Gravitational Waveforms.* arXiv:2010.04131. Phys. Rev. D, 103, 124023.

5. **Charlie Hoy** and Vivien Raymond. *PESummary: the code agnostic Parameter Estimation Summary page builder.* arXiv:2006.06639. Chapter 8 in this thesis.

6. I. M. Romero-Shaw and others inc. **Charlie Hoy**. *Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue.* arXiv:2006.00714. Mon. Not. Roy. Astron. Soc., 499, 3295–3319.

7. Stephen Fairhurst, Rhys Green, **Charlie Hoy**, Mark Hannam, and Alistair Muir. *Two-harmonic approximation for gravitational waveforms from precessing binaries.* arXiv:1908.05707. Phys. Rev. D, 102, 024055, Chapter 2 in this thesis.

8. Stephen Fairhurst, Rhys Green, Mark Hannam, and **Charlie Hoy**. *When will we observe binary black holes precessing?* arXiv:1908.00555. Phys. Rev. D, 102, 041302, Chapter 4 in this thesis.

9. R. Abbott and others inc. **Charlie Hoy**. *GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object.* arXiv:2006.12611. Astrophys. J. Lett., 896, L44, Chapter 5 in this thesis.

10. B.P. Abbott and others inc. **Charlie Hoy**. *GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs.* arXiv:1811.12907. Phys. Rev. X, 9, 031040, Chapter 5 in this thesis.

I am also an author on many other collaboration papers not listed here for which my contributions vary.

Acknowledgments

Firstly, I would like to thank my supervisor Stephen Fairhurst. You always had the answer to every question (even when I asked it multiple times!) and never ceased to amaze me with your knowledge. Whenever things were not going well, you always had a suggestion and nine times out of ten it always worked. I will always be thankful that your "one good idea every 10 years" happened during my PhD. Without it, who knows where we would have ended up! It has been an absolute pleasure working with you.

Thanks also go to Mark Hannam for answering my numerous questions and queries, even over a panicked WhatsApp... I hope you don't regret giving me your phone number! Without you I would still be adding flavourless cherries to my cocktails, I wouldn't know the in and outs of Gimlets and I would never have thought to mix Gin and pasta. But of course, you're right Dude. 100% certain.

I would also like to thank Dave and Jonathan. Whisky night has been a constant in the whirlwind of changes this past year. It has been a necessary release and one that I have thoroughly enjoyed. Thank you Dave for being a first-class housemate. I have never laughed as much as when we were inventing the "Marktini" or researching the "Slippery Lionel".

Special thanks also goes to everyone at 53. It may no longer be ours, but the conversations and procrastinations are as firmly dented in my mind as the impactful paper aeroplane. Ed, thank you for all of your help over the course of my PhD; you have taught me so much and I am forever grateful. Rhys, thank you for being there on my first day and sharing your Master's project with me, even though it did take over 2 years longer than what we initially anticipated! Phil, Virginia, thank you for listening to me rant and making me laugh, especially over the last year. Cam, thank you for putting up with my relentless requests over the last couple of weeks. Eleanor and 2-I'd Iain thank you for making my PhD fun and memorable.

Thank you to all the incredible students, postdocs and faculty that I have collaborated with over the years, particularly Vicky. Your guidance during GW190814 was invaluable and has shaped me into the researcher I am today.

Of course, thank you to my family. Mum, Dad, Danielle, Lee, thank you for always supporting me and trying to understand this research. After many years, I can now finally get a 'proper job'. Himani and Rusty, thank you for welcoming me into your home throughout the pandemic; I am aware that I am not the easiest person to live with! Kimi, you are my best friend and future wife. You have always been there every step of the way and I would not been able to do this without you.

Finally, a special thank you to Flair, Hario, Jonathan, Phil and Vaibhav. You all went above and beyond what I could ever have imagined.

Chapter 1

Introduction

Albert Einstein revolutionised our understanding of gravity in his groundbreaking theory of General Relativity [1]. Einstein demonstrated that the equations proposed by Isaac Newton over 200 years earlier [2] were actually a special case of a more general theory; a theory which predicted phenomena that could not be explained in the Newtonian picture. One such example are *gravitational waves* [3]. Gravitational waves are disturbances to space-time caused by the acceleration of matter. Almost 100 years after Einstein first predicted the existence of gravitational waves, these cosmic ripples were first detected in 2015 by the LIGO Scientific and Virgo collaborations with two of the most advanced experimental setups ever built [4, 5, 6]. This observation marked the beginning of the era of gravitational wave astronomy [7, 8].

1.1 Gravitational Waves

1.1.a Gravitational Wave Sources

Similar to electromagnetic radiation, different sources emit gravitational waves at different frequencies (see the *gravitational wave spectrum* [9]). Typically, gravitational waves are grouped into 3 main categories: high, low and very low frequency. High frequency gravitational waves ($10\text{Hz} < f < 10^3\text{Hz}$), detectable with current ground based gravitational wave detectors such as GEO [10], LIGO [4, 5], Virgo [6] and KAGRA [11], originate from the collision of stellar mass compact objects (such as black holes and neutron stars), supernovae [12] and rotating neutron stars [13, 14]. Potential space-based gravitational wave detectors, such as the Laser Interferometer Space Antenna (LISA) [15], are designed to be sensitive to low frequency gravitational waves: $10^{-4}\text{Hz} < f < 1\text{Hz}$. These detectors are likely to observe the coalescence of super massive black holes [16]. Gravitational waves emitted at even lower frequencies are generated by non-stellar objects and may be detected through pulsar timing arrays [17] and/or their imprint on the cosmic microwave background [18]. In this thesis, I will focus on high frequency gravitational waves; specifically those

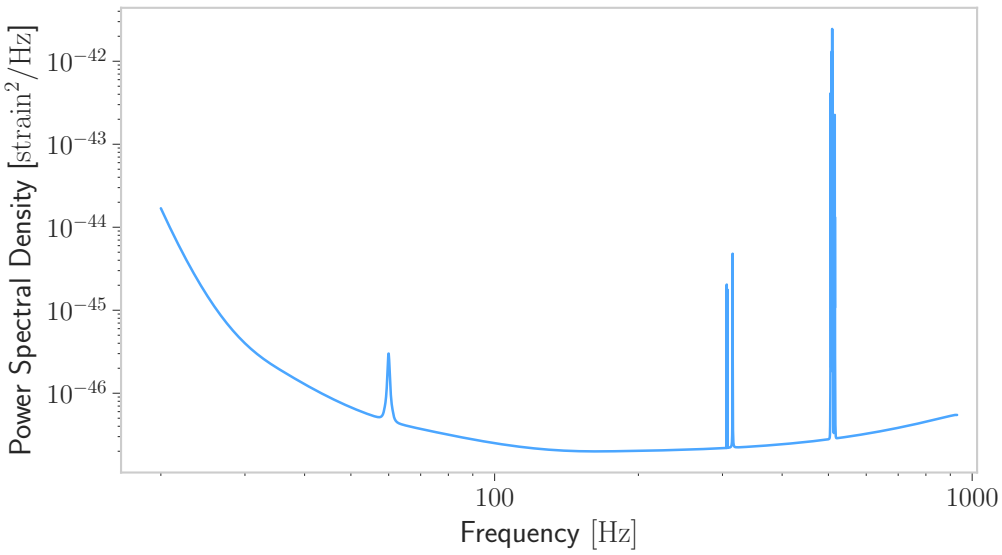

Figure 1.1: Typical sensitivity curve for the LIGO gravitational wave detector located in Livingston, Louisiana. The bump at 60Hz is a consequence of the detectors AC power supply vibrating electronic equipment. The forest of peaks at $\sim$ 300Hz and $\sim$ 500Hz are 'violin modes' arising from the vibration of wires suspending the beam splitter and mirrors in each arm [25].

generated by the coalescence of stellar mass compact objects.

1.1.b Detecting Gravitational Waves

High frequency gravitational waves can be detected on Earth with Michelson interferometers [19] (although other types of detectors have been proposed [20]). Here, a laser is split by a beam splitter into two beams which are sent down two orthogonal arms of the interferometer. The beams are reflected by mirrors at the ends of each arm, and recombined at the beam splitter where they are observed by a photodiode. Any passing gravitational wave periodically changes the relative difference in arm length and as such, modifies the interference pattern observed at the photodiode [see e.g. 21].

After travelling cosmological distances, the gravitational wave amplitude, expressed as the dimensionless strain, is very small: $O(10^{-21})$. This means that a gravitational wave passing through a Michelson interferometer with arm length 1m periodically changes the relative difference in arm length by 10^{-21}m [22]. Ground-based gravitational wave detectors are therefore more complex than traditional Michelson interferometers in order to detect these very small perturbations. For instance, in order to amplify the effect of the gravitational wave signal, the orthogonal arms are kilometres long [5, 23] and the laser is reflected back and forth many times prior to recombining to increase the effective path length of each laser beam [24].

Ground based gravitational wave detectors are subject to many sources of un-

wanted noise; for example ground noise caused by the motion of cars nearby and strong winds hitting the sides of the orthogonal arms [26]. Consequently, sophisticated techniques are employed to isolate and reduce external noise as far as possible [see e.g. 27, 28, 29].

In the absence of a passing gravitational wave, the output from the photodiode can be used to generate a one-sided power spectral density (PSD) representing the detectors sensitivity [30]. We show a typical PSD belonging to the gravitational wave detector located in Livingston Louisiana in Figure 1.1, produced using the `BayesLine` algorithm [31]. Although in general we see a relatively smooth function, the PSD includes narrow lines arising from noise sources in the instrument. We see a bump at 60Hz originating from the detectors AC power supply vibrating electronic equipment and a *forest of peaks* at $\sim$ 300Hz and $\sim$ 500Hz arising from resonances in the wires suspending the beam splitter and mirrors in each arm [25].

The Amplitude Spectral Density (ASD), defined as $\sqrt{\text{PSD}}$, describes the amplitude of noise in the detector. Any gravitational wave signal originating from the coalescence of stellar mass compact objects with amplitude lower than the ASD is therefore too quiet to detect. We see that the noise amplitude is lowest within the frequency range $10\text{Hz} < f < 10^3\text{Hz}$ with amplitude $\text{O}(10^{-23}\,\text{strain}/\sqrt{\text{Hz}})$.

1.1.c Observed Gravitational Waves

During the first gravitational wave observing run (O1), September 2015 to January 2016, the LIGO Scientific and Virgo collaborations witnessed the first set of gravitational waves from binary black hole mergers [7, 32, 33]. The second gravitational wave observing run (O2), November 2016 to August 2017, saw the first gravitational wave from a binary neutron star coalesence [34] and a further seven gravitational wave observations from binary black hole mergers [35, 36, 37, 38]. Groups outside of the LIGO Scientific, Virgo and KAGRA collaborations have reported on additional gravitational wave candidates found during O1 and O2 [39, 40, 41, 42]. The first half of the third gravitational wave observing run (O3a), April 2019 to October 2019, saw a second binary neutron star coalescence [43], numerous binary black hole mergers [44, 45, 46] and potentially the first set of observations from neutron star black hole mergers [44, 47]. We show the population of gravitational wave candidates detected by the LIGO Scientific, Virgo and KAGRA collaborations during O3a in Figure 1.2. This plot is split into three regions of parameter space; each region is quantified by the mass of the primary and secondary compact objects in the binary, m_1 and m_2 respectively ($m_1 > m_2$). Systems with $m_1, m_2 > 3M_\odot$ are most likely binary black holes as neutron stars are unlikely to have masses above $3M_\odot$ [48, 49].

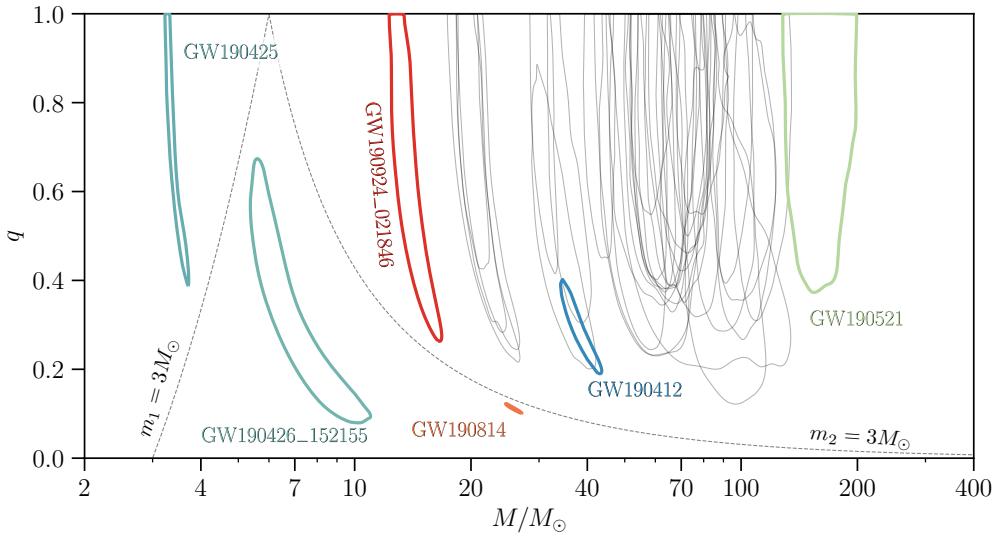

Figure 1.2: High frequency gravitational wave sources from binary mergers observed by the LIGO Scientific, Virgo and KAGRA collaborations within the total binary mass M and binary mass ratio $q = m_2/m_1 \leq 1$[1]parameter space. Each contour shows the 90% credible region for a given event. The dashed lines delineate regions where the primary/secondary compact object can have a mass below $3M_\odot$. The following *exceptional* event candidates are highlighted: GW190412 [45], GW190425 [43], GW190426_152155 [44], GW190521 [46], GW190814 [47], and GW190924_021846 [44]. Plot taken from Ref. [44].

1.2 Spin-Induced Orbital Precession

Compact stellar objects are parameterised by their mass m and spin $\mathbf{S}$[2]. The spin of each compact object is often parametrised by the dimensionless spin $\chi = \mathbf{S}/m^2$, which obeys the Kerr limit $|\chi| \leq 1$ [52].

Denoting the primary by 1 and the secondary by 2, a binary of compact stellar objects has total spin $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$, orbital angular momentum $\mathbf{L}$ and total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$, as shown in Figure 1.3. If the spins of the compact objects are parallel/anti-parallel with $\mathbf{L}$, the direction of $\mathbf{S}_1$, $\mathbf{S}_2$ and $\mathbf{L}$ all remain fixed as the emission of gravitational waves decreases the amplitude of the orbital angular momentum and causes the system to spiral inwards. It is convenient to describe this binary as an (anti-)aligned-spin system and we will use this notation throughout this thesis. For a non-eccentric (anti-)aligned-spin binary, the amplitude and frequency of the emitted gravitational wave increases as the orbital separation decreases. The system eventually merges and forms a single perturbed black hole that emits gravitational radiation as a superposition of quasinormal ringdown multipoles, until the

[1]This plot defines the mass ratio to be $q = m_2/m_1 \leq 1$. The rest of this thesis defines the mass ratio to be $q = m_1/m_2 \geq 1$

[2]A neutron star has additional degrees of freedom quantifying its ability to deform, see Chapter 5 (Sec. 5.3.a) and e.g. [50, 51]

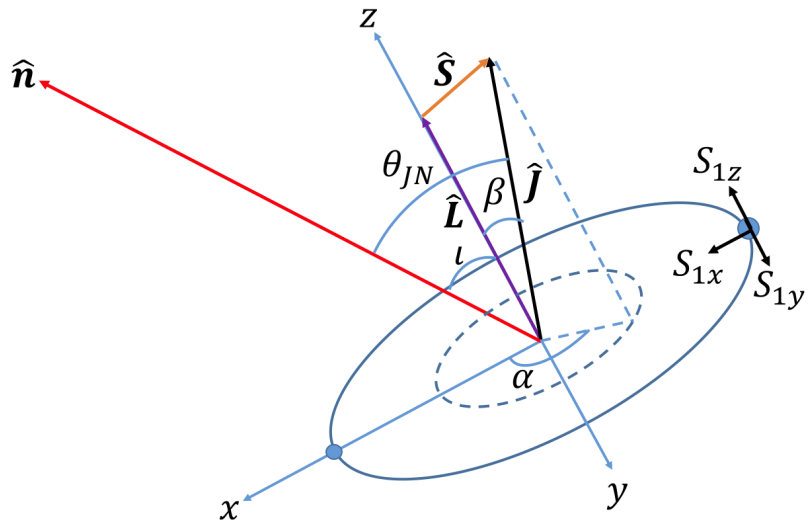

Figure 1.3: A plot showing the orbital angular momentum **L**, total spin **S** and total angular momentum **J** of the binary in the co-precessing frame. The normal vector $\hat{\mathbf{n}}$ here indicates the line of sight of the observer, $\hat{\mathbf{L}}$, $\hat{\mathbf{S}}$ and $\hat{\mathbf{J}}$ are the orbital angular momentum, spin angular momentum and total angular momentum vectors respectively. S_{1x}, S_{1y} and S_{1z} are the x, y and z components of the spin on the larger object. Note that $\mathbf{J}, \mathbf{L}$ and $\hat{\mathbf{n}}$ are shown to be co-planar only for ease of illustration. It is not true in general.

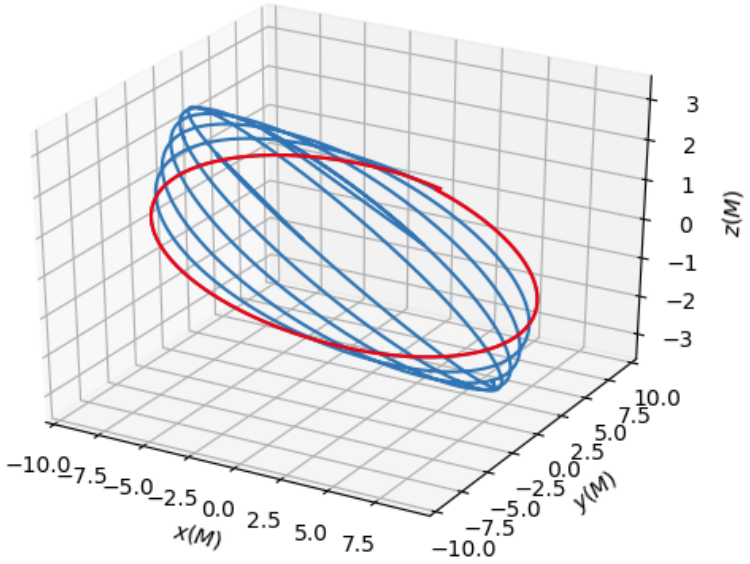

Figure 1.4: A plot showing the trajectory of a binary undergoing simple precession. For simplicity, only the path of the secondary compact object is shown. The red line shows the first initial orbital plane of the binary. This plot was produced using the BAM Numerical Relativity software [55, 56, 57]. As is customary in Numerical Relativity, the x, y and z coordinates of the orbital plane are expressed in terms of the total mass of the system M. The conversion factor is G/c^2.

system settles down to its final state [53]. If the spins of the compact objects are not parallel with $\mathbf{L}$, i.e. the spins have a non-zero component within the plane of the binary, the direction of $\mathbf{S}_1, \mathbf{S}_2$ and both the amplitude and direction of $\mathbf{L}$ vary in time.[3]. This causes the orbital plane to undergo a General Relativistic phenomenon known as spin-induced orbital precession [54].

For an (anti-)aligned-spin/non-precessing system, the orbital plane remains fixed throughout the coalescence, as shown by the red track in Figure 1.4. This means that the inclination angle of the source ι, the angle between the line of sight $\hat{\mathbf{n}}$ (here a caret denotes a unit vector, e.g. $\hat{\mathbf{a}} = \mathbf{a}/|\mathbf{a}|$) and $\mathbf{L}$, remains fixed. For a system where $\mathbf{L} \nparallel \mathbf{S}$ and $L \ll J$, $\mathbf{L}$, and therefore the orbital plane, precesses around (the approximately constant) $\mathbf{J}$ with the spins precessing such that $\dot{\mathbf{S}} = -\dot{\mathbf{L}}$ [54, 58]. This is shown by the blue tracks in Figure 1.4 and is known as *simple precession*. Since precession causes the inclination angle of the system to become time dependent, it is often convenient to denote the inclination of a precessing binary by the approximately constant angle between $\mathbf{J}$ and $\hat{\mathbf{n}}$, denoted by $\theta_{\rm JN}$. As the maximum instantaneous energy emission is approximately in the direction of $\hat{\mathbf{L}}$, the energy emitted in gravitational waves from a precessing binary system in the $\hat{\mathbf{n}}$ direction is also time dependent. This introduces characteristic amplitude and phase modulations in the observed gravitational wave which are modulated on the precession timescale [54, 58]. For a small minority of cases where $\mathbf{L} \approx -\mathbf{S}$, we see *transitional precession* [54]. As $\mathbf{L}$ is radiated away during inspiral through gravitational wave emission (recall that in the Newtonian limit $L \propto \sqrt{r}$, where r is the orbital separation), $\mathbf{J}$ crosses $\mathbf{J} = 0$ resulting in a change of sign and thus direction. In this thesis, I focus on systems that undergo *simple precession* since detecting high frequency gravitational waves exhibiting the effects of transitional precession are expected to be very rare [54].

For a system undergoing simple precession, the observed amplitude and phase modulations depends on $\theta_{\rm JN}$, the polarisation ψ of the observed signal, β, the polar angle between $\mathbf{L}$ and $\mathbf{J}$ and α, the azimuthal angle of $\mathbf{L}$ around $\mathbf{J}$. These angles are illustrated in Figure 1.3. For now we note several well-known features of precessing waveforms [54, 58], which will be further sharpened throughout this thesis. The strength of spin-induced orbital precession is characterised by β and by the precession frequency $\Omega_{\rm p}$. β is determined primarily by the total spin within the plane of the binary, the binary's mass ratio $q = m_1/m_2 \geq 1$, and separation r. At leading order the magnitude of $\mathbf{L}$ is given by $L = \mu\sqrt{Mr}$, meaning that we can write,

$$\cos\beta = \frac{\mu\sqrt{Mr} + S_{\parallel}}{\left[\left(\mu\sqrt{Mr} + S_{\parallel}\right)^2 + S_{\perp}^2\right]^{1/2}}, \tag{1.1}$$

[3] Although the amplitude of $\mathbf{S}_1$ and $\mathbf{S}_2$ can vary in time [54], this effect is significantly smaller than the change in $\mathbf{L}$ and therefore it is common to assume that it is constant. To a good approximation, the orbit-averaged spin components parallel and perpendicular to L, $S_{1,||}, S_{2,||}$ and $S_{1,\perp}, S_{2,\perp}$ also remain fixed

where $M = m_1 + m_2$ is the total mass of the source, $\mu = m_1 m_2/M$ is the reduced mass and $S_{\parallel}$ and $S_{\perp}$ are the total spin parallel and perpendicular to $\mathbf{L}$ respectively. In general, the larger the "opening angle" β, the more prominent the precession effects. Although in simple precession cases β slowly increases during inspiral [54], it typically varies very little over the portion of a binary's inspiral that is visible in a gravitational wave detector. This means that a) it is often possible to make the approximation that β is constant and b) if the binary is viewed along $\hat{\mathbf{J}}$ (i.e $\hat{\mathbf{n}}$ is aligned with $\hat{\mathbf{J}}$), then $\iota \approx \beta$ and therefore ι varies slowly and with minimal oscillations due only to orbital nutation. This approximation has been used to good effect in Ref. [59], and we will use it in some of the discussion in this thesis.

The precession frequency $\Omega_{\rm p}$ of $\mathbf{L}$ around $\mathbf{J}$ is given by $\Omega_p = \dot{\alpha}$, which, at leading order, can be written as,

$$\Omega_p \approx \left(2 + \frac{3}{2q}\right) \frac{J}{r^3}. \tag{1.2}$$

This means that to first approximation it *does not* depend on the spins (or therefore the opening angle β), but only on the binary's total mass, mass ratio, and separation (or equivalently orbital frequency). The number of precession cycles over a certain time or frequency range (e.g., over the course of an observation), depends on the total mass and mass ratio of the binary. In a gravitational wave observation there is a partial degeneracy between the mass ratio and the aligned spin $S_{||}$ [60, 61, 62], meaning that one of the chief effects for a measurement of precession will be to improve the measurement of these two physical properties [63].

It is often convenient to quantify the binary's in-plane spin by the scalar quantity $\chi_{\rm p}$ [64] (although see Refs. [65, 66] for alternative measures). $\chi_{\rm p}$ is obtained by averaging the in-plane spin components that drive precession over a precession cycle, meaning that it is more accurate for a system that undergoes many precession cycles. Motivated by the leading-order post-Newtonian (PN) precession equations [54, 58],

$$\chi_{\rm p} = \frac{1}{A_1} \max\left(A_1 S_{1\perp}, A_2 S_{2\perp}\right), \tag{1.3}$$

where $A_1 = 2 + 3/(2q)$, $A_2 = 2 + 3q/2$ and $S_{i\perp}$ is the component of the spin perpendicular to $\mathbf{L}$ on the ith object in the binary. $\chi_{\rm p}$ takes values between 0 (non-precessing system) and +1 (maximally precessing system).

Whether or not a binary of compact stellar objects undergoes spin-induced orbital precession depends primarily on how the system came into existence. The two favoured formation mechanisms for merging compact stellar objects are isolated binary evolution and dynamical formation in clusters [see e.g. 67, for a review]. In isolated models (also referred to as field models), two stars are born in a relatively wide binary and coevolve. The radius of the binary is then rapidly reduced through dynamically unstable mass transfer [68]. It is expected that binaries formed through this channel have spins distributed about $\mathbf{L}$ with a some small unknown misalign-

ment angle [see e.g. 69, 70, 71]; little to no precession is expected in these binaries. In dynamical models, two massive stars collapse to compact objects independently within a dense stellar environment, such as globular clusters, and then form a binary through interactions with other objects. It is expected that binaries formed through this channel have spins isotropically orientated [72]; precession is expected in these binaries.

1.2.a Modelling a precessing gravitational wave signal

Throughout this thesis, I primarily use the `IMRPhenomPv2` phenomenological model presented in Ref. [73] to model the gravitational wave emitted from a precessing binary. This model exploits the phenomenology of simply precessing binaries, described above, with the additional approximation that a precessing waveform can be factorised into an underlying non-precessing waveform and the precessional dynamics [74]. `IMRPhenomPv2` uses the `IMRPhenomD` [75, 76] waveform model as the underlying non-precessing waveform. In `IMRPhenomD` both spin components aligned with $\mathbf{L}$ are used to generate an approximate PN phasing and amplitude, with corrections provided by fits to numerical-relativity waveforms, that are parameterised by two different combinations of the two spin components. Although `IMRPhenomD` has been found to model well two-spin systems [77], its dominant spin dependence can be characterised well by the effective spin,

$$\chi_{\rm eff} = \frac{1}{M}\left(\frac{\mathbf{S_1}}{m_1} + \frac{\mathbf{S_2}}{m_2}\right)\cdot\hat{\mathbf{L}}, \tag{1.4}$$

which takes values between -1 (both maximal anti-aligned spins) and $+1$ (both maximal aligned spins) to describe the magnitude of the spin aligned with $\mathbf{L}$. For a given configuration `IMRPhenomPv2` a) modifies the final spin of the remnant object calculated by `IMRPhenomD` to account for the additional in-plane spin components, and b) applies a frequency dependent rotation to the modified `IMRPhenomD` waveform model to introduce the precession dynamics, which are modelled by frequency-domain PN expressions for the precession angles for an approximately equivalent single-spin system [64, 73], where the large black hole has spin $\chi_{\rm p}$.

There are several important features which are not incorporated in the `IMRPhenomPv2` waveform. These include two-spin effects [78, 79, 80], gravitational wave multipoles other than the leading 22 mode [81], significant precession during merger [82], and spin alignment due to spin-orbit resonances during inspiral [83, 84]. Some of these effects will have an impact upon the distributions of black hole spin orientations when the binaries enter the LIGO or Virgo sensitivity band while others can leave imprints on the waveform which may be observable, particularly close to the merger. As an example, by neglecting two-spin effects, $\chi_{\rm eff}$ ceases to be a constant of motion. Nonetheless, the `IMRPhenomPv2` has been used in the analysis of all

LIGO-Virgo observations during the first three observing runs [7, 32, 85, 38, 44], and it captures much of the dominant phenomenology of precessing-binary waveforms.

1.3 Summary of Thesis

This thesis is split into three parts. The first explores the general relativistic phenomenon of spin-induced orbital precession and introduces an elegant new representation. The second applies learnt techniques to observations made during the first, second and third gravitational wave observing runs. The third presents a new and innovate software package.

Chapter 2 presents a new formalisation for modelling a precessing gravitational wave as a series of beating non-precessing harmonics. In Chapter 3 I use the formalisation introduced in Chapter 2 to identify which binaries are most likely to emit gravitational waves with measurable precession while at the same time verifying the concepts in Chapter 2.

In Chapter 4 I estimate the probabilitity for observing a precessing gravitational wave and make a testable prediction for how many precessing binaries we expect to observe in the third gravitational wave observing run. Chapter 5 presents the inferred properties of potentially the first neutron star black hole candidate ever detected and investigates whether the large mass ratio allows for an informative precession measurement. In Chapter 6 I consider all of the gravitational wave candidates observed in the third gravitational wave observing run and identify if spin-induced orbital precession has been measured in the population. Chapter 7 then builds upon Chapter 6 and presents a study which identifies the mostly likely spin distribution of black holes detected through gravitational waves.

Finally, Chapter 8 presents a new and innovative software package that has been used to analyse all gravitational wave candidates to date and is central to making gravitational wave data open and easily reproducible.

Chapter 2

Two-harmonic approximation for gravitational waveforms from precessing binaries

This chapter is based upon the text of Fairhurst *et al.*[86] published as Phys. Rev. D, 102, 024055. This work was led by Stephen Fairhurst, Rhys Green and Charlie Hoy. My main contributions were to the code development, the analysis and figure formatting. All lead authors jointly contributed to the writing of the paper.

2.1 Introduction

When the spins of black holes in a binary system are mis-aligned with the binary's orbital angular momentum, both the spins and orbital angular momentum will precess [54, 58, 87, 88]. We therefore expect that most astrophysical binaries will undergo precession, but to date there has been no evidence of precession in gravitational-wave (GW) observations from the Advanced LIGO and Virgo detectors [38, 85]. This is not necessarily surprising, because precession often leaves only a weak imprint on the observable signal, particularly when the black holes are of comparable mass and the binary's orbit is face-on to the detector, which are the most likely configurations that have been observed so far. Despite this heuristic picture, there is no simple means to estimate the measurability of precession of a given binary configuration, and as such it is difficult to predict when precession effects will be conclusively observed in GW events.

Detailed parameter estimation techniques have been developed, which enable the reconstruction of the parameters of observed signals [89, 90, 91, 92, 93], in addition to approximate Fisher-matrix methods [61, 94]. In parallel, techniques have been developed that provide an intuitive understanding of the measurement accuracy of certain parameters (or parameter combinations) [95, 62, 96, 97, 98, 99]. These have typically involved either approximations (such as leading order, Fisher Matrix type calculations), restriction to a subset of system parameters (for example masses and spins; timing and sky location; binary orientation). Combined, these give an understanding of the accuracy of parameter estimation for non-precessing systems.

In parallel, there have been significant developments in understanding the implications of precession, starting with the early work in Refs. [54, 87, 88] which provided insights into the impact of precession on the gravitational waveform emitted during the inspiral of compact binaries. Subsequently, black hole binary waveforms which incorporate precession through merger have been developed [73, 78, 81, 100, 101]; large scale parameter estimation studies of precession have been performed to identify the regions of parameter space where precession will be observable [102, 103, 104, 63, 105, 106, 107, 108]; and new theoretical insights into the impact of precession on both detection and parameter estimation have been obtained [59, 109, 110]. Complementary to this, there have been several efforts to understand the impact of precession on searches [111, 59], and to implement searches for precessing signals [87, 112, 113, 114, 88]. This has led to an increasingly clear picture of the impact of precession: it is most significant for binaries with large mass ratios, where the in-plane spin components are large and for systems where the total angular momentum is mis-aligned with the line of sight.

At leading order, the gravitational waveform emitted by a precessing binary is composed of five harmonics, which are offset by multiples of the precession frequency [73, 110]. We show that these harmonics form a natural hierarchy with the ampli-

tude of the sub-leading harmonics suppressed by a factor that depends upon the opening angle (the angle between the orbital and total angular momenta). Using this approximation, and restricting to the two leading harmonics, we are able to obtain relatively simple expressions for the precession waveform. Each harmonic takes the form of a non-precessing-binary waveform (i.e., with monotonic amplitude and frequency evolution during the inspiral of non-eccentric systems), and the amplitude and phase modulations of the complete precessing-binary waveform arise as beating between the two harmonics.

The purpose of this chapter is to introduce this decomposition (Sec. 2.2), and the two-harmonic approximation (Sec. 2.3), and to identify its range of validity and accuracy (Sec. 2.4). We then discuss a proposed search for precessing binaries using the two-harmonic approximation (Sec. 2.5) and finally introduce the notion of a "precession SNR" that can be used to determine whether precession effects are observable in a given system (Sec. 2.6).

2.2 Harmonic decomposition of the waveform from a precessing binary

The gravitational waveform emitted by a precessing system, as observed at a gravitational wave detector, can be expressed approximately as [88, 59]

$$h(t) = -\left(\frac{d_o}{d_L}\right) A_o(t)\,\mathrm{Re}\Big[e^{2i\Phi_S(t)}\big(F_+(C_+ - iS_+) + F_\times(C_\times - iS_\times)\big)\Big]\,. \tag{2.1}$$

Here, $A_o(t)$ denotes the amplitude of the gravitational wave signal in a (time-varying) frame aligned with the orbital angular momentum of the binary, and depends upon the masses and spins of the binary.[1] Since the amplitude scales linearly with the luminosity distance, we have chosen to introduce a fiducial normalization $A_o(t)$ for a waveform at a distance d_o and explicitly extract the distance dependence.[2] Φ_S is the phase evolution in the frame aligned with the total angular momentum $\mathbf{J}$ of the binary. The phase evolution, Φ_S, is related to the orbital phase, $\phi_{\rm orb}$, as

$$\Phi_S(t) = \phi_{\rm orb}(t) - \epsilon(t) \tag{2.2}$$

where [115]

$$\dot{\epsilon}(t) := \dot{\alpha}(t)\cos\beta(t) \tag{2.3}$$

[1] Note that this expression, and what follows, does not include additional gravitational wave multipoles beyond the leading order quadrupole. We will investigate including these additional terms in a future work.

[2] Of course, the observed waveform is also affected by the redshifting of frequencies. For the calculation discussed here, we work in the detector frame and consider the *observed* masses, which are $(1+z)$ times the source frame masses.

and, as before, β is the opening angle and α gives the phase of the precession of $\mathbf{L}$ around $\mathbf{J}$ as shown in Fig. 1.3. F_+ and $F_\times$ give the detector response relative to the $\mathbf{J}$-aligned frame and $C_{+,\times}$, $S_{+,\times}$ encode the time-varying response to the gravitational wave due to the evolution of the binary's orbit relative to the detector. They depend upon the three angles introduced previously: the precession opening angle β and phase α and the angle between the total orbital angular momentum and the line of sight $\theta_{\rm JN}$. See Ref. [53] for the evolution equations of β and α accurate to 2PN order. In terms of these angles, we can express $C_{+,\times}$ and $S_{+,\times}$ as[3]

$$\begin{aligned}
C_+ &= -\left(\frac{1+\cos^2\theta_{\rm JN}}{2}\right)\left(\frac{1+\cos^2\beta}{2}\right)\cos 2\alpha \\
&\quad -\frac{\sin 2\theta_{\rm JN}}{2}\frac{\sin 2\beta}{2}\cos\alpha - \frac{3}{4}\sin^2\theta_{\rm JN}\sin^2\beta, \\
S_+ &= \left(\frac{1+\cos^2\theta_{\rm JN}}{2}\right)\cos\beta\sin 2\alpha + \frac{\sin 2\theta_{\rm JN}}{2}\sin\beta\sin\alpha, \\
C_\times &= -\cos\theta_{\rm JN}\left(\frac{1+\cos^2\beta}{2}\right)\sin 2\alpha - \sin\theta_{\rm JN}\frac{\sin 2\beta}{2}\sin\alpha, \\
S_\times &= -\cos\theta_{\rm JN}\cos\beta\cos 2\alpha - \sin\theta_{\rm JN}\sin\beta\cos\alpha.
\end{aligned} \tag{2.4}$$

The non-precessing expressions can be recovered in the limit of $\beta \to 0$ and $\alpha \to$ constant (which is then degenerate with the polarization of the system). When β is non-zero, the effect of precession is to modulate the detector response at frequencies $\Omega_{\rm p}$ and $2\Omega_{\rm p}$. To make the harmonic content of $C_{+,\times}$ and $S_{+,\times}$ more explicit, we first introduce the parameter,

$$b = \tan\left(\beta/2\right), \tag{2.5}$$

and write the response functions in terms of it. The terms involving β can be expressed as

$$\begin{aligned}
\frac{1+\cos^2\beta}{2} &= \frac{1+b^4}{(1+b^2)^2}, \\
\cos\beta &= \frac{1-b^4}{(1+b^2)^2}, \\
\frac{\sin 2\beta}{2} &= \frac{2b(1-b^2)}{(1+b^2)^2}, \\
\sin\beta &= \frac{2b(1+b^2)}{(1+b^2)^2}, \\
\sin^2\beta &= \frac{4b^2}{(1+b^2)^2}\,.
\end{aligned} \tag{2.6}$$

[3]We have re-written the C_+ term relative to what is normally given in the literature, e.g. [88, 59], to group terms with the same α dependence.

Substituting the trigonometric identities from Eq. (2.6) into the expressions for C_+ and S_+ in Eq. (2.4) we obtain,

$$\begin{aligned}\left(\frac{d_o}{d_L}\right)(C_+ - iS_+) &= -e^{2i\alpha}\sum_{k=0}^{4}\mathcal{A}_k^+\left[\frac{b^k e^{-ik\alpha}}{(1+b^2)^2}\right],\\ \left(\frac{d_o}{d_L}\right)(C_\times - iS_\times) &= ie^{2i\alpha}\sum_{k=0}^{4}\mathcal{A}_k^\times\left[\frac{b^k e^{-ik\alpha}}{(1+b^2)^2}\right],\end{aligned} \tag{2.7}$$

where we have introduced $\mathcal{A}_k^+$ and $\mathcal{A}_k^\times$ as

$$\begin{aligned}\mathcal{A}_0^+ = \mathcal{A}_4^+ &= \frac{d_o}{d_L}\left(\frac{1+\cos^2\theta_{\rm JN}}{2}\right),\\ \mathcal{A}_0^\times = -\mathcal{A}_4^\times &= \frac{d_o}{d_L}\cos\theta_{\rm JN},\\ \mathcal{A}_1^+ = -\mathcal{A}_3^+ &= 2\frac{d_o}{d_L}\sin\theta_{\rm JN}\cos\theta_{\rm JN},\\ \mathcal{A}_1^\times = \mathcal{A}_3^\times &= 2\frac{d_o}{d_L}\sin\theta_{\rm JN},\\ \mathcal{A}_2^+ &= 3\frac{d_o}{d_L}\sin^2\theta_{\rm JN},\\ \mathcal{A}_2^\times &= 0\,.\end{aligned} \tag{2.8}$$

In the approximation where the direction of total angular momentum is constant, the $\mathcal{A}_k^{+,\times}$ are time independent amplitudes, and the time dependence of the amplitude functions is captured as a power series in the parameter $b = \tan(\beta/2)$.

Finally, we can use the harmonic decomposition in Eq. (2.7) to obtain a decomposition of the waveform, Eq. (2.1),

$$h(t) = {\rm Re}\left[\left(\frac{A_o(t)e^{2i(\Phi_S+\alpha)}}{(1+b^2)^2}\right)\sum_{k=0}^{4}(be^{-i\alpha})^k(F_+\mathcal{A}_k^+ - iF_\times\mathcal{A}_k^\times)\right]. \tag{2.9}$$

This allows us to clearly identify the impact of precession on the waveform. First, precession leads to an additional phase evolution at frequency $2\Omega_{\rm p}$ and a decrease in the amplitude by a factor $(1+b^2)^2$. The precessing waveform contains five harmonics that form a power series in b, whose amplitude depends upon the detector response, distance and viewing angle of the binary. The frequency of each harmonic is offset from the next by the precession frequency $\Omega_{\rm p}$. Similar results have been obtained previously, by manipulating the spin-weighted spherical harmonic decomposition of the waveform, e.g. [109, 110]. However, it was not previously observed that the relative amplitudes of the harmonics were related in a straightforward manner.

As a final step, we would like to explicitly extract three more time-independent angles that characterize the waveform, namely the polarization angle ψ, the initial

phase ϕ_o and the initial polarization phase α_o.[4]

The unknown polarization ψ is currently folded into the detector response functions $F_{+,\times}$. It is more useful to extract ψ and then consider the detector response to be a *known* quantity dependent upon only the details of the detector and the direction to the source. Thus, we write the detector response as,

$$\begin{aligned} F_+ &= w_+ \cos 2\psi + w_\times \sin 2\psi, \\ F_\times &= -w_+ \sin 2\psi + w_\times \cos 2\psi, \end{aligned} \tag{2.10}$$

where w_+ and $w_\times$ are the detector response functions in a fixed frame — for a single detector it is natural to choose $w_\times = 0$ and for a network to work in the dominant polarization for which w_+ is maximized [116]. The unknown polarization of the source relative to this preferred frame is denoted ψ.

To isolate the initial orbital and precession phases, we explicitly extract them from the binary's phase evolution by introducing,

$$\begin{aligned} \Phi(t) &:= \Phi_S(t) - \phi_o + \alpha(t) - \alpha_o \\ &= \phi_{\rm orb}(t) - \phi_o + \int_{\alpha_o}^{\alpha(t)} \frac{2b^2}{1+b^2}\, d\alpha\,. \end{aligned} \tag{2.11}$$

Thus $\Phi(t)$ vanishes at $t = 0$ and evolves as the sum of the orbital phase and an additional, precession dependent, contribution.

We then substitute the expressions for $F_{+,\times}$, Eq. (2.10), and Φ, Eq. (2.11), into the expression for $h(t)$ given in Eq. (2.9), and isolate the time-varying terms from the constant, orientation dependent angles. The waveform can be written as the sum of five precessing harmonics, the amplitudes of which are constants that depend upon the binary's sky location, distance and orientation:

$$h = \sum_{k=0}^{4} w_+ (h_0^k \mathcal{A}_k^1 + h_{\frac{\pi}{2}}^k \mathcal{A}_k^3) + w_\times (h_0^k \mathcal{A}_k^2 + h_{\frac{\pi}{2}}^k \mathcal{A}_k^4), \tag{2.12}$$

where $h_{0,\frac{\pi}{2}}^k$ are the waveform harmonics and $\mathcal{A}_k^\mu$ are constants. The waveform harmonics are

$$\begin{aligned} h_0^k(t) &= {\rm Re}\left[A_o(t) e^{2i\Phi} \left(\frac{b^k e^{-ik(\alpha-\alpha_o)}}{(1+b^2)^2} \right) \right], \\ h_{\frac{\pi}{2}}^k(t) &= {\rm Im}\left[A_o(t) e^{2i\Phi} \left(\frac{b^k e^{-ik(\alpha-\alpha_o)}}{(1+b^2)^2} \right) \right]. \end{aligned} \tag{2.13}$$

The harmonics form a simple power series in $be^{-i\alpha}$, so the amplitude of each successive harmonic is reduced by a factor of b, and the frequency is reduced by $\Omega_{\rm p}$.

[4] The initial polarization phase α_o is sometimes denoted in the literature as ϕ_{JL}.

The amplitudes for the harmonics are given by

$$\begin{aligned}
\mathcal{A}^1_k &= \mathcal{A}^+_k \cos\phi_k \cos 2\psi - \mathcal{A}^\times_k \sin\phi_k \sin 2\psi, \\
\mathcal{A}^2_k &= \mathcal{A}^+_k \cos\phi_k \sin 2\psi + \mathcal{A}_\times \sin\phi_k \cos 2\psi, \\
\mathcal{A}^3_k &= -\mathcal{A}^+_k \sin\phi_k \cos 2\psi - \mathcal{A}^\times_k \cos\phi_k \sin 2\psi, \\
\mathcal{A}^4_k &= -\mathcal{A}^+_k \sin\phi_k \sin 2\psi + \mathcal{A}^\times_k \cos\phi_k \cos 2\psi,
\end{aligned} \tag{2.14}$$

where the $\mathcal{A}^{+,\times}$ were introduced in Eq. (2.8), ψ is the polarization and the phase angle for each harmonic is,

$$\phi_k = 2\phi_o + (2-k)\alpha_o\,. \tag{2.15}$$

These amplitudes form a generalization of the $\mathcal{F}$-statistic decomposition of the non-precessing binary waveform (see e.g. [116]). In the limit that $b \to 0$, the precessing decomposition reduces to the standard expression for the non-precessing waveform as the amplitude for all harmonics other than $k = 0$ vanish.

The precessing waveform can equally well be written in the frequency domain by performing a Fourier transform of the time-domain expressions given above [117]. In this case, Eq. (2.12) is unchanged, as are the constant amplitude terms in Eq. (2.14). The frequency dependent harmonics are simply the Fourier transform of the time-domain modes given in Eq. (2.13), and naturally satisfy $h^k_{\frac{\pi}{2}} = i h^k_0$.

The expansion above is most natural when $b < 1$, which corresponds to opening angles of $\beta < 90°$. In cases where the opening angle is greater than 90° it is natural to re-express the waveform in terms of $c = b^{-1} = \cot(\beta/2)$ in which case the waveform can be expressed as a power series in c. We will not discuss the large opening angle calculation further in this chapter, but note that many of the arguments presented below would extend in a straightforward manner to this case.

2.2.a Obtaining the harmonics

Here, we give an explicit prescription to obtain the five harmonics for the waveform, introduced in Eq. (2.12). To do so, we generate waveforms for orientations that contain only a subset of the harmonics, and combine them to isolate a single harmonic. For simplicity, we restrict attention to the + polarization by fixing $w_+ = 1$, $w_\times = 0$ and consider a binary at a distance $d_L = d_o$.

Harmonics $k = 0$ and $k = 4$. When the viewing angle of the signal is aligned with the total angular momentum, $\theta_{\rm JN} = 0$, the observed waveform contains only the zeroth and fourth harmonics as $\mathcal{A}^{+,\times}_{1,2,3}$ vanish for $\theta_{\rm JN} = 0$. We also fix $\alpha_o = 0$,

to obtain,

$$\begin{aligned} h_{\phi_o=0;\psi=0} &= h_0^0 + h_0^4, \\ h_{\phi_o=\frac{\pi}{4};\psi=\frac{\pi}{4}} &= -h_0^0 + h_0^4. \end{aligned} \tag{2.16}$$

From these, we can extract the $k = 0$ and 4 harmonics,

$$\begin{aligned} h_0^0 &= \tfrac{1}{2}\left(h_{\phi_o=0,\psi=0} - h_{\phi_o=\frac{\pi}{4},\psi=\frac{\pi}{4}}\right), \\ h_0^4 &= \tfrac{1}{2}\left(h_{\phi_o=0,\psi=0} + h_{\phi_o=\frac{\pi}{4},\psi=\frac{\pi}{4}}\right). \end{aligned} \tag{2.17}$$

The $\frac{\pi}{2}$ phases of the harmonics can be obtained in an identical way.

Harmonics $k = 1$ and $k = 3$. When the signal is edge on, the $\times$ polarization contains only the first and third harmonics. Then, fixing $\theta_{\rm JN} = \frac{\pi}{2}$ and $\psi = \frac{\pi}{4}$, we have,

$$\begin{aligned} h_{\alpha_o=0;\phi_o=\frac{\pi}{4}} &= -2\left(h_0^1 + h_0^3\right), \\ h_{\alpha_o=\frac{\pi}{2};\phi_o=0} &= -2\left(h_0^1 - h_0^3\right), \end{aligned} \tag{2.18}$$

so that,

$$\begin{aligned} h_0^1 &= -\tfrac{1}{4}\left(h_{\alpha_o=0;\phi_o=\frac{\pi}{4}} + h_{\alpha_o=\frac{\pi}{2},\phi_o=0}\right), \\ h_0^3 &= -\tfrac{1}{4}\left(h_{\alpha_o=0;\phi_o=\frac{\pi}{4}} - h_{\alpha_o=\frac{\pi}{2},\phi_o=0}\right). \end{aligned} \tag{2.19}$$

Harmonic $k = 2$. Finally, from the + polarization of the edge-on waveform, we can extract the second harmonic — in principle we could also get $k = 0$ and $k = 4$, but we have already described a method to obtain them. Fixing $\theta_{\rm JN} = \frac{\pi}{2}$ and $\psi = 0$ we have,

$$\begin{aligned} h_{\alpha_o=0,\phi_o=0} &= \tfrac{1}{2}h_0^0 + 3h_0^2 + \tfrac{1}{2}h_0^4, \\ h_{\alpha_o=\frac{\pi}{2},\phi_o=0} &= -\tfrac{1}{2}h_0^0 + 3h_0^2 - \tfrac{1}{2}h_0^4, \end{aligned} \tag{2.20}$$

so that,

$$h_0^2 = \tfrac{1}{6}\left(h_{\alpha_o=0,\phi_o=0} + h_{\alpha_o=\frac{\pi}{2},\phi_o=0}\right). \tag{2.21}$$

Throughout the remainder of this work, we use the `IMRPhenomPv2` signal model to generate h. Since this model only incorporates leading order spin-orbit couplings in the precession equations (see Sec. 1.2.a), the evolution of α and β are accurate to 1.5PN order. However, since our harmonic decomposition is in no way tied to the particular waveform model used, we could use other waveform models for precessing binaries which, for example, incorporate 2PN spin-spin effects and precession during

merger. We expect that the broad features of many of the results presented in the remainder of the chapter are relatively unaffected by the specific waveform choice, but the details for any specific signal could change.

2.2.b Precession with varying orientation

The observable effect of precession will vary significantly with the binary orientation, as has been discussed in many previous works, for example [54, 59]. Interestingly, both the amplitude and frequency of the observed precession depends upon the viewing angle. The harmonic decomposition derived above provides a straightforward way to understand this effect. The observed amplitude and phase modulations can be understood as the beating of the different harmonics against each other, with the amplitude of the composite waveform being maximum when the harmonics are in phase and minimum when they are out of phase.

In Fig. 2.1, we show the waveform for four different orientations: a) along $\mathbf{J}$, b) $\times$ polarization at 45° to $\mathbf{J}$, c/d) $+/\times$ polarization orthogonal to $\mathbf{J}$. In all cases, we show the last two seconds of the waveform (from around 25 Hz) for a $40M_{\odot}$ binary, with $q = 6$, and in-plane spin on the larger black hole of $\chi_{\rm p} = 0.6$. This configuration gives an opening angle of $\beta \approx 45°$ (and $b \approx 0.4$) which leads to significant precession effects in the waveform.

When viewed along $\mathbf{J}$, there is minimal precession as only the $k = 0$ and 4 harmonics are present in the system and the $k = 4$ harmonic is down-weighted by a factor of $b^4 \approx 0.03$ relative to the leading harmonic. Furthermore, the modulation comes from the beating of the $k = 0$ and $k = 4$ harmonics and occurs at four times the precession frequency. When the line of sight is orthogonal to the total angular momentum, the $k = 0, 2, 4$ harmonics are present in the $+$ polarized waveform and $k = 1, 3$ in the $\times$ polarization. The $k = 0$ and 2 harmonics have close to equal amplitude (although $k = 2$ is down-weighted by $b^2 \approx 0.17$, the amplitude as given in Eq. (2.8) is maximal). Consequently the observed waveform has maximal amplitude and phase modulation due to precession. For the $\times$ polarized signal, it is the $k = 1, 3$ harmonics that contribute, with $k = 3$ a factor of $b^2 \approx 0.17$ smaller than $k = 1$. Consequently, precession effects are less significant. In both cases, precession occurs at twice the precession frequency as it is from the beating of $k = 0$ and $k = 2$ ($+$ polarization) or $k = 1$ and $k = 3$ ($\times$ polarization). For the $\times$ polarized signal with $\theta_{\rm JN} = 45°$, the $k = 0, 1, 3, 4$ harmonics are present, with $k = 0, 1$ dominating and having approximately equal amplitude. For this signal, the binary precesses from a face-on orientation, $\iota = 0$ to edge-on, $\iota = 90°$, and the waveform amplitude oscillates from the maximum to zero. Here, modulations occur at the precession frequency.

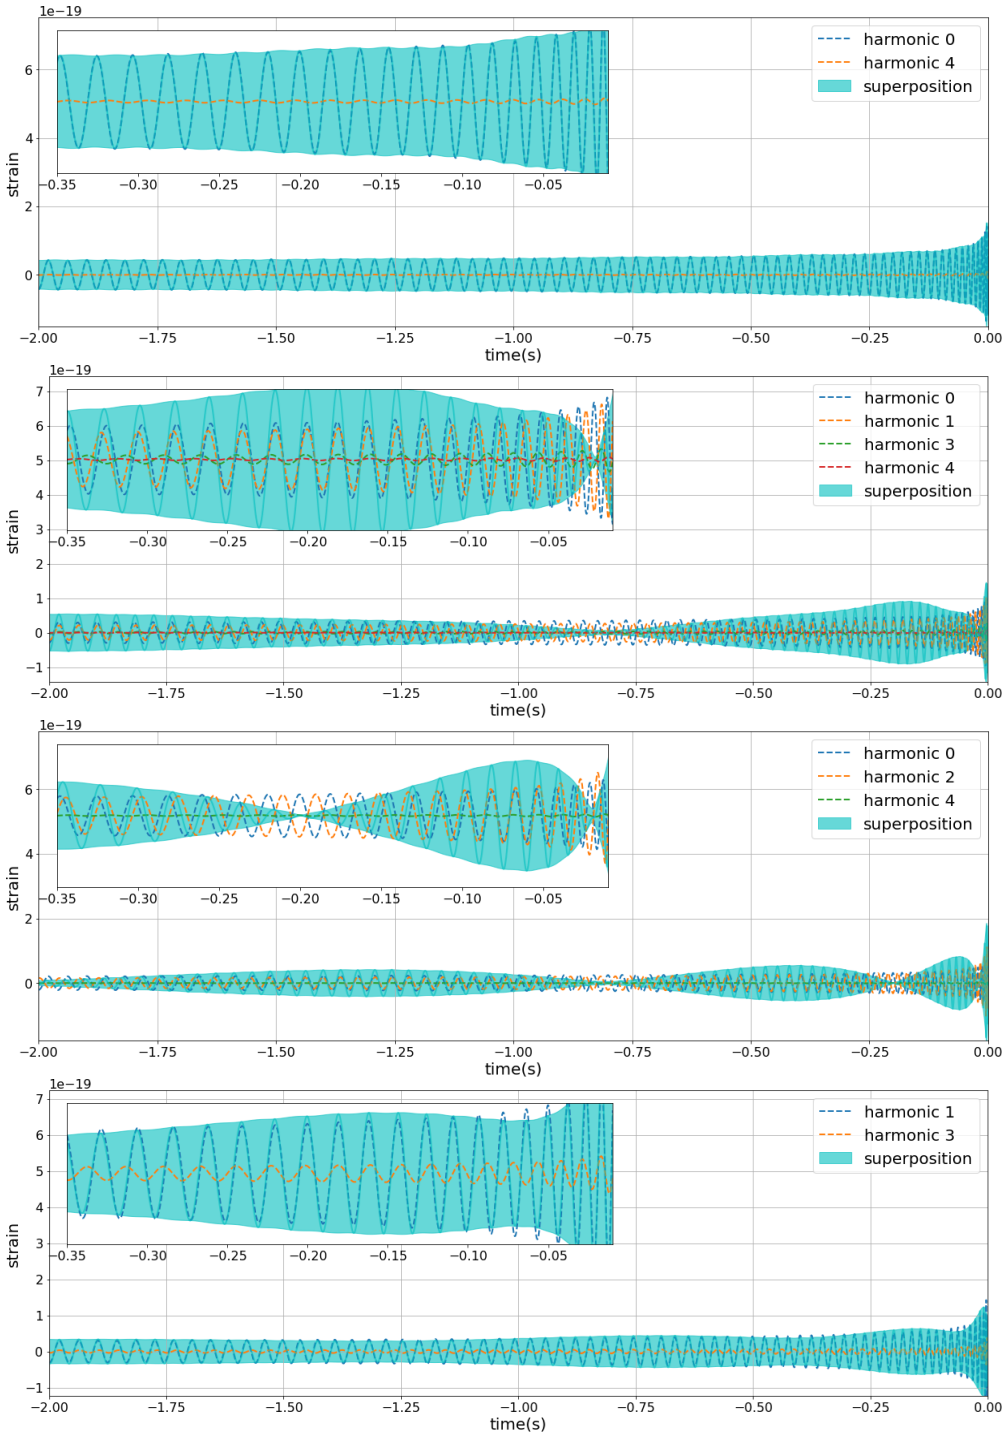

Figure 2.1: The observed waveform from a $40M_\odot$ binary with mass ratio $q = 6$, $\chi_{\rm eff} = 0$ and $\chi_{\rm p} = 0.6$. The waveform is shown for four different binary orientations: $\theta_{JN} = 0$ (first row); $\theta_{JN} = 45°$, $\times$ polarization (second row); $\theta_{JN} = 90°$, $+$ polarization (third row); $\theta_{JN} = 90°$, $\times$ polarization (fourth row). For each waveform, the harmonics that contribute to the signal, their sum and the envelope of the full precessing waveform are shown. The insets show a zoom of a portion of the waveform to more clearly demonstrate that precession arises as a beating between the different harmonics.

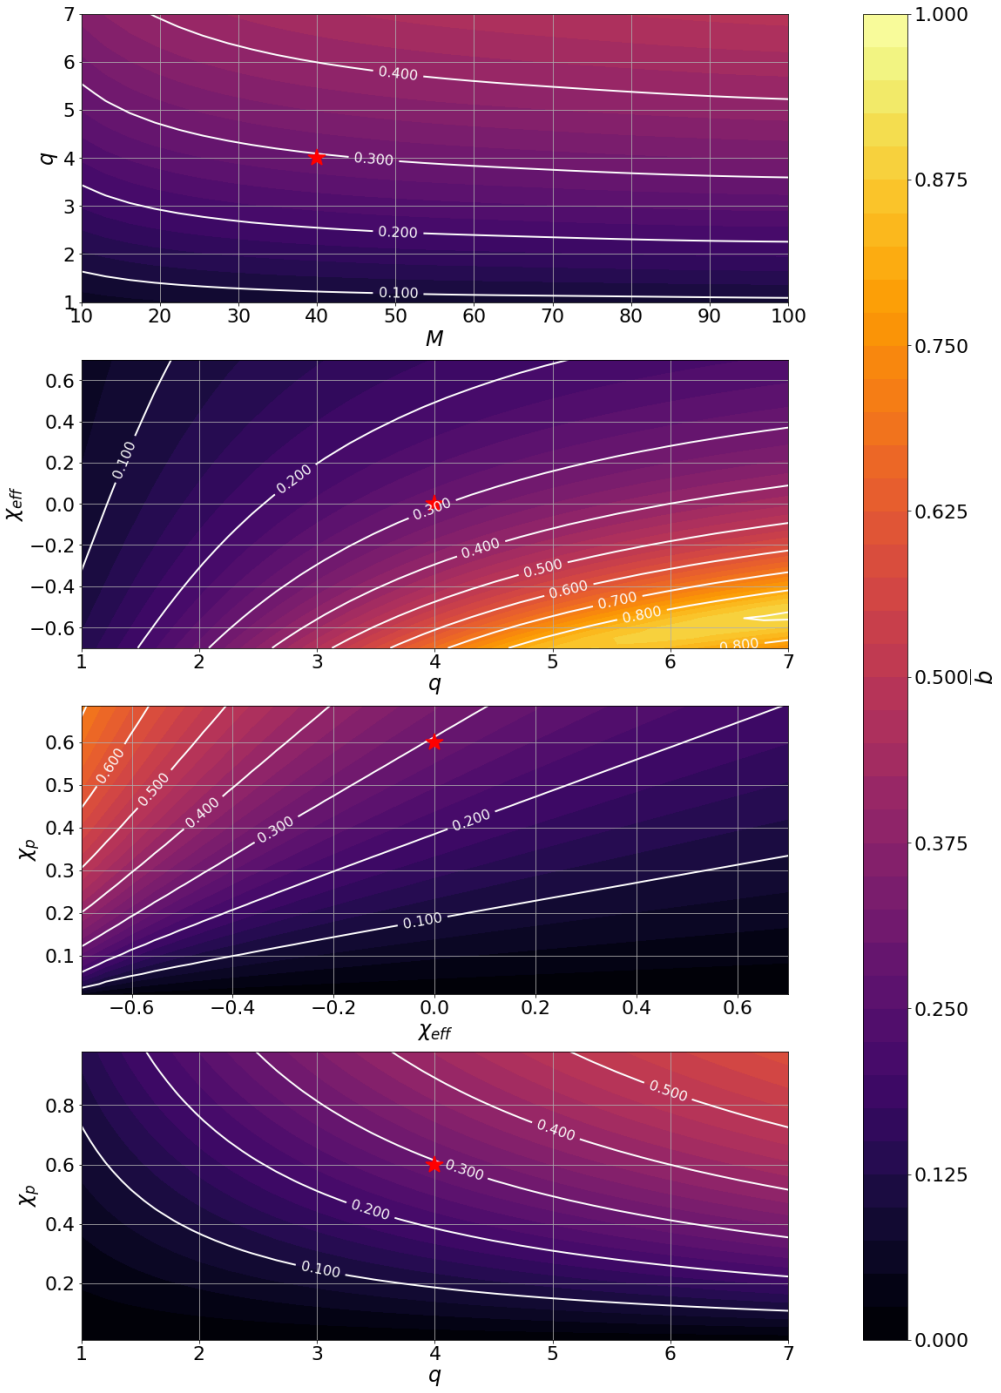

Figure 2.2: The value of $\bar{b}$ across the parameter space of total mass, mass ratio, $\chi_{\rm eff}$ and $\chi_{\rm p}$. In each figure, two of the parameters are varied while the other two are fixed to their fiducial values of $M = 40M_{\odot}$, $q = 4$, $\chi_{\rm eff} = 0$, $\chi_{\rm p} = 0.6$ (this point is marked with a $\star$ in all the plots). The total mass has a limited impact on the value of $\bar{b}$, for masses over $M \approx 40M_{\odot}$; below this the $\bar{b}$ increases with mass, as the later parts of the merger are brought into the most sensitive band of the detector. The value of $\bar{b}$ is seen to increase as the mass ratio or precessing spin $\chi_{\rm p}$ are increased and decrease as the aligned component of the spin $\chi_{\rm eff}$ increases. Thus, the value of b is largest for a binary with unequal masses, a large spin on the more massive component which has significant components both in the plane of the orbit and anti-aligned with the orbital angular momentum.

2.2.c Importance of precession over parameter space

From the intuitive discussion of precession presented in [54, 88, 59] and summarized in Section 1.2, it is straightforward to identify regions of parameter space where precession is most likely to have a significant impact upon the binary dynamics and, consequently, the observed waveform. Specifically, we expect that higher mass ratios, larger in-plane spins and negative aligned spin components will all lead to a larger opening angle and more significant precession [59]. Here we briefly revisit this discussion, framing our results in terms of the parameter b introduced earlier. Explicitly, we introduce the waveform-averaged value of b as,

$$\overline{b} := \frac{|h^1|}{|h^0|} = \sqrt{\frac{\int df \frac{|h_1|^2}{S_n(f)}}{\int df \frac{|h_0|^2}{S_n(f)}}}\,, \tag{2.22}$$

where $h^{0,1}$ are the harmonics of the waveform introduced in Eq. (2.13) and $S_n(f)$ is the noise power spectrum of the detector. For this work, we choose $S_n(f)$ to be the design-sensitivity Advanced LIGO noise curve [38] and evaluate the integral over the frequency range $f \in [20, 1024]\,\mathrm{Hz}$ [5]. For binaries where the opening angle β is approximately constant, $\overline{b} \approx \tan(\overline{\beta}/2)$.

Fig. 2.2 shows the value of $\overline{b}$ on several two-dimensional slices through the four dimensional parameter space of total mass M, mass ratio q, effective spin $\chi_{\rm eff}$ and precessing spin $\chi_{\rm p}$. Keeping other quantities fixed, the value of $\overline{b}$ increases with total mass. For higher masses, the late inspiral and merger occur in the sensitive band of the detectors and, close to merger, the opening angle increases as orbital angular momentum is radiated. For masses above $40M_\odot$ the mass dependence of $\overline{b}$ is small, with only a 10% decrease from $40M_\odot$ to $100M_\odot$. Thus, for the other figures, we fix $M = 40M_\odot$ and investigate the dependence of $\overline{b}$ on $q, \chi_{\rm eff}$ and $\chi_{\rm p}$. The dependence of $\overline{b}$ follows directly from Eq. (1.1). The opening angle will increase with mass ratio, as the orbital angular momentum decreases. The opening angle, and also $\overline{b}$, increase with $\chi_{\rm p}$. It follows directly from the definition that $\tan\beta$ scales linearly with $\chi_{\rm p}$, and hence approximately linearly for $b = \tan(\beta/2)$. Finally, the opening angle decreases as the effective spin $\chi_{\rm eff}$ increases, so that the largest value of $\overline{b}$ is obtained with significant spin anti-aligned with $\mathbf{J}$.

Over much of the parameter space we have explored, $\overline{b} \lesssim 0.3$. This includes binaries with mass ratio up to 4:1, with precessing spin $\chi_{\rm p} \lesssim 0.6$, and zero or positive aligned spin, $\chi_{\rm eff} \geq 0$. Only a small part of parameter space has $\overline{b} > 0.4$, the value used in generating the waveforms in Figure 2.1, and $b > 0.5$ is only achieved with at least two of: a) close to maximal $\chi_{\rm p}$, b) high mass ratio, $q \gtrsim 5$ or c) significant spin anti-aligned with the orbital angular momentum $\chi_{\rm eff} \lesssim -0.4$.

[5] Using a realistic noise curve similar to the observed curves during O1 and O2 would change the reported values slightly, as these noise curves are less sensitive than design, particularly at low

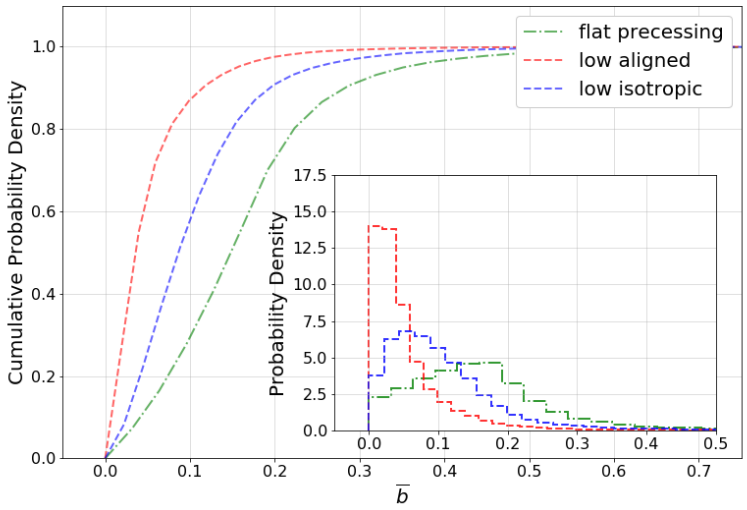

Figure 2.3: The distribution of $\overline{b}$ for a 3 different populations of binary black holes. Each population assumes either a low-isotropic, low-aligned or a flat precessing spin distribution. A power-law distribution in masses is assumed in all cases (see text for details).

Next, we consider the importance of precession for an astrophysically motivated population. In Fig. 2.3, we show the distribution of $\overline{b}$ for three distributions of black hole masses and spins. For each population, we generate 100,000 binaries uniformly in co-moving distance, with masses drawn from a power law distribution — $p(m_1) \propto m_1^{-\alpha}$, with $\alpha = 2.35$ — and different spin distributions, which are the same as those used in Refs. [118, 119, 120]. We consider populations where the spins are preferentially low and aligned with the binary orbit; low and isotropically aligned or drawn from a flat distribution and preferentially leading to precession. A low spin distribution is a triangular distribution peaked at zero spin and dropping to zero at maximal spin while a flat distribution is a uniform between zero and maximal spin. The *aligned* distribution is strongly peaked towards aligned spins, while the *isotropic* distribution assumes a uniform distribution of spin orientations over the sphere. The *precessing* distribution is strongly peaked towards spins orthogonal to the orbital angular momentum, i.e., with significant orbital precession [121, 122]. To account for observational biases, we keep only those signals that would be observable above a fixed threshold in a gravitational wave detector. We find that even for the most extreme precessing population considered, the mean value of $\overline{b}$ is 0.15 with over 90% of binaries having $\overline{b} < 0.3$. This result is obviously sensitive to the assumptions on the mass and spin distribution. In Chapter 4 we investigate a larger number of spin distributions, including ones which allow for large spin magnitudes, and we find that the peak of the $\overline{b}$ distribution is below 0.2 and that over 90% of binaries have

frequencies. The qualitative patterns seen in the figure would remain the same however

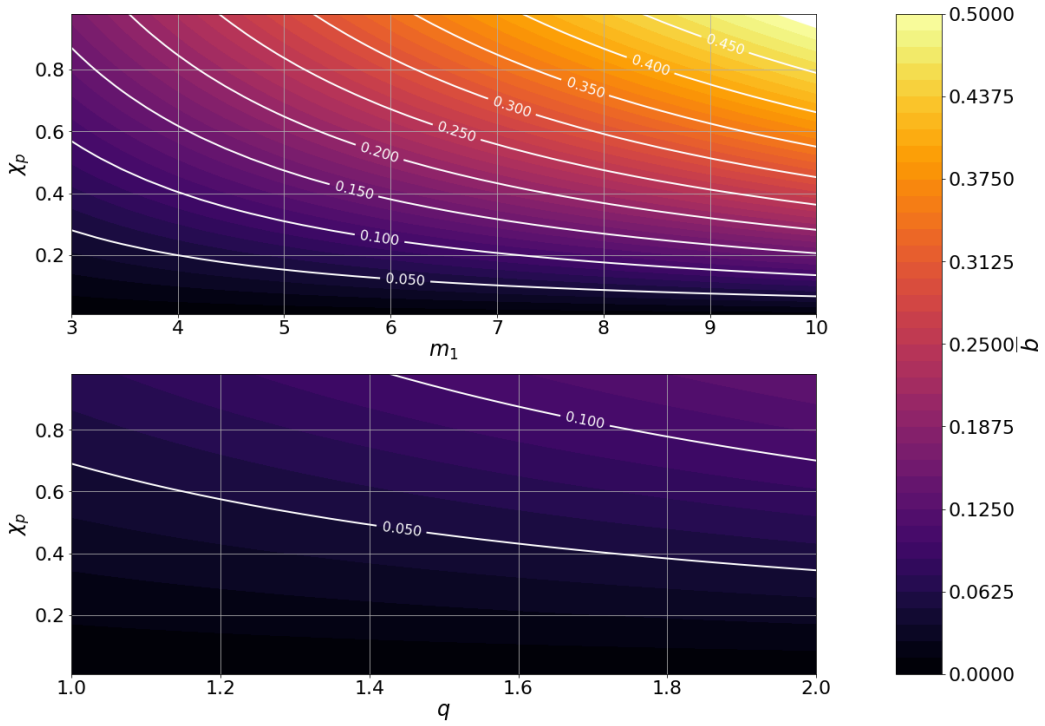

Figure 2.4: The value of $\bar{b}$ across the binary neutron star and neutron-star–black-hole space. The top figure shows the variation of $\bar{b}$ for an NSBH system with a $1.4M_\odot$ neutron star, $\chi_{\rm eff} = 0$ and varying black hole mass and $\chi_{\rm p}$. The bottom figure shows the variation of $\bar{b}$ against mass ratio and $\chi_{\rm p}$ for a binary neutron star system of total mass $2.7M_\odot$ and $\chi_{\rm eff} = 0$.

$\bar{b} < 0.4$ in all cases.

Fig. 2.4 shows $\bar{b}$ for a range of neutron star–neutron star and neutron star–black hole binaries. For neutron star–black hole binaries, the picture is similar to that for black hole binaries, with large values of $\bar{b}$ observed for high mass ratios and large $\chi_{\rm p}$. However, as an earlier part of the waveform is in the detector's sensitive band, the impact of precession is less observable at fixed mass ratio than for higher mass black hole binaries. For neutron star binaries, the value of $\bar{b}$ remains below 0.15 across the parameter space, and is less than 0.05 for reasonable neutron star spins, $\chi \lesssim 0.4$.

2.3 The two-harmonic approximation

The precessing waveform can be expressed as the sum of five harmonics whose amplitudes form a power series in $b = \tan(\beta/2)$. Furthermore, over the majority of the space of binary mergers, the value of b is less than 0.3. In addition, for $\bar{b} \leq 0.4$ the dominant harmonic — the one containing the most power — must be either $k = 0$ or 1. Thus, for the vast majority of binary mergers, we expect that these two harmonics will be the most significant.

This motivates us to introduce the *two-harmonic approximation*, in which we

generate a waveform containing only the $k = 0$ and $k = 1$ harmonics, i.e.,

$$h = \sum_{k=0,1} w_+(h_0^k \mathcal{A}_k^1 + h_{\frac{\pi}{2}}^k \mathcal{A}_k^3) + w_\times(h_0^k \mathcal{A}_k^2 + h_{\frac{\pi}{2}}^k \mathcal{A}_k^4)\,. \tag{2.23}$$

We note that Eq. 2.23 assumes a constant opening angle since we surmise that $\overline{b^N} = \bar{b}^N$ when dropping higher order terms in the harmonic decomposition. Although the opening angle must continually increase throughout the inspiral [54], it typically varies very little across the sensitive frequency range of a GW detector [59]. This means that it is often possible to make the approximation that β (and hence b) remains approximately constant. We refer the reader to Ref. [59] for a detailed discussion of the constant opening angle approximation. We note that all previous equations are independent of the constant opening angle approximation.

The expression for the two-harmonic waveform can be simplified by restricting to the single detector case (i.e., setting $w_+ = 1$ and $w_\times = 0$), explicitly working with the waveform in the frequency domain, for which $h_{\frac{\pi}{2}}^k(f) = ih_0^k(f)$, and dropping the subscript 0 on the zero-phase waveform, so that $h^k(f) := h_0^k(f)$. The two harmonics of interest are,

$$h^0(f) = A_o(f)e^{2i\Phi(f)}\left(\frac{1}{(1+b(f)^2)^2}\right), \tag{2.24}$$

$$h^1(f) = A_o(f)e^{2i\Phi(f)}\left(\frac{b(f)e^{-i(\alpha(f)-\alpha_o)}}{(1+b(f)^2)^2}\right), \tag{2.25}$$

and the two-harmonic waveform then becomes,

$$h_{\rm 2harm} = \mathcal{A}_0 h^0 + \mathcal{A}_1 h^1\,, \tag{2.26}$$

where,

$$\begin{aligned}\mathcal{A}_0 =& \frac{d_0}{d_L}\left(\frac{1+\cos^2\theta_{\rm JN}}{2}\cos 2\psi - i\cos\theta_{\rm JN}\sin 2\psi\right)\times e^{-i(2\phi_o+2\alpha_o)},\\ \mathcal{A}_1 =& \frac{d_0}{d_L}\left(\sin 2\theta_{\rm JN}\cos 2\psi - 2i\sin\theta_{\rm JN}\sin 2\psi\right)\times e^{-i(2\phi_o+\alpha_o)}.\end{aligned} \tag{2.27}$$

Thus, the two-harmonic waveform is composed of two components that have frequencies offset by $\Omega_{\rm p}$, and any observed amplitude and phase modulation of the waveform is caused by the beating of one waveform against the other. The relative amplitude and phase of the two harmonics is encoded by

$$\begin{aligned}\zeta \;&:=\; \frac{\bar{b}\mathcal{A}_1}{\mathcal{A}_0} \\ &=\; \bar{b}e^{i\alpha_o}\left(\frac{\sin 2\theta_{\rm JN}\cos 2\psi - 2i\sin\theta_{\rm JN}\sin 2\psi}{\frac{1}{2}(1+\cos^2\theta_{\rm JN})\cos 2\psi - i\cos\theta_{\rm JN}\sin 2\psi}\right).\end{aligned} \tag{2.28}$$

The value of ζ depends upon the viewing angle, encoded in $\theta_{\rm JN}$ and ψ, and the initial precession phase α_o. It is not difficult to show that ζ can take any value as the parameters $\theta_{\rm JN}$, ψ, α_o are varied. For example, at $\theta_{\rm JN} = 0$, $\mathcal{A}_1$ vanishes and so does ζ, while at $\theta_{\rm JN} = \pi/2$ and $\psi = \pi/4$, $\mathcal{A}_0$ vanishes and $\zeta \to \infty$. Since the initial precession phase α_o is a free parameter, the phase of ζ also can take any value. The overall amplitude and phase of the signal also depends upon the distance and coalescence phase so that any values of the amplitude and phase of the signal in the two harmonics are consistent with a signal.

2.4 Validity of the two-harmonic waveform

To investigate the validity of the two-harmonic approximation, we compare the approximate waveform with the full, five-harmonic, precessing waveform across the parameter space. The error will be of order b^2, which is small over much of the parameter space, and for the majority of orientations.

Fig. 2.5 shows the overlap between the full waveform and a subset of the harmonics for a binary with $M = 40M_\odot$, $q = 4$ and $\chi_{\rm eff} = 0$, while varying the orientation and value of $\chi_{\rm p}$. In each case, we calculate,

$$O(h, h') = \frac{\max_{\phi_o}(h|h')}{|h||h'|}, \tag{2.29}$$

where,

$$(a|b) = 4\,\mathrm{Re}\int_{f_o}^{f_{\rm max}} \frac{a^\star(f)b(f)}{S(f)}df, \tag{2.30}$$

and $S(f)$ is the power spectral density of the detector data. Thus the overlap is maximized over the phase, but not over time or any of the mass and spin parameters. An overlap of close to unity shows that the two waveforms are very similar, while a lower value of overlap implies significant deviations between the waveforms. As a rule of thumb, an overlap $O(h, h') \lesssim 1 - 3/\rho^2$ will be observable at a signal to noise ratio ρ [123, 124, 62].

We calculate the overlap of the full waveform, h, against

1. the leading order waveform in the precession expansion, h^0;
2. the dominant harmonic, i.e. the harmonic of h^0 and h^1 which contains the largest fraction of the power in the full waveform;
3. the two-harmonic waveform with the appropriate values of $\mathcal{A}_0$ and $\mathcal{A}_1$.

For the + polarized waveform (left column), the $k = 0$ harmonic is dominant for all values of $\theta_{\rm JN}$ and $\chi_{\rm p}$, so that the observed overlap with the full waveform is above 0.8 across the parameter space. For $\theta_{\rm JN} \approx 0$ or small values of $\chi_{\rm p}$, the other harmonics make a minimal contribution and the overlap is close to unity. For larger

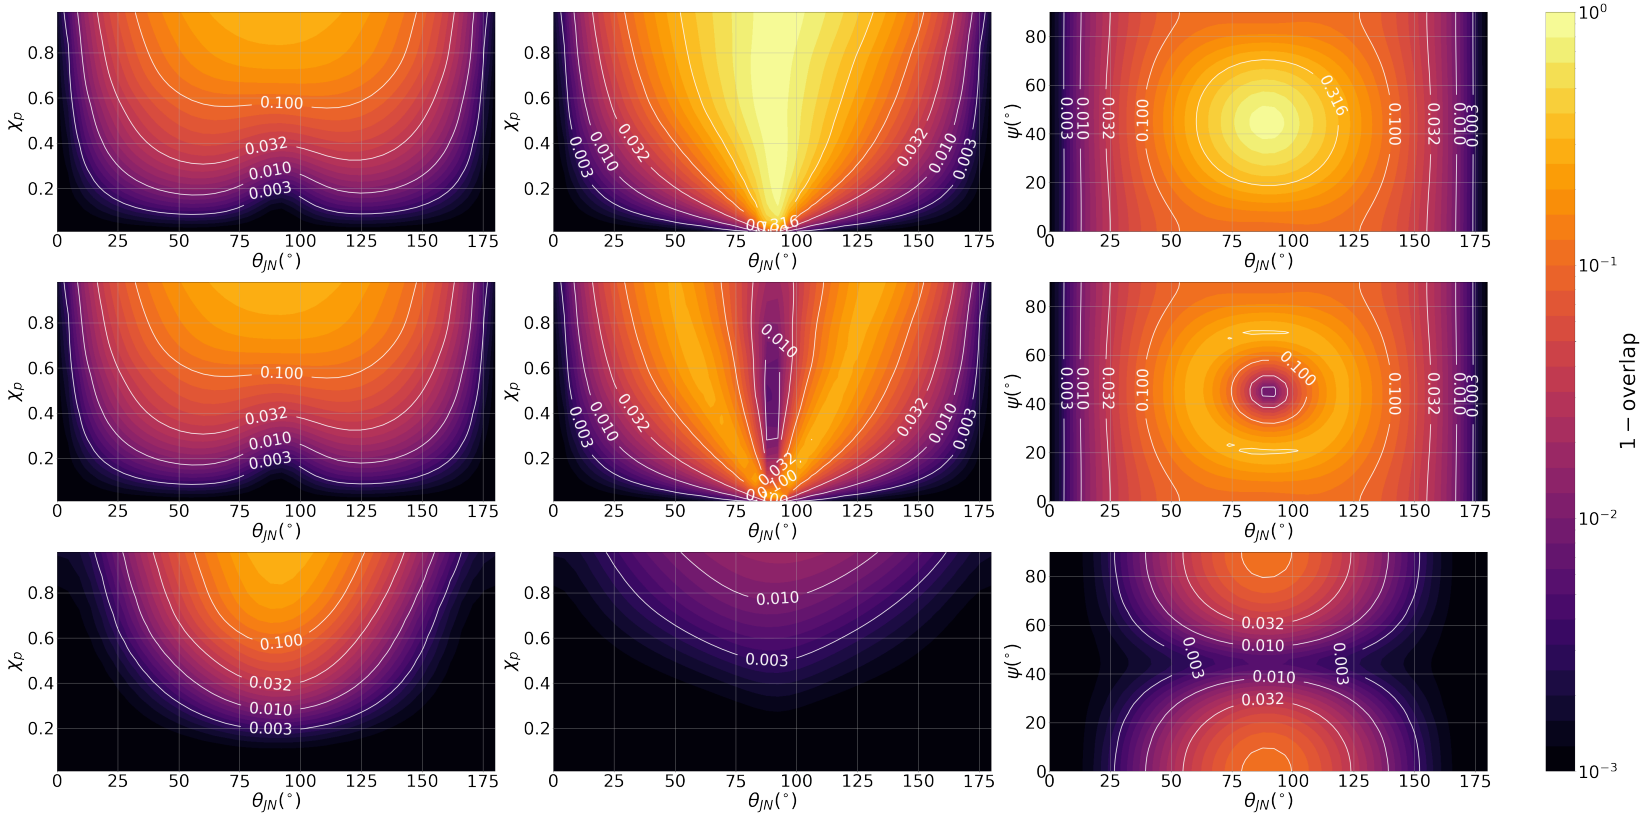

Figure 2.5: The overlap between a precessing waveform and a subset of the harmonics, as a function of the precessing spin and binary orientation for a $40M_\odot$ binary with mass ratio $q = 4$ and $\chi_{\rm eff} = 0$. The top row shows the overlap between the leading, $k = 0$, harmonic and the full waveform; the second row shows the overlap between the dominant harmonic and the full waveform; the bottom row shows the overlap between our two-harmonic precessing waveform and the full waveform. The first column is for the + polarization, second for × and third for fixed $\chi_{\rm p} = 0.6$ and varying polarization.

values of $\theta_{\rm JN}$ and $\chi_{\rm p}$ the other harmonics are more significant and the overlap drops to 0.9 or less. The two-harmonic waveform is a significantly better match to the full waveform, with an overlap greater than 0.99 for much of the parameter space, and only below 0.9 for edge-on systems with high $\chi_{\rm p}$ where the $k = 2$ harmonic contributes most strongly (and the $k = 1$ contribution vanishes).

For the $\times$ polarized waveform (center column), the effect of incorporating the $k = 1$ harmonic is dramatic. For $\theta_{\rm JN} = 90^\circ$ the $k = 0$ contribution vanishes and only the $k = 1, 3$ harmonics are present. Thus, the overlap with harmonic $k = 0$ is essentially zero. Using the best of $k = 0, 1$ provides a good overlap with the edge-on waveform, but there is still a poor overlap at $\theta_{\rm JN} \approx 60^\circ$ where both the $k = 0$ and 1 harmonics contribute significantly to the waveform. This effect has been observed previously, for example in [59, 109] and a geometric understanding of its origin provided. The two-harmonic waveform matches remarkably well to the full waveform, with the largest differences for $\theta_{\rm JN} = 90^\circ$ and $\chi_{\rm p} \approx 1$ where the overlap drops to 0.99 due to the contribution from the $k = 3$ harmonic.

The right column shows the overlap as the orientation of the binary changes. As expected, at points where the $k = 0$ harmonic vanishes ($\theta_{\rm JN} = 90^\circ$ and $\psi = 45^\circ$), the overlap with this harmonic drops to zero. The dominant harmonic is a good match to the waveform, except for orientations where two harmonics contribute significantly. As discussed in detail in Ref. [59], this corresponds to configurations where the binary orientation passes through the null of the detector response (i.e. the signal goes to zero) once per precession cycle. Thus, the radius of the circle with poor overlaps is approximately equal to the opening angle of the binary. The two-harmonic approximation provides an excellent fit to the full waveform over the majority of orientations, only dropping below 0.95 for orientations where $\theta_{\rm JN} \to 90^\circ$ and $\psi \approx 0, 90^\circ$, where the $k = 2$ harmonic is most significant.

Next, we investigate the validity of the two-harmonic approximation for a population of binaries. To begin with, let us fix the masses and spins and just consider the effect of binary orientation. As before, we choose $M = 40 M_\odot$, $q = 4$, $\chi_{\rm eff} = 0$ and $\chi_{\rm p} = 0.6$, corresponding to $\bar{b} \approx 0.3$, with the binary orientation distributed uniformly over $\cos(\theta_{\rm JN}), \phi_o, \alpha_o, \psi$. Fig. 2.6 shows the distribution of the overlap between the full waveform and 1) the $k = 0$ harmonic, 2) the dominant harmonic and 3) the two-harmonic approximation. The results are shown for both a uniformly distributed population, and a population of signals observable above a fixed threshold in the detector — thereby favoring orientations that produce the largest amplitude gravitational wave. The median overlap with either the $k = 0$ or dominant harmonic is $\lesssim 0.9$, while the two-harmonic approximation improves the median overlap to 0.99. Using the dominant harmonic, there are a small fraction of signals with overlaps of 0.7 or lower (and for the $k = 0$ harmonic, this tail extends to overlaps of 0.2), while for the two harmonic approximation, the worst overlap is 0.88.

We can use these results to obtain a *rough* sense of the benefits of performing

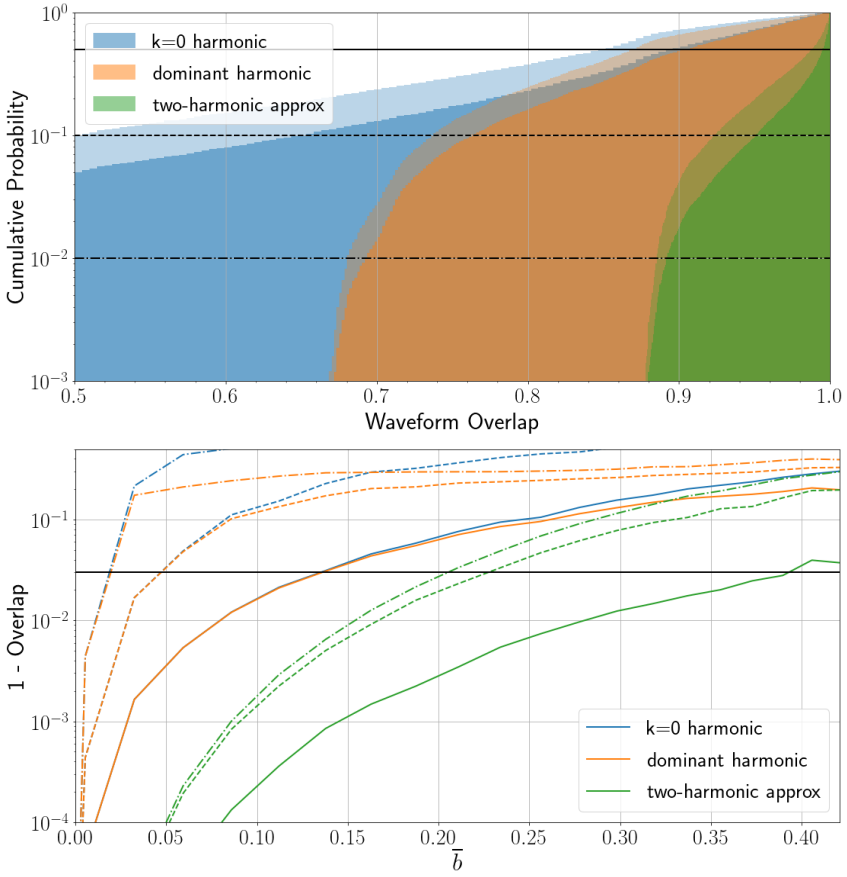

Figure 2.6: The distribution of the overlap of the precessing waveform with the $k = 0$, dominant and two-harmonic waveforms for a population of signals with $M = 40M_\odot$, $q = 4$, $\chi_{\rm eff} = 0$. The top plot shows the overlap distribution for $\chi_{\rm p} = 0.6$, with random orientation of the signal. The lighter shaded regions give the distribution for a randomly oriented population of sources and the darker regions for the expected observed distribution (for a uniform-in-volume source). The lower plot shows the overlap between full and approximate waveforms as a function of $\overline{b}$. The lines on the plot show the value of the overlap for the median (solid line), worst 10% (dashed) and worst 1% (dot-dashed) of signals.

a search using the two-harmonic approximation. Previous, more detailed, investigations of this question have been carried out in, e.g. [114, 59, 125]. Current gravitational wave searches make use of spin-aligned waveforms [126, 127], and a precessing waveform will naturally be identified by a spin-aligned waveform which matches well the dominant harmonic. Thus, we can use the overlaps between the precessing waveforms and dominant harmonics as a proxy for the performance of an aligned spin search. Since the median overlap is 0.9 we would expect to recover approximately 70% as many signals ($\approx 0.9^3$ for a population uniform in volume) as with a full precessing search, above a fixed threshold. A search based upon the two-harmonic approximation would recover around 97% of these signals, indicating an improvement of over 30% in sensitivity to such systems.

We also show how the distribution of overlaps varies across the mass and spin parameter space, as encoded by the parameter $\overline{b}$ and plotted for three choices of spin distribution in Figure 2.3.[6] For $\overline{b} \lesssim 0.13$ — accounting for three quarters of signals in the low-isotropic population — the median overlap between the dominant harmonic and the full waveform is above 0.97. Thus, for the majority of expected signals, the spin-aligned search will have good sensitivity. However, even for low values of $\overline{b}$ there will be some orientations of signals where two dominant harmonic will not match the waveform well, while the two-harmonic waveform still provides an essentially perfect representation of the waveform for all orientations. At $\overline{b} \approx 0.25$ the median overlap with the dominant harmonic waveform drops to 0.9, and it is here that a search with the two-harmonic approximation could provide a 30% improvement. We note, however, that for the low-isotropic distribution this accounts for only 5% of systems. While systems with such significant precession may be rare they would come from interesting areas of parameter space, with high mass ratios and spins. It is only at $\overline{b} = 0.4$ that the median overlap for the two harmonic waveform drops to 0.97, indicating a 10% loss relative to an ideal search, but also 70% improvement over a spin-aligned search.

2.5 Searching for precessing binaries

The two-harmonic approximation provides an ideal basis to develop a search for binaries with precession. The typical approach to searching for binary coalescences has been to generate a template-bank of waveforms that covers the parameter space [128, 129, 130]. These templates comprise discrete points in the mass and spin space chosen so that the waveform produced by a binary anywhere in the parameter space of interest has a match of at least 97% with one of the templates. The waveform for each template is then match-filtered against the data to identify peaks of high SNR, and various signal consistency and coincidence tests are used to differentiate signals from non-stationary noise transients [131, 132, 133, 126, 127]. Current searches make use of a template bank covering the four dimensional mass and aligned-spin space [134, 135].[7] The search takes advantage of the fact that changing the sky location, distance and orientation of the binary only changes the overall amplitude and phase

[6] While these plots were made with fixed masses and $\chi_{\rm eff}$, they should give a reasonable indication of the accuracy of the two-harmonic waveform across the mass and spin parameter space, as a function of $\overline{b}$. For different masses and spins, the evolution of the precession angle during the coalescence can have a slight impact upon the relative importance of the modes but, as b typically does not change significantly over the observable waveform, this effect is likely to be small. Furthermore, as different modes are not perfectly orthogonal, the degree to which they are not will also have a small effect upon the results. As shown in Section 2.6, the harmonics are close to orthogonal for $M \lesssim 40 M_{\odot}$ so that the results shown here will be representative, at least at lower masses.

[7] As we have discussed, the most significant effect on the observed waveform arises due to the effective spin $\chi_{\rm eff}$, which is a combination of the aligned spin components of the two waveforms. Thus, although the template space is four dimensional, one of the spin directions provides limited variation to the waveforms, and thus is relatively straightforward to cover.

of the signal, and these quantities can be maximized over in a simple manner.

When developing a search for precessing binaries, the search becomes more challenging due to the increasing number of parameters. In principle, it is necessary to search over two masses and six spin components, although, in practice it will probably be sufficient to restrict to the masses, $\chi_{\rm eff}$ and $\chi_{\rm p}$. The second complication is that the observed morphology of the waveform varies as the orientation of the binary changes, and it becomes necessary to search over binary orientation $\theta_{\rm JN}$, polarization ψ and precession phase α_o, although methods have been developed to straightforwardly handle a subset of these parameters [136, 114].

The two-harmonic waveform can be used to maximize the SNR over the binary orientation in a simple way. The two complex amplitudes $\mathcal{A}_0$ and $\mathcal{A}_1$, defined in Eq. (2.27), are dependent upon five variables: the distance, d_L, binary orientation, $\theta_{\rm JN}$, ψ, and the initial orbital and precession phases, ϕ_o, α_o. Since $\mathcal{A}_0$ and $\mathcal{A}_1$ can take any value in the complex plane, it is possible to construct the two-harmonic SNR by filtering the two harmonics h_0 and h_1 against the data and then freely maximizing the amplitudes so that,

$$\rho^2_{\rm 2harm} = \rho^2_0 + \rho^2_1 \,. \tag{2.31}$$

If the harmonics are not orthogonal, the two-harmonic SNR should be calculated using h^0 and $h^1_\perp$ — the $k = 1$ harmonic with any component proportional to h^0 removed. The extrinsic parameters of the binary (distance, sky location, orientation, orbital and precession phase) can be searched over through maximization over the amplitudes of the two harmonics, leaving only the masses and spins as dimensions to search using a bank of waveforms.

We must still construct a bank of waveforms to cover the four-dimensional parameter space of masses, the effective aligned $\chi_{\rm eff}$ and precessing $\chi_{\rm p}$ components of the spins. The amplitude and phase evolution of a single harmonic does not carry the tell-tale amplitude and phase modulation caused by precession, but does have a different phase evolution due to precession [110, 73]. Since the phase evolution of each precessing harmonic is degenerate with a non-precessing waveform with different mass-ratio or effective spin, the bank of templates will essentially be a bank of non-precessing waveforms. This may allow us to reduce the size of the template bank.

The k=0 harmonic of the precessing waveform has an additional phase (see Eq. (2.11)) of,

$$\delta\phi_0(t) = \int_{t_o}^{t} \frac{2b^2}{1+b^2}\,\dot{\alpha}\,dt' \,. \tag{2.32}$$

For systems in which orbital angular momentum dominates over spin angular momentum, the precession frequency is inversely proportional to orbital frequency, $\Omega_{\rm p} = \dot{\alpha} \propto f^{-1}$ [54, 88, 59]. This is the same frequency dependence as the 1PN

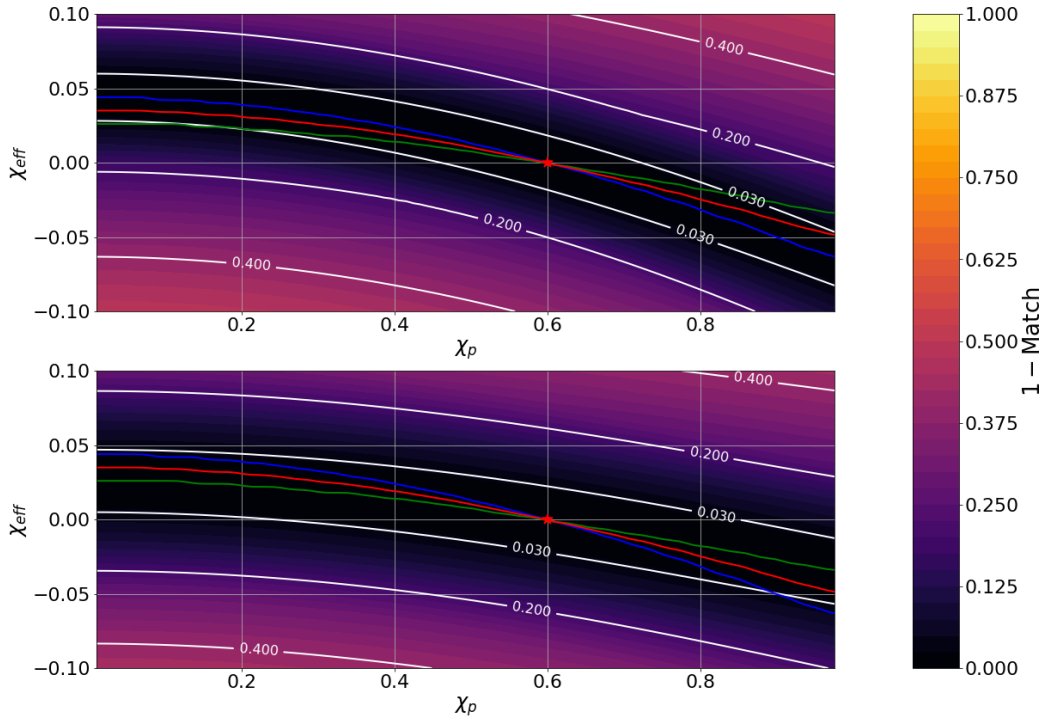

Figure 2.7: The mismatch between the $k = 0$ (top) and $k = 1$ (bottom) harmonic of two precessing signals as the effective spin $\chi_{\rm eff}$ and precessing spin $\chi_{\rm p}$ are varied. For all waveforms, the total mass is fixed to $40M_\odot$ and the mass ratio to 4. One waveform has $\chi_{\rm eff} = 0$ and $\chi_{\rm p} = 0.6$ (the point marked by a star), while the spins of the second waveform are varied. The blue and green lines show the value of $\chi_{\rm eff}$, for the $k = 0$ and $k = 1$ harmonics respectively, which gives the largest match with the fiducial waveform; the red line is the average of these values.

contribution to the waveform, whose amplitude depends upon the mass ratio. Consequently, it is reasonable to expect that the precession-induced phase will be indistinguishable from a systematic offset in the binary mass ratio, or the effective spin [96]. Similarly, the $k = 1$ harmonic has essentially the same amplitude evolution as the non-precessing waveform, but with a phase difference of,

$$\delta\phi_1(t) = -\int_{t_o}^{t} \frac{1-b^2}{1+b^2}\,\dot{\alpha}\,dt'\,, \tag{2.33}$$

which will also, in many cases, be degenerate with a change in the mass ratio or aligned spin.

In Figure 2.7, we investigate the degeneracy in the spin ($\chi_{\rm eff}$–$\chi_{\rm p}$) space of the two leading precession harmonics. We consider a system with masses, $M = 40M_\odot$ and $q = 4$, and spins $\chi_{\rm eff} = 0$, $\chi_{\rm p} = 0.6$ and investigate how the two waveform harmonics vary as the spins are changed. The figure shows the match — the overlap maximized over time-offsets — between our fiducial waveform and one with the same masses but different spins. For both harmonics, there is a band in the $\chi_{\rm eff}$–$\chi_{\rm p}$ plane where

the is mismatch is small — the different phase evolution of each harmonic caused by varying $\chi_{\rm p}$ can be offset by a suitable change in $\chi_{\rm eff}$. The relation is approximately quadratic, $\Delta\chi_{\rm eff} \propto (\Delta\chi_{\rm p})^2$, which is to be expected. Recall, from Eq. (2.32), that change in phase due to precession is quadratic in b, and therefore also in $\chi_{\rm p}$ at least for small values of b. Meanwhile the phasing of the waveform varies, at leading order, linearly with $\chi_{\rm eff}$.

This degeneracy in the $\chi_{\rm eff}$–$\chi_{\rm p}$ plane suggests that a single template waveform could be used to search over an extended region corresponding, for example, to the region of mismatch < 0.03 in Figure 2.7. However, this will only work if the degenerate region for the $k = 0$ and $k = 1$ harmonics is the same. It is clear from Equations (2.32) and (2.33) and Figure 2.7 that they are not identical. Nonetheless,[8] for the example we have considered, the two degenerate regions are similar, and along the line that traces the mid-point between best fit values of $\chi_{\rm eff}$ for the two harmonics, both harmonics have a match above 0.97 with the initial point. Thus, to an accuracy appropriate for generating a template bank, we can use the two harmonics from a single waveform to cover a band in the $\chi_{\rm eff}$–$\chi_{\rm p}$ plane which spans all values of $\chi_{\rm p}$. This effectively reduces the dimensionality of the parameter space to three dimensions: mass, mass ratio and one spin parameter.

Our proposal to develop a precessing search is as follows. First, generate a bank of templates to cover the space of non-precessing binaries. At each M, q, $\chi_{\rm eff}$ point in the template bank, construct the two-harmonic waveform for a fixed value of $\chi_{\rm p}$. Then, filter the data against the two harmonics and calculate the two-harmonic SNR, as defined in Eq. (2.31) to identify candidate events in a single detector. It will be necessary to extend the existing χ^2 signal consistency test [132] to each harmonic, taking into account the presence of the other harmonics, to reduce the impact of non-stationarity in the data. Next, perform coincidence between detectors by requiring a signal in the same template at the same time, up to the allowed time delays based upon speed of propagation. For a non-precessing signal observed in two detectors, the relative amplitude and phase of the SNR in each detector can take any value, even though some are astrophysically more likely [137] (and this can be used to increase search sensitivity). However, for the two-harmonic waveform not every signal observed in two detectors will be compatible with an astrophysical source. This can be seen through simple parameter counting: there are ten measured quantities (two complex amplitudes and a time of arrival in each detector), which depend upon eight parameters, the five orientation parameters (d_L, $\theta_{\rm JN}$, ψ, ϕ_o, α_o), sky location and merger time. An additional coincidence test to check for consistency between parameters will likely be necessary to reduce the search background. A similar problem arises already in extending the amplitude and phase consistency of

[8] Strictly, when doing this comparison, we must use the same time offset for the two harmonics, whereas the figure allows for an independent maximization of the time delay for each harmonic. Fixing a single time delay does slightly decrease the matches, but not significantly enough to change the conclusions.

[137] to three or more detectors and methods developed for that purpose may be helpful for the precessing search.

We can estimate the likely sensitivity improvement from a precessing search, as we have briefly discussed in Section 2.4. A non-precessing search will typically find the dominant harmonic of the waveform. Thus, for signals where two harmonics provide a significant contribution, a search based on the two-harmonic waveform has the potential to out-perform the non-precessing search. The two-harmonic waveform has four degrees of freedom, encoded in $\mathcal{A}_0$ and $\mathcal{A}_1$, compared to two for the non-precessing search. Thus, the noise background is higher for the two-harmonic search and, based upon a comparison of the tails of the χ^2 distribution with 2 and 4 degrees of freedom, an increase of around 5% in SNR is required to obtain the same false alarm rate (see e.g., Ref. [114] for a discussion of this issue). Thus, a signal will be observed as more significant in the two-harmonic search than a non-precessing search if the SNR can be increased by 5% or more. Fig. 2.6 shows that this occurs for $\overline{b} \gtrsim 0.15$, and for binaries with $\overline{b}$ above this value the two-harmonic search has the potential to outperform a non-precessing search. We note, however, that a given template will cover a range of spin values and consequently a range of $\overline{b}$, so it may be more appropriate to deploy the two-harmonic search for templates with an *average* of $\overline{b}$ which is greater than 0.15.

Another challenge of searches for precessing systems is the associated computational cost [114], which can be prohibitive. The maximum computational cost for the two-harmonic search would be double that of a comparable non-precessing search: it becomes necessary to filter both the $k = 0$ and 1 harmonics, and computational time is dominated by this matched filtering. However, since both the $k = 0$ and $k = 1$ harmonics are essentially non-precessing waveforms, there may be waveforms associated with the $k = 1$ harmonics are *already* in the set of $k = 0$ waveforms, but associated with different parameters. If so, this could further reduce the computational cost.

2.6 Observability of precession

The two-harmonic approximation allows us to easily identify regions of the binary merger parameter space for which precession will leave an observable imprint on the waveform. Since the amplitude and phase evolution of a single harmonic is generally consistent with that of a non-precessing waveform (see above and [110, 109]), it is only when two harmonics can be observed that we are able to clearly identify precession in the system. We are therefore interested in deriving an expression for the *precession SNR*, $\rho_{\rm p}$, defined as the SNR in the second most significant harmonic, and determining when it will be observable. If the two harmonics h^0 and h^1 in

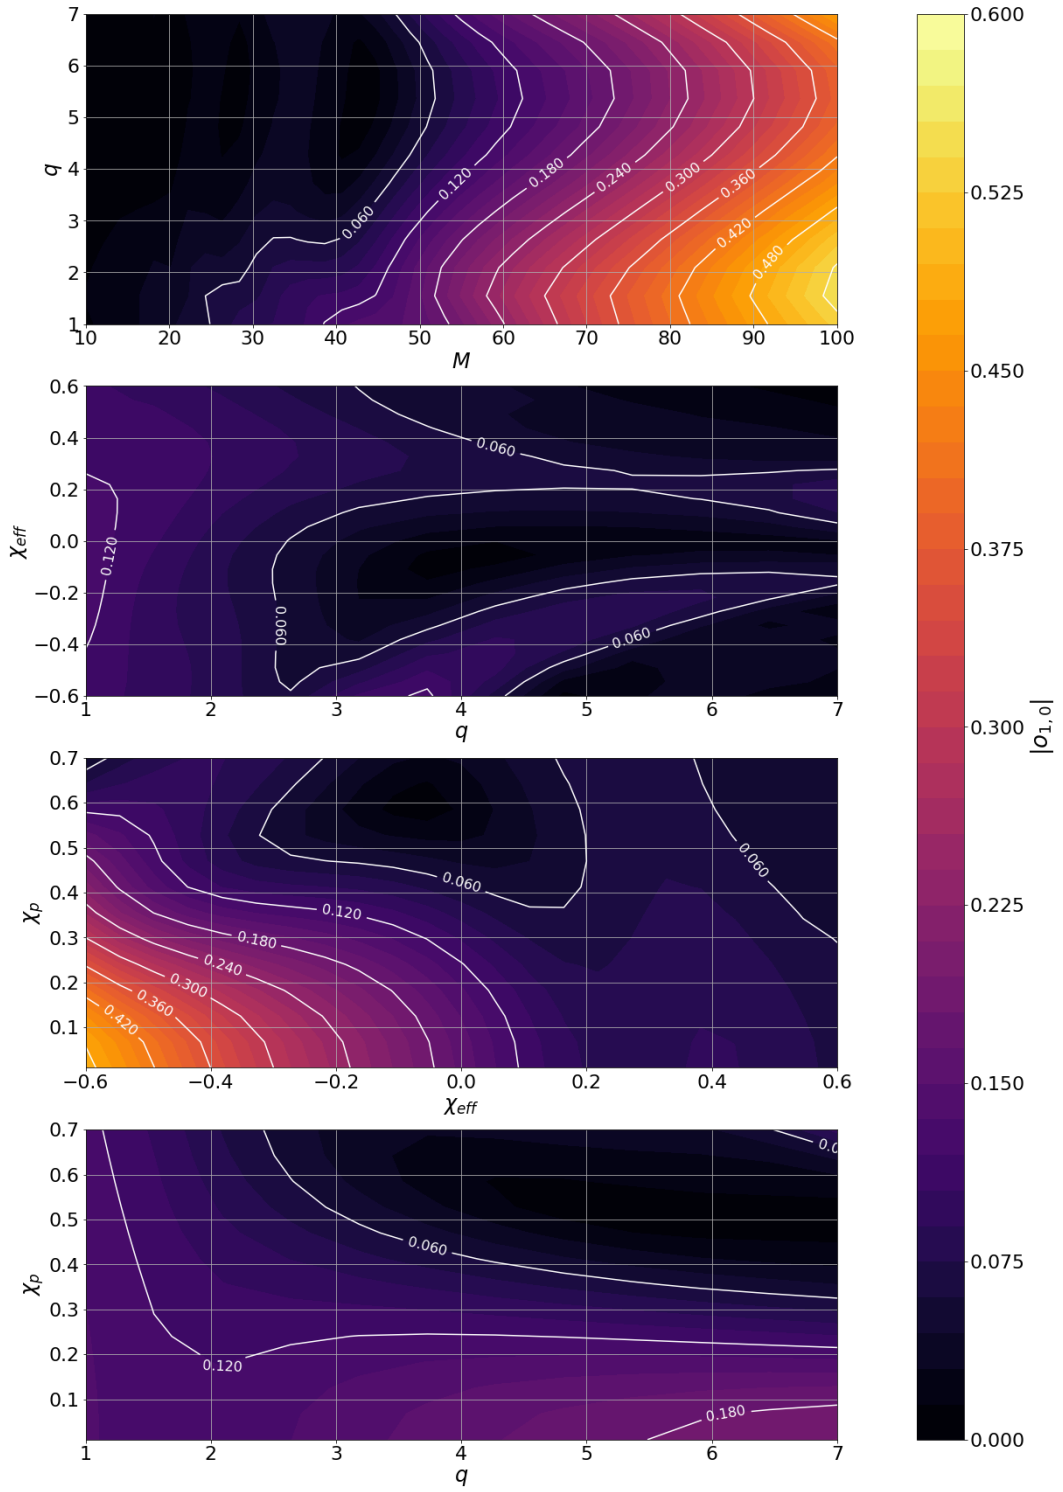

Figure 2.8: The overlap $O(h_0, h_1)$ between the $k = 0$ and $k = 1$ harmonics across two-dimensional slices in the parameter space of total mass, mass ratio, $\chi_{\rm eff}$ and $\chi_{\rm p}$. In each plot, two of the parameters are varied while the other two are fixed to their fiducial values of $M = 40M_\odot$, $q = 4$, $\chi_{\rm eff} = 0$, $\chi_{\rm p} = 0.6$.

Eq. (2.26) are orthogonal, then the precession SNR is simply,

$$\begin{aligned}\rho_{\rm p} &= \min(|\mathcal{A}_0 h^0|, |\mathcal{A}_1 h^1|), \\ &= \rho_{\rm 2harm} \left(\frac{\min(1, |\zeta|)}{\sqrt{1+|\zeta|^2}} \right), \end{aligned} \tag{2.34}$$

where ζ, defined in Eq. (2.28), gives the ratio of the SNR in the $k = 1$ and $k = 0$ harmonics and $\rho_{\rm 2harm}$ is the total SNR in the two-harmonic waveform.

Let us briefly examine where in parameter space the two harmonics are close to orthogonal. Where there are sufficient precession cycles we expect the two harmonics, h^0 and h^1, will be close to orthogonal, and the overlap to be close to zero [110]. The overlap between the two harmonics for various two-dimensional slices through the parameter space is shown in Fig. 2.8. At higher masses, where the binary completes one, or fewer, precession cycles in the detector's sensitive band, there is a larger overlap between the harmonics. At negative $\chi_{\rm eff}$ and minimal $\chi_{\rm p}$, the overlap is also significant. However, providing the mass of the system is below $50 M_\odot$, for the much of the parameter space the overlap is less than 0.1 and simple expression in Eq. (2.34) will be applicable.

Taking into account the overlap between harmonics, the total power in the two-harmonic waveform is,

$$\rho^2_{\rm 2harm} = |\mathcal{A}_0 h^0|^2 \left(1 + 2{\rm Re}[\zeta\, o_{1,0}] + |\zeta|^2\right) . \tag{2.35}$$

where $o_{1,0}$ is complex overlap between the two harmonics:

$$o_{1,0} = \frac{(h^1|h^0) + i(h^1|ih^0)}{|h^1||h^0|} . \tag{2.36}$$

We can project the SNR onto directions parallel and perpendicular to the h^0 waveform to obtain the SNR in these two directions as,

$$\begin{aligned}\rho^2_0 &= |\mathcal{A}_0 h^0|^2 \left(1 + 2{\rm Re}[\zeta\, o_{1,0}] + |\zeta\, o_{1,0}|^2\right), \\ \rho^2_{\perp,0} &= |\mathcal{A}_0 h^0|^2 |\zeta|^2 \left(1 - |o_{1,0}|^2\right) . \end{aligned} \tag{2.37}$$

Similarly, the power parallel to and perpendicular to the $k = 1$ harmonic is,

$$\begin{aligned}\rho^2_1 &= |\mathcal{A}_0 h^0|^2 \left(|o_{1,0}|^2 + 2{\rm Re}[\zeta\, o_{1,0}] + |\zeta|^2\right), \\ \rho^2_{\perp,1} &= |\mathcal{A}_0 h^0|^2 \left(1 - |o_{1,0}|^2\right) . \end{aligned} \tag{2.38}$$

The precession SNR is defined as the power orthogonal to the dominant har-

monic,[9]

$$\begin{aligned}\rho_{\rm p} &:= \min(\rho_{\perp,0}, \rho_{\perp,1}), \\ &= \rho_{\rm 2harm} \min(1, |\zeta|) \left(\frac{1 - |o_{1,0}|^2}{1 + 2{\rm Re}[\zeta\, o_{1,0}] + |\zeta|^2} \right)^{\frac{1}{2}} .\end{aligned} \tag{2.39}$$

As expected, the precession SNR scales with the total SNR of the signal, so that precession will be more easily observed for louder events. If there is significant degeneracy between the harmonics, the numerator will be reduced, making the observation of precession more difficult. Finally, in the limit that $o_{1,0} \to 0$, the expression simplifies to the one given earlier for orthogonal harmonics in (2.34), as expected.

What value of $\rho_{\rm p}$ will be required to observe precession? This will happen if the evidence for a signal with $\chi_{\rm p} \neq 0$ in the data is greater than that for a non-precessing source. This can be evaluated through Bayesian model selection, by considering the Bayes factor between the hypotheses. However, such a calculation requires a full exploration of the parameter space. We can, instead, obtain an approximate answer by considering the maximum likelihood. Since the two-harmonic waveform is more general than the non-precessing waveform, it will always give a larger maximum likelihood *even in the absence of precession* do to its ability to fit the detector noise. Thus, we are interested in examining the expected increase in SNR due to the inclusion of the second harmonic, in the absence of any power in it.

The two-harmonic SNR can be written as

$$\rho^2_{\rm 2harm} = \rho^2_{\rm np} + \rho^2_{\rm p} \,. \tag{2.40}$$

where ρ_{np} is the non-precessing SNR or, equivalently, the SNR in the dominant harmonic. In the absence of precession, $\rho_{\rm p}$ will be χ^2 distributed with 2 degrees of freedom, as we are able to freely maximize over the amplitude and phase of the two harmonics independently [131, 133]. Consequently, in 90% of cases, noise alone will give a value of $\rho_{\rm p} < 2.1$. Therefore, as a simple criterion, we require that,

$$\rho_{\rm p} \geq 2.1, \tag{2.41}$$

for precession to be observable. In Chapter 3 we use this definition to investigate in detail the observability of precession over the parameter space.

As can be seen from Eq. 2.34, $\rho_{\rm p}$ relies upon the two harmonic approximation and therefore inherently assumes that the $k = 0, 1$ harmonics dominate the harmonic decomposition of a precessing gravitational wave. While this assumption is valid for the vast majority of parameter space (see e.g. Figs. 2.2 and 2.5), there are regions where either the $k = 2, 3$ or 4 harmonics contain more power than the $k = 1$

[9] In exceptional circumstances, where the overlap is large and $\zeta o_{1,0}$ is close to -1, there can be more power in $\rho_{\perp,i}$ than ρ_i. In such cases, it is natural to use ρ_i to determine if precession is present, although this is not ideal as $\rho_{\perp,i}$ need not resemble a non-precessing waveform.

harmonic, e.g. for systems that are viewed close to edge on with large in-plane spins. This means that $\rho_{\rm p}$ will be underestimated for these systems. If we assume an astrophysically motivated population of binary black holes, where the systems are distributed uniformly in binary orientation, masses drawn from a power law distribution and spins drawn from a low isotropic distribution (see Chapter 4 for details), we find that 3 in every 1000 binaries detected by LIGO–Virgo contains more power in either the $k = 2, 3$ or 4 harmonics than the $k = 1$ harmonic. This means that on average $\rho_{\rm p}$ will be underestimated for 0.3% of binaries detected by LIGO–Virgo. Therefore, while there are potential systematics in computing $\rho_{\rm p}$ using only the $k = 0$ and $k = 1$ harmonics, this only affects a small population of binaries which are unlikely to be observed with LIGO–Virgo.

2.7 Discussion

We have presented a new, intuitive way to understand the observability of precession in GW observations. By keeping only the leading precession term, we have derived a precession SNR and argued that this can be used to determine when precession will be observable. Before discussing applications we point out the main limitations of this analysis. As is clear from the formulation, this analysis works best for binaries where $b = \tan(\beta/2)$ is small. This typically corresponds to situations where the masses are comparable, the precessing spin is small and any aligned component of the spin is aligned (rather than anti-aligned) with the orbital angular momentum. We have shown above that this assumption is valid for a reasonable population.

We now point to several advantages and applications of this formulation:
First, it gives new understanding of the observability of precession, and also of the origin of precession as the beating of two waveform components with slightly differing frequencies (also discussed in [110]). It is difficult to identify the presence of precession in a GW observation directly from $\chi_{\rm p}$, since the prior astrophysical expectation disfavours $\chi_{\rm p} = 0$. While the deviation from the prior can be determined through the Bayes factor, the results in this chapter suggest that the precession SNR $\rho_{\rm p}$ could provide a direct measure of whether precession has been measured in a signal. The potential applications of $\rho_{\rm p}$ are discussed in Chapter 4, and will be investigated in more detail in Chapter 3, where we probe the measurability of precession across the gravitational wave parameter space.

There exist a number of detailed population analyses which extract the features of the underlying population of gravitational waves from the set of observed gravitational wave events, for example [138, 139, 140, 85]. These typically use the full posterior distributions recovered from the gravitational wave signal [91, 141] to infer the population and, as such, naturally account for precession effects in the observed signals when inferring the black hole mass and spin populations. Nonetheless, there have been a number of studies performed which investigate the population proper-

ties using a subset of the recovered parameters, see e.g. [142, 118, 143, 144, 119, 85], and have been successfully used to infer interesting properties of the mass and spin distributions. The majority of these studies have restricted attention to the aligned components of the spins. The precession SNR provides a straightforward method to determine the significance of precession, and provides a way to probe observability of precession in populations of binaries. In using this method we have been able to derive constraints on the preferred spin distribution including precession effects (see Chapter 4).

Both of the applications highlighted above are currently possible using other more sophisticated but computationally expensive methods such as Bayesian model comparison. This is, of course, a more general method that makes fewer assumptions than we do in computing $\rho_{\rm p}$, however the computational costs associated with calculating the marginal likelihood over multiple, e.g. precessing and non-precessing, models per binary are not feasible for a large number of binaries. For example the analysis in Chapter 4 involved calculating $\rho_{\rm p}$ for 1 million binaries, and computing the Bayes factor for 1 million binaries would certainly not be practical. Similar, lightweight analyses, could also be developed using the formalism introduced in, e.g. [59], and if this is done,it would be interesting to compare them with the results from the two harmonic analysis.

Finally, we have outlined a method by which the two-harmonic approximation could be used to develop a search for precessing binaries. We have shown that in principle that this approach could result in a significant increase in sensitivity without the computational overheads associated with other precessing search methods. In addition, the formalism should provide a way to identify the parts of parameter space where a precessing search is likely to increase sensitivity. We plan a detailed investigation into the feasibility of a precessing search based upon the two-harmonic approximation in future work.

Chapter 3

Identifying when Precession can be Measured in Gravitational Waveforms

This chapter is based upon the text of Green and Hoy et al. [145] published as Phys. Rev. D, 103, 124023. This work was led by Rhys Green and Charlie Hoy. My main contributions were generating the workflow, managing the runs and paper writing. Significant contributions were made to all sections but I specifically led Sections 3.4.a, 3.4.d, 3.4.g. All authors edited the text.

3.1 Introduction

In September 2015, the first direct detection of GWs marked the beginning of GW astronomy [7]. Another 14 detections have been announced by the LIGO Scientific and Virgo collaborations (LVC), the vast majority of which were due to black hole (BH) mergers [38, 43, 45, 47, 146, 46]. Additional events have also been reported by independent groups [39, 40, 41, 42]. These GW observations have already provided significant insights into gravitational physics, cosmology, astronomy, nuclear physics and fundamental physics (see e.g. Refs. [147, 148, 149, 150, 151, 152, 153, 85]). With an order of magnitude more observations expected over the next 5-10 years, as the sensitivities of the LIGO [4, 5], Virgo [6] and KAGRA [11] detectors improve and additional detectors come online, GW astronomy from compact-binary mergers has the potential to transform our understanding of gravitational and fundamental physics [154, 155, 156].

Everything we learn from GW BBH observations is a consequence of a detailed parameter estimation analysis that extracts the source parameters of the binary. While some parameters are extracted with good precision, inspiral dominanted signals show strong correlations between certain parameters which means that they cannot be measured so accurately, for example correlations between the binary's distance and inclination [60, 8, 99], the two masses [60, 61], and the mass-ratio and spin components aligned to the binary's orbital angular momentum [61, 62, 157, 158]. As well as studies of the inspiral, work has been done to extract the source properties for high mass signals dominated by the merger ringdown, see e.g. [159, 160, 104, 161].

Spin components misaligned with the binary's orbital angular momentum, leading to a precession of the binary's orbital plane and hence modulations of the amplitude and phase, have not yet been unambiguously measured in GW observations [38], see Fig. 3.1. Precession effects and correlations with other parameters are understood in principle [54, 58] but since theoretical signal models of precessing binaries that include the merger and ringdown date from only shortly before the first detections [73, 162], we have less experience of when precession will be measurable, and what the impact will be on other parameter measurements.

The purpose of this chapter is to explore when precession will be measurable, and its impact on other parameter measurements, in the kind of configurations that are representative of expectations from binary populations based on LIGO-Virgo-KAGRA observations to date [38]. By utilizing the precession SNR $\rho_{\rm p}$, introduced in Chapter 2, as a quantifier for the measurability of precession, we also verify that $\rho_{\rm p}$ is indeed a good metric for the measurability of precession across the vast majority of the parameter space, and relate it to the standard means to identify the presence of precession, the Bayes factor. In doing so, we show that computationally expensive parameter estimation runs can be avoided by simply calculating the precession SNR.

Previous work has explored the general phenomenology of precession effects:

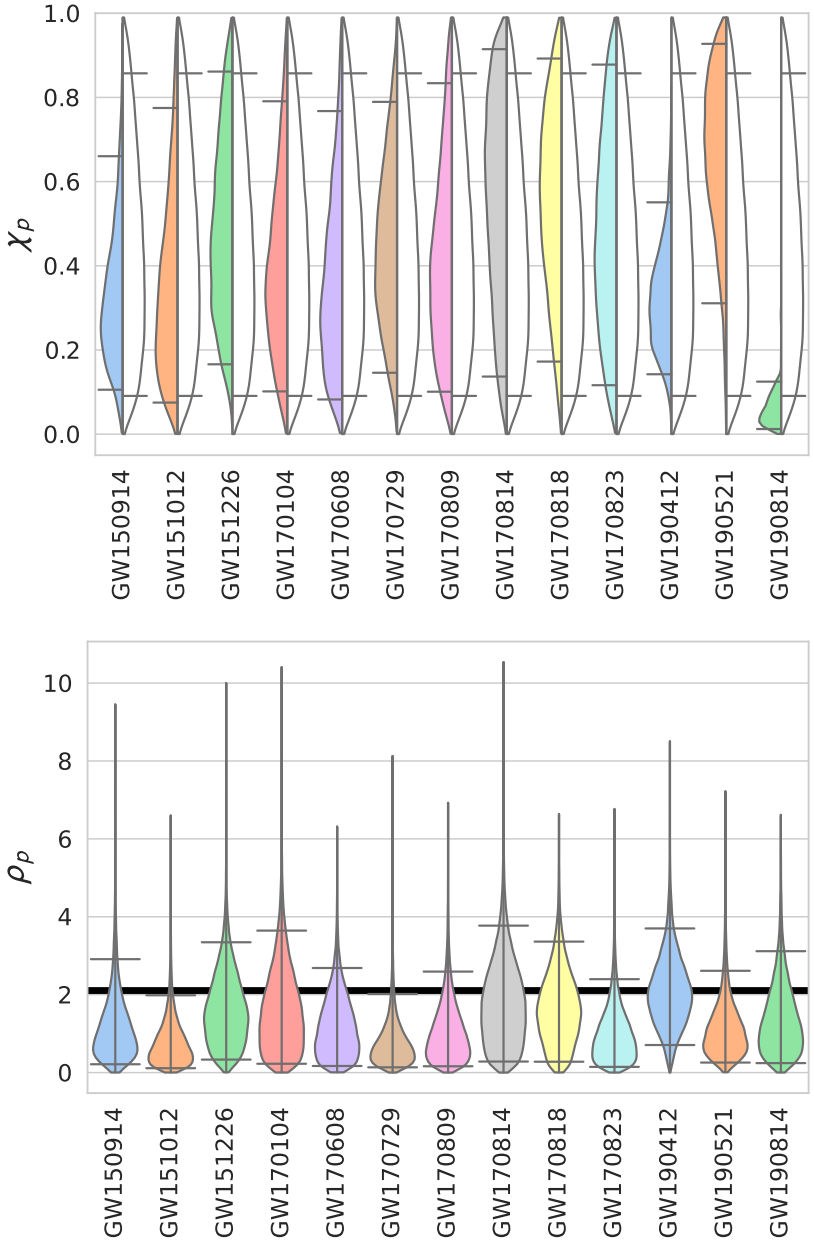

Figure 3.1: Plot showing the posterior distributions for χ_{p} and ρ_{p} for all LIGO/Virgo BBH observations. The χ_{p} posterior distribution (left hand side, colored) is compared to its prior (right hand side, white) in the form of a split violin plot. The ρ_{p} posterior distribution is shown as a single violin. Horizontal grey lines show the 90% symmetric credible interval. The solid black line shows the $\rho_{\mathrm{p}} = 2.1$ threshold. Bounded KDEs are used for estimating the probability density. Data obtained from the Gravitational Wave Open Science Center [163].

its increased measurability with large in-plane spins [164, 94, 165], large mass ratios [164, 94], high inclination [54, 59, 166, 167, 120, 104], and of course high SNR [164, 94, 168]. Beyond these general expectations, the *quantitative* behaviour of parameter measurements in the presence of precession has not been studied in great detail for typical LIGO-Virgo-KAGRA observations. The measurability of precession for high mass ratio LIGO-Virgo-KAGRA observations like GW190814 has been investigated in recent work [169].

In this chapter, we focus on the region of parameter space most likely to yield binaries with observable precession: binaries of comparable mass, with moderate in-plane and aligned-spin components [120]. We perform a series of one-dimensional investigations of the parameter space, in which we vary one parameter at a time: total mass, mass ratio, spins (both in-plane as characterized by χ_p, and the aligned spin combination $\chi_{\rm eff}$), the binary orientation (both the inclination of the orbit and also binary polarization), and the sky location and show the impact of varying each of the binary parameters individually. These investigations serve to confirm that much of the known phenomenology is apparent even at relatively low SNR, while also demonstrating that the precession SNR can be effectively used across a significant fraction of the parameter space to *predict* the observable consequences of precession *without* the need for computationally costly parameter estimation analyses.

This chapter is structured as follows: Sec. 3.2 provides a brief recap of the two-harmonic approximation that allows us to define $\rho_{\rm p}$, and a summary of the importance of precession across the parameter space. Sec. 3.3 provides an introduction to the parameter estimation techniques used here, and parameter estimation results and interpretation for our fiducial system. In Sec. 3.4 we perform a series of one-dimensional explorations of the parameter space. In Sec. 3.6 we compare the predicted precession SNR with observations and in Sec. 3.5 we compare precession SNR with the Bayes factors between precessing and non-precessing runs. We conclude with a summary and discussion of future directions.

3.2 Two harmonic approximation

In Chapter 2 we introduced a method for decomposing a precessing waveform into a series of five *non-precessing harmonics*, where the characteristic modulations of a precessing signal are caused by the beating of these *harmonics*. These harmonics form a power series in the parameter $b = \tan(\beta/2)$. We found that in most regions of parameter space, the two leading harmonics are sufficient to capture the significant precession features in the waveform and the other harmonics can be neglected. As discussed in detail in Sec. 2.4, these other precession harmonics can be ignored for binaries where $\beta \lesssim 45^\circ$. Thus, for almost all signals we expect to observe, the

waveform can faithfully be expressed as,

$$h(f) \approx \mathcal{A}_0 h^0(f) + \mathcal{A}_1 h^1(f)\,, \tag{3.1}$$

where $\mathcal{A}_0$ and $\mathcal{A}_1$ are complex, orientation dependent amplitudes, and $h^0(f)$ and $h^1(f)$ are the waveforms of the two leading harmonics.

The observability of precession can then be characterised by the precession SNR $\rho_{\rm p}$, defined as the SNR in the weaker of the two harmonics. $\rho_{\rm p}$ is a function of ρ, the total signal SNR, the overlap between the two harmonics and ζ, the relative significance of the two harmonics, see Eq. 2.39. ζ is defined as Eq. 2.28. However, this expression is restricted to the special case where the detector response is only a function of the polarization angle ψ, $F_+ = \cos 2\psi$ and $F_\times = -\sin 2\psi$. In this chapter, we use a generalized form of Eq. 2.28 which accounts for the detector responses dependence on the sky location,

$$\begin{aligned} \zeta &:= \frac{\bar{b}\mathcal{A}_1}{\mathcal{A}_0} \\ &= \bar{b}e^{i\phi_{JL}} \left(\frac{F_+ \sin 2\theta_{\rm JN} + 2iF_\times \sin\theta_{\rm JN}}{\frac{1}{2}F_+(1+\cos^2\theta_{\rm JN}) + iF_\times \cos\theta_{\rm JN}} \right), \end{aligned} \tag{3.2}$$

where the detector response $F_{+,\times}$ is calculated using the polarization angle appropriate for a co-ordinate system defined with the z-axis along the direction of total angular momentum $\mathbf{J}$ and the sky location. $\bar{b}$ is defined in Eq. 2.22, ϕ_{JL} is the reference precession phase[1] and $\theta_{\rm JN}$ is the angle between $\mathbf{J}$ and the line of sight (see Fig. 1.3).

The quantity $\rho_{\rm p}$ parameterises the *observable* precession, it is therefore the *measured* quantity in the data. By considering what we actually measure in the data we are able to understand many of the correlations and degeneracies in the physical parameters that are *caused* by the presence of (or lack of) measurable precession.

As discussed in Sec. 2.6, in the absence of precession, $\rho_{\rm p}^2$ will be χ^2 distributed with two degrees of freedom. Consequently, if there is no observable precession in the system, $\rho_{\rm p} \geq 2.1$ in only 10% of cases. Thus far we have used $\rho_{\rm p} = 2.1$ as a simple threshold to determine if there is any observable precession in the system. We revisit this in more detail in Sec. 3.6.b.

3.2.a Observability of precession

The strength of the modulations in the GW signal depend primarily on the opening angle, β, and this is reflected in the expansion parameter b in the two-harmonic approximation; the precession frequency $\dot{\alpha}$ also plays a role. The strength of the

[1] This is equivalent to α_0 in Chapter 2. We use ϕ_{JL} in this chapter to be consistent with other parameter estimation studies.

modulations in the *observed* signal also depend on the binary's inclination to the observer, $\theta_{\rm JN}$, and the detector polarisation ψ, and these are all incorporated into the precession SNR $\rho_{\rm p}$, through Eqs. (3.2) and (2.39). From these we can draw immediate conclusions about the scenarios in which precession will be most easily measured. These observations are in general not new (see, as always, the pioneering discussions in Refs. [54, 58]), but we summarise them here and, where salient, present them in terms of the two harmonic formalism, which highlights the insights and intuition that are simplified in this formulation. We then compare these expectations with the quantitative results that we find in our full parameter estimation study.

Our first basic picture of the strength of precession effects comes from Eq. (1.1), which gives the dominant effect on β during the inspiral. If we first consider cases where the spin is entirely in the orbital plane, i.e., $S_{||} = 0$, we see that the opening angle β will be zero if $S_\perp = 0$ (as we would expect), and increases linearly for small $S_\perp$. The opening angle also increases as μ decreases, i.e., as the mass ratio is increased. Eq. (1.1) is no longer accurate near merger, and for equal-mass systems β does not become large, but for large mass ratios the opening angle can approach 90°.

If we now consider non-zero $S_{||}$, we see that the level of precession will be reduced for systems with a positive aligned-spin component, and will be increased for systems with a negative aligned-spin component. The importance of this effect will depend on the other terms, but we can see that for a high-mass-ratio system where μ is very small, and close to merger, so rM is also small, the aligned-spin component will have a strong effect on β, and therefore the measurability of precession. A negative $S_{||}$ is necessary to achieve $\beta > 90^\circ$, and for large mass-ratio systems near merger (small μ and rM) and large negative $S_{||}$, β can approach 180°, but such systems will be rare.

The measurability of precession also depends on the orientation of the binary with respect to the detector, $\theta_{\rm JN}$. As we see in Eq. (3.2), precession effects will be minimal if $\theta_{\rm JN} \sim 0^\circ$ or 180°, i.e., the observer views the system from the direction of $\hat{\mathbf{J}}$. We expect precession to be strongest in the observed waveform for orientations close to $\theta_{\rm JN} \sim 90^\circ$. Additionally, when the detector, or network is primarily sensitive to the $\times$ polarization, precession effects will be more significant. The amplitude of the $k = 1$ harmonic vanishes in the $+$ polarization for both face on $\theta_{\rm JN} = 0^\circ$ and 180° and edge-on $\theta_{\rm JN} = 90^\circ$ systems, while the $\times$ polarization is maximal for edge-on systems. Additionally, the $\times$ polarization for the $k = 0$ harmonic vanishes for edge on systems, while the $+$ polarization is only reduced by a factor of two. Thus, even when b is small, there can be observable precession when the system is close to edge on and the network is preferentially sensitive to the $\times$ polarization. For a given choice of masses and spins, the maximum precession SNR is $\rho_{\rm p} = \rho/\sqrt{2}$.

3.3 Parameter Estimation Results

3.3.a Standard configuration

We begin by describing the results of the parameter recovery routine for a specific simulated signal. The details of the signal are given in Tab. 3.1. These parameters were chosen so that precession effects would be significant in the observed waveform while still being consistent with the observed population of BBHs. In the following sections, we vary over the parameters of the signal one-by-one to investigate the impact of each parameter on the observability of precession and the accuracy of parameter recovery. For each parameter, we are able to both increase and decrease the significance of precession.

By taking the inferred properties of the BBHs observed in the first, second and third observing runs [44], it is predicted that 90% of detected binaries will have mass ratios $q < 4$ and $\sim 97\%$ of BHs in these binaries will have masses less than $45M_{\odot}$ [170]. Our "standard" simulated signal was chosen to have total mass $M = 40M_{\odot}$ and mass ratio $q = 2$ inclined at an angle of $\theta_{JN} = 60^{\circ}$. This corresponds to component masses of $26.7M_{\odot}$ and $13.3M_{\odot}$. This mass ratio and inclination was chosen to increase the observability of precession.

Of the 50 events reported by the LIGO/Virgo, 13 exclude the aligned-spin measure $\chi_{\rm eff} = 0$ at 90% confidence [33, 38, 44]. The other 37 observations peak at $\chi_{\rm eff} = 0$ [38, 44]. Based on this, studies have shown that it is likely BHs in binaries have low spin magnitudes [85, 118, 119, 120]. For this reason, in our standard configuration the BH spins were chosen such that there is zero spin aligned with the binary's orbital angular momentum, $\chi_{\rm eff} = 0$. We introduce precession by giving the more massive BH a spin of 0.4 in-plane and leaving the second BH with zero spin; two-spin effects are generally far weaker than the dominant precession effect, which exhibits the same phenomenology as a single-spin system [171, 64]. From Eq. (1.3) we see that this gives us a system with $\chi_{\rm p} = 0.4$. The opening angle for the binary when the signal enters the detector's sensitivity band is 10° and the average value of the parameter $b = \tan(\beta/2)$ is $\overline{b} = 0.11$, from Eq . (2.22). The signal is simulated using the `IMRPhenomPv2` waveform model that incorporates precession effects, but not higher harmonics ($\ell > 2$) in the signal [73, 172].

Our "standard" simulated signal was chosen to be more favourable to precession measurements than typical LIGO-Virgo observations. Assuming systems are distributed uniformly in binary orientation, masses drawn from a power law distribution and spins drawn from a low isotropic distribution (see Ref. [120] for details), we expect that 4 in every 100 binaries detected by LIGO-Virgo will be inclined at angles greater than 60° and have $\overline{b} > 0.11$.

The sky location of the binary was chosen to have RA $= 1.88\,\mathrm{rad}, \mathrm{DEC} = 1.19\,\mathrm{rad}$. The coalescence time is $t = 1186741861$ GPS (corresponding to the merger

time of GW170814 [37]). The polarization angle, defined by the orientation of the orbital plane when entering the sensitive band at 20Hz, is $\psi = 40°$. The two harmonic approximation is calculated in the J-aligned frame ($\hat{z} = \hat{\mathbf{J}}$). In this frame, the polarization angle is $\psi_J = 120°$, which gives antenna factors for H1 of $F_+ = 0.34$ and $F_\times = 0.53$ and for L1 of $F_+ = -0.45$ and $F_\times = -0.30$, thus both detectors are roughly equally sensitive to the two GW polarizations.

We injected the signals into zero noise. The zero-noise analysis results will be similar to those obtained from the average results of multiple identical injections in different Gaussian noise realisations. The simulated signal is recovered using the LIGO Livingston and Hanford detectors with sensitivities matching those achieved in the second observing run (O2) [38]. A low frequency cut-off of 20Hz was used for likelihood evaluations, this frequency is also used as the reference frequency when defining all frequency dependent parameters such as θ_{JN}. Both the LIGO Livingston and Hanford sensitivities improved prior to the third observing run [173] and are expected to improve further prior to the fourth observing run (O4) [4]. The results presented in this work are unlikely to be affected significantly by these changes and therefore we expect the main conclusions to be valid for O4 and beyond.

The SNR of the signal is fixed to be 20, corresponding to a moderately loud signal for aLIGO and AdV observations [4]. This sets the distance to $d_L = 223\,\mathrm{Mpc}$. The simulated SNR in the two detectors is 16.2 in L1 and 11.7 in H1. The simulated precession SNR in each of the detectors is 3.7 and 3.4 respectively, giving a network precession SNR of 5.0. Thus, we expect that precession will be clearly observable in this signal.

3.3.b Parameter Estimation Techniques

We will adopt a parameter estimation methodology that uses matched filtering with phenomenological gravitational waveforms and Markov Chain Monte Carlo (MCMC) techniques to sample the posterior.

We begin by introducing the matched filtering formalisation for parameter estimation. We assume that the time series received from the GW detectors can be decomposed as a sum of the GW signal, $h(t)$, plus noise, $n(t)$, which is assumed stationary and Gaussian with zero mean,

$$d(t) = h(t) + n(t). \tag{3.3}$$

Under the assumption of Gaussian noise, the probability of observing data d given a signal $h(\boldsymbol{\lambda})$ parameterised by $\boldsymbol{\lambda} = \{\lambda_1, \lambda_2, ..., \lambda_N\}$, otherwise known as the likelihood, is [174],

$$p(d|\boldsymbol{\lambda}) \propto \exp\left(-\frac{1}{2}\langle d - h(\boldsymbol{\lambda})|d - h(\boldsymbol{\lambda})\rangle\right), \tag{3.4}$$

where $\langle a|b\rangle$ denotes the inner product between two waveforms a and b and is defined in Eq. 2.30.

The posterior probability density function can then be computed through a simple application of Bayes' theorem,

$$\begin{aligned} p(\boldsymbol{\lambda}|d) &= \frac{p(\boldsymbol{\lambda})p(d|\boldsymbol{\lambda})}{p(d)}, \\ &\propto p(\boldsymbol{\lambda}) \exp\left(-\frac{1}{2}\langle d - h(\boldsymbol{\lambda})|d - h(\boldsymbol{\lambda})\rangle\right), \end{aligned} \tag{3.5}$$

where $p(\boldsymbol{\lambda}|d)$ is the posterior distribution for the parameters λ, $p(\boldsymbol{\lambda})$ is the prior probability distribution where $\int p(\boldsymbol{\lambda})d\boldsymbol{\lambda} = 1$, and $p(d)$ is the marginalised likelihood where $p(d) = \int p(\lambda_i)p(d|\lambda_i)d\lambda_i$. Posterior distributions for specific parameters can then be found by marginalising over all other parameters,

$$p(\lambda_i|d) = \int p(\boldsymbol{\lambda}|d)d\lambda_1...d\lambda_{i-1}d\lambda_{i+1}...d\lambda_N. \tag{3.6}$$

In the idealised situation of zero noise, Eq. (3.4) has a maximum at $h(\boldsymbol{\lambda}) = h(\boldsymbol{\lambda}_0)$. However, as can be seen in Eq. (3.5) the posterior also includes priors, this means that, as well as effects due to noise, certain priors may cause the maxima to be deflected away from $h(\boldsymbol{\lambda}) = h(\boldsymbol{\lambda}_0)$. This would then lead to Eq. (3.6) recovering a biased posterior. In this work, we consider the effect of three closely related priors,

- *Global*: the prior used during the parameter estimation analysis. This reflects our prior belief before observing any data,
- *Conditioned*: the global prior conditioned upon the posterior distributions of other parameters from the same analysis. For example since $\chi_{\rm eff}$ and $\chi_{\rm p}$ are correlated, any informative measurement of $\chi_{\rm eff}$ modifies our prior beliefs about $\chi_{\rm p}$. This prior has been used in previous LVC publications, see e.g. [38],
- *Informed*: the global prior conditioned upon the posterior distributions from a different analysis. Here, we use this to inform our expectations of the degree of precession given the results from a non-precessing analysis. See Section 3.6 for details.

3.3.c Parameter recovery

We performed parameter estimation on the signal using the `LALInference` [91] and `LALSimulation` libraries within `LALSuite` [175]. Parameter recovery was performed with the `IMRPhenomPv2` model [172, 73], which matches the simulated signal to remove any systematic error caused by waveform uncertainty, and the corresponding `IMRPhenomD` aligned-spin waveform model [76, 75], which does not include any precession effects. Additionally, all analyses used exactly the same

	Simulated	Median		maxL	
	Precessing	Non-Precessing	Precessing	Non-Precessing	Precessing
Total mass $M/M_\odot$	40.0	40^{+3}_{-2}	40^{+4}_{-2}	40.161	40.507
Chirp mass $\mathcal{M}/M_\odot$	16.22	$16.5^{+0.3}_{-0.2}$	$16.3^{+0.3}_{-0.3}$	16.459	16.113
Mass ratio q	2.0	$1.8^{+0.8}_{-0.7}$	$1.9^{+0.9}_{-0.7}$	1.895	2.191
Inclination angle $\theta_{JN}/^\circ$	60.0	110^{+50}_{-100}	120^{+40}_{-90}	30.0	40.0
Precession phase $\phi_{JL}/^\circ$	45.0	–	200^{+100}_{-200}	–	80.0
Effective aligned spin, $\chi_{\rm eff}$	0.0	$0.044^{+0.099}_{-0.084}$	$-0.005^{+0.098}_{-0.092}$	0.06	−0.011
Effective precessing spin, $\chi_{\rm p}$	0.4	–	$0.5^{+0.4}_{-0.3}$	–	0.554
Right ascension RA/rad	1.88	3^{+3}_{-3}	3^{+3}_{-3}	1.418	1.325
Declination DEC/rad	1.19	$0.2^{+1.0}_{-1.2}$	$0.2^{+1.0}_{-1.2}$	1.229	1.221
Luminosity distance d_L/Mpc	223	500^{+200}_{-200}	400^{+200}_{-200}	451.834	372.706
Network SNR ρ	20.0	$19.3^{+0.1}_{-0.2}$	$19.7^{+0.2}_{-0.2}$	19.52	19.936
Precessing SNR $\rho_{\rm p}$	5.05	–	4^{+2}_{-2}	–	4.649

Table 3.1: Table showing the simulated and inferred parameters for the "standard" injection when recovered by a non-precessing (`IMRPhenomD`) and a precessing (`IMRPhenomPv2`) waveform model. We report the median values along with the 90% symmetric credible intervals and the maximum likelihood (maxL) value.

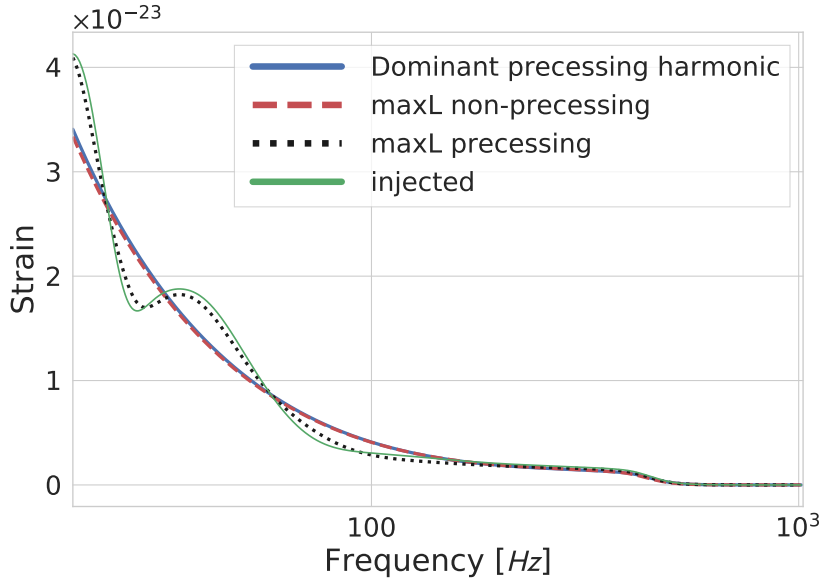

Figure 3.2: Comparison of the simulated precessing (green), non-precessing maximum likelihood (red), precessing maximum likelihood (black) and dominant precessing harmonic (blue) waveforms as a function of frequency. Waveforms are projected onto the LIGO Hanford detector.

priors as those used in the LIGO-Virgo discovery papers, for details, see Appendix B.1 of [38]. All post-processing was handled by the `PESummary` python package (see Chapter 8).

Tab. 3.1 summarises the key results for the standard configuration. All uncertainties are the 90% symmetric credible intervals.

We begin by comparing the overall differences between parameter recovery with the precessing, `IMRPhenomPv2`, and non-precessing, `IMRPhenomD`, runs. From the table, we see that the maximum likelihood SNR for the non-precessing model is, as expected, lower than for the precessing waveforms. This can be easily understood from the two-harmonic approximation. Since the precessing waveform is well approximated by the sum of two non-precessing harmonics, we would expect the non-precessing recovery to accurately recover the more significant of these two. If that were the case, the we would expect that,

$$\rho_D^2 \approx \rho^2 - \rho_{\rm p}^2, \tag{3.7}$$

and this is indeed the case, as $\rho_D = 19.52$, $\rho = 19.94$ and the recovered power in the second harmonic is $\rho_{\rm p} = 4.6$. Furthermore, we see that the recovered waveforms confirm this expectation: the recovered waveform when we include precession matches well with the simulated signal, while the non-precessing run recovers a waveform that matches the dominant harmonic, as show in Fig. 3.2.

We first consider the accuracy with which the masses and (aligned) spins are

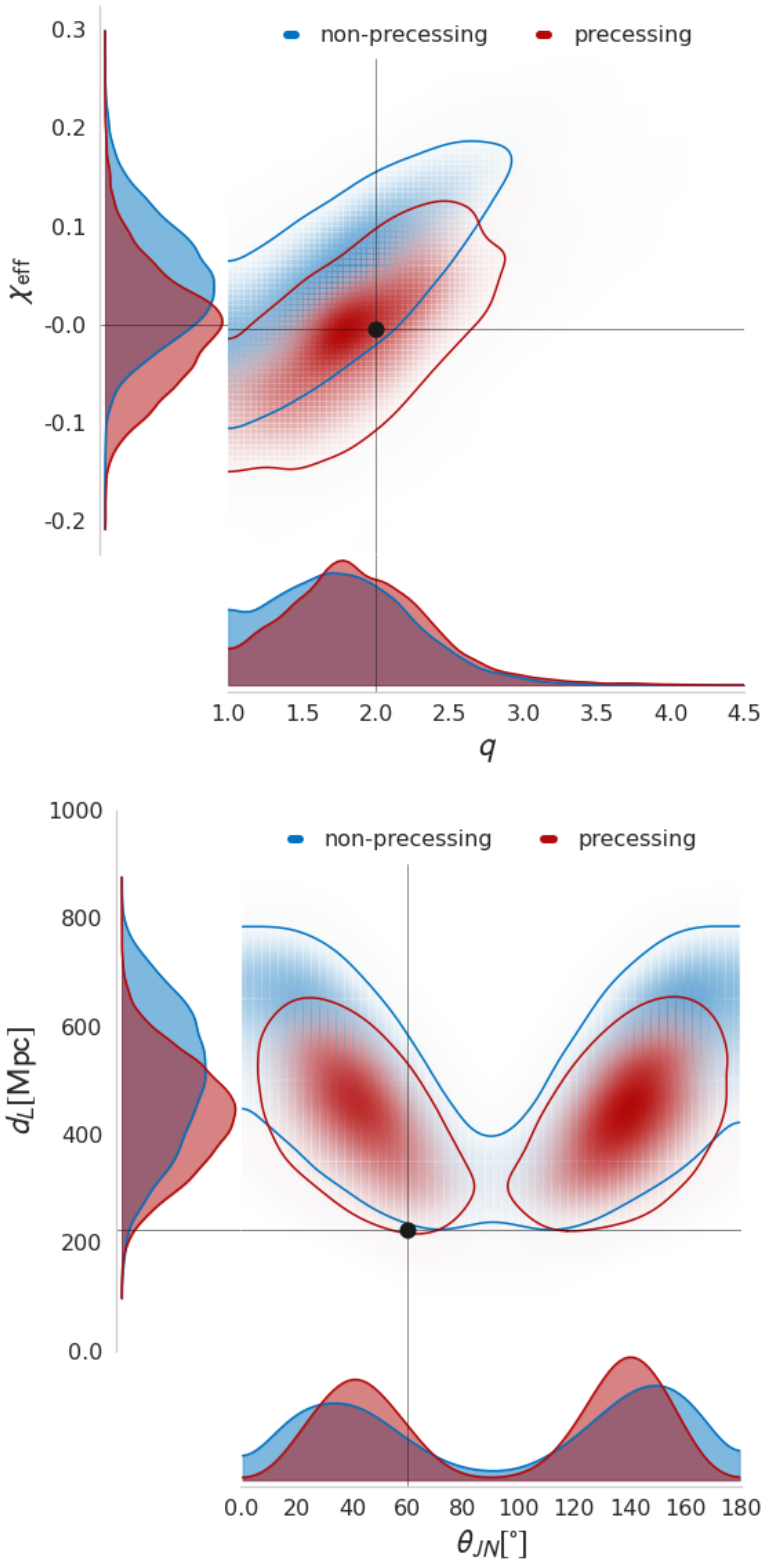

Figure 3.3: 2d contour comparing q–$\chi_{\rm eff}$ (top) and distance–inclination (bottom) degeneracies when precession effects are included. Contours show the 90% confidence interval. Bounded two-dimensional KDEs are used for estimating the joint probability density. The black circle indicates the simulated values.

recovered. As expected, the chirp mass of the system is well recovered, in that it matches the simulated value with only a 2% uncertainty, which remains constant for both precessing and non-precessing runs. As is well known, there is a degeneracy between mass-ratio and spin, particularly during the inspiral part of the waveform [61, 62, 157, 158], which leads to significant uncertainty in both parameters. In Fig. 3.3 we show the recovery of the mass ratio and spin, for both precessing and non-precessing runs. When the model used to recover includes precession effects, the peak of the posteriors is located close to the simulated value ($\chi_{\rm eff} = 0$ and $q = 2.0$) and, while the degeneracy leads to significant uncertainty in both parameters, the mass-ratio distribution is clearly peaked away from $q = 1$. Interestingly, when we recover with a non-precessing waveform model, the inferred *aligned* spin component is systematically offset, with a peak at $\chi_{\rm eff} \approx 0.05$. This can be understood by recalling that precession induces a secular drift in the phase evolution of the binary, and this can be mimicked by a change in the value of the aligned spin (see Chapter 2 and Ref [54]). This discrepancy has not been seen in LIGO/Virgo observations [38] as we have not observed any systems with significant $\rho_{\rm p}$ (see Fig. 3.1). We investigate this further in Sec. 3.4.d, where we study the effect of varying the mass ratio.

For non-precessing binaries, it is generally not possible to accurately recover the distance and orientation of the source, due to a well known degeneracy (see e.g., Ref. [99] for details), although the observation of higher signal harmonics can break this degeneracy through an independent measurement of the source inclination [60, 99, 176]. Similarly, the observation of precession can break this degeneracy [177]. Precession causes an oscillation of the orbital plane leading to a time-dependence of the orientation of the orbital plane relative to the line of sight. Equivalently, in the two-harmonic picture, precession leads to the observation of a second harmonic and, consequently, additional constraints on the binary orientation as the amplitudes of the harmonics depend upon the viewing angle. In Fig. 3.3, we show the inferred two-dimensional distance and inclination posteriors for the precessing and non-precessing runs. As expected, the precessing run constrains the source to be away from face-on, while the non-precessing run simply returns the prior. However, even with observable precession, the simulated distance and orientation are not accurately recovered — a significant fraction of the posterior support is for a system at a greater distance and oriented closer to face-on. We will see how these measurements improve with stronger precession in Sec. 3.4.b.

The sky location of the source is not well recovered. The analysis was performed with only the two LIGO detectors, and therefore we expect to recover the source restricted to a ring on the sky, which corresponds to a fixed time delay between the detectors [95, 97]. The location along the ring cannot be well constrained and, as expected the inferred location is preferentially associated with sky positions where the detector network is more sensitive. Thus, while the simulated sky location is within the 90% region, it is not at or close to the peak. This impacts the recovery

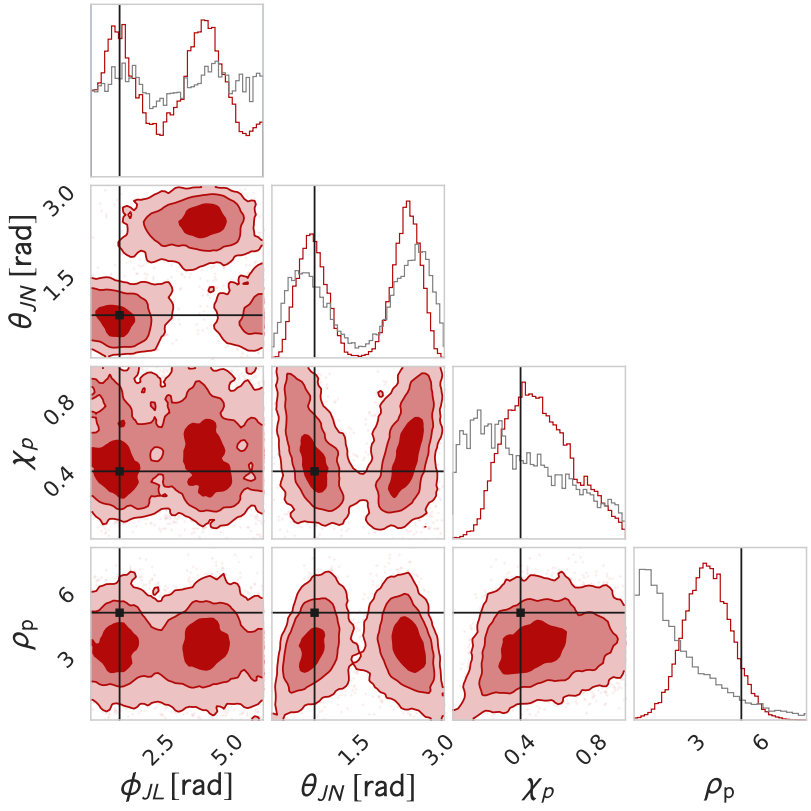

Figure 3.4: A corner plot showing the recovered values of binary orientation θ_{JN}, precessing spin $\chi_{\rm p}$, precession phase ϕ_{JL} and precession SNR $\rho_{\rm p}$. Shading shows the 1σ, 3σ and 5σ confidence intervals. Black dots show the simulated values. The grey histograms show the *informed* prior, see Sec. 3.6. There is a clear correlation between the binary orientation and inferred precession spin, with signals which are close to face on ($\cos\theta \approx \pm 1$) having larger values of precessing spin, while those which are more inclined having less precessing spin. The precession SNR only weakly correlated with $\chi_{\rm p}$.

of the distance, with the signal being recovered at larger distances, although the simulated distance remains within the 90% range. In Section 3.4.g, we show results from a set of runs with varying sky location, and verify that at sky locations where the network is more sensitive, the distance posterior is more consistent with the simulated value.

Lastly, we turn to measurement of precession. In Fig. 3.4 we show the recovered distributions for binary orientation, θ_{JN}, precessing spin $\chi_{\rm p}$, initial precession phase, ϕ_{JL}, and precession SNR, $\rho_{\rm p}$. There is a clear correlation between the inferred orientation and $\chi_{\rm p}$, with binaries that are more inclined having lower values of $\chi_{\rm p}$. Neither of these quantities are directly observable, it is only the amount of observable precession in the system, encoded by $\rho_{\rm p}$, that can be measured. Thus the orientation and spin must combine to give the right amount of power in precession, and we see that this is the case — there is little correlation between the recovered values of

$\rho_{\rm p}$ and the precessing spin $\chi_{\rm p}$. The inferred value of the precessing spin $\chi_{\rm p}$ and precession SNR $\rho_{\rm p}$ are both consistent with the simulated values. Specifically, the signal has $\chi_{\rm p} = 0.4$ and this is consistent with the recovered value, although the posterior distribution is broad, with support over essentially the entire range from 0 to 1. The precession SNR peaks well away from zero, giving clear indication of precession in the system. However, the peak of the distribution occurs at 3.5, while the simulated value is 5.0. We have deliberately chosen an event with significant observable precession. Only a small fraction of the parameter-space volume leads to such significant precession as shown by the *informed* prior on Fig. 3.4. This is calculated by estimating the allowed values of $\rho_{\rm p}$ conditioned on the measurements from a non-precessing analysis. See Sec. 3.6 for further details.

The precession phase, ϕ_{JL}, while not measured with great accuracy, does show two peaks, which are consistent with the simulated value of 45° (0.8 rad). The precession phase can be inferred from the relative phase of the two precessing harmonics using Eq. (3.2), provided the binary orientation is well measured. There is a clear dependence with the binary orientation: if $\theta_{JN} < 90°$ then the peak is in ϕ_{JL} at the simulated value and if it is greater then ϕ_{JL} is offset by 180°, to compensate for the change in sign of the $\cos\theta_{JN}$ terms in Eq. (3.2).

3.4 Impact of Varying Parameters

We now look at the effect of varying individual parameters one at a time on the recovered posteriors, in particular focusing on the measurement of precession as described by the posterior distributions of $\rho_{\rm p}$ and $\chi_{\rm p}$. All subsequent one-dimensional investigations of the parameter space maintain a constant SNR (except for Sec.3.4.a where the effect of the SNR is investigated). This is achieved by varying the distance to the source.

Primary results presented in this section will be displayed in the form of violin plots. We show the $\chi_{\rm p}$ posterior distribution (left hand side, colored) compared to the global prior (right hand side, white) unless otherwise stated. We show the $\rho_{\rm p}$ posterior distribution as a single violin. Horizontal grey lines show the 90% symmetric credible interval. Horizontal red lines show the simulated value. A solid black line corresponds to the $\rho_{\rm p} = 2.1$ threshold. Bounded KDEsare used for estimating the probability density. We use the same 2d contour plots and multi-dimensional corner plots as described in Sec. 3.3.c. Plots were generated with the PESummary python package (see Chapter 8).

3.4.a SNR

We start with the fiducial run configuration described above and vary the SNR of the simulated signal.

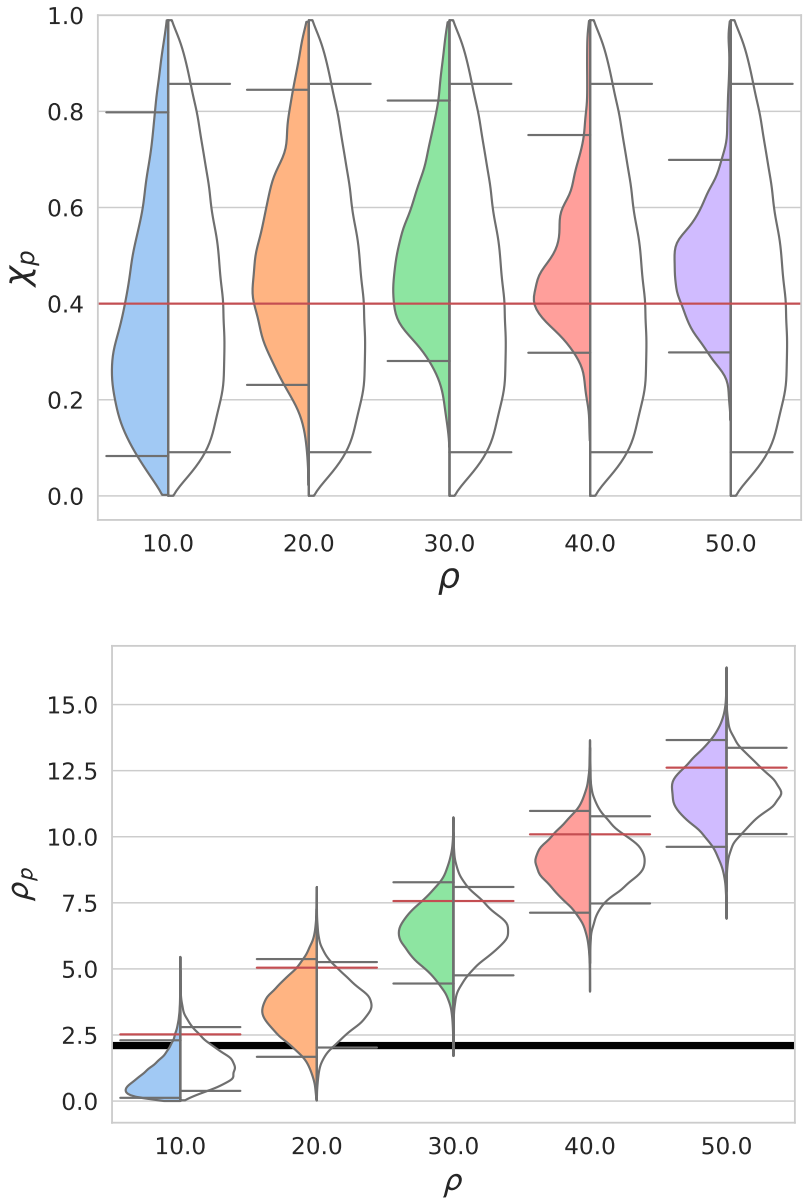

Figure 3.5: Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (top) and ρ_p compared to a non-central χ distribution with 2 degrees of freedom and non-centrality equal to the median of the ρ_p distribution (bottom). Distributions are plotted for varying SNR. Parameters other than the SNR of the signal match the "standard injection" (see Table 3.1).

In the strong-signal limit, where the likelihood surface can be well approximated by a multivariate gaussian, it is well known that the accuracy with which parameters can be measured is generally inversely proportional to the SNR [60, 61]. However, this is not always the case due to, for example, degeneracies between parameters (see Ref. [178] for a discussion of the limits of this approximation).

Fig. 3.5 shows that as the SNR of the simulated signal increases, the accuracy and precision of the inferred $\chi_{\rm p}$ posterior distribution improves. As expected the width of the 90% credible interval decreases approximately linearly with increasing SNR. The improvement in the $\chi_{\rm p}$ posterior distribution can be mapped to a linear increase in $\rho_{\rm p}$.

When the simulated signal has low SNR ($\rho = 10$), the recovered $\chi_{\rm p}$ posterior distribution resembles the prior, implying that there is no information about precession in the data. For this case, $\rho_{\rm p}$ matches the expected distribution in the absence of any measurable precession — a χ distribution with 2 degrees of freedom. As the SNR increases (ρ = 20-30), the 5th percentile of the the $\rho_{\rm p}$ distribution is comparable or greater than the $\rho_{\rm p} = 2.1$ threshold. This maps to the $\chi_{\rm p}$ posterior distribution removing all support for near-zero $\chi_{\rm p}$ ($\chi_{\rm p} \lesssim 0.1$). For larger SNRs ($\rho > 40$), the entire $\rho_{\rm p}$ distribution is greater than the 2.1 threshold. This implies significant power from precession. For these cases, we remove support for maximal precession $\chi_{\rm p} \sim 1$.

As expected we find good agreement between $\rho_{\rm p}$ and a non-central χ distribution with 2 degrees of freedom and non-centrality equal to the inferred power in the second harmonic (median of the $\rho_{\rm p}$ distribution).

3.4.b In-plane spin components

We now look at the effect of varying the amount of precession in the system, varying $\chi_{\rm p}$ from 0 to 1 in steps of 0.25. At $\chi_{\rm p} = 1$ we have maximal spin, all in the plane of the binary. The inferred values of precessing spin and precession SNR are shown in Fig. 3.6. We observe, as expected, that increasing the in-plane spin leads to an increase in the magnitude of precession effects observable in the system. With zero precessing spin, there is no evidence for precession in the system; the recovered $\chi_{\rm p}$ is consistent with zero[2]. Similarly, there is no support for significant precession SNR, with $\rho_{\rm p}$ constrained near zero. As $\chi_{\rm p}$ increases, the amount of precession in the system grows and the measurement of $\chi_{\rm p}$ becomes both more accurate and more precise. Fig. 3.6 shows the relationship between $\rho_{\rm p}$ and $\chi_{\rm p}$, and a larger value for $\rho_{\rm p}$ enables a better measurement for $\chi_{\rm p}$.

Fig. 3.7 shows how the inferred mass ratio–aligned spin and distance–orientation contours change as the magnitude of the in-plane spins change. When there is no observable precession in the system, there is a clear degeneracy in both cases. However, as precession effects become stronger the degeneracy between both pairs

[2]We do not expect the $\chi_{\rm p}$ posterior to contain $\chi_{\rm p} = 0$ as there is no prior support there, however the posterior is relatively well constrained at low precession.

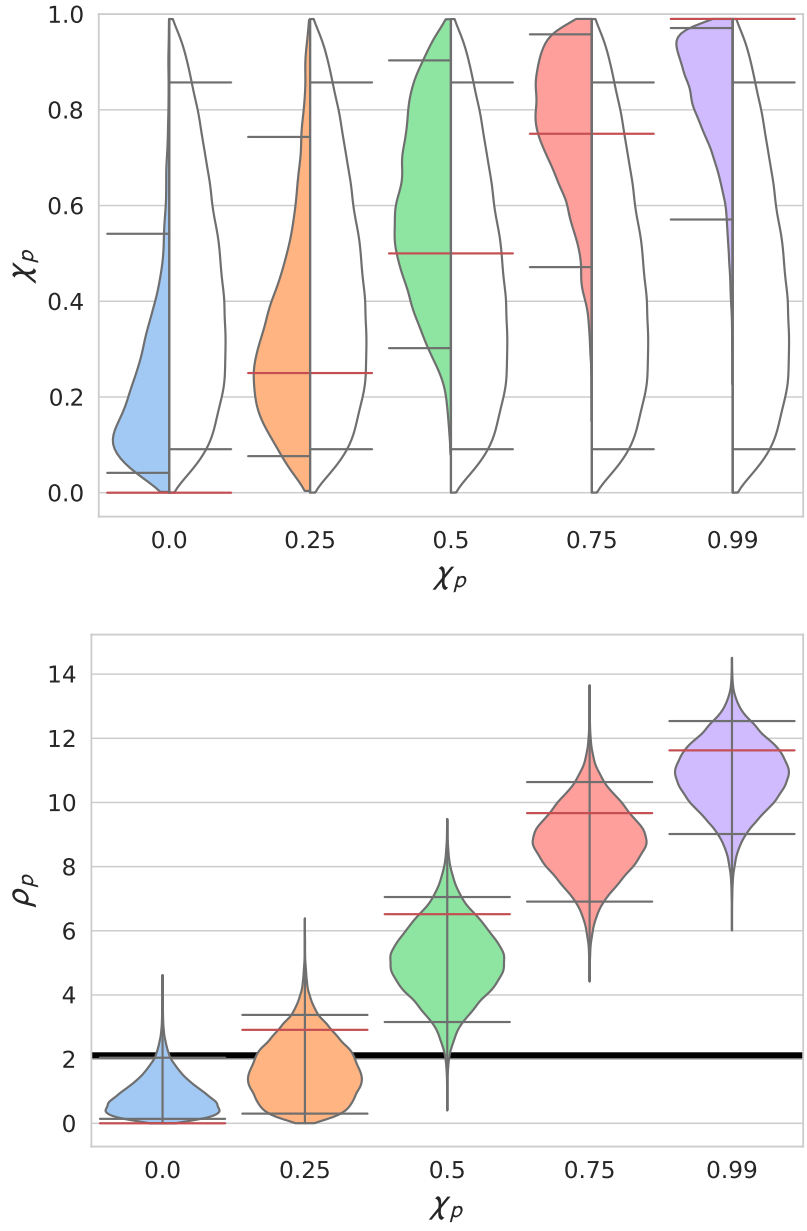

Figure 3.6: Violin plots showing the recovered posterior distributions distributions for χ_{p} compared to its prior (top) and ρ_{p} (bottom). Distributions are plotted for varying χ_{p}. Parameters other than χ_{p} match the "standard injection" (see Table 3.1)

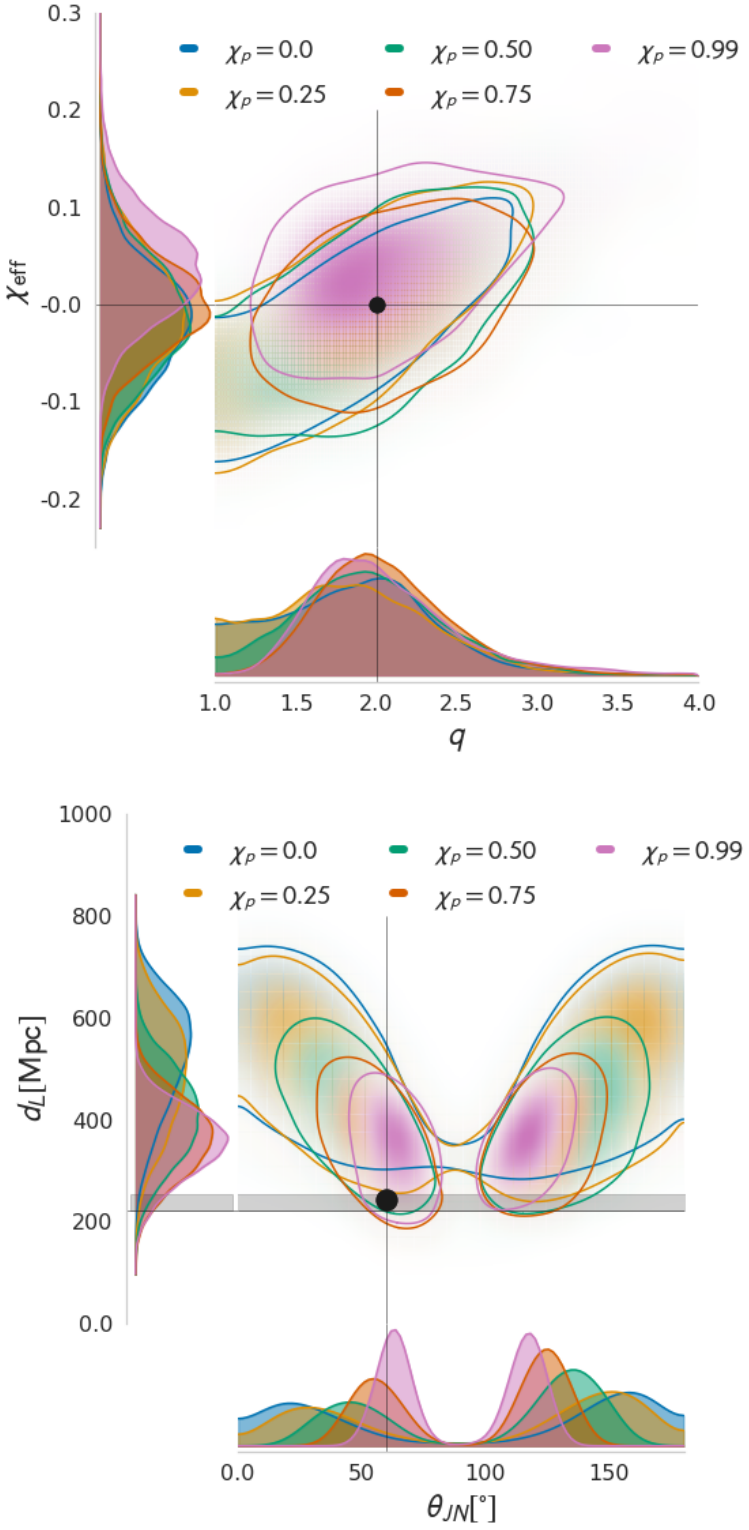

Figure 3.7: Two dimensional posteriors for (top) mass ratio and aligned spin, $\chi_{\rm eff}$, (bottom) binary orientation and distance. Contours show the 90% confidence interval. Bounded two-dimensional KDEs are used for estimating the joint probability density. The black circle with corresponding horizontal and vertical lines indicates the simulated values. For the simulated distance, a solid horizontal band indicates the maximum and minimum simulated values.

of parameters is broken. If $\rho_{\rm p}$ is small then this can be explained by both a small amount of precession observed at almost any inclination angle, or a large $\chi_{\rm p}$ observed close to face on, as seen in Fig. 3.4. Since precession effects are not strong enough to provide an accurate measurement of the orientation, the degeneracy between distance and θ_{JN} persists. When $\rho_{\rm p}$ clearly excludes small values, there is *no support* for close to face-on signals, allowing a more precise measurement of the inclination angle θ_{JN}, breaking the degeneracy with distance.

Stronger precession also allows for improved measurement of the mass ratio. The opening angle β, and consequently the precession parameter $\bar{b}$, increases as the mass-ratio is increased, as can be seen from Eq. (1.1). Thus, when strong precession effects are observed, the signal is inconsistent with an equal mass system. In addition, the difference in frequency between the two leading precession harmonics depends upon the mass-ratio (see Chapter 2), and this may also improve our measurement of q. This can also be seen from the precession dynamics, where the precession rate of L around J, $\dot{\alpha}$, depends the mass ratio, and the number of observable precession cycles corresponds to improved accuracy in the measurement of the mass ratio [63].

As $\chi_{\rm p}$ is increased, the peak of the recovered $\rho_{\rm p}$ distribution is closer to the simulated value. This is likely due to a better measurement of the binary orientation as shown in Fig. 3.7.

3.4.c Inclination

It is well known that the inclination angle will affect our ability to measure precession, as outlined in the discussion in Sec. 1.2. In particular, from Eq. (3.2) we see that in the two-harmonic approximation the second harmonic vanishes when $\theta_{\rm JN} = 0^\circ$ or 180°. In this section we consider the effect of changing the orientation of our standard configuration, which allows us to quantify how it will manifest in realistic LIGO-Virgo signals. A related study has looked at the effect at higher mass ratios [169].

The effect of varying $\theta_{\rm JN}$ is shown in Fig. 3.8. For binaries where the total angular momentum is nearly aligned with the line of sight, precession effects are not observable, as is clear from both the $\rho_{\rm p}$ and $\chi_{\rm p}$ posteriors. It is not until $\theta_{JN} \geq 40^\circ$ that we begin to be able to measure precession. Although the accuracy of the measurement clearly improves as we increase θ_{JN}, the uncertainty in the measurement of $\chi_{\rm p}$ remains large and even at $\theta_{JN} = 90^\circ$ the posterior is very broad. This can be understood by considering the degeneracies shown in Fig. 3.4 for the standard signal and in Fig. 3.9 for the $\theta_{JN} = 90^\circ$ signal. In both cases, the measured quantity, $\rho_{\rm p}$, is relatively well constrained but neither the binary orientation nor $\chi_{\rm p}$ are accurately measured. The observed precession is consistent with both a highly inclined system with lower precessing spin (i.e., low $\chi_{\rm p}$ and large θ_{JN}) or by a less inclined system with higher precessing spin (i.e., high $\chi_{\rm p}$ and small θ_{JN}). Both of

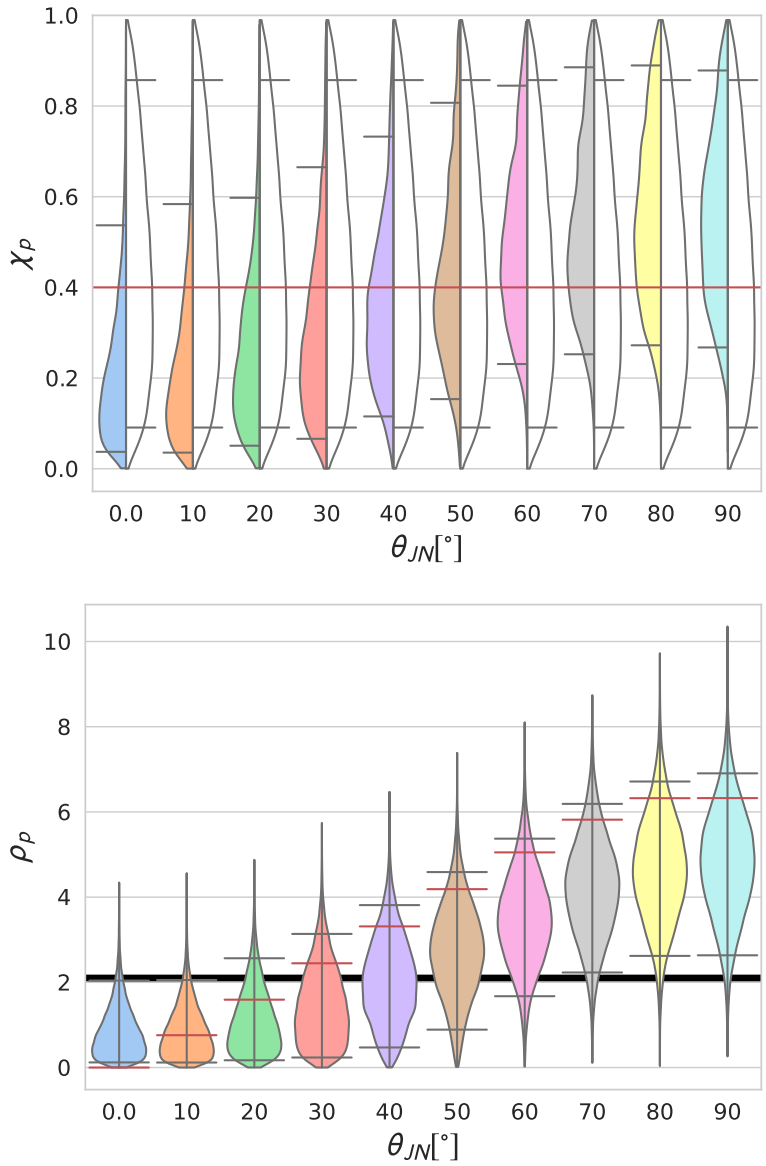

Figure 3.8: Violin plots showing the recovered posterior distributions distributions for $\chi_{\rm p}$ compared to its prior (top) and $\rho_{\rm p}$ (bottom). Distributions are plotted for varying θ_{JN}. Parameters other than θ_{JN} match the "standard injection" (see Table 3.1)

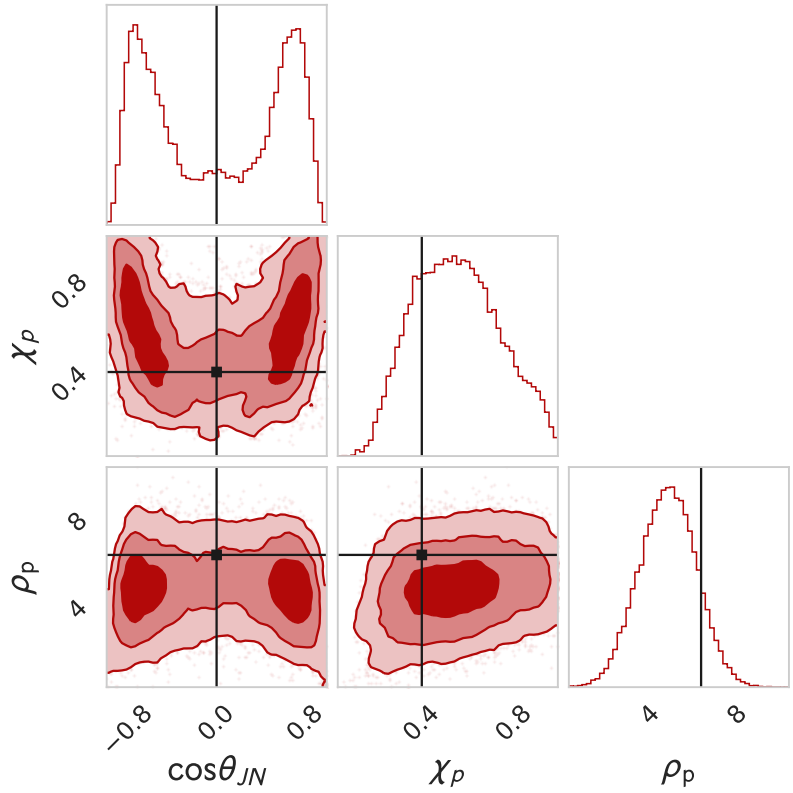

Figure 3.9: A corner plot showing the recovered values of binary orientation θ_{JN}, precessing spin χ_p and precession SNR ρ_p for a system simulated at edge on. Shading shows the 1σ, 3σ and 5σ confidence intervals. Black dots show the simulated values, We see the strong correlation between θ_{JN} and χ_p reflecting the measurement of a certain ρ_p

these will produce similar observable effects in the waveform.

This allows us to explain the measured posterior for χ_p. At low inclination the posterior is consistent with small values of χ_p. While we are unable to rule out large χ_p, there is limited support as it would require the system to be observed very close to face-on, otherwise precession effects become significant. At large values of θ_{JN}, when precession is clearly observable in the signal, $\chi_p = 0$ is excluded but the distribution remains broad and extends to $\chi_p = 1$.

3.4.d Mass ratio and aligned spin

Fig. 3.10 shows how the inferred precessing spin and precession SNR varies with the mass ratio of the system. As expected from the general considerations presented in Sec. 1.2, as the mass ratio increases, an in-plane spin on the larger BH leads to a larger opening angle and more significant precession effects. For near equal-mass systems ($q \lesssim 1.5$), the inferred χ_p posterior distribution resembles its prior, and there is not significant power in precession, as shown by the value of ρ_p. As the

mass ratio increases, the inferred power in precession also increases and for $q \gtrsim 2.5$, the 90% credible interval of the inferred $\rho_{\rm p}$ distribution is entirely above $\rho_{\rm p} = 2.1$. At this stage, precession is clearly identified and $\chi_{\rm p} \approx 0$ is clearly excluded. In addition, the maximum value of χ_p is also bounded away from maximal.

Fig. 3.11 shows how varying $\chi_{\rm eff}$ affects our ability to measure precession. A system with a large negative $\chi_{\rm eff}$ results in a larger opening angle compared to an equivalent system with a large positive $\chi_{\rm eff}$. Thus, based upon Eq. (1.1), we expect the observable impact of precession to be greater for negative values of $\chi_{\rm eff}$ and smaller for positive values. The results are consistent with this expectation, in that the precession SNR decreases with increasing $\chi_{\rm eff}$ and the width of the recovered χ_p distribution increases. However, for the $\chi_{\rm eff} = 0.4$ analysis, we find that the range of $\chi_{\rm p}$ is restricted, with both $\chi_{\rm p} = 0$ and $\chi_{\rm p} = 1$ excluded. This is *not* due to the measurement of precession, but is actually due to the measured non-zero aligned-spin component.

A non-zero measurement of $\chi_{\rm eff}$ forces $\chi_{\rm p} < 1$ as the primary and secondary spin magnitudes must be less than unity. For example, in the $\chi_{\rm eff} = 0.4$ analysis, we measure $\chi_{\rm eff} = 0.38^{+0.07}_{-0.07}$. Under the single spin assumption, this limits $\chi_{\rm p} < 0.95$. Similarly, since we are using prior distributions that are uniform in spin magnitude and orientation, the observation of a large aligned spin component leads to greater support for a large in-plane spin component. This is shown in Fig. 3.12, where we plot both the uninformed prior on the primary spin as well as the prior conditioned on $\chi_{\rm eff} = 0.4$, which removes all support for $\chi_p \approx 0$.

The $\chi_{\rm p}$ measurement for the $\chi_{\rm eff} = 0.27$ and 0.4 analyses are similar to the conditional prior but do restrict the lower $\chi_{\rm p}$ bound beyond prior effects. Although the distribution for $\rho_{\rm p}$ does extend to zero, it still peaks above $\rho_{\rm p} = 2.1$ indicating some evidence, although not particularly strong, for precession.

As we vary the mass ratio and aligned spin, the length of the waveform will change. In particular, the aligned spin and high mass ratio configurations produce longer waveforms than those with anti-aligned spins and equal masses [179]. In principle, this will impact the measurability of precession, as longer waveforms allow for a greater number of precession cycles in the detectors' sensitive band. For very short signals, with less than one precession cycle in band, the two leading harmonics are no longer orthogonal (or even approximately so), which make it more challenging to unambiguously identify the second harmonic. This is not an issue for the signals considered here, but does become important when we vary the mass of the binary in Section 3.4.e. With a greater number of precession cycles, we will also be able to more accurately measure the precession frequency (the frequency difference between the harmonics), which may improve the measurement of mass ratio [63]. However, it is still the precession SNR that determines the observability of precession. Finally, we note that changing the mass ratio and aligned spin will change the overall amplitude of the waveform. Since our study is performed at a *fixed SNR*, this simply leads to

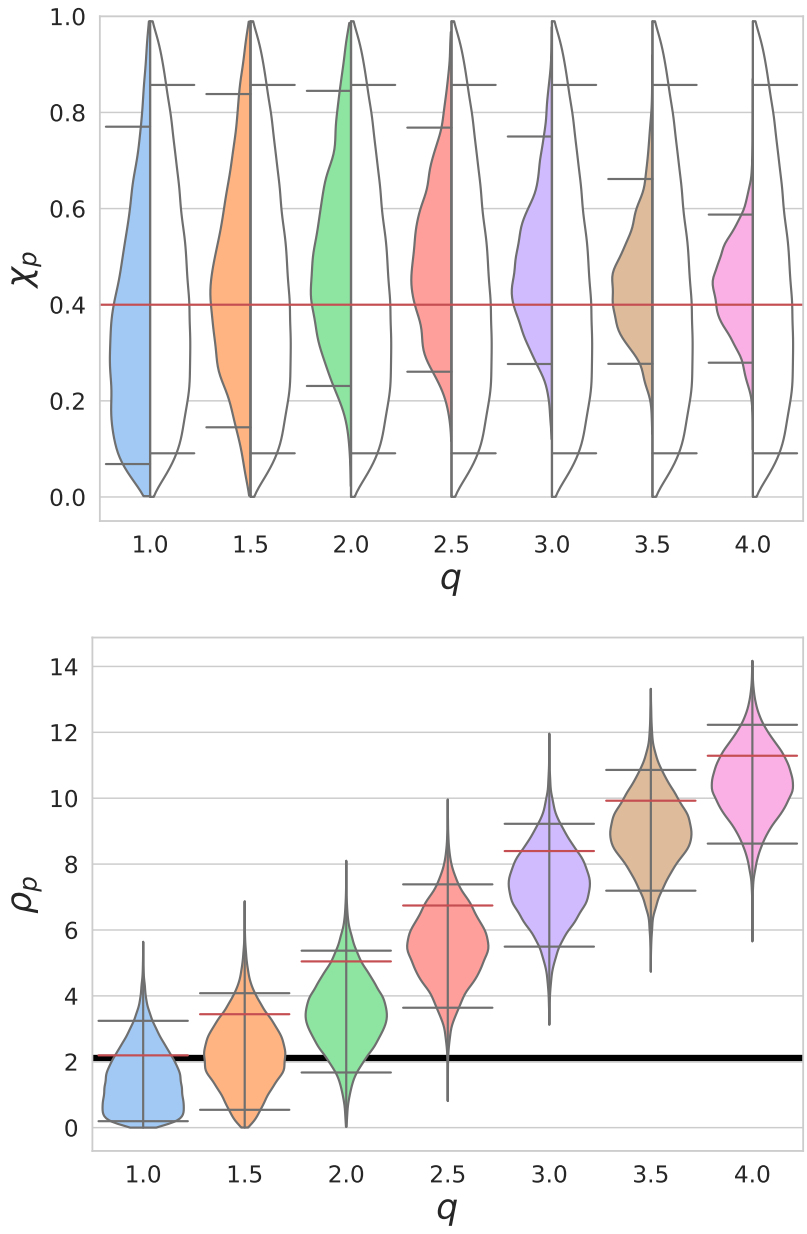

Figure 3.10: Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (top) and ρ_{p} (bottom). Distributions are plotted for varying mass ratio. Parameters other than the mass ratio of the signal match the "standard injection" (see Table 3.1).

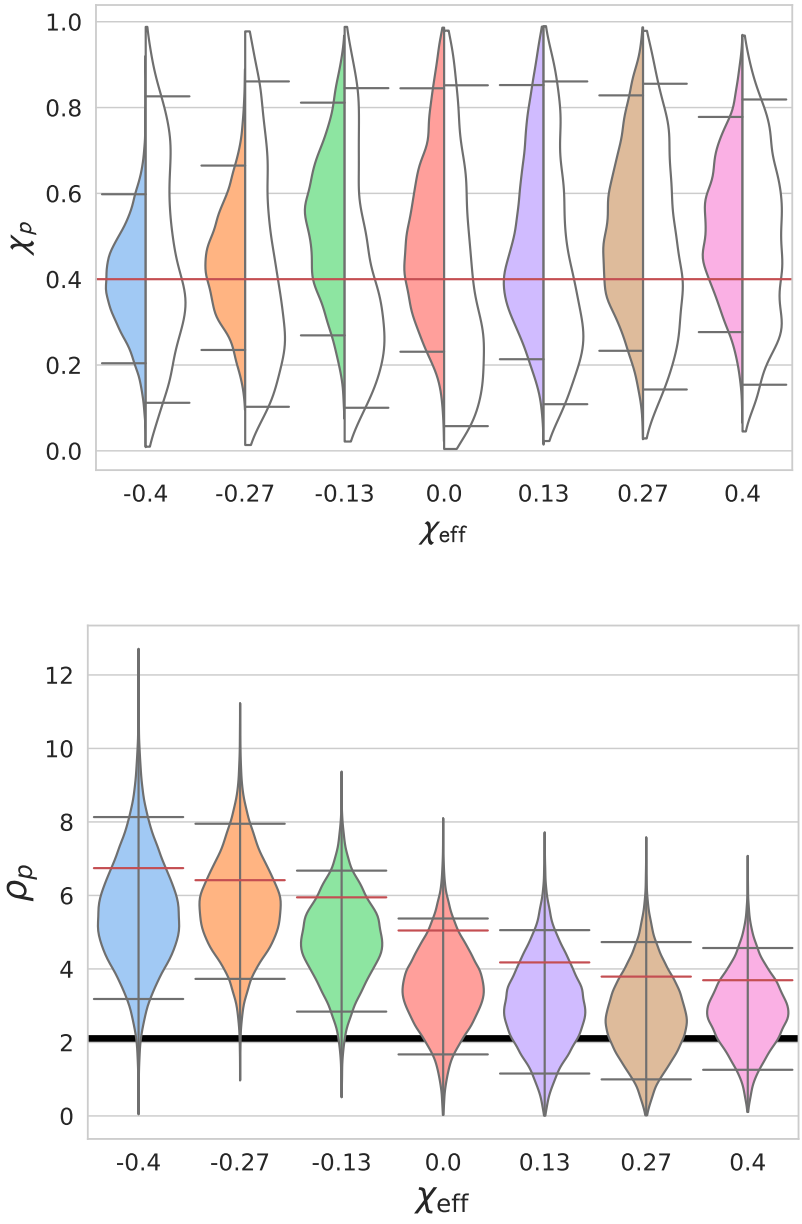

Figure 3.11: Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior conditioned on the $\chi_{\rm eff}$ and mass ratio posterior distributions (top) and $\rho_{\rm p}$ (bottom). Distributions are plotted for varying $\chi_{\rm eff}$. Parameters other than the $\chi_{\rm eff}$ of the signal match the "standard injection" (see Table 3.1).

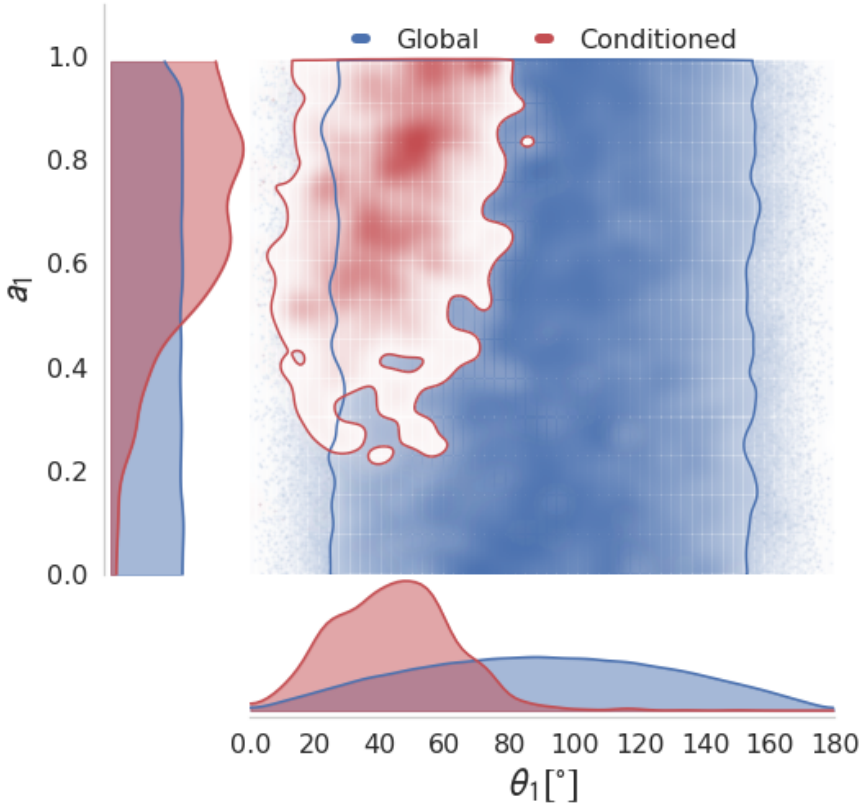

Figure 3.12: 2d contours showing the prior 90% credible interval over the primary spin magnitude and spin direction parameter space. Blue shows the global prior and red shows the global prior conditioned on the $\chi_{\rm eff} = 0.4$ mass ratio and $\chi_{\rm eff}$ posterior distributions

the signals being placed at a larger or smaller distance and therefore doesn't impact the results presented here.

3.4.e Total mass

We now vary the total mass of the system, keeping all other parameters including mass ratio fixed, in steps of $20\,M_\odot$. As before, we keep the SNR of the system constant at 20, so the higher mass systems are generated at a greater distance. The inferred distributions for $\chi_{\rm p}$ and $\rho_{\rm p}$ are shown in Fig. 3.13.

As the total mass of the source increases, the length of the waveform decreases, as does the number of precession cycles, with the number scaling approximately inversely to the total mass (see Eq. (45) of [54]). From the two-harmonic perspective, a small number of precession cycles leads to a large overlap between the harmonics. Specifically, for the $M = 100M_\odot$ system the overlap between the normalised harmonics is $\langle \hat{h}_0 | \hat{h}_1 \rangle = 0.77$, where $\hat{h} = h/|h|$ and the inner product is

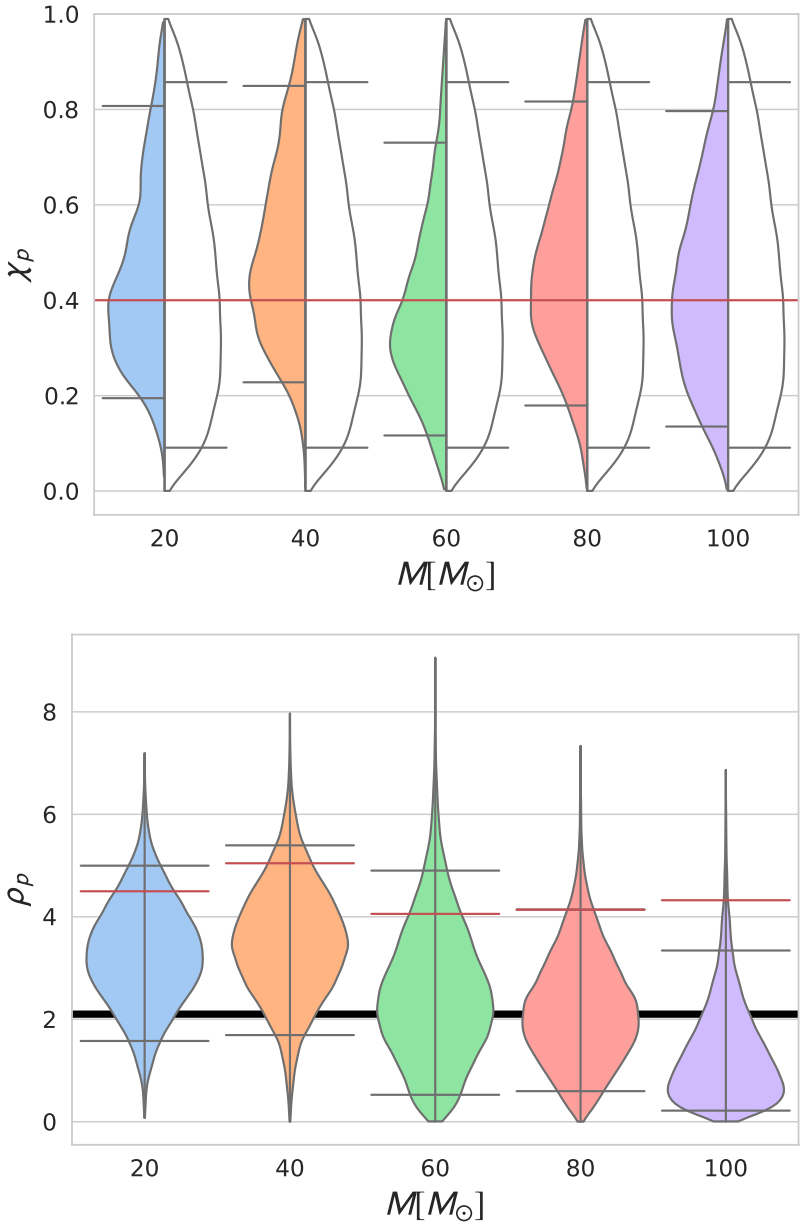

Figure 3.13: Violin plots showing the recovered posterior distributions distributions for $\chi_{\rm p}$ compared to its prior (top) and $\rho_{\rm p}$ (bottom). Distributions are plotted for varying total mass. Parameters other than the total mass of the signal match the "standard injection" (see Table 3.1)

defined in Eq. (2.30). At $M = 20M_\odot$, the harmonics are close to orthogonal with $\langle \hat{h}_0 | \hat{h}_1 \rangle = 0.15$. The opening angle doesn't change significantly, with $\bar{b} = 0.14$ at $M = 20M_\odot$ and $\bar{b} = 0.21$ at $M = 100M_\odot$.

At lower masses, $M \leq 40M_\odot$, while the precessing spin is not tightly constrained, it is clearly restricted to be non-zero and the precession SNR has essentially no support for $\rho_{\rm p} = 0$. For the $60M_\odot$ and $80M_\odot$ mergers, the precessing spin is still peaked close to the simulated value while $\rho_{\rm p}$ peaks above 2.1 showing evidence for observable precession, although both $\rho_{\rm p}$ and $\chi_{\rm p}$ distributions do extend to zero.

For the high-mass system, $M = 100M_\odot$, the $\chi_{\rm p}$ posterior more closely matches the prior and we are unable to exclude $\chi_{\rm p} = 0$. The inferred $\rho_{\rm p}$ distribution peaks close to zero, and is consistent with no precession, even though the precession SNR in the simulated signal is similar to the lower mass signals. This is likely due to the breakdown of the two-harmonic approximation for this short signal. In particular, for a high-mass system, the power orthogonal to the leading harmonic will depend sensitively upon the initial precession phase ϕ_{JL}. The fact that the recovered value of $\rho_{\rm p}$ is inconsistent with the simulated value may be due to this fact: the value of $\phi_{JL} = 45^\circ$ used in the simulation leads to maximal observable precession. Across the full parameter space there are very few configurations with significant precession, so this observation is dis-favoured by our priors. We explore the prior effects such as this in detail in Sec. 3.6.b.

3.4.f Polarization

The effect of changing the relative sensitivity to the two GW polarizations is clear from Eq. (3.2). Recalling that $\bar{b} = 0.11$ and $\theta_{JN} = 60^\circ$, we can express ζ (the ratio of the amplitudes of the two harmonics) as

$$|\zeta| = 0.15 \left| \frac{F_+ + 2iF_\times}{1F_+ + 0.8iF_\times} \right| ,$$

Thus, ζ, and consequently the imprint of precession on the waveform, will be maximized when the detector network is primarily sensitive to the $\times$ polarization and minimized when the network is sensitive to the + polarization. We can investigate this by varying the polarization angle of the simulated signal, in steps of 10° from the "standard" value of 40°. At $\psi = 40^\circ$, the sensitivity to the two polarizations is approximately equal, $|F_\times|/|F_+| = 0.9$. It is largest for $\psi = 20^\circ$ where $|F_\times|/|F_+| = 25$ and smallest for $\psi = 60^\circ$ where $|F_\times|/|F_+| = 0.04$. This leads to a variation in the precession SNR from $\rho_{\rm p} \approx 3$ to $\rho_{\rm p} \approx 7$.

In Fig. 3.14 we show the recovered posteriors for $\chi_{\rm p}$ and $\rho_{\rm p}$ for a set of runs where the precession is varied. The precession SNR varies in accordance with expectation — it is largest at $\psi = 20^\circ$, where the median of the posterior is at $\rho_{\rm p} = 6$ and there is no support for non-precessing systems, and smallest at 60° where the posterior

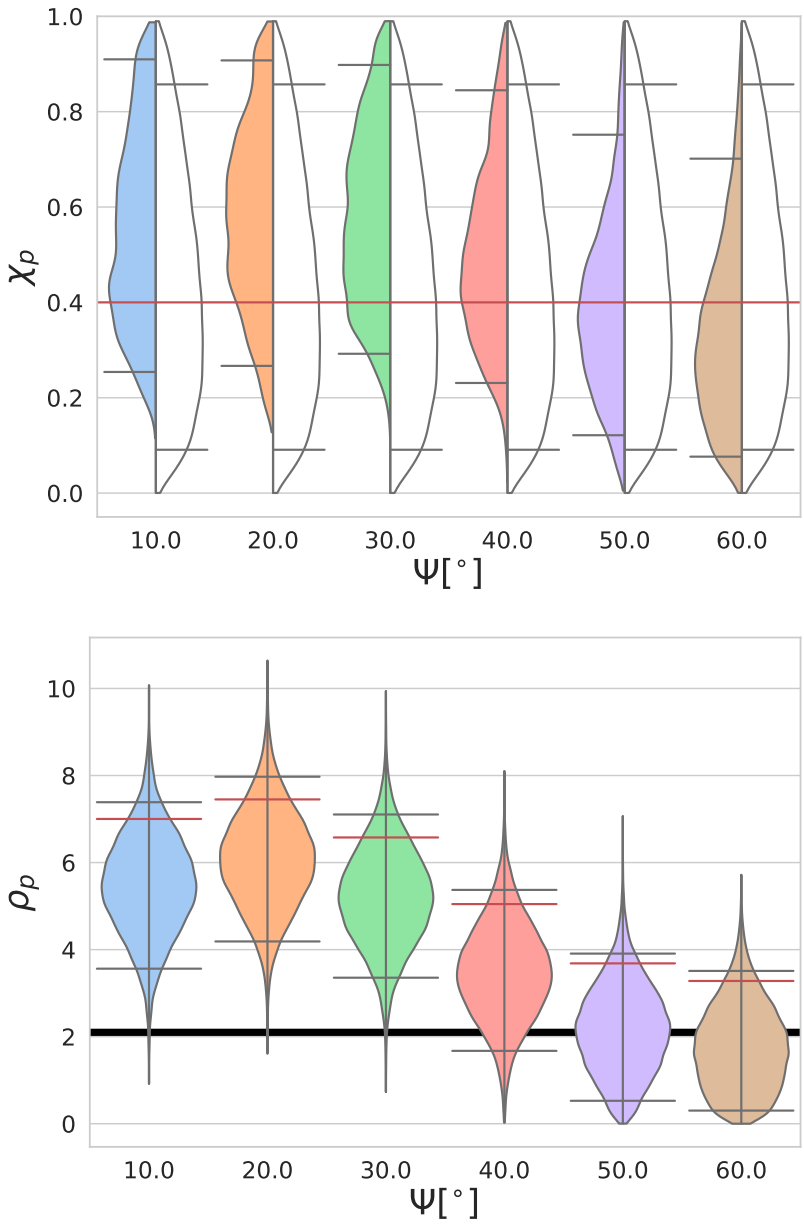

Figure 3.14: Violin plots showing the recovered posterior distributions distributions for $\chi_{\rm p}$ compared to its prior (top) and $\rho_{\rm p}$ (bottom). Distributions are plotted for varying ψ_J. Parameters other than ψ_J match the "standard injection" (see Table 3.1)

extends down to $\rho_{\rm p} = 0$. The amount of observable precesssion directly impacts the inferred distribution for $\rho_{\rm p}$. For the $\psi = 60°$ signal, the posterior for $\chi_{\rm p}$ is consistent with zero, or small in-plane spins, and large values are excluded. Meanwhile for $\psi = 20°$, $\chi_{\rm p} < 0.1$ is excluded while extremal in-plane spins are consistent with the observation.

It is well known that precession leaves a stronger imprint upon the $\times$ polarization. However, we are not aware of previous results showing how simply changing the polarization of the system can so dramatically change the observable consequences of precession — from being barely observable when the observed signal is primarily the + polarization to being strongly observed in $\times$. Using the two-harmonic approximation, we are able to straightforwardly predict this effect and then verify it with detailed parameter estimation studies.

3.4.g Sky Location

We performed a series of runs where we altered the sky location of the signal, keeping the masses and spins of the components fixed. We also maintained the binary orientation $\theta_{\rm JN} = 60°$, but varied the distance and polarization of the source to ensure that the SNR remained constant and that the relative contribution of the + and $\times$ polarizations was consistent with the standard run. Furthermore, sky locations were restricted to those for which the relative time of arrival between the Hanford and Livingston detectors remains the same (i.e., we were sampling from the nearly degenerate ring on the sky of constant time delays). Details of the runs are given in Tab. 3.2.

Table 3.2 shows that the inferred luminosity distance remains approximately constant despite the simulated luminosity distance varying by almost a factor of two. In addition, the recovered $\rho_{\rm p}$ distribution remains consistent with the "standard" injection. Fig. 3.15 shows that the inferred sky position of the source remains essentially unchanged, and consistent with locations of the detectors' greatest sensitivity. We note here that for this study we only considered the two detector LIGO network. Including VIRGO would likely have considerably improved the precision of the inferred sky location. We do not expect that this would affect any of the inferred physical parameters or any of the main conclusions in this work.

3.5 Relating $\rho_{\rm p}$ posteriors to Bayes Factors

An alternative method for identifying evidence for precession can be calculated within the Bayesian framework. We can calculate the Bayes factor, $\mathcal{B}$, by comparing the marginalized likelihoods (see Eq. (3.5)) from two competing hypotheses (A, B) [180],

$$\ln \mathcal{B} = \ln p(d_A) - \ln p(d_B). \tag{3.8}$$

Label	RA/rad	DEC/rad	$\psi/^\circ$	d_L/MPc	$\rho_{\rm p}$	d_L/MPc
A	0.31	0.92	320	370	5.02	480^{+130}_{-180}
B	0.80	1.15	345	320	5.09	470^{+140}_{-160}
C	1.31	1.22	10	280	5.11	450^{+150}_{-160}
D	1.88	1.19	40	220	5.05	430^{+160}_{-160}
E	6.11	0.21	40	310	5.09	440^{+150}_{-170}

Table 3.2: Table showing the simulated parameters for the sky location set (see Sec. 3.4.g). All other parameters match the "standard injection" (see Table 3.1). The recovered luminosity distance (far right column) is also shown.

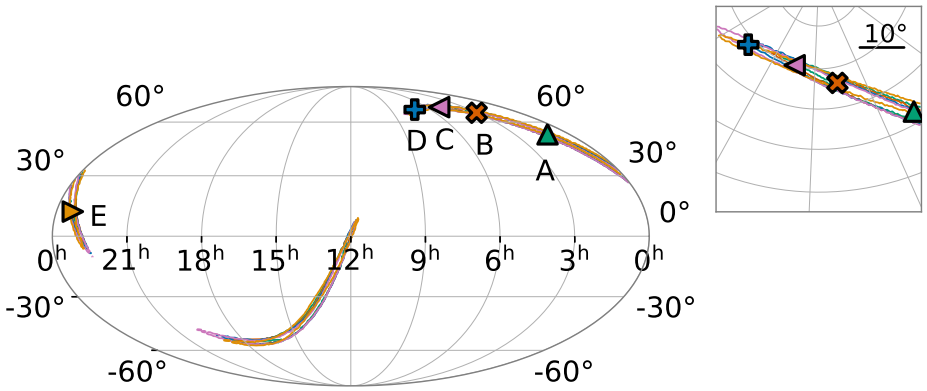

Figure 3.15: Skymap showing the different simulated sky positions, see Table 3.2. The solid lines show the 90% credible intervals and the markers show the simulated sky position. Their respective colors matches their corresponding credible intervals. We vary the distance and polarization of the source to ensure that the SNR remains consistent with the standard injection in Table 3.1.

Bayes factors have thus far been the gold standard for identifying evidence for precession within the GW community and have been used extensively in previous works, see e.g., Ref. [169].

In the same way that Bayes factors can be used to quantify evidence for precession, it is also possible to quantify the significance of a GW signal by calculating the Bayes factor for signal verses noise [90]. It has been shown that the log Bayes factor for signal versus noise scales approximately with ρ^2 [181]. Here, we investigate the relationship between the Bayes factor in favour of precession and the precession SNR $\rho_{\rm p}$. Both of these quantities have been used together in recent works when assessing the evidence for observable precession [45, 47, 169]

For a subset of the runs described in Section. 3.4.c, we reran the analysis using the aligned-spin waveform model `IMRPhenomD`. Bayes factors in favour of precession could then be calculated and compared to the derived $\rho_{\rm p}$ posterior distributions.

Fig. 3.16 shows an approximately linear relationship between the log Bayes factor ($\ln \mathrm{BF}$) and the square of the precession SNR ($\rho_{\rm p}^2$). This is expected given that the likelihoods recovered from the precessing waveform model will be larger than the

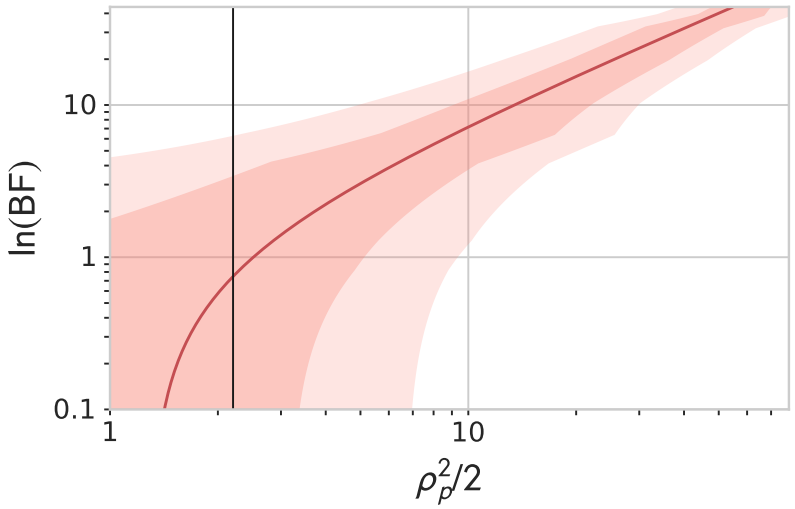

Figure 3.16: Plot comparing the Bayes factor in favour of precession to the inferred $\rho_{\rm p}$ distribution. Bayes factors were calculated by comparing the evidences for a precessing analysis and a non-precessing analysis. The uncertainties on the Bayes factors are calculated by taking the 90% confidence interval across multiple LALInferenceNest chains. The solid line uses the median of the $\rho_{\rm p}$ distribution. The shading gives the 1σ and 2σ uncertainties on the $\rho_{\rm p}$ measurement. The solid black lines shows the $\rho_{\rm p} = 2.1$ threshold.

likelihoods recovered from the aligned-spin waveform model by a factor of $\exp(\rho_{\rm p}^2/2)$.

The commonly used heuristic when assessing the strength of evidence using Bayes factors is that $1 \leq \ln \mathrm{BF} \leq 3$ is marginal evidence and $\ln \mathrm{BF} > 3$ is strong evidence in favour of a hypothesis. From the plots above we conclude that if 90% (50%) of the $\rho_{\rm p}$ posterior distribution is above the $\rho_{\rm p} = 2.1$ threshold, this corresponds to a $\ln \mathrm{BF} \approx 3.5$ ($\ln \mathrm{BF} \approx 0.8$) and is therefore very strong (marginal) evidence for precession. The posterior distribution on $\rho_{\rm p}$ can therefore be approximately mapped to the commonly used $\ln \mathrm{BF}$. Assessing the strength of evidence for precession using $\rho_{\rm p}$ would also reduce the need for additional parameter estimation runs using non-precessing models, which are *necessary* to compute the Bayes factor. This reduction in computational cost will not be significant for a single event, but for population analyses and large scale PE studies this alternative metric could be extremely useful.

3.6 Predicting the Precession SNR Posterior

For the majority of simulations presented in this chapter, the distribution for the precession SNR, $\rho_{\rm p}$, has been peaked significantly below the simulated value, although in nearly every case the simulated value does lie within the 90% confidence region. While the naive expectation is that the recovered posterior will peak at the simulated value, for complex parameter recovery where there are dependencies and

degeneracies between the different parameters, this is often not the case. We have already seen that the distance is typically over-estimated in the simulations we have performed — this is a well-known effect and arises for two reasons, first that the network is less sensitive to sources from the chosen sky location than from other locations consistent with the observed signal (as discussed in Sec. 3.4.g), and second that the signal was simulated significantly inclined from face-on, yet preferentially recovered close to face-on (as discussed in Sec. 3.4.c). Similarly, it seems likely that the signals we have simulated have more significant precession effects (deliberately, as we wish to understand the observability of precession) than the vast majority of possible sources. Thus, our conjecture is that the likelihood peaks at the simulated value of $\rho_{\rm p}$ but the posterior distribution will be biased to recover a smaller value owing to the much larger volume of parameter space consistent with low $\rho_{\rm p}$. To demonstrate this, we calculate a prior distribution for $\rho_{\rm p}$ which uses the information gleaned from a non-precessing analysis to take into consideration the much larger volume of parameter space consistent with low $\rho_{\rm p}$. We then show that when multiplying the likelihood by the prior, the *predicted* posterior for $\rho_{\rm p}$ agrees well with the *inferred* posterior from a fully precessing parameter estimation analysis.

Let us first show that the likelihood peaks at the simulated value of $\rho_{\rm p}$. The two-harmonic approximation allows us to factorize the likelihood in Eq. (3.9) into two terms: a non-precessing component (dependent on h_0) $\Lambda_{\rm np}(\boldsymbol{\lambda})$ and precessing component (dependent on h_1) $\Lambda_{\rm p}(\boldsymbol{\lambda})$,

$$\begin{aligned} p(d|\boldsymbol{\lambda}) &\propto \exp\left(-\frac{1}{2}\langle d-(\mathcal{A}_0h^0+\mathcal{A}_1h^1)|d-(\mathcal{A}_0h^0+\mathcal{A}_1h^1\rangle\right) \\ &\propto \exp\left(\langle d|\mathcal{A}_0h^0\rangle-\frac{|\mathcal{A}_0|^2}{2}\langle h^0|h^0\rangle\right)\times\exp\left(\langle d|\mathcal{A}_1h^1\rangle-\frac{|\mathcal{A}_1|^2}{2}\langle h^1|h^1\rangle\right) \\ &\propto \Lambda_{\rm np}(\boldsymbol{\lambda})\times\Lambda_{\rm p}(\boldsymbol{\lambda}), \end{aligned} \tag{3.9}$$

For simplicity we use the approximations that $\langle h^0|h^1\rangle = 0$ and that h^0 is the dominant harmonic, i.e., that the SNR in the h^0 harmonic is larger than in h^1. The calculation proceeds analogously when h^1 is dominant, and can be extended to the general case by replacing h^1 by its projection onto the space orthogonal to h^0.

We can re-express the precessing contribution to the likelihood Λ_p in terms of the precession SNR using Eq. (2.34). To do so, we introduce $\hat{\rho}_{\rm p}$ which is the simulated value of $\rho_{\rm p}$, and $\rho_{\rm p}(\boldsymbol{\lambda})$ which is the precession SNR for the set of parameters $\boldsymbol{\lambda}$. Furthermore, we define the simulated phase (as given in Eq. (2.27)) of the precession harmonic as $\hat{\phi}_1$ and the phase associated with the parameters $\boldsymbol{\lambda}$ as $\phi_1(\boldsymbol{\lambda})$. Following the procedure described in, e.g. Ref. [182], we can rewrite the precession likelihood as

$$\Lambda_{\rm p}(\rho_{\rm p},\phi_1)\propto\exp\left(-\tfrac{1}{2}\left(\rho_{\rm p}^2(\boldsymbol{\lambda})-2\hat{\rho}_{\rm p}\rho_{\rm p}(\boldsymbol{\lambda})\cos(\hat{\phi}_1-\phi_1)+\hat{\rho}_{\rm p}^2\right)\right). \tag{3.10}$$

In general, we have no prior knowledge of the precession phase, so it is natural to assume a uniform prior on ϕ_1. We may then analytically marginalise $\Lambda_{\rm p}(\rho_{\rm p}, \phi_1)$ over ϕ_1 to obtain,

$$\begin{aligned} \Lambda_{\rm p}(\rho_{\rm p}) &\propto \int_0^{2\pi} \Lambda_{\rm p}(\rho_{\rm p}, \phi_1)\, p(\phi_1)\, d\phi_1 \\ &\propto I_0(\hat{\rho}_{\rm p}\, \rho_{\rm p}) \exp\left(-\frac{\hat{\rho}_{\rm p}^2 + \rho_{\rm p}^2}{2}\right). \end{aligned} \tag{3.11}$$

We therefore see that the precession likelihood peaks at $\hat{\rho}_{\rm p}$. We may then calculate the posterior distribution for $\rho_{\rm p}$ using Bayes' Theorem,

$$p(\rho_{\rm p}|d) \propto p(\rho_{\rm p}) \Lambda_{\rm p}(\rho_{\rm p})\,, \tag{3.12}$$

where $p(\rho_{\rm p})$ is the prior for the precession SNR.

Previously, in Chapter 2, we obtained a distribution for $p(\rho_{\rm p}|d)$ by maximising the likelihood over $\mathcal{A}_1$. This is equivalent to assuming uniform priors for the real and imaginary components of $\mathcal{A}_1$, and leads to a prior $p(\rho_{\rm p}) \propto \rho_{\rm p}$. It follows from Eq. 3.12 that this results in a χ^2 distribution with 2 degrees of freedom. Here, we instead use a prior for $\rho_{\rm p}$ which is informed by the information obtained from a non-precessing analysis, we refer to this as the *informed* prior. This *informed* prior better represents our prior knowledge about $\rho_{\rm p}$ before explicitly accounting for precession in our analysis.

The majority of parameters required to calculate the *informed* prior are already given in the non-precessing results. The two exceptions are the amplitude of the precessing spin $\chi_{\rm p}$ and the initial precession phase ϕ_{JL}. As discussed in Section 3.4.d, we can obtain a prior for $\chi_{\rm p}$ *conditioned* upon the other parameters, specifically the mass ratio and aligned spin $\chi_{\rm eff}$, and this can be used to generate the informed prior on $\rho_{\rm p}$. The initial precession phase is unconstrained by the non-precessing parameter recovery, this then allows us to assume it to be uniformly distributed. By calculating the *predicted* posterior distribution for $\rho_{\rm p}$ based upon a set of non-precessing samples, we may examine the effect of other measured parameters on the final $\rho_{\rm p}$ distribution. For example, if the aligned-spin run favours a binary that is close to equal mass and an orientation consistent with a face-on system, then our prior belief will be that the precessing SNR will be low — it is only with unequal masses and systems misaligned with the line of sight that there are significant precession effects in the observed waveform. A prior belief of $\rho_{\rm p}$ peaking at low values will cause the predicted $\rho_{\rm p}$ to peak at values lower than the simulated one and consequently so too will the inferred posterior distribution for $\rho_{\rm p}$ inferred from a full 15-dimensional parameter estimation analysis.

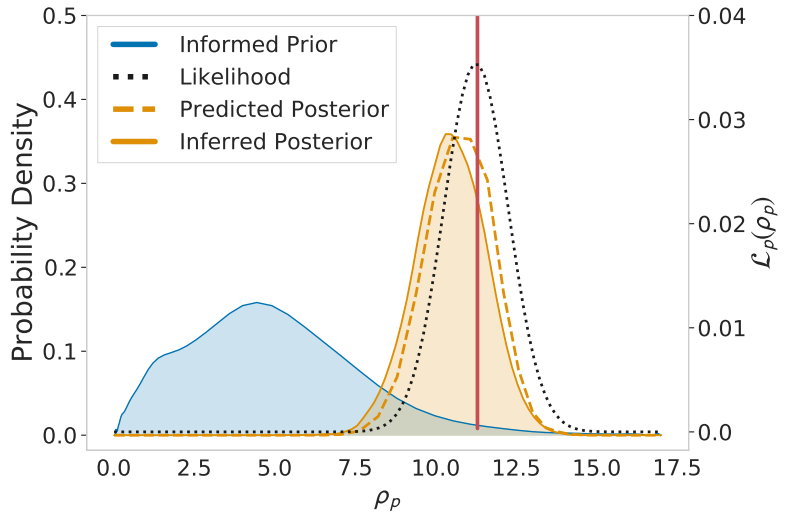

Figure 3.17: The predicted distribution for the precession SNR $\rho_{\rm p}$ (dashed orange) calculated as the product of the precessing contribution to the likelihood (black dotted line) and the informed prior of $\rho_{\rm p}$ (blue) for the $q = 4$ simulation presented in Sec. 3.4.d. For comparison, we show the inferred $\rho_{\rm p}$ posterior distribution from the full 15 dimensional parameter estimation analysis (solid orange) and $\rho_{\rm p}$ for the injection (red line). The informed prior is peaked at low values of $\rho_{\rm p}$ causing the peak of the posterior to be smaller than the maximum likelihood value.

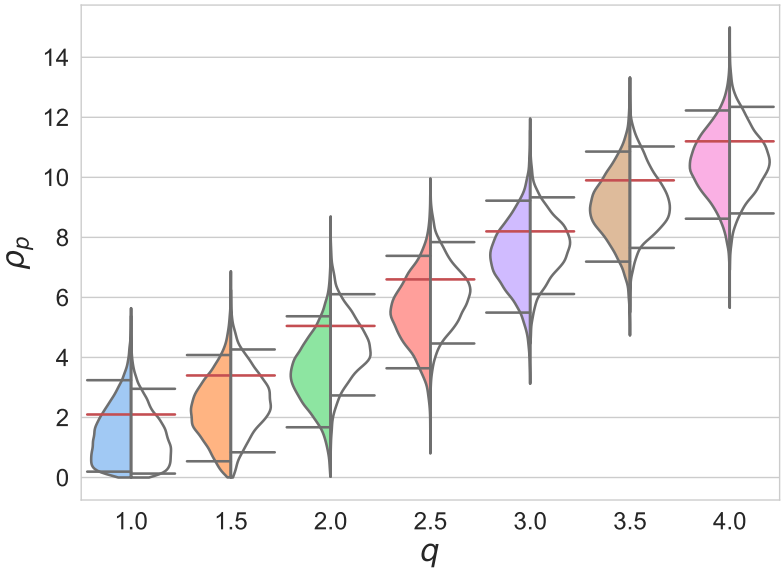

Figure 3.18: Violin plot comparing the observed $\rho_{\rm p}$ distribution (colored) from a precessing analysis, and the predicted distribution (white) based on the aligned-spin results and simulated value of $\rho_{\rm p}$ for the set of varying mass ratio simulations presented in Sec. 3.4.d. The predicted and observed distributions for precession SNR are in good agreement, even though the $\rho_{\rm p}$ in the simulated signal (red lines) lies above the peak of either distribution.

3.6.a Precessing signal

We now apply this conjecture to a precessing signal by attempting to *predict* the posterior distributions for $\rho_{\rm p}$. This allows us to investigate how much our recovered posterior distributions may differ from the idealised case of a precession likelihood function distributed about the simulated (true) value. In Fig. 3.17 we show the results of this for the $q = 4$ simulation presented in Sec. 3.4.d. This specific simulation was chosen since this case has the largest $\rho_{\rm p}$ and corresponds to a simulation where a non-precessing analysis is less justified. It is therefore a good case to show how the combination of the informed prior and the additional likelihood from precession $\Lambda_{\rm p}$ correctly estimates the large $\rho_{\rm p}$. In Fig. 3.18, we show how the predicted posterior distribution compares to the inferred distribution over the full range of mass ratio simulations presented in Sec. 3.4.d.

In Fig. 3.17 we show this predicted distribution, the informed prior, the χ^2 likelihood function and the posterior distribution obtained from a full parameter estimation analysis. By explicitly calculating the informed prior and likelihood terms separately for $\rho_{\rm p}$, we can see the effect of the prior on the $\rho_{\rm p}$ posterior. The prior strongly disfavours large observable precession and therefore pulls the posterior towards *smaller* values than the simulated value i.e. where the likelihood function peaks.

In Fig. 3.18, we show a comparison between the predicted and measured $\rho_{\rm p}$ distributions for the set of runs with varying mass ratio presented in Sec. 3.4.d. When we calculate the posterior, explicitly accounting for the parameter space weighting encoded in the informed prior on $\rho_{\rm p}$, we find good agreement between the predicted and the inferred $\rho_{\rm p}$ distributions and note that neither predicted nor inferred are centred around the *true* value for the set of signals that we have simulated. Of course, if we were to draw signals uniformly from the prior distribution, we would expect to observe the inferred distributions of $\rho_{\rm p}$ matching with the simulated values.

3.6.b Non-precessing signal

We now look at the expected posterior distribution for $\rho_{\rm p}$ when there is no precession in the the signal. As explained in Sec. 2.6, previously a χ^2 distribution with two degrees of freedom was used to model the $\rho_{\rm p}$ distribution in the absence of any precession. This then led to the natural heuristic that $\rho_{\rm p} = 2.1$ should be the threshold for observable precession. Using Eq. (3.12) we can now use a more informative prior on $\tilde{\rho}_{\rm p}$ and obtain a more accurate estimate of the expected posterior distribution in the absence of precession. We do this by using parameter estimation samples from an aligned-spin model and setting the simulated precession SNR to be 0, this then allows us to account for the effects of priors and different noise realisations.

In Fig. 3.19 we show the predicted and observed distributions for the precession SNR for a non-precessing signal. We use a non-precessing equivalent of the

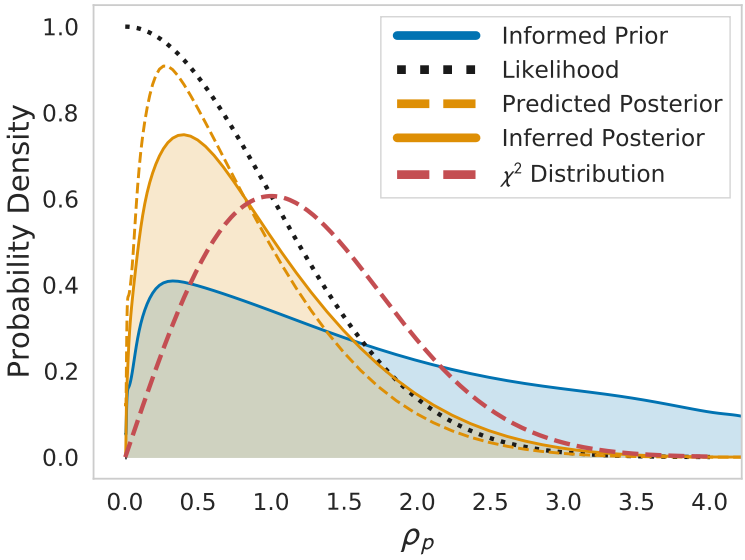

Figure 3.19: Distribution of $\rho_{\rm p}$ in the absence of precession for the "standard injection". The inferred $\rho_{\rm p}$ distribution using the IMRPhenomPv2 approximant for recovery is shown by the solid orange line. The dashed orange line shows the predicted distribution using samples collected from an aligned-spin analysis and setting the simulated precession SNR to be 0. We also shows the χ^2 distribution derived previously in Chapter 2 as a red dashed line

"standard" injection as our simulated signal (i.e., we set $\chi_{\rm p} = 0$ while ensuring all other parameters match those in Tab. 3.1). We inject with zero noise and use the `IMRPhenomPv2` model for parameter recovery.

The inferred $\rho_{\rm p}$ distribution is peaked at lower values that the χ^2 distribution as shown in Fig. 3.19. However using the prediction from the likelihood (Eq.3.9) and the *informed prior* we are able to obtain a better estimate of the posterior in the absence of precession. This estimate can be obtained without performing parameter estimation incorporating precession, this therefore allows for a better metric for determining whether or not there is measurable precession in the system.

The distribution for the informed prior on precession SNR will depend upon the details of the signal. In particular, it will be strongly peaked near zero for events that are likely to have small opening angle (eqivalently $\bar{b}$), i.e., events that are close to equal mass and have significant spin aligned with the orbital angular momentum, while high mass-ratio events and those with large anti-aligned spins will lead to greater support for large values of $\rho_{\rm p}$. Furthermore, for binaries where the orientation can be well measured, *without* precession information, for example where higher modes are important, those that are close to face-on will lead to predictions of smaller $\rho_{\rm p}$ while those that are edge-on will give larger values. Given that the majority of signals observed to date are consistent with equal mass binaries, in most cases the prior on $\rho_{\rm p}$ will tend to be peaked at low values. Consequently, the simple

threshold of $\rho_{\rm p} \gtrsim 2.1$ as evidence for precession, remains appropriate and is likely more stringent than suggested by the simple likelihood calculation.

3.7 Discussion

In most candidate astrophysical binary distributions, precession is likely to be first measured in a comparable-mass binary (see Chapter 4). We have considered a fiducial example of such a possible signal (mass-ratio $q = 2$, SNR $\rho = 20$, and in-plane spin $\chi_p = 0.4$, such that the precession contribution to the total SNR is $\rho_{\rm p} = 5$), and performed an extensive parameter-estimation study that has systematically explored the impact on parameter measurements of changes in each of the key source parameters: the SNR, the in-plane spin magnitude, binary inclination, the binary mass ratio and aligned-spin contribution, the binary's total mass, the polarisation, and sky location. These examples illustrate well-known features of precession signals [164, 94, 54, 59, 166, 167, 120, 168], and quantify their effect on both the measurement of precession, and their impact on the measurement accuracy and precision of other parameters.

We have also verified that $\rho_{\rm p}$ provides a suitable and intuitive metric for determining whether or not we have measured precession, and shown that there is an approximate mapping between $\rho_{\rm p}$ and the use of the Bayes factor to assess the evidence of precession. We suggest that given these results, future large scale studies of precession can be made considerably computationally cheaper by computing $\rho_{\rm p}$, rather than a full Bayesian analysis.

We note that as $\rho_{\rm p}$ captures precession by identifying additional power beyond a simple non-precessing waveform model, it could therefore be effected by phenomena such as eccentricity and higher order multipoles. As BFs simply compare the evidence for two models, one precessing and one non-precessing, using BFs as the sole metric would also be biased by properties like eccentricity and higher order multipoles.

However, a similar approach to the 2-harmonic decomposition for precessing signals has recently been applied to GWs including the effects of higher harmonics [182]. In future work, we will combine these approaches and explore the measurability of precession in systems with significant evidence for higher harmonics, and the impact of the combination of higher modes and precession upon parameter accuracy. It may also be possible to account for eccentricity through a similar decomposition.

As highlighted in section 3.6 these decompositions provide powerful insights into how the addition of physical phenomena introduce information into the analysis. Here we show that the likelihood can be simply factored into precessing and non-precessing contributions. This then allows us quantify the extra information that can be gained from a precessing analysis and even predict the recovered $\rho_{\rm p}$ distribution with or without these effects taken into consideration in the analysis.

The current study does not include higher harmonics, and uses a signal model (`IMRPhenomPv2`) that neglects two-spin precession effects, mode asymmetries that lead to out-of-plane recoil [183], and detailed modelling of precession effects through merger and ringdown. Although these effects are typically small, so is the imprint of precession on the signal, and it would be interesting in future to investigate the impact of these additional features on our results. We also emphasize that, although we consider it to be extremely useful to provide quantitative examples of the effects of each of the binary parameters, these will necessarily depend on the location in parameter space of our fiducial example. However, having chosen a configuration from amongst what we expect to be the most likely signals, we hope that these examples will act as a useful guide in interpreting precession measurements when they arise in future gravitational-wave observations.

Chapter 4

When will we observe binary black holes precessing?

This chapter is based upon the text of Fairhurst *et al.*[120] published as Phys. Rev. D, 102, 041302. This work was led by Stephen Fairhurst, Rhys Green, Mark Hannam and Charlie Hoy. My main contributions were writing the code which performed the population analysis described in Sections 4.3 and 4.4 and paper writing. All authors contributed equally to writing.

4.1 Introduction

The Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) [5] and Advanced Virgo (AdV) [23], provide a unique method of observing mergers of black holes and/or neutron stars. Observations to date already provide insights into the mass and spin distributions of black holes [38, 85].

One important general relativistic effect that has *not yet* been observed is orbital precession. This arises when the black-hole spins are not aligned with the binary's orbital angular momentum. In contrast to Newtonian mechanics, where all angular momenta are individually conserved, in general relativity the binary's total angular momentum is (approximately) conserved, and the orbital angular momentum (and hence the orbital plane) and spins precess around it [54, 87]. This leads to modulations in the amplitude and phase of the gravitational wave (GW) signal. These are in general small effects and, in addition, whether they can be measured depends not only on the black-hole masses and spin magnitudes and directions, but also on the binary's orientation relative to the detector, and the observed GW polarization. For this reason, until now there was no straightforward way to determine how significantly precession would be imprinted onto a given waveform. The usual approach is to perform computationally expensive Bayesian analyses (see e.g. [38, 91]), but even then, the misaligned spin components (which signify whether the binary is precessing) are degenerate with other parameters, and do not provide a direct measure of precession features in the signal. This makes it difficult to infer the impact of precession measurements on the properties of astrophysical binary populations and their formation mechanisms.

4.2 Observability of precession

In Chapter 2, we demonstrated that a precessing waveform can be decomposed into a series of five *non-precessing harmonics*, where the characteristic modulations of a precessing signal are caused by the beating of these *harmonics*. Since the individual harmonics are indistinguishable from non-precessing waveforms, it is only when two precession harmonics can be independently observed that precession can be unambiguously identified. For precession to be observable, we therefore require that the expected signal-to-noise ratio (SNR) in *both* of the harmonics is above some threshold.

The expected SNR for a signal h embedded in data from a detector with a noise power spectral density $S(f)$ is given as $\hat{\rho}^2 = (h|h)$, where $(h|h)$ is the inner product defined in Eq. 2.30. In cases with more than one precession cycle between f_o and f_{max}, the two harmonics will be close to orthogonal, $(h^0|h^1) \approx 0$, and the precession SNR is simply defined as the expected SNR in the weaker harmonic. When there is less than one precession cycle, it is necessary to also subtract the power aligned

with the dominant harmonic before determining whether the second harmonic is observable. This is explained in detail in Chapter 2.

It remains to determine a threshold above which $\rho_{\rm p}$ can be considered as evidence for precession. Consider the situation where an event has been observed, so there is significant SNR in at least one harmonic. In the absence of measurable precession, and assuming well-modelled Gaussian noise, the SNR in the second harmonic will be χ^2 distributed with two degrees of freedom, where the two degrees of freedom correspond to the real and imaginary parts of the complex amplitude. Therefore, in the absence of precession, $\rho_{\rm p} > 2.1$ is expected in less than 10% of cases, and $\rho_{\rm p} > 3$ in approximately 1% of cases. We therefore use these simple thresholds as a measure of the strength of evidence for observable precession.[1]

In Fig. 4.1 we show the recovered distribution of $\chi_{\rm p}$ and $\rho_{\rm p}$ for a number of signals, both real and simulated. For each signal, we use a nested sampling routine within the LALInference code [90, 91] to obtain posterior probability distributions for the parameters. This is the same infrastructure that was used to measure the properties of the LIGO-Virgo observations, and we present our results in the same form as in, for example, the GWTC-1 catalogue [38], by using the PESummary library (see Chapter 8). The new feature is our calculation of $\rho_{\rm p}$.

First, we show the recovered $\chi_{\rm p}$ and $\rho_{\rm p}$ distributions for a set of simulated signals, generated using the `IMRPhenomPv2` model [73], each with the same choices of masses and spins — total mass $M = 40M_{\odot}$, mass ratio 2:1, and an in-plane spin of $\chi_{\rm p} = 0.4$ on the large black hole only — but varying orientation, encoded by the angle θ between the total angular momentum and the line of sight. The distance to each signal is chosen to ensure a *fixed* expected SNR of 20 in the aLIGO detectors at the sensitivity of the second observing run (O2) [38], resulting in a distance variation by a factor of ≈ 3.5 between the least and most inclined systems.

For binaries with total angular momentum closely aligned with the line of sight, $\theta < 45^{\circ}$, the precessing SNR is consistent with no power in the h^1. The posterior on $\chi_{\rm p}$ is consistent with the prior at low $\chi_{\rm p}$ but excludes $\chi_{\rm p} \gtrsim 0.7$. When $\theta > 45^{\circ}$, the angular momentum is significantly mis-aligned with the line of sight and there is significant power in both harmonics, leading to a value of $\rho_{\rm p}$ inconsistent with noise alone and little support for values of $\chi_p \lesssim 0.1$. However, using $\chi_{\rm p}$ alone, *even after performing the parameter recovery* there are no simple criteria to determine when precession is observed. A natural choice might require that the 90% confidence interval for χ_p exclude zero, but this will *always* be the case, primarily due to the shape of the prior. Furthermore, even though we know all of the parameters *a priori*, it is impossible to determine whether precession will be observable without generating the waveform and performing the parameter recovery.

[1] A more detailed analysis would consider the volume of the binary parameter space consistent with the non-precessing and precessing parameters and use these to appropriately weight the likelihoods of the precessing and non-precessing signals. This will have some impact on the required $\rho_{\rm p}$ threshold, but is unlikely to change it significantly.

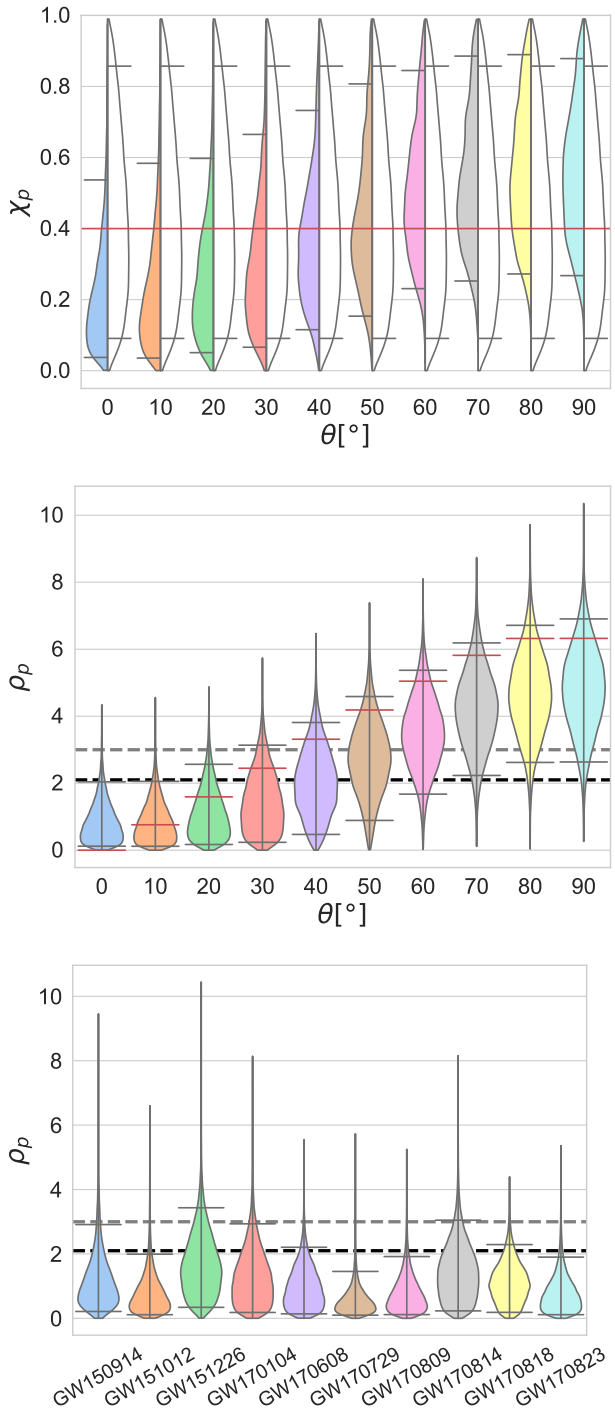

Figure 4.1: For a set of simulated signals with fixed masses and spins (see text), we show the posterior and prior (white) distributions for χ_p (top), and posterior distributions for $\rho_{\rm p}$ (middle) for a range of different binary orientations, θ. The grey lines show the 90% confidence regions, the solid red lines show the *true* values of $\chi_{\rm p}$ and $\rho_{\rm p}$ respectively and the dashed black and grey lines indicates the thresholds for observable precession at $\rho_{\rm p} = 2.1$ and $\rho_{\rm p} = 3$. The bottom panel shows the $\rho_{\rm p}$ distribution for the ten binary-black-hole observations in O1 and O2 [38].

The precession SNR solves these problems. A value of $\rho_{\rm p} > 2.1$ tells us immediately that there is evidence of observable precession. Most significantly, the expected precession SNR $\rho_{\rm p}$ (red lines on the middle plot) can be calculated directly from the signal parameters; no detailed parameter estimation analysis is necessary. Thus, for the first time, we are able to identify immediately whether precession would be measurable in a given configuration. We see in Fig. 4.1 that for each inclination, the true value for $\rho_{\rm p}$ lies within the recovered 90% credible interval however the posterior is not centred around the true value. This is due to selection and prior effects. In Chapter 3, we investigated these selection effects as well as providing a detailed exploration of the observability of precession over the parameter space of masses, spins and binary orientation.

Fig. 4.1 also shows the distribution of $\rho_{\rm p}$ for the BBH merger signals that were observed in the aLIGO and AdV O1 and O2 runs [38]. No evidence of precession was found in these signals [85], as is made clear from the recovery of $\rho_{\rm p}$. There are several cases where the distribution extends to higher values, but the median never exceeds the 2.1 threshold. These results demonstrate the efficacy of $\rho_{\rm p}$.

4.3 When will we observe precession?

We can use the observation of precession to distinguish different binary formation scenarios. The precession SNR makes it straightforward to perform an in-depth investigation of various models and identify the fraction of signals for which precession effects will be observable. Such a study was not previously possible, due to the difficulty in classifying observability of precession. Instead, limited investigations of the parameter space have been performed [102], or inferences of the distributions for the spin magnitudes and orientations obtained [140, 138], again with a limited sample size.

We investigate nine astrophysically-motivated populations of black hole binaries, comprised of three distributions of spin magnitude, and three distributions of spin orientation for the individual black holes in the binary. The spin-magnitude distributions are those used in Refs. [118, 142, 119]: *low* and *high* are triangular, peaked either at zero or extremal spin, and *flat* is a uniform distribution between zero and one. The spin-orientation is characterized by the distribution for the angle σ between each black hole's spin and the orbital angular momentum: *aligned* is a triangular distribution in $\cos\sigma$, which peaks at 1 and can take values $0.85 < \cos\sigma < 1.0$, ($\sigma \lesssim 30^\circ$); *precessing* is triangular in $\cos\sigma$ peaked at 0, with values $-0.15 < \cos\sigma < 0.15$, ($80^\circ \lesssim \sigma \lesssim 100^\circ$); *isotropic* is uniform in $\cos\sigma$ between -1 and 1. For each population, we generate 10^5 binaries with masses drawn from a power law mass distribution with $p(m_1) \propto m^{-2.35}$, and $p(m_2)$ uniform in m_2 between $5\,M_\odot$ and m_1 (as in [119]), and distributed uniformly in volume and binary orientation.

	Aligned		Isotropic		Precessing	
Low	0.043	0.644	0.151	0.194	0.173	0.150
Flat	0.077	0.448	0.276	0.040	0.327	0.019
High	0.105	0.331	0.354	0.013	0.412	0.005

Table 4.1: The probability of observing precession, $\rho_{\rm p} > 3$, for an observed binary (white) from each spin distribution and the probability of *not* observing precession in 10 random draws (grey) from each spin distribution.

Table 4.1 shows the probability of observing precession in a single event drawn from each of the nine populations, observed with O2 sensitivity while assuming zero noise. For this study, we use the higher threshold of $\rho_{\rm p} > 3$, corresponding to a 1% false rate, to indicate strong evidence for observed precession. When observing a population of events, the number of events exceeding this threshold when there is no precession in the system remains low.[2]

As expected, we are most likely to observe precession when the black holes have high spins that lie preferentially in the orbital plane (*high-precessing* configurations) and least likely for black holes with low spins, or with spins preferentially aligned with the orbital angular momentum (*low-aligned* configurations). Given that precession has not been observed in GW detections to date, we are able to restrict the spin distribution. Table 4.1 shows the probability of detecting ten signals with no observable precession from each of the nine spin distributions. Based on precession measurements alone, we strongly disfavour all *precessing* distributions. Although these are already considered astrophysically unlikely, there are models that predict preferentially in-plane spins [121, 122]. We also disfavour *isotropic* spins with *flat* or *high* magnitudes. Thus, the lack of observed precession points towards low spins, or spins preferentially aligned with the orbital angular momentum.

Previous constraints on spins have primarily been provided by considering the measurable aligned-spin component [118, 142, 119, 38] and provide strong evidence against all but *low aligned* or *iostropic* distributions, with *low isotropic* spins preferred. Combining the aligned spin and precession results will further restrict the spin distribution consistent with GW observations, and will likely require spin magnitudes even smaller than our *low* distribution (see also Ref. [85]).

[2] As our analysis assumes zero noise, the fraction of binaries with observable precession will be slightly underestimated. At a threshold of $\rho_{\rm p} > 2.1$, the effect would be significant while at a threshold of $\rho_{\rm p} > 3$, the difference is small.

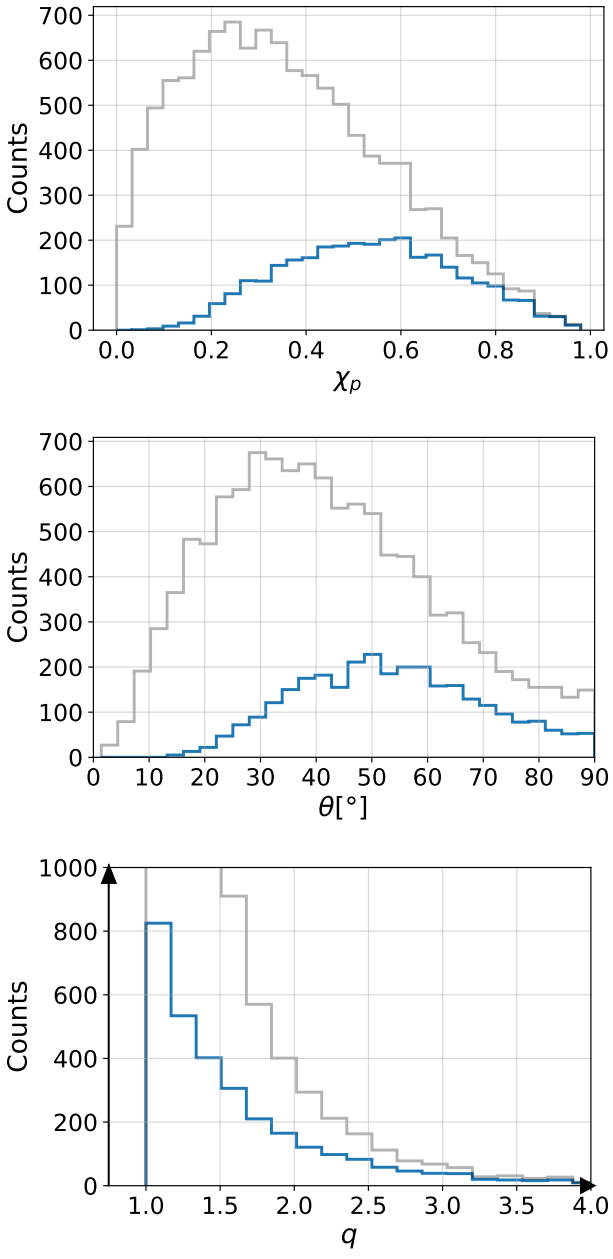

Figure 4.2: The distribution of $\chi_{\rm p}$, θ and q for observable binaries (grey), and those with measurable precession (blue), assuming a low isotropic spin distribution. θ is the inclination angle folded to $[0, \pi/2]$. The y-axis labels the number of observed events in each bin, out of 10^5 simulated signals with low isotropic spins.

4.4 Where will we observe precession?

We are able to identify, for the first time, the regions of parameter space that lead to signals with observable precession. In Fig. 4.2, we show the expected distribution of the precessing spin $\chi_{\rm p}$, binary orientation θ and mass ratio q for observable binaries and binaries with observable precession, $\rho_{\rm p} > 3$, assuming a *low isotropic* distribution of spins. We identify clear regions of the parameter space where precession is more likely to be observed: large values of $\chi_{\rm p}$, binaries that are close to edge-on, $\theta > 45^\circ$, and systems with high mass ratio. Regions where the chance of observing precession is close to zero include binaries with $\chi_{\rm p} < 0.2$ or where the total angular momentum is within 20° of the line of sight. These results are consistent with expectations based upon smaller studies using detailed parameter estimation techniques [102]. We also note that most observations of precession will be in comparable-mass binaries, i.e., $q \leq 2$. This is a surprising, new result. It is well known that precession is more easily measured at higher mass ratios [54], which is confirmed by our study: precession is observed in <12% of detections with $q < 2$, but >35% for $q > 2$. However, with $\sim 90\%$ of observations expected to have $q < 2$, these vastly outnumber the higher-mass-ratio observations, and we find that ∼75% of precession observations will come from detections of binaries with $q < 2$.

4.5 Discussion

In this chapter we have used a simple method to identify when precession is measurable in a compact binary GW signal. The gravitational waveform is well approximated by the first two harmonics in a power series expansion in the tangent of the half-opening angle (see Chapter 2), and the unambiguous observation of precession requires the identification of both of these harmonics in the data. The precession SNR $\rho_{\rm p}$ is a simple measure of this observability. We have demonstrated the efficacy of $\rho_{\rm p}$ through parameter estimation studies and also provided the distributions of $\rho_{\rm p}$ for the aLIGO-AdV observations to date. Using our definition of precession SNR, we have identified how often precession will be observed for a variety of potential astrophysical spin distributions. For the most likely distribution, based on current observations (*low-aligned*) there is a 83% chance that precession will be measured after ∼40 observations, and is therefore likely to be observed during the current third aLIGO-AdV observing run (O3). The non-measurement of precession by the end of O3 would place much stronger constraints on spin orientations and magnitudes.

The precession SNR has many applications. Most immediately, it allows us to determine the measurability of precession in a system *without performing computationally expensive parameter estimation.* This allows us to, e.g., easily fold precession information into population analyses of black-hole binaries. In future work, we will explore whether the value of $\rho_{\rm p}$ can be used to predict the measured $\chi_{\rm p}$ distribution.

The precessing SNR also gives us a simple way to identify regions of the parameter space where precession is important, a necessary first step in extending existing GW searches to explicitly use precessing waveforms [114].

Chapter 5

GW190814: a potential neutron star black hole observation

In this chapter I explore neutron star-black hole (NSBH) binaries. In Section 5.2 I introduce how the lack of a clear NSBH binary observation in the first and second gravitational wave observing runs allowed us to place upper limits on the inferred merger rates of NSBH binaries in the universe. This work was published by the LIGO Scientific and Virgo collaborations in Phys. Rev. X, 9, 031040. I produced Figure 5.1 and this demonstrates my contribution to Figure 14 in Ref. [38]. In Sec. 5.3 I discuss the properties of a potential NSBH binary detected in the first half of the third gravitational wave observing run and identify if precession has been measured in this asymmetric binary. Section 5.3.c was published by the LIGO Scientific and Virgo collaborations in Astrophys. J. Lett., 896, L44. I was one of the 8 lead authors for this paper and was a leading contributor to Section 4: Properties of GW190814. I also assisted with other sections of this paper for which my contributions varied. I produced Figures 5.7, 5.10, 5.14 and the top panel of Figure 5.9 for the Astrophys. J. Lett., 896, L44 publication. The other plots in this chapter were generated by PESummary (see Chapter 8) using public data from the Gravitational Wave Open Science Center [184]. These plots have not previously been published. All text and figures in Section 5.3.e are unpublished.

DISCLAIMER: This chapter presents results and text that was previously published by the LIGO Scientific and Virgo collaborations. Some of the plots in this chapter have not previously been published and therefore have not been through a rigorous review process. Consequently, these plots and associated text do not reflect the scientific opinion of the LIGO Scientific and Virgo collaborations. I would, however, like to gratefully acknowledge LIGO, the LIGO Scientific and Virgo collaborations and their funding agencies.

5.1 Introduction

During the first two observing runs (O1 and O2), Advanced LIGO [5] and Advanced Virgo [6] observed the first detection of a binary black hole (BBH) (GW150914; [7]) and a binary neutron star (BNS) (GW170817; [34]) coalesence. Throughout O1 and O2, the LIGO–Virgo collaboration (LVC) announced, in total, 10 BBH systems and 1 BNS system [38]. Additional events have been reported by independent groups [39, 40, 41, 42].

The first six months of the third observing run (O3a) saw a further 39 gravitational-wave (GW) candidates announced [44]. These additional candidates include GW190425, the coalescence signal of what is most likely a BNS with unusually large chirp mass and total mass [43], GW190412, the first BBH coalescence with an unequivocally unequal mass ratio $q = m_1/m_2$ of $3.61^{+1.06}_{-1.11}$ [45] and GW190521, the heaviest BBH coalescence detected by the LVC [46] (all measurements are reported as symmetric 90% credible intervals around the median of the marginalized posterior distribution, unless otherwise specified). Although the astrophysical origin of most GW candidates in O1, O2 and O3a are known, two remain uncertain: GW190426 [44] with masses $m_1 = 5.7^{+3.9}_{-2.3}\,M_\odot, m_2 = 1.5^{+0.8}_{-0.5}\,M_\odot$ and GW190814 [47] with masses $m_1 = 23.2^{+1.1}_{-1.0}\,M_\odot, m_2 = 2.59^{+0.08}_{-0.09}\,M_\odot$. Both of these candidates are consistent with originating from either a BBH or neutron star-black hole (NSBH) binary, hereafter denoted as "potential NSBH" candidates.

One of the methods for understanding if a GW originated from a BBH or NSBH is through examining the secondary mass of the binary. If the secondary mass is heavier than the maximum allowed neutron star (NS) mass $M_{\rm max}$, the GW cannot have originated from an NSBH. Unfortunately, although theoretical estimates allow for masses up to $\sim 3M_\odot$ [48, 49], the maximum NS mass is unknown as it is determined by the unidentified NS equation of state (EOS). We must therefore compare the secondary mass to the observed population of NS masses: $[1.0 - 2.14]M_\odot$ [185, 186][1]. If the secondary mass is heavier than the maximum allowed NS mass and the lowest observed black hole (BH), the GW likely originated from a BBH system. Electromagnetic observations have highlighted that black holes are unlikely to have masses less than $5M_\odot$ [188, 189, 190, 191]. If the secondary mass lies within the *mass gap* between the known NS and BH populations, understanding the origin of the GW remains a challenge.

The observation of a confirmed NSBH would be the first detection of a brand new subclass of compact binary objects. Owing to the lack of a direct observation, the event rates for NSBH mergers is highly uncertain [38, 192]. Based on estimates from Ref. [192], the fourth observing run (O4) will likely see 1 to 92 confirmed NSBH

[1]The estimate for the most massive known neutron star in the galaxy has recently been revised in Ref. [187]. Since the work from Fonseca *et al.* [187] was circulated after the GW190814 discovery paper [47] was published, we continue to use the estimate from Ref. [186] for the rest of this chapter to ensure consistency

mergers [4].

NSBH binaries are an astrophysically interesting class of systems. Their GW signal contains signatures of the physical conditions of matter at nuclear densities (see e.g. [193]). A direct observation may therefore provide insight into the highly uncertain NS EOS and for models of stellar evolution and core collapse [194]. Although typical NSBH mass ratios are unknown, population synthesis models suggest that mass ratios $q < 10$ are astrophysically most likely [195, 196]. The unequal mass ratio implies that NSBH binaries are also one of the most promising sources for detecting precession (see Chapter 3 and Ref. [169]) and multipoles beyond the dominant quadrupole moment [197].

In this chapter, we present a calculation for the upper bound of the merger rate of NSBH systems based on the detector sensitivity network operating at O1 and O2 sensitivity. We then provide details about one of the potential NSBH systems observed during the third gravitational wave observing run, including a full bayesian analysis to determine its source properties.

5.2 NSBH merger rates during O1 and O2

O1 and O2 saw no clear NSBH mergers, although GW data alone cannot exclude the possibility that GW170817 is an NSBH [198, 199, 200]. We therefore wish to place upper limits on the rate of NSBH mergers in the local universe. This requires knowing the volume over which the gravitational wave detectors are sensitive to NSBH signals. We estimate this by injecting a large set of simulated waveforms sampled from an astrophysical population of NSBH sources into the gravitational wave strain data and identifying how many of these sources are found/missed at a detection threshold of FAR (false alarm rate) $= 0.01\text{yr}^{-1}$. The sensitive volume can then be calculated as,

$$\langle V \rangle = V_0 \frac{N_{\text{found}}}{N_{\text{inj}}}, \tag{5.1}$$

where N_{found} is the total number of found injections, N_{inj} is the total number of injections performed and V_0 is the astrophysical volume. We define V_0 as,

$$V_0 = \int_0^{z_{\text{max}}} \frac{dV_c}{dz} \frac{1}{1+z} dz, \tag{5.2}$$

where z_{max} is the maximum redshift used in the injection campaign, dV_c/dz is the differential co-moving volume and $1 + z$ in the denominator accounts for the time-dilation in the intrinsic rate caused by the expanding universe. We evaluate this integral through Monte Carlo integration methods. This procedure is explained in detail in Ref. [201]. We use a total of six NSBH populations which differ by their combination of BH masses and spin distributions. We fix the black hole mass

to be $5M_\odot, 10M_\odot$ and $30M_\odot$ and use 2 different spin distributions: isotropic and aligned/anti-aligned. The BH spin magnitudes are uniform in $[0, 1]$ and the NS spin magnitudes uniform in $[0, 0.05]$. We fix the neutron star mass to the canonical value $M_{\rm NS} = 1.4M_\odot$, use an isotropic distribution of sky location and source orientation and choose distances assuming a uniform distribution in volume. We use the *pycbc* search pipeline [202, 127] to identify the sources and model the waveforms using precessing BBH models [79] (since at the time of this analysis NSBH models were not available [203, 204]). This may introduce a small error especially for the low black hole mass injections were tidal effects are more prominent (see Section 5.3.a for details).

The astrophysical rate of NSBH coalescences R is then calculated by,

$$R = \frac{\xi}{\langle VT \rangle} \tag{5.3}$$

where ξ is the expected number of observed NSBH events in a given analysis, $\langle VT \rangle$ is the product of the population averaged sensitive volume $\langle V \rangle$ and the time T over which ξ observations have been made. The likelihood for finding zero observations in the data s follows the Poisson distribution for zero events $p(s|\xi) = e^{-\xi}$. From Bayes' theorem, the posterior for ξ is,

$$p(\xi|s) \propto p(\xi)e^{-\xi}, \tag{5.4}$$

where $p(\xi)$ is the prior on ξ, chosen to be uniform as it yields a more conservative upper limit. It follows from Eq. 5.3 that the joint posterior distribution on the astrophysical rate and the sensitive volume is,

$$p(R, \langle VT \rangle | s) \propto p(R, \langle VT \rangle)e^{-R\langle VT \rangle}, \tag{5.5}$$

where $p(R, \langle VT \rangle)$ is the joint prior. A posterior for the rate is then obtained by marginalising over $\langle VT \rangle$,

$$\begin{aligned} p(R|s) &= \int d\langle VT \rangle \, p(R, \langle VT \rangle | s) \\ &\propto \int d\langle VT \rangle \, p(R, \langle VT \rangle)e^{-R\langle VT \rangle}. \end{aligned} \tag{5.6}$$

The upper limit R_c on the rate with confidence C is then given by the solution to,

$$\int_0^{R_c} dR p(R|s) = C \tag{5.7}$$

To evaluate Eq. 5.6, we expand the joint prior $p(R, \langle VT \rangle) = p(R|\langle VT \rangle)p(\langle VT \rangle)$ where $p(R|\langle VT \rangle)$ is chosen to be uniform on R. As with Ref. [205], we use a log-normal prior on $\langle VT \rangle$,

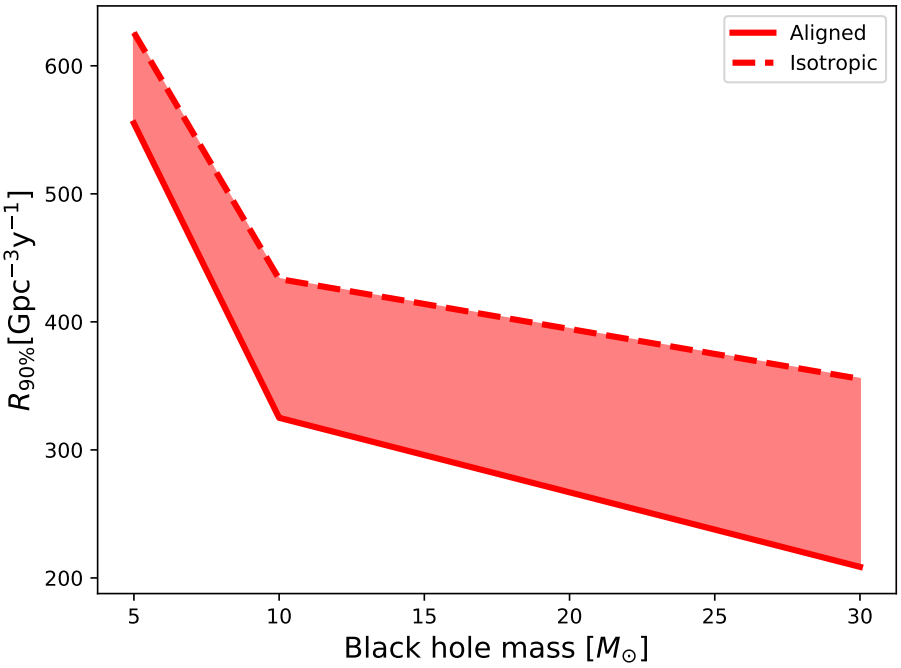

Figure 5.1: The 90% upper limit for the rate of NSBH mergers based on a LIGO-Hanford, LIGO-Livingston detector network operating at O1 and O2 sensitivity. The rate is measured at a set of three discrete black hole masses (5, 10, and 30$M_\odot$) and the neutron star mass is fixed at 1.4$M_\odot$. The rate is calculated for two spin distribution choices: isotropic (dashed) and aligned/anti-aligned (solid).

$$p(\langle VT \rangle) = \ln \mathcal{N}(\mu, \sigma^2) \tag{5.8}$$

where μ is the calculated value of $\ln \langle VT \rangle$ and σ represents the fractional uncertainty in $\langle VT \rangle$ which we set to be 18% to account primarily for calibration error.

In Fig. 5.1 we show how the 90% upper limit on the rate of NSBH binaries varies for different black hole masses and spin distributions during O1 and O2. We see that the rate is larger for isotropic spins compared to (anti-)aligned spins and for smaller BH mass systems, and that all upper limits are below 610Gpc^{-3}y^{-1}. The rates are larger for isotropic spins because the search uses a template bank of (anti-)aligned systems and consequently loses sensitivity when searching for systems with misaligned spins. The rates are larger for small BH mass systems since lower mass NSBH sources produce quieter signals and therefore cannot be observed to the same distances as higher mass NSBH sources. Our results are of course dependent on the prior chosen for R and consequently ξ. An alternative prior which has been used previously in literature [206] is the Jeffreys prior, defined as $p(\xi) \propto \xi^{-0.5}$. This prior suppresses larger rates and therefore we would expect to observe fewer mergers per year.

5.3 GW190814

On 2019 August 14 at 21:11:00 UTC, GW190814 was observed in coincidence by the LIGO-Livingston and Virgo gravitational wave detectors by the low-latency GstLAL matched-filtering search pipeline for coalescing binaries [207, 208, 126, 209, 210] [2]. Although LIGO-Hanford was in a stable operating configuration at the time of GW190814, the detector was not in observing mode due to a routine procedure to minimize angular noise coupling to the strain measurement [213] (this same procedure took place at LIGO-Hanford around the time of GW170608; we refer the reader to [36] for details of this procedure). Within a 5 minute window around GW190814, this routine procedure was not taking place. Therefore upon re-examination, LIGO-Hanford data for GW190814 are usable in the nominal range of analysed frequencies. GstLAL and PyCBC reanalysed the data and found a coincident gravitational-wave signal in all three detectors. Alerts were then issued notifying the astrophysical community that based on the outputs from GstLAL and PyCBC, GW190814 was likely to originate from an NSBH source with more than $> 99\%$ probability [214].

The source properties of GW190814 were inferred by performing a coherent Bayesian analysis on 16 s of data from LIGO-Livingston, LIGO-Hanford and Virgo. Both the LALInference stochastic sampling software [91] and a parallelised version of the parameter estimation software Bilby (pBilby; [215, 141]) were used to sample the posterior distribution. We used a low-frequency cutoff of 20 Hz for LIGO-Hanford and Virgo and 30 Hz for LIGO-Livingston for all likelihood evaluations. LIGO-Livingston used a higher low-frequency cutoff owing to the presence of scattered light, a common source of noise in all three interferometers [216].

Our analysis revealed a merger signal with the most unequal mass ratio observed with gravitational waves ($1/q = 0.112^{+0.008}_{-0.009}$) at a signal-to-noise ratio (SNR) of $\simeq 25$ with a primary component conclusively a black hole with mass $m_1 = 23.2^{+1.1}_{-1.0} M_\odot$. The nature of the $2.59^{+0.08}_{-0.09} M_\odot$ secondary component is unclear. Forming coalescing compact binaries with this unusual combination of masses at such a rate challenges our current understanding of astrophysical models.

5.3.a Difference between NSBH and BBH signal models

GW190814 was first classified as a potential NSBH source with more than $> 99\%$ probability [214], based on the secondary mass of the best matching (highest SNR) template from the GstLAL search pipeline. This template had a secondary mass which lies in the observed *mass gap* between known NS and BH populations. Excluding the possibility that GW190814's source includes an exotic object (such as

[2] Other low-latency searches, including the matched-filtering based MBTA [211] and PyCBC [127, 137, 212, 202] pipelines, could not detect the event at the time as its SNR in Virgo data was below their single-detector detection thresholds

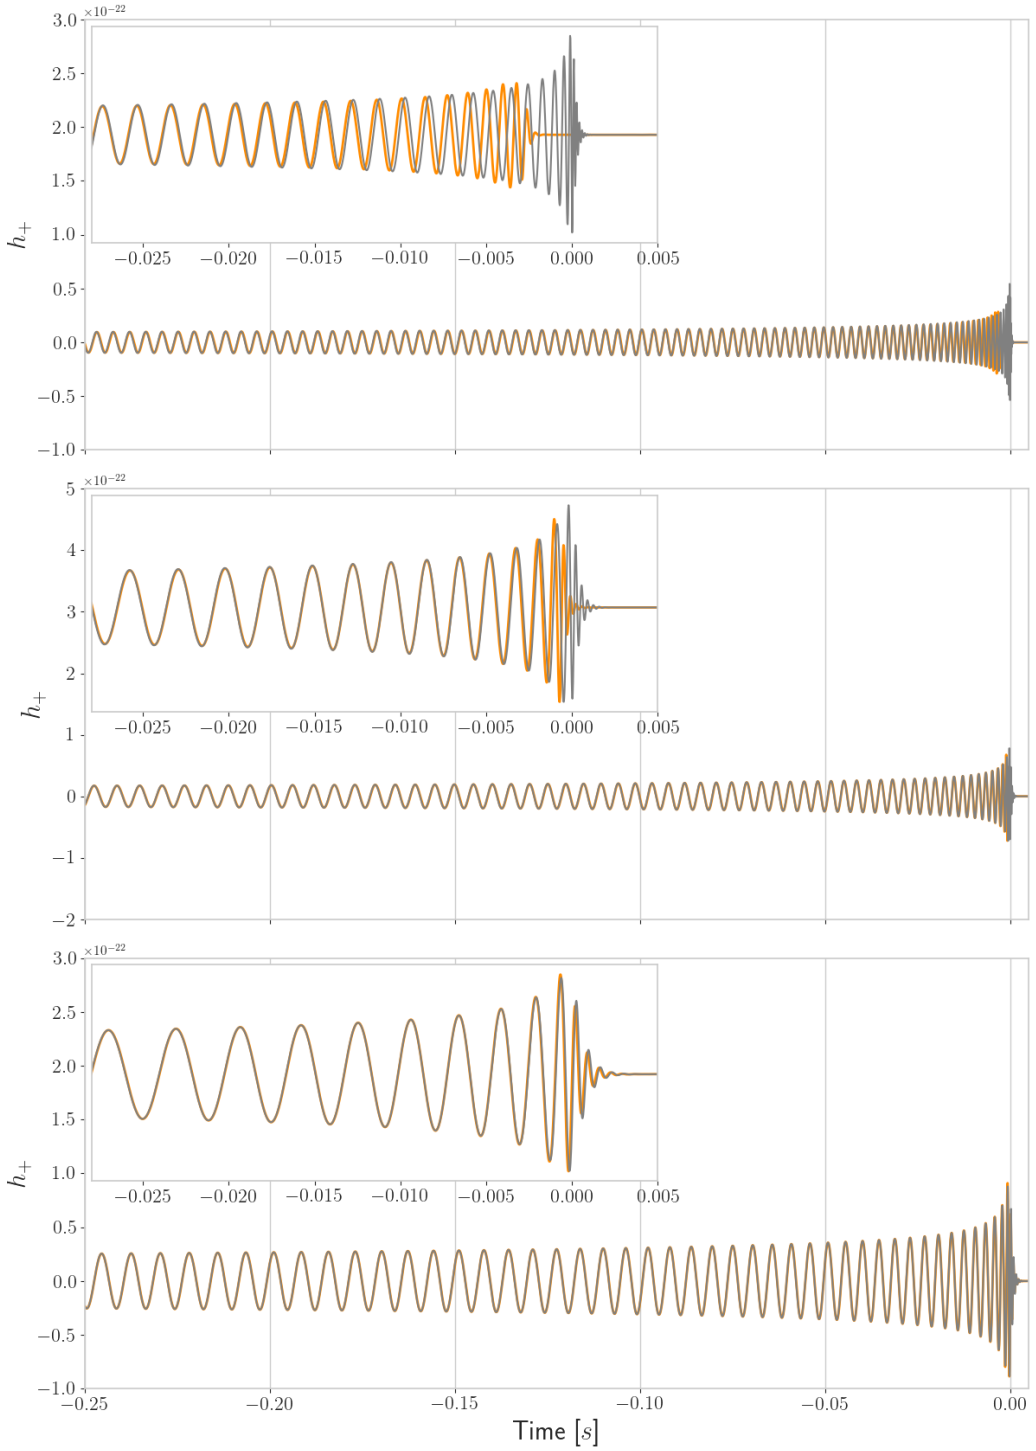

Figure 5.2: Comparison of the NSBH signal model EOBNR NSBH [204] (orange) and the BBH signal model SEOBNRv4 [219] (grey) for parameters matching the NSBH numerical relativity waveform of *top*: SXS:BHNS:0004, *middle*: SXS:BHNS:0003 and *bottom*: SXS:BHNS:0001 [220, 221]. These parameters were chosen such that the top panel has $f_{\rm td} < f_{\rm ISCO}$, middle has $f_{\rm ISCO} < f_{\rm td} < f_{\rm ringdown}$ and bottom has $f_{\rm td} > f_{\rm ringdown}$. The inset in each plot shows a zoomed in portion of the signal around merger. We apply time and phase shifts to align the waveforms.

a boson star [217] or a gravastar [218]), GW190814's source is therefore either an NSBH or a BBH.

In order to avoid making any unnecessary assumptions, we wish to analyse GW190814 using a single waveform model which remains a good approximation for both systems. Given that there have been some major developments at including precession and subdominant multipole moments into BBH signal models in recent years [222, 80, 81, 79, 223, 224], but not yet NSBH signal models [204, 203], it is preferable to analyse GW190814 assuming the system is either a BBH or can be well approximated by a BBH system. This is because both precession and subdominant multipole moments are likely to leave strong imprints in the gravitational wave given the high mass ratio estimate (see Chapter 3 and Refs [169, 176]. Below we discuss the differences between NSBH and BBH signal models and then describe under what conditions an NSBH signal model can be approximated by an equivalent BBH model.

When in the presence of another object's gravitational field, a star is subject to tidal forces. If these two objects are at a close enough separation $\mathcal{R}$ that the tidal forces are larger than the star's self-gravity, the star is tidally disrupted and will loose mass in a process known as mass shedding (depending on the duration of exposure). By assuming a quasi-circular orbit, we can relate this separation to an orbital frequency $f_{\rm td}$ [225, 226].

Mass shedding modifies both the amplitude and phase of the emitted gravitational wave [204, 203]. However, depending on the value of $f_{\rm td}$, these modulations may or may not be measurable. For certain $f_{\rm td}$ they may not even occur at all. The observability of mass shedding depends on $f_{\rm td}$, the frequency at the Innermost Stable Circular Orbit (ISCO), $f_{\rm ISCO}$, and the frequency at which the binary merges, often approximated by the ringdown frequency of the binary $f_{\rm ringdown}$.

For the case where $f_{\rm td} < f_{\rm ISCO}$, the NS undergoes mass shedding with potentially some of the NS matter remaining outside of the BH in the form of an accretion disk [225]. For such cases, the late inspiral, merger and ringdown present for BBH signals is exponentially suppressed. This results in the emitted GW exhibiting a truncated chirp-like shape with a phase drift compared to the GW emitted by a BBH system, see the top panel of Figure 5.2.

For the case where $f_{\rm ISCO} < f_{\rm td} < f_{\rm ringdown}$, the NS undergoes mass shedding during late inspiral but forms no accretion disk around the remnant black hole. This is because all of the matter that has been shed from the NS lies inside of the ISCO and thus cannot form a bound orbit around the BH. Since most of the NS remains intact as it plunges into the black hole, the emitted GW will be similar to a BBH up until merger where both a mild phase shift and amplitude suppression are seen. For this case the ringdown portion of the signal is also suppressed, see the middle panel of Figure 5.2.

For $f_{\rm td} > f_{\rm ringdown}$, it is assumed that the NS remains completely intact as it

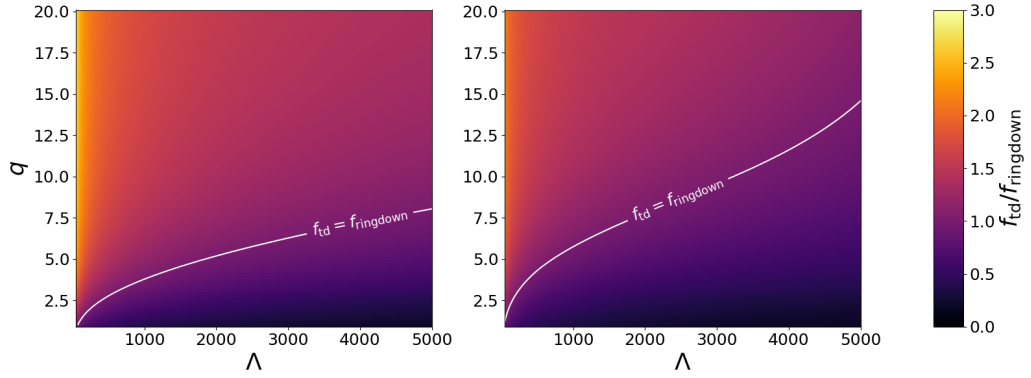

Figure 5.3: Plot showing $f_{\rm td}/f_{\rm ringdown}$ across the mass ratio tidal deformability parameter space for 2 different black hole spins, *left*: $\chi_{1z} = -0.5$, *right*: $\chi_{1z} = 0.5$. The contours shows the region where $f_{\rm td} = f_{\rm ISCO}$. Light (dark) colours show the region of the parameter space where $f_{\rm td} > (<) f_{\rm ringdown}$. We used the IMRPhenomNSBH [203] model to calculate $f_{\rm td}$ and $f_{\rm ringdown}$.

plunges into the BH. Here, the GW amplitude (early and late inspiral, merger and ringdown) is almost indistinguishable from a GW emitted by an equivalent BBH, with often a very small phase drift, see the bottom panel of Figure 5.2. For this case, BBH waveform models are able to approximate the emitted GW. This discussion excludes the effects of precession and subdominant multipole moments. We briefly discuss the impact of these phenomena at end of Sec. 5.3.b.

Whether or not the GW emitted by an NSBH binary can be approximated by a BBH signal model is therefore heavily dependent on the ratio of $f_{\rm td}$ and $f_{\rm ringdown}$. Let us now briefly summarise the dependencies of $f_{\rm td}$ and $f_{\rm ringdown}$ and quantify the region of the NSBH parameter space for which $f_{\rm td} > f_{\rm ringdown}$.

Both the tidal disruption and ringdown frequencies are a function of the BH's mass and spin and the NS's mass and radius [227, 228, 229, 230, 231]. Given that the NS radius is dependent on the unknown NS EOS, it is customary to replace this dependency with a single EOS-invariant parameter: the tidal deformability parameter Λ [232]. Λ encodes how easily an object can be deformed in the presence of an external tidal field [50]. For instance, a large, less compact star is easier to deform than a small, more compact star. This means that a large less compact star has a larger tidal deformability [51]. Λ is bounded to be greater than or equal to 0 where a value of $\Lambda = 0$ corresponds to an object which cannot be deformed under tidal forces. It has been shown that this limit corresponds to a black hole [see e.g. 233, 234, 235].

In Figure 5.3 we show how $f_{\rm td}/f_{\rm ringdown}$ varies across the mass ratio – tidal deformability – black hole spin χ_{1z} parameter space. χ_{1z} can take values between $[-1, 1]$, where positive (negative) values imply that the black hole spin is (anti-)aligned with the orbital angular momentum. We see that $f_{\rm td} > f_{\rm ringdown}$ for smaller tidal deformability and for larger mass ratios [236, 237] with the black hole spin

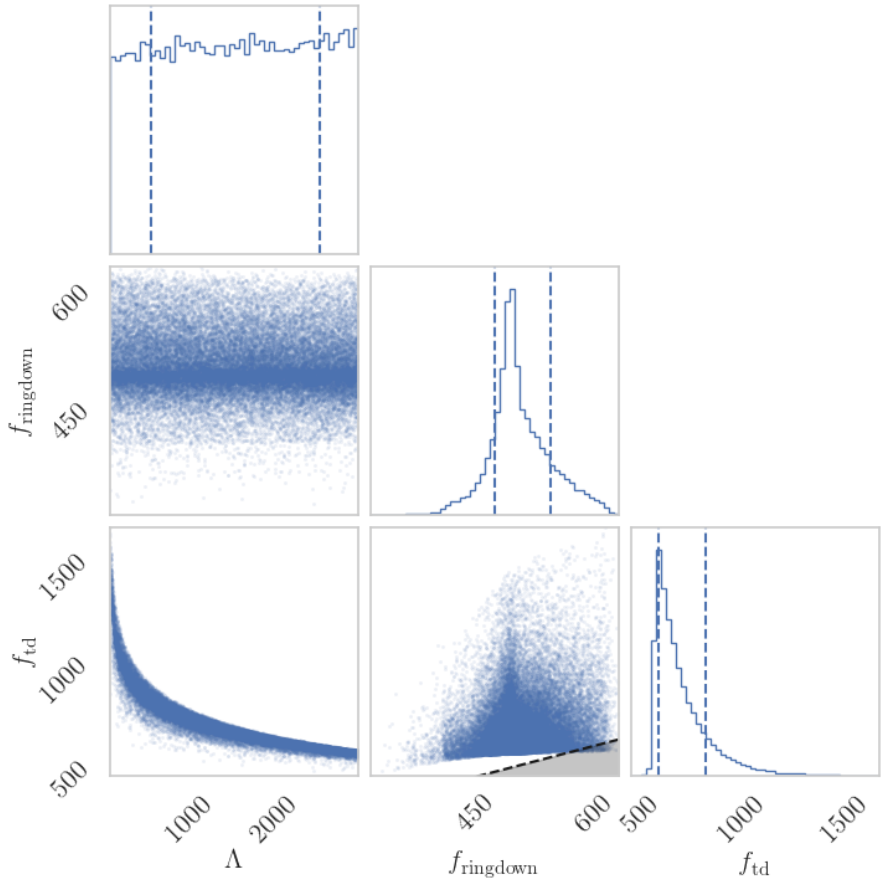

Figure 5.4: Corner plot showing the posterior distribution of the tidal deformability Λ, mass shedding starting frequency $f_{\rm td}$ and ringdown frequency $f_{\rm ringdown}$ for a combined set of samples collected with two NSBH signal models [203, 204]. The histograms along the diagonal show the marginalized 1d posterior distributions. In the $f_{\rm ringdown}$–$f_{\rm td}$ panel, the black dashed line shows $f_{\rm td} = f_{\rm ringdown}$ and the grey region shows $f_{\rm td} < f_{\rm ringdown}$.

having little effect in comparison to Λ and q. We understand this because a) tidal effects scale inversely with mass ratio ($\sim q^{-4}$) [203, 204] meaning that for unequal mass ratio systems mass shedding starts at higher orbital frequencies and b) NSs with smaller tidal deformabilities are smaller, more compact and have a larger self-gravity which requires larger tidal forces (smaller separations) to overcome.

5.3.b Motivating the use of BBH signal models for the analysis of GW190814

If we wish to describe GW190814 with BBH signal models, we require $f_{\rm td} > f_{\rm ringdown}$. From Figure 5.3, we understand this requires large mass ratios and/or small tidal deformabilites. To obtain estimates for these parameters, we performed parameter

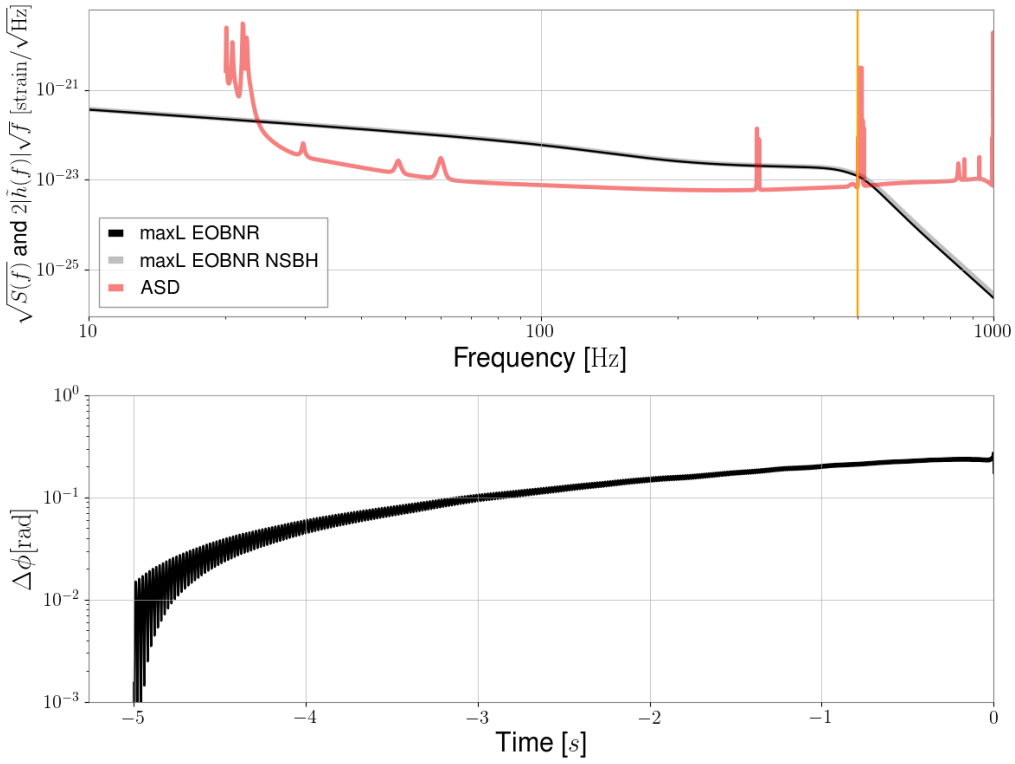

Figure 5.5: Comparison between the maximum likelihood waveforms for GW190814 inferred from two different analyses: one where GW190814's source is assumed to be a BBH (EOBNR) and another where GW190814's source is assumed to be an NSBH (EOBNR NSBH). The top panel shows both maximum likelihood waveforms projected onto the LIGO-Hanford detector within the frequency domain. For comparison we also show the amplitude spectral density (ASD) of LIGO-Hanford at the time of the detection in red. The vertical orange line shows an estimate for the merger frequency $f_{\rm ringdown}$. We apply a normalisation factor ($2\sqrt{f}$) to all waveforms such that the area between the signal and ASD is indicative of the SNR of GW190814 [32, 30]. The bottom panel shows the accumulated time-domain phase error between the two maximum likelihood waveforms over the length of the EOBNR NSBH signal. We apply time and phase shifts to aligned the waveforms.

estimation using the EOBNR NSBH [204] and Phenom NSBH [203] signal models[3]; two complementary models representing two different approaches to model the same binary system, one based on the effective-one-body approach (EOBNR; [219, 79, 223, 224]) and the other on a phenomenological approach (Phenom; [76, 75, 222, 80, 81]).

Figure 5.4 shows the three-dimensional joint posterior distribution for $f_{\rm td}$, $f_{\rm ringdown}$ and the tidal deformability Λ for the combination of EOBNR NSBH and Phenom NSBH samples. We see that $f_{\rm td} > f_{\rm ringdown}$ at more than 99% probability. This implies that if GW190814's source is an NSBH binary, it is extremely likely that the NS plunged into the BH without a) being tidally disrupted, although we are unable

[3]In LALSimulation [175] the full name of the EOBNR model is SEOBNRv4_ROM_NRTidalv2_NSBH and Phenom model is IMRPhenomNSBH

	EOBNR	EOBNR NSBH
Primary mass $m_1/M_\odot$	24.43	23.76
Secondary mass $m_2/M_\odot$	2.46	2.49
Inverted mass ratio $1/q$	0.1	0.1
Chirp mass $\mathcal{M}/M_\odot$	6.04	6.02
Total mass $M/M_\odot$	26.89	26.25
Primary spin magnitude χ_1	0.08	0.04
Effective inspiral spin parameter $\chi_{\rm eff}$	0.04	0.03
Tidal deformability parameter Λ	-	635.37
Luminosity distance $D_{\rm L}$/Mpc	284.14	302.93
Source redshift z	0.06	0.07
Inclination angle θ_{JN}/rad	0.64	0.52
Network Signal to noise ratio $\rho_{\rm HLV}$	23.97	23.97

Table 5.1: Table showing the maximum likelihood values for GW190814 inferred from two different analyses: one where GW190814's source is assumed to be a BBH (EOBNR) and another where GW190814's source is assumed to be an NSBH (EOBNR NSBH). The inclination angle is folded to $[0, \pi/2]$.

to exclude the possibility of tidal disruption or b) the SNR of the signal is not large enough measure the very small phase drift (see Fig. 3 in Ref [203]). This implies that any differences in the amplitude and phase compared to a BBH signal model are at magnitude that cannot be measurable at current detector sensitivities. This explains why our analysis returned an uninformative posterior distribution for the tidal deformability.

Interestingly, our analysis reveals a sharp cutoff in $f_{\rm td}$ at $\sim 500\,$Hz. From Figure 5.4 we see that this cutoff is a result of limiting the tidal deformability $\Lambda < 3000$ (prior railing). If Λ were allowed to take values larger than 3000, the NS would be larger and less compact for a given mass. This implies that a lower tidal force would be required to initiate mass shedding, which, given that tidal forces scale inversely with binary separation (for fixed configuration), occurs earlier in the inspiral and therefore at lower frequencies. Our analysis adopted a boundary at $\Lambda = 3000$ due to limitations with the available NSBH signal models [203]. This restriction is reasonable given that GW170817 constrained the NS tidal deformabilities $\Lambda < 3000$ at more than 99% probability.

Given that $f_{\rm td} > f_{\rm ringdown}$ at more than 99% probability, it is likely that GW190814 can be well approximated by a BBH signal model. To confirm this, we analysed GW190814 using a set of aligned spin quadrupole only BBH signal models and compared maximum likelihood samples (maxLs) with the NSBH signal models. By doing so, we compare waveforms which best match the data and neglect any differences in prior volumes. These BBH models, SEOBNRv4_ROM [219] (EOBNR) and IMRPhenomD [76, 75] (Phenom), are close to, and if not exactly equivalent to, the BBH limit of the NSBH signal models used previously.

Figure 5.5 plots the recovered maxL waveforms from the EOBNR and EOBNR NSBH analyses in the frequency domain. We see that the two maxL waveforms agree very well. The amplitudes are almost indistinguishable and the phase difference is comparable to the error between EOBNR NSBH and numerical relativity waveforms [204]. We note that over time, the phase error increases. This is consistent with the fact that the tidal-phase correction accumulates over many cycles. We also highlight that a small difference in amplitude is seen for frequencies post merger $f \gtrsim f_{\text{ringdown}}$. However, since these differences lie below the ASD, they cannot be detected at current detector sensitivities.

We also show the maxL values for each parameter in table 5.1. Although in general we see very good agreement between the two analyses, we see differences in the inferred distance and inclination angle. This is expected given that the distance-inclination angle degeneracy [60, 8] is only broken with aligned spin quadrupole only models for a small fraction of binaries: those that are viewed close to edge-on $75^\circ \lesssim \theta_{JN} \lesssim 105^\circ$ [99]. We do observe that the waveform amplitudes ($\cos\theta_{JN}/D_{\text{L}}$, $(1+\cos\theta_{JN}{}^2)/2D_{\text{L}}{}^4$) remain approximately constant between the models. Based on this comparison, we can conclude that if GW190814 originated from an aligned spin NSBH binary, it is exceptionally well modelled by an aligned spin BBH merger. This confirms the conclusions in Ref. [236].

So far, we have excluded the effects of precession and subdominant multipole moments in our discussion. However, it has been shown previously that a) tidal disruption of the NS is less likely in a precessing NSBH binary [see e.g. 238, 239] and b) higher order tidal corrections caused by subdominant multipole moments introduces a $\leqslant 0.7\%$ correction [240]. Consequently, if the effects of precession and subdominant multipole moments were included in our analysis, a GW190814-like NSBH merger would still be exceptionally well modelled by a BBH merger. Given that precessing higher order multipole moment models currently only exist for BBH systems, these results show that it is reasonable to perform a full Bayesian analysis of GW190814 using these BBH signal models which contain more physics.

5.3.c Analysing GW190814 with BBH waveform models

Our primary analyses include the effects of subdominant multipole moments (also referred to as higher order multipole moments HM) in precessing waveform template models (PHM): IMRPhenomPv3HM (Phenom PHM; [80, 81]) from the phenomenological family and SEOBNRv4PHM (EOBNR PHM; [79, 224]) from the EOBNR family. We find no significant evidence that one waveform family is preferred over the other as the Bayes factor between Phenom PHM and EOBNR PHM is $\log_{10}\mathcal{B} \simeq 1.0$. As a result, we combine the posterior samples with equal weight, in

[4]The waveform amplitudes are defined in terms of the angle between the orbital angular momentum and line of sight ι, however since both EOBNR and EOBNR NSBH signal models are non-precessing, $\iota = \theta_{JN}$ and therefore we are free to use either variable

	EOBNR PHM	Phenom PHM	Combined
Primary mass $m_1/M_\odot$	$23.2^{+1.0}_{-0.9}$	$23.2^{+1.3}_{-1.1}$	$23.2^{+1.1}_{-1.0}$
Secondary mass $m_2/M_\odot$	$2.59^{+0.08}_{-0.08}$	$2.58^{+0.09}_{-0.10}$	$2.59^{+0.08}_{-0.09}$
Inverted mass ratio $1/q$	$0.112^{+0.008}_{-0.008}$	$0.111^{+0.009}_{-0.010}$	$0.112^{+0.008}_{-0.009}$
Chirp mass $\mathcal{M}/M_\odot$	$6.10^{+0.06}_{-0.05}$	$6.08^{+0.06}_{-0.05}$	$6.09^{+0.06}_{-0.06}$
Total mass $M/M_\odot$	$25.8^{+0.9}_{-0.8}$	$25.8^{+1.2}_{-1.0}$	$25.8^{+1.0}_{-0.9}$
Final mass $M_\mathrm{f}/M_\odot$	$25.6^{+1.0}_{-0.8}$	$25.5^{+1.2}_{-1.0}$	$25.6^{+1.1}_{-0.9}$
Upper bound on primary spin magnitude χ_1	0.06	0.08	0.07
Effective inspiral spin parameter χ_eff	$0.001^{+0.059}_{-0.056}$	$-0.005^{+0.061}_{-0.065}$	$-0.002^{+0.060}_{-0.061}$
Upper bound on effective precession parameter χ_p	0.07	0.07	0.07
Final spin χ_f	$0.28^{+0.02}_{-0.02}$	$0.28^{+0.02}_{-0.03}$	$0.28^{+0.02}_{-0.02}$
Luminosity distance D_L/Mpc	235^{+40}_{-45}	249^{+39}_{-43}	241^{+41}_{-45}
Source redshift z	$0.051^{+0.008}_{-0.009}$	$0.054^{+0.008}_{-0.009}$	$0.053^{+0.009}_{-0.010}$
Inclination angle θ_{JN}/rad	$0.9^{+0.3}_{-0.2}$	$0.8^{+0.2}_{-0.2}$	$0.8^{+0.3}_{-0.2}$
Signal to noise ratio in LIGO-Hanford ρ_H	$10.6^{+0.1}_{-0.1}$	$10.7^{+0.1}_{-0.2}$	$10.7^{+0.1}_{-0.2}$
Signal to noise ratio in LIGO-Livingston ρ_L	$22.21^{+0.09}_{-0.15}$	$22.16^{+0.09}_{-0.17}$	$22.18^{+0.10}_{-0.17}$
Signal to noise ratio in Virgo ρ_V	$4.3^{+0.2}_{-0.5}$	$4.1^{+0.2}_{-0.6}$	$4.2^{+0.2}_{-0.6}$
Network Signal to noise ratio ρ_HLV	$25.0^{+0.1}_{-0.2}$	$24.9^{+0.1}_{-0.2}$	$25.0^{+0.1}_{-0.2}$

Table 5.2: Source properties of GW190814: We report the median values along with the symmetric 90% credible intervals for the SEOBNRv4PHM (EOBNR PHM) and IMRPhenomPv3HM (Phenom PHM) waveform models. The primary spin magnitude and the effective precession is given as the 90% upper limit. The inclination angle is folded to $[0, \pi/2]$. The last column is the result of combining the posteriors of each model with equal weight.

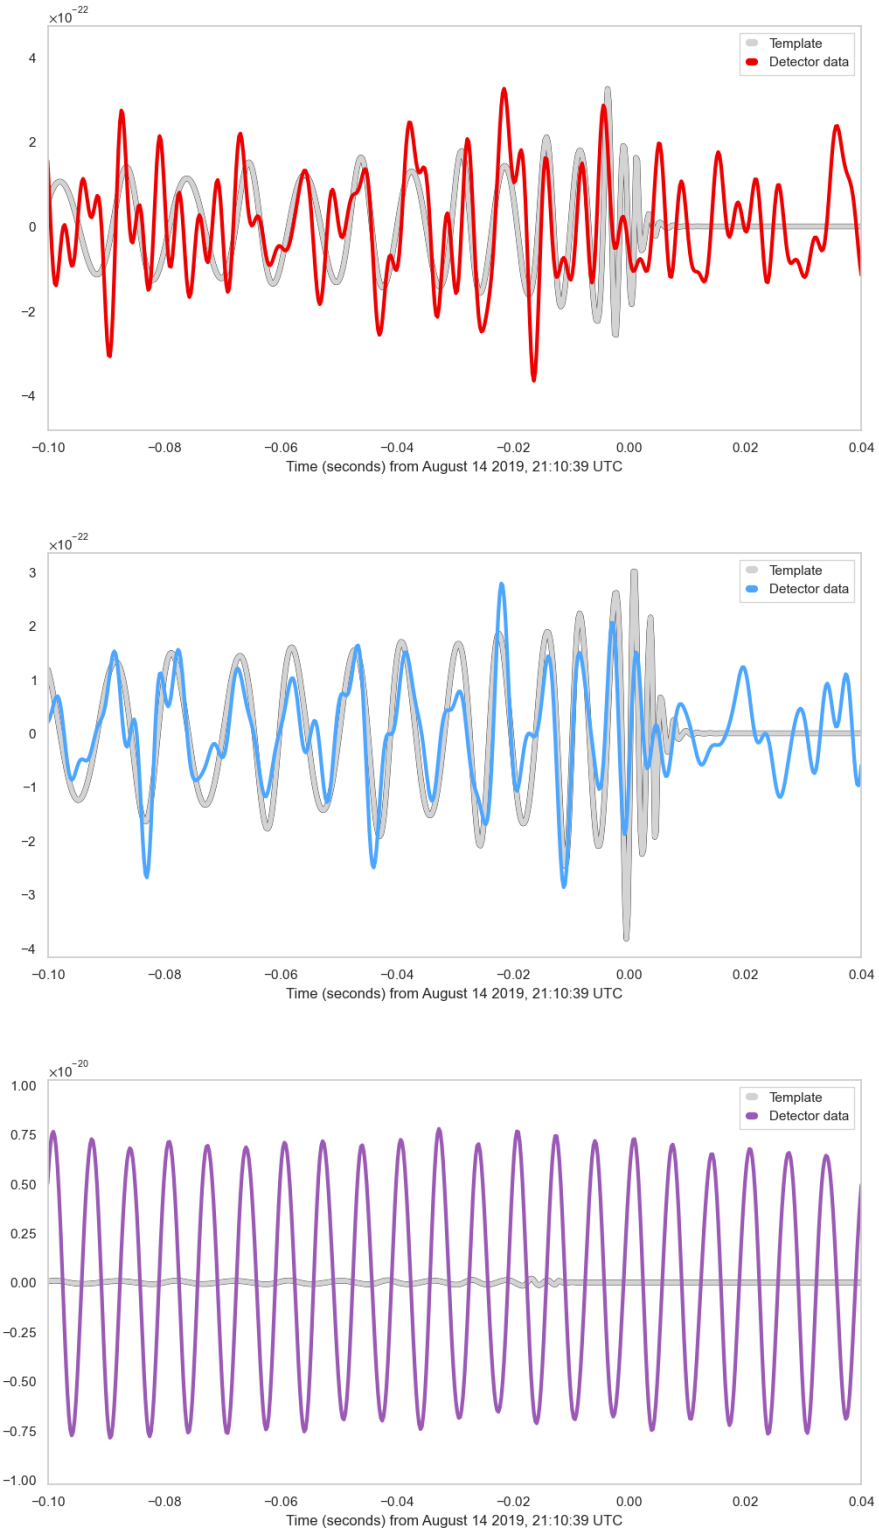

Figure 5.6: Time-domain data (sampled at 4096 Hz) and reconstructed waveforms of GW190814, for the *top*: LIGO-Hanford, *middle*: LIGO-Livingston and *bottom*: Virgo GW detectors. Times are shown relative to August 14 2019 at 21:10:39 UTC.

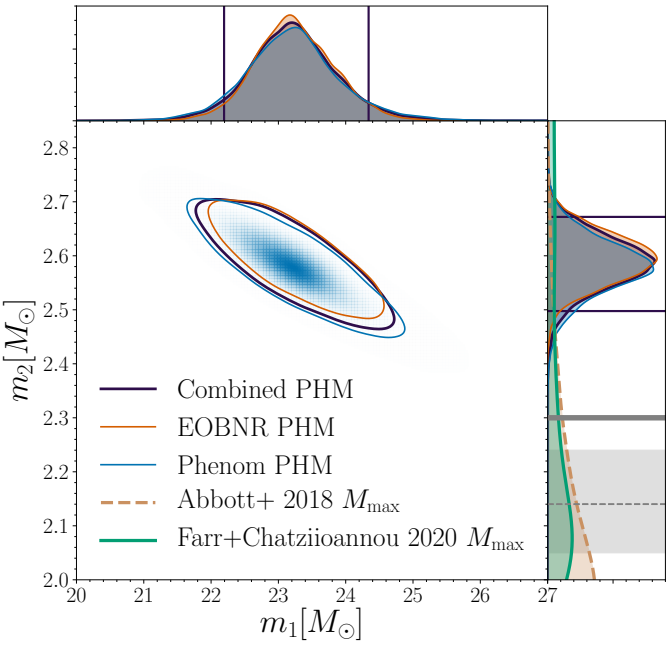

Figure 5.7: Posterior distribution of the primary and secondary source masses for the SEOBNRv4PHM and IMRPhenomPv3HM waveform models. The posterior distribution resulting from the combination of their samples is also shown. Each contour, as well as the coloured horizontal and vertical lines, shows the 90% credible intervals. The right panel compares m_2 to predictions for the maximum NS mass, $M_{\rm max}$ from a) studies of the remnant of GW170817 (solid grey band) b) theoretical estimates (orange) and c) fitting the known population of NSs in binaries (green). The grey dashed line and shading represent the measured mass of the heaviest pulsar in the Galaxy (median and 68% confidence interval; [186]).

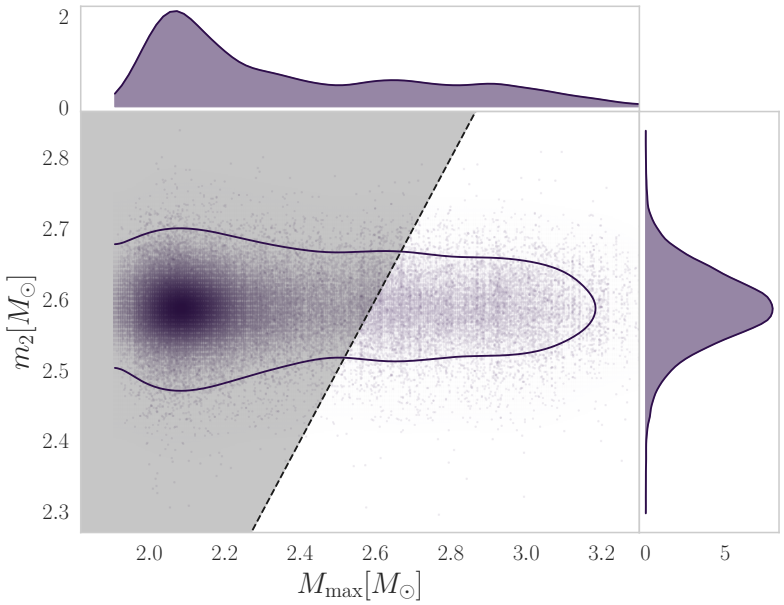

Figure 5.8: Posterior distribution of the secondary source mass compared to the posterior distribution for the maximum NS mass estimated by fitting the known population of NSs in binaries [241] (green curve in the right hand panel of Figure 5.7). The black dashed line shows $m_2 = M_{\rm max}$ and the grey region shows $m_2 > M_{\rm max}$.

effect marginalizing over a discrete set of signal models with a uniform probability. A summary of the inferred source properties of GW190814 is given in Table 5.2. Figure 5.6 compares the time-domain data and the reconstructed waveforms for GW190814 for the 3 GW detectors used in this analysis.

We found that the secondary mass lies in the range 2.50–2.67 $M_\odot$, placing it in the hypothesised lower mass gap between known NSs and BHs, 2.0–5 $M_\odot$ [188, 189, 190, 191], see Figure 5.7. The secondary object is heavier than the most massive known pulsar in the Galaxy ($2.14^{+0.10}_{-0.09}\,M_\odot$ at 68.3% credible interval; [186]), the primary component of GW190425 (1.61–2.52 $M_\odot$; [43]), and the bounds on maximum NS mass from a) studies of the remnant of GW170817, b) theoretical estimates [242] and c) fitting the known population of NSs in binaries [241]. It is, however, comparable to the $\sim 2.7\,M_\odot$ putative BH remnant mass of GW170817 [243]. The primary object is most likely a BH with mass $23.2^{+1.1}_{-1.0} M_\odot$.

In Figure 5.8 we directly compare the posterior distribution for the secondary mass with the maximum NS mass estimated by fitting the known population of NSs in binaries [241]. We identify that the secondary mass is less than the maximum NS mass at a probability $\mathrm{P}(m_2 \leq M_{\mathrm{max}}) \sim 0.29$. We therefore favour the $m_2 > M_{\mathrm{max}}$ scenario, albeit not very strongly because of the distribution's long tail up to $\sim 3\,M_\odot$. However, the empirical M_{max} prediction is sensitive to selection effects that could potentially bias its posterior [244]. This posterior also does not take into consideration the discovery of GW190425 [43] which could impact the M_{max} prediction.

The time delay of a signal across a network of gravitational wave detectors, together with the relative amplitude and phase at each detector, allows us to measure the location of the GW source on the sky [245]. We localize GW190814's source to within $18.5\,\mathrm{deg}^2$ at 90% probability, as shown in Figure 5.9. Despite this tight constraint, we fail to localise GW190814 to a single point on the sky. As shown in Figure 5.9, the secondary sky position is caused by a small secondary peak in the arrival time at Virgo. This is consistent with a secondary peak in the Virgo SNR time series reported by the search pipelines.

Spins are a fundamental property of BHs. Their magnitude and orientation carry information regarding the evolution history of the binary. We parameterise the spin components parallel to the orbital angular momentum by the effective inspiral spin parameter χ_{eff} [246, 247, 248, 249]. We infer that $\chi_{\mathrm{eff}} = -0.002^{+0.060}_{-0.061}$. The tight constraints are consistent with being able to measure the phase evolution from the long inspiral.

Orbital precession occurs when there is a significant spin component perpendicular to the orbital angular momentum (within the orbital plane of the binary) [54]. We parameterize precession by the effective precession spin parameter $0 \leq \chi_{\mathrm{p}} \leq 1$ [64]. It is well known that orbital precession is easier to measure for asymmetric mass ratios (see Chapter 3 and Ref. [169]) and edge-on systems [54, 88, 102, 104, 120,

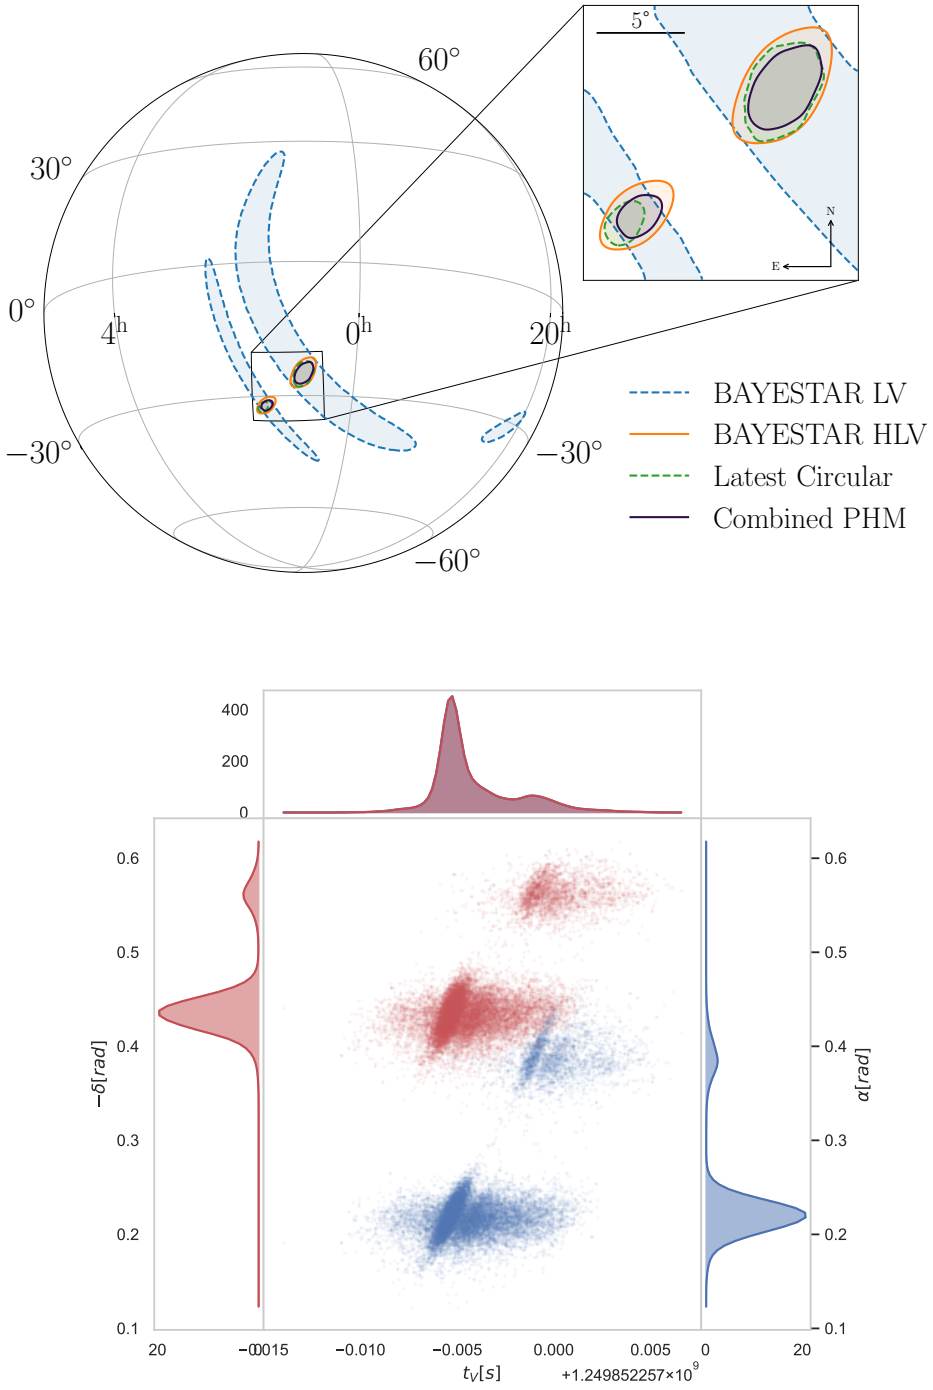

Figure 5.9: Posterior distributions for the sky location of GW190814. The top panel shows the 90% credible interval for a LIGO Livingston–Virgo (blue) and LIGO Hanford–LIGO Livingston–Virgo (orange) detector network based on the rapid localization algorithm BAYESTAR [97]. The sky localization circulated 13.5 hours after the event, based on a LIGO Hanford–LIGO Livingston–Virgo bayesian analysis is shown in green. The purple contour indicates the final sky localization. The bottom panel plots two 2D probability distributions on the same axis. In blue we show the 2D probability distribution for the merger time as reported by the Virgo gravitational wave detector t_V and the right ascension α and in red we show the 2D probability distribution for t_V and declination δ (red).

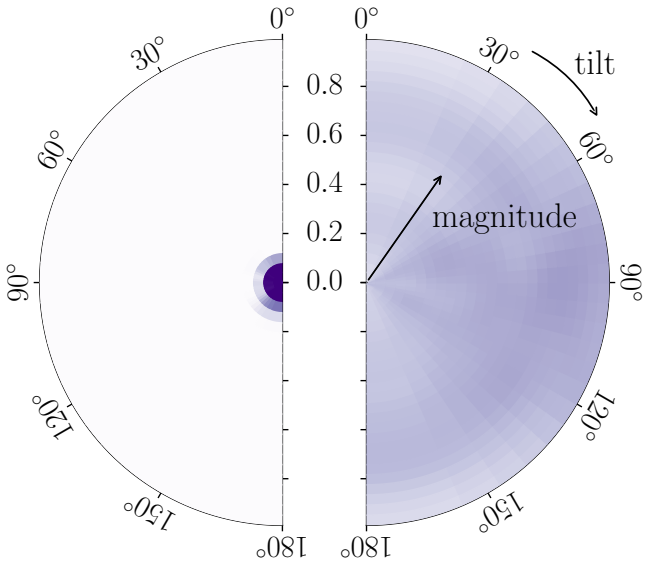

Figure 5.10: Two-dimensional posterior probability for the tilt-angle and spin-magnitude for the primary object (left) and secondary object (right) based on the Combined samples. The tilt angles are 0° for spins aligned and 180° for spins anti-aligned with the orbital angular momentum. The tiles are constructed linearly in spin magnitude and the cosine of the tilt angles such that each tile contains identical prior probability. The color indicates the posterior probability per pixel. The probabilities are marginalized over the azimuthal angles.

86, 145]. Given that GW190814 constrains the inclination of the binary to be $\theta_{JN} = 0.8^{+0.3}_{-0.2}$ rad (folded to $[0, \pi/2]$) and has an inferred inverted mass ratio of $1/q = 0.112^{+0.008}_{-0.009}$, GW190814 provides the tightest constraint on precession from any GW to date: $\chi_{\rm p} = 0.04^{+0.04}_{-0.03}$. By computing the Bayes factor between a precessing and non-precessing signal model ($\log_{10}\mathcal{B} \sim 0.5$ in favor of precession), we find inconclusive evidence for in-plane spin.

The asymmetry in the masses of GW190814 means that the spin of the more massive object dominates contributions to $\chi_{\rm eff}$ and $\chi_{\rm p}$. As both $\chi_{\rm eff}$ and $\chi_{\rm p}$ are tightly constrained, we are able to bound the primary spin of GW190814 to be $\chi_1 \leq 0.07$ as shown in Figure 5.10. This is the strongest constraint on the primary spin for any gravitational-wave event to date [38, 43, 45].

The joint posterior probability of the magnitude and orientation of χ_1 and χ_2 are shown in Figure 5.10. This plot is constructed such that the prior has uniform shading and therefore any deviations from uniform shading indicates a distribution which differs from the prior. We see that the primary spin is tightly constrained to small magnitudes, but its orientation is indistinguishable from the prior distribu-

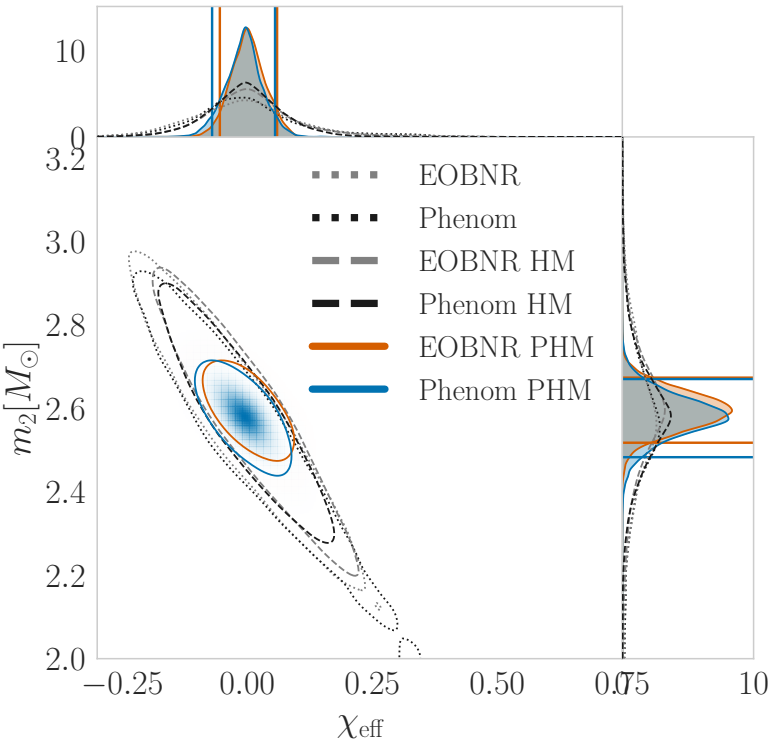

Figure 5.11: Posterior distribution of the effective spin and secondary source mass using a suite of waveform models. Each contour shows the 90% credible intervals. The top and right panels show the marginalized posterior distribution for effective spin and secondary source mass respectively. The labels Phenom/EOBNR PHM (generic spin directions + higher multipoles), Phenom/EOBNR HM (aligned-spin + higher multipoles) and Phenom/EOBNR (aligned-spin, quadrupole only) indicate the different physical content in each of the waveform models.

tion. The spin of the less massive object, χ_2, remains unconstrained; the posterior distribution is broadly consistent with the prior.

5.3.d Near-zero precession impacting accuracy of source properties

Our coherent Bayesian analysis has shown that GW190814 is consistent with a $23.2^{+1.1}_{-1.0}M_\odot - 2.59^{+0.08}_{-0.09}M_\odot$ binary merger with near-zero in-plane and aligned spin components. Given that GW190814 lies in a region of the parameter space that has previously been unexplored via gravitational-wave emission, we tested the accuracy of our gravitational wave models by comparing the Phenom and EOB waveform families. Differences in the inferred secondary mass and effective spin are shown in Figure 5.11. The results indicate that both the Phenom and EOB waveform families agree exceptionally well with the inferred secondary mass and effective spin magnitude robust to possible waveform systematics.

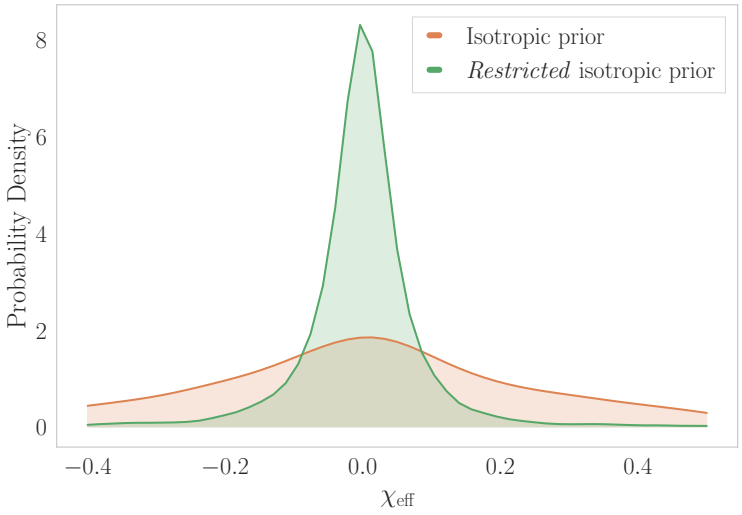

Figure 5.12: Two different prior distributions for $\chi_{\rm eff}$. The $\chi_{\rm eff}$ prior arising from isotropic priors on the component spins (orange) and the *restricted* $\chi_{\rm eff}$ prior (green) which is equivalent to the orange curve but with the constraint that $\chi_{\rm p} < 0.05$.

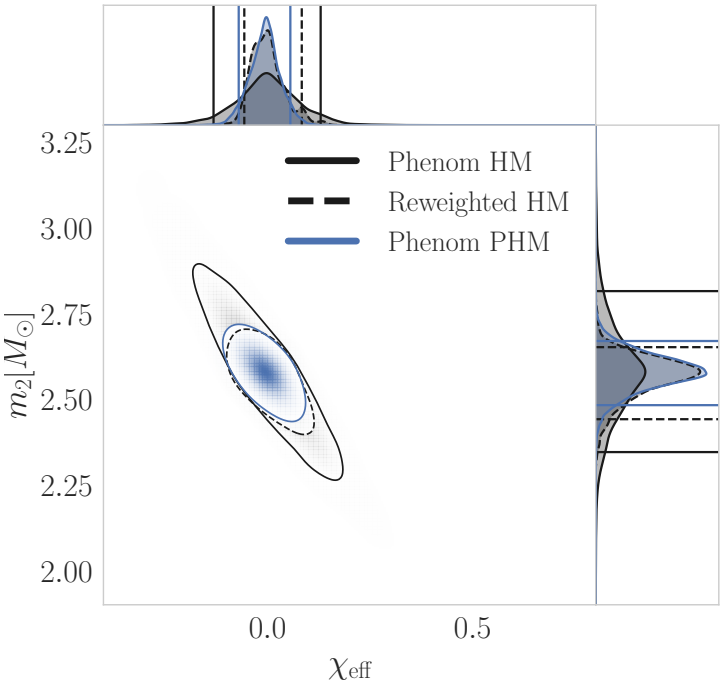

Figure 5.13: Posterior distribution of the effective spin and secondary source mass for an aligned spin and a precessing analysis. Each contour shows the 90% credible intervals. The top and right panels show the marginalized posterior distribution for effective spin and secondary source mass respectively. The Reweighted HM analysis reweights the Phenom HM analysis to use the *restricted* isotropic prior as described in Figure. 5.12.

Surprisingly, Figure 5.11 also shows that the non-measurement of precession impacts the inferred mass ratio and effective spin of the binary, with precessing signal models constraining the mass ratio and effective spin more than aligned spin signal models despite measuring near-zero $\chi_{\rm p}$. This arises because inferences between in-plane and aligned spin components are correlated owing to the prior distribution used. To remain agnostic, we often use an isotropic distribution for spin orientations and a uniform distribution for spin magnitudes [see Appendix B.1 of 38]. This prior leads to significantly larger support for small values of aligned spin upon an inference of small in-plane spin. To demonstrate this, Figure 5.12 compares 2 priors for $\chi_{\rm eff}$: an isotropic prior described above, and a *restricted* isotropic prior which is equivalent to the isotropic prior with the constraint that $\chi_{\rm p} < 0.05$. We see that an inference of low in-plane spin causes $\sim 4\times$ as much support for $\chi_{\rm eff} = 0$ compared to the isotropic case, with close to zero prior support for $|\chi_{\rm eff}| \gtrsim 0.2$. This means that as GW190814 is consistent with near-zero $\chi_{\rm p}$, precessing signal models effectively use the *restricted* isotropic prior for $\chi_{\rm eff}$, whereas because aligned spin signal models cannot make any inferences about the in-plane spin, they use an equivalent of the wider isotropic prior [see equation A7 in 250]. Since the mass ratio and aligned spin are degenerate [62] the tighter constraint on $\chi_{\rm eff}$ for the precessing signal models leads to a tighter constraint on the mass ratio, enabling for a more precise measurement of the secondary mass.

To demonstrate that the improved measurement in $\chi_{\rm eff}$ and m_2 is a result of a near-zero measurement on $\chi_{\rm p}$, we can simply re-weight the aligned spin results to use the *restricted* isotropic prior. In Figure 5.13 we perform this re-weighting with the Phenom HM analysis and we now see excellent agreement between the 90% credible intervals of the Phenom PHM and the reweighted Phenom HM analysis. We have therefore shown that the non-measurement of precession allows for tighter constraints on the effective spin and secondary mass when using the isotropic prior. However, it can be shown that this effect is robust and does not depend on the selected prior. For instance, we see a similar effect if we use a more astrophysically motivated distribution for $\chi_{\rm eff}$ (for example the distribution in Ref. [170]). This effect has also been studied for NSBH systems in Ref. [251].

5.3.e Minimal assumption analysis of GW190814

As we have seen, GW190814 is unlike any other gravitational wave signal observed to date. Not only does our analysis indicate that GW190814 was detected with the second largest signal-to-noise ratio (SNR) of any gravitational wave detection $\simeq 25$, but its source has the most unequal mass ratio observed with gravitational waves: $1/q = 0.112^{+0.008}_{-0.009}$. Given that relative importance of a subdominant multipole moment increases with mass ratio, GW190814 exhibited the strongest evidence for higher-order multipoles compared to any previous gravitational wave observa-

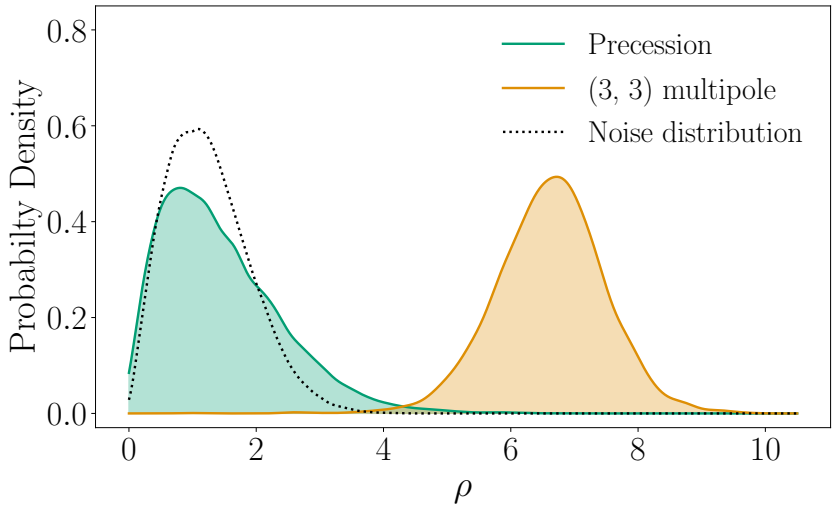

Figure 5.14: Posterior distributions for the precessing SNR, $\rho_{\rm p}$ (green) and the optimal SNR in the (3,3) sub-dominant multipole moment, ρ (orange). The grey dotted line shows the expected distribution for Gaussian noise.

tion with a Bayes Factor in favour of a higher-multipole vs. a pure quadrupole model of $\log_{10}\mathcal{B} \simeq 9.6$. Through computing the optimal SNR in each subdominant multipole [182], we found that the $(\ell, m) = (3, 3)$ multipole was strongest with an orthogonal optimal SNR of $\rho_{33} = 6.6^{+1.3}_{-1.4}$, as shown in Figure 5.14.

GW190814 also provides the tightest constraint on precession from any GW to date. This, combined with the constraint on the inclination angle, unequal mass ratio and large SNR, meant that the inferred power from precession (see Chapter 2) was $\rho_{\rm p} = 1.27^{+1.9}_{-1.0}$. This resembles the expected distribution in the absence of any precession in the signal, as shown in Figure 5.14.

Given the strength of both ρ_{33} and $\rho_{\rm p}$, only a limited region of parameter space is consistent with these SNRs. For instance, as demonstrated in Ref. [169], if GW190814's source had $\chi_{\rm p} > 0.1$, we would recover a significant Bayes factor with precession SNR distinctly different from the noise distribution. We therefore wish to perform a minimal assumption analysis of GW190814 and limit the parameter space consistent with ρ_{33} and $\rho_{\rm p}$ without performing an expensive Bayesian analysis. We choose to focus our attention on the orientation, mass ratio, and $\chi_{\rm p}$ parameter space, although in principle it could be extended for any parameter choice. We define the orientation to be the angle between the total angular momentum and the line of sight θ_{JN} such that it remains approximately constant for precessing systems. The technique described below is independent of the distance to the source and will be useful if search pipelines are capable of computing ρ_{33} and $\rho_{\rm p}$ (see Sec. 2.5 for details about how $\rho_{\rm p}$ could be returned by the search pipelines).

We start with the chirp mass, merger time and SNR for the best matched template from the search pipelines [207, 208, 126, 209, 210, 211, 127, 137, 212]

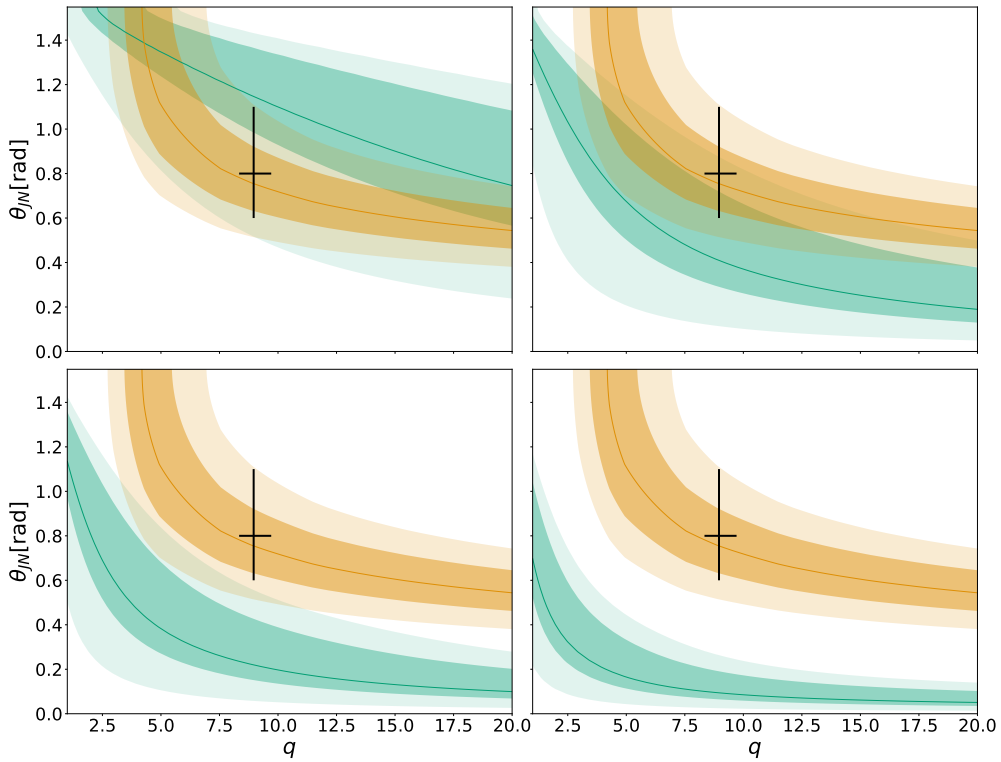

Figure 5.15: Region of the inclination θ_{JN}, mass ratio $q = m_1/m_2 > 1$ and in-plane spin $\chi_{\rm p}$ parameter space that is consistent with the chirp mass and merger time from the best matched template and the reported SNR in the (3, 3) multipole ρ_{33} (orange) and the SNR from precession $\rho_{\rm p}$ (green). Reading from left to right and down, $\chi_{\rm p}$ increases; in the *top left*: plot we fix $\chi_{\rm p} = 0.01$, *top right*: $\chi_{\rm p} = 0.05$, *bottom left*: $\chi_{\rm p} = 0.1$ and *bottom right*: $\chi_{\rm p} = 0.25$. The solid green and orange lines show the region of the parameter space consistent with the peak of the inferred distribution for $\rho_{\rm p}$ and ρ_{33}. The darker and lighter colored regions encase 68% and 95% of the ρ_{33} and $\rho_{\rm p}$ distributions. The single data point with associated errors corresponds to the reported inclination and mass ratio values in Table 5.2.

and the sky location from the skymap produced by the rapid Bayesian algorithm BAYESTAR [97], all of which are available through Gravitational-Wave Candidate Event Database [252]. We then assume that the peaks of the inferred distributions for $\rho_{\rm p}$ and ρ_{33} calculated from a full Bayesian analysis are equivalent to what would be returned from by the search pipeline. We then associate an uncertainty for each of these SNRs according to non-central χ^2 with 2 degrees of freedom. This was shown to be a reasonable approximation in Ref [182]. Next we grid the orientation, mass ratio, and $\chi_{\rm p}$ parameter space using 6400 bins. We restrict the orientation to take values between $0 < \theta_{JN} < \pi/2$, mass ratio between $1 < q < 20$ and $\chi_{\rm p} \in [0.01, 0.05, 0.1, 0.25]$. We then calculate $\rho_{\rm p}$ and ρ_{33} for each bin and isolate the region consistent with the reported values. We also assume that $\chi_{\rm eff} = 0$ throughout this analysis, however, it is trivial to include this as an extra degree of freedom without significantly increasing the computational cost.

First we consider the region of parameter space consistent with ρ_{33}. Ref. [182] provides a relationship between ρ_{33}, the inclination angle[5], the intrinsic properties of the source and the SNR in the dominant quadrupole ρ_{22}, valid only for non-precessing systems,

$$\rho_{33} = 2\rho_{22}\,\alpha_{33}\,\sin\theta_{JN}, \tag{5.9}$$

where α_{33} encodes the relative amplitude of (3, 3) multipole and is a function of the intrinsic properties of the source[6]. Given that ρ_{33} is dependent on the luminosity distance to the source, the GW detector network and its sensitivity, we consider the detector and distance invariant quantity ρ_{33}/ρ_{22}. Since $\rho_{33} = 6.6^{+1.3}_{-1.4}$ and $\rho_{22} = 24.5$, we bound $\rho_{33}/\rho_{22} \in [0.19, 0.35]$ ($\rho_{33}/\rho_{22} \in [0.23, 0.31]$) at 95% (68%) confidence.

In Figure 5.15 we show the region of the parameter space consistent with the ρ_{33}/ρ_{22} measurement. Given we also vary the in-plane spin of the binary $\chi_{\rm p}$, we make the assumption that Eq. 5.9 holds for precessing systems with $\chi_{\rm p} \in [0.01, 0.05, 0.1, 0.25]$. This assumption is valid for small values of $\chi_{\rm p}$ but becomes invalid as $\chi_{\rm p}$ increases. This is why the region consistent with ρ_{33}/ρ_{22} remains constant in all panels. Focusing on the top left panel, we see that as θ_{JN} decreases, q must increase to maintain constant ρ_{33}. We understand this because the significance of the (3, 3) multipole (α_{33}) increases with mass ratio [see e.g. Refs. 182, 253, 254, 255, 256] and therefore $\sin\theta_{JN}$ must decrease to compensate. We also see that for $q \lesssim 3$ there is no region of the parameter space consistent with ρ_{33}/ρ_{22}. This because for $\rho_{33}/\rho_{22} = 0.19$ (lower bound of the 95% credible interval), $\alpha_{33} > 0.095$. For GW190814's chirp mass, this is only achievable for $q \gtrsim 2.6$.

Next we consider the region of parameter space consistent with $\rho_{\rm p}$. In Chapter 2, we defined $\rho_{\rm p}$ as,

$$\rho_{\rm p} = \rho_{22}\left(\frac{\min(1, |\zeta|)}{\sqrt{1+|\zeta|^2}}\right) \tag{5.10}$$

where ζ is a function of the inclination angle, detector response functions $F_{+,\times}$, the reference precession phase ϕ_{JL} and the binary's opening angle $\overline{b} \approx \tan\left(\overline{\beta}/2\right)$ (this is valid when β is approximately constant over the observed waveform),

$$\zeta \;=\; \overline{b}e^{i\phi_{JL}}\left(\frac{F_+\sin 2\theta_{JN} + 2iF_\times\sin\theta_{JN}}{\frac{1}{2}F_+(1+\cos^2\theta_{JN}) + iF_\times\cos\theta_{JN}}\right). \tag{5.11}$$

Similar to ρ_{33}, we see that $\rho_{\rm p}$ is also dependent on the luminosity distance,

[5]Ref. [182] defines the equation for ρ_{33} in terms of the angle between the orbital angular momentum and the line of sight ι. Since this equation is only valid for non-precessing systems, ι is equivalent θ_{JN}

[6]α_{33} also depends on the detector sensitivity but since this is a weak dependence, we ignore this for our analysis

GW detector network and its sensitivity. We therefore consider the ratio $\rho_{\rm p}/\rho_{22}$ which is generally detector invariant (see Eq. 2.34 and surrounding text for details). Consequently, we bound $\rho_p/\rho_{22} \in [0.01, 0.14]$ ($\rho_p/\rho_{22} \in [0.04, 0.10]$) at 95% (68%) confidence. Unlike ρ_{33}/ρ_{22}, $\rho_{\rm p}$ is also dependent on the in-plane spin (through $\overline{b}$). We therefore consider the region of $\theta_{JN} - q - \chi_{\rm p}$ parameter space which is consistent with $\rho_{\rm p}/\rho_{22}$. Figure 5.15 shows that as θ_{JN} decreases, the mass ratio must increase to ensure constant $\rho_{\rm p}$ (for fixed $\chi_{\rm p}$). We understand this because there is a positive relationship between q, θ_{JN} and $\rho_{\rm p}$, i.e. an increase in q and/or θ_{JN}, causes an increase in $\rho_{\rm p}$ (see Chapter 3). This can be seen from Eqs. 5.10, 5.11 and the equation for the opening angle in Chapter 1 (Eq. 1.1). For instance, when θ_{JN} increases we see that ζ increases. Likewise, when the mass ratio increases, the binary's opening angle increases and as such ζ also increases. Thus, to maintain constant ζ, and therefore $\rho_{\rm p}/\rho_{22}$, a decrease in θ_{JN} is compensated by an increase in q. We also see that to ensure fixed $\rho_{\rm p}/\rho_{22}$, θ_{JN} and q both decrease for increasing $\chi_{\rm p}$. This is because an increase in $\chi_{\rm p}$ translates to a larger binary opening angle and consequently a larger ζ. This means that the calculated values of ζ in each bin for the $\chi_{\rm p} = 0.25$ case will be greater than the corresponding values for the $\chi_{\rm p} = 0.1$ case. To ensure constant ζ and therefore $\rho_{\rm p}/\rho_{22}$, θ_{JN} and q both decrease.

By combining the constraints from the ρ_{33} and $\rho_{\rm p}$ measurements we are able to limit the parameter space consistent with the reported chirp mass, merger time, $\rho_{\rm p}/\rho_{22}$ and ρ_{33}/ρ_{22} of GW190814. We are able to constrain $\chi_{\rm p} < 0.1$ with $\chi_{\rm p} < 0.05$ more likely, $3 \lesssim q < 20$, $0.5 \lesssim \theta_{JN}$. We see that based on the results from Table 5.2, this minimal non-Bayesian analysis has captured $\chi_{\rm p}$ and θ_{JN} to a good estimate but it has failed to reduce the parameter space consistent with q. It is possible that we can reduce the allowed mass ratio by considering the orthogonal optimal SNR in multipoles other than $(\ell, m) = (3, 3)$, since each have a different dependence on the mass ratio and inclination angle [see e.g. Fig.2 in 182]. However, we leave this to future work.

5.4 Conclusion

Neutron star black hole binaries are an interesting subclass of compact binary objects, which if detected could provide valuable insights into the highly uncertain NS EOS models of stellar evolution and core collapse [194]. In this chapter we have described how the data from the first and second gravitational wave observing runs combined with the lack of a direct NSBH binary allow us to place limits on the rate of NSBH mergers. We then described how during their third observing run, the LIGO and Virgo detectors observed a novel source unlike any other known compact binary coalescence detected so far: GW190814. We verified that in this region of parameter space there is little to gain from performing this analysis under the assumption that the source is a BBH or NSBH. We therefore used the more advanced

BBH signal models to extract the source properties of GW190814. Thanks in part to the observation of significant power in the subdominant multipoles, and the conclusive measurement of little to no spin precession, we obtained precise measurements of its physical source properties.

We also demonstrated that, owing to the unique combination of masses, large SNR and low primary spin magnitude, we are able to bound GW190814's source properties to within a good estimate without using expensive Bayesian techniques. We suggest that given this result, future search pipelines should return an estimate for $\rho_{\rm p}$ and ρ_{33} as it would then be possible to provide a quick estimate for the source properties which could then be used as jump proposals to reduce computational cost in a full Bayesian analysis.

Chapter 6

Precession in the first half of the third gravitational wave observing run

In this chapter I explore whether precession has been observed in the first half of the third gravitational wave observing run. This work presents my contribution to Hoy and Mills *et al.* (*in preparation*). This work was produced by Charlie Hoy.

6.1 Introduction

Between 2015 and 2017, the Advanced LIGO [5] (aLIGO) and Advanced Virgo [6] (AdV) gravitational wave (GW) detectors performed their first and second GW observing runs (O1 and O2). During this time, the LIGO Scientific and Virgo collaboration (LVC) announced GWs originating from a binary neutron star [34] and 10 binary black holes [38]. Independent groups also reported on additional GW candidates [39, 40, 41, 42].

An important General Relativistic effect that was not clearly observed during O1 and O2 was spin-induced orbital precession (see Chapters 3 and 4). Spin-induced orbital precession arises when there is a misalignment between the orbital angular momentum and the spins of each compact object [54]. The importance of spin-induced orbital precession increases as the binary's mass ratio ($q = m_1/m_2$) increases [54, 164, 94, 169, 145]. Clear evidence for asymmetric masses was absent in the binaries detected during O1 and O2 [38], making the observation of precession challenging.

In 2019, aLIGO and AdV conducted the first half of their third GW observing run (O3a). Within 6 months of observing time, the LVC revealed a further 39 GW candidates in the second gravitational wave catalog (GWTC-2) [44]. In contrast to O1 and O2, several events in O3a had unequivocally unequal masses. First among these is GW190412 [45], with a mass ratio of ∼4:1. GW190412 was the first event where the amount of precession in the system was constrained away from the prior. Several months later GW190814 was detected with highly asymmetric component masses (∼9:1) and a secondary component with a mass larger than any previously discovered neutron star and lighter than any black hole [47]. This led to the most precise precession measurement of any event observed to date. It was demonstrated that the precession measurement improved parameter estimates, and in particular reduced the uncertainty on the mass of the smaller object. It has previously been reported that no single event in O3a unambiguously exhibits spin-induced orbital precession [45, 47, 44, 170].

In this chapter, we take advantage of $\rho_{\rm p}$, first introduced in Chapter 2, to build upon these statements. We calculate the signal-to-noise ratio (SNR) from precession for every event in O3a and compare it to the expected distribution in the absence of precession. Unlike Refs. [44, 170], we show that several events in O3a exhibit spin-induced orbital precession. Namely, GW190412, GW190929_012149 and GW190915_235702.

6.2 Setup

To calculate the $\rho_{\rm p}$ for each GW candidate, we use the publicly available data files made available through the Gravitational Wave Open Science Center [184]. We use

Event	ρ_p	$D_{\rm JS}$
GW190408_181802	$1.0^{+1.8}_{-0.9}$	0.03
GW190412	$3.0^{+1.6}_{-1.5}$	0.36
GW190413_134308	$0.7^{+1.5}_{-0.6}$	0.04
GW190413_052954	$0.6^{+1.4}_{-0.5}$	0.01
GW190421_213856	$0.7^{+1.4}_{-0.6}$	0.03
GW190424_180648	$0.6^{+1.4}_{-0.5}$	0.01
GW190425	$0.6^{+1.6}_{-0.5}$	-
GW190503_185404	$0.8^{+1.8}_{-0.7}$	0.03
GW190512_180714	$0.8^{+1.6}_{-0.7}$	0.01
GW190513_205428	$0.8^{+1.6}_{-0.6}$	0.01
GW190514_065416	$0.5^{+1.2}_{-0.4}$	0.03
GW190517_055101	$1.0^{+2.0}_{-0.8}$	0.02
GW190519_153544	$1.0^{+1.9}_{-0.7}$	0.07
GW190521	$0.7^{+1.4}_{-0.6}$	-
GW190521_074359	$1.6^{+2.5}_{-1.2}$	0.09
GW190527_092055	$0.7^{+1.7}_{-0.6}$	0.01
GW190602_175927	$0.5^{+1.0}_{-0.4}$	0.01
GW190620_030421	$0.8^{+1.7}_{-0.6}$	0.01
GW190630_185205	$1.0^{+1.8}_{-0.8}$	0.02
GW190701_203306	$0.5^{+1.0}_{-0.4}$	0.0
GW190706_222641	$0.5^{+1.1}_{-0.4}$	0.01
GW190707_093326	$0.8^{+1.4}_{-0.6}$	0.0
GW190708_232457	$0.7^{+1.5}_{-0.6}$	0.0
GW190719_215514	$0.6^{+1.5}_{-0.5}$	0.01
GW190720_000836	$0.7^{+1.2}_{-0.6}$	0.01
GW190727_060333	$0.7^{+1.6}_{-0.6}$	0.01
GW190728_064510	$0.8^{+1.3}_{-0.6}$	0.01
GW190731_140936	$0.5^{+1.3}_{-0.4}$	0.0
GW190803_022701	$0.6^{+1.4}_{-0.5}$	0.01
GW190814	$1.8^{+1.6}_{-1.2}$	0.03
GW190828_063405	$0.9^{+1.6}_{-0.8}$	0.01
GW190828_065509	$1.0^{+1.9}_{-0.8}$	0.03
GW190909_114149	$0.6^{+1.4}_{-0.5}$	0.05
GW190910_112807	$0.8^{+1.6}_{-0.7}$	0.02
GW190915_235702	$1.7^{+2.5}_{-1.4}$	0.17
GW190924_021846	$0.6^{+1.2}_{-0.5}$	0.0
GW190929_012149	$2.1^{+3.0}_{-1.9}$	0.23
GW190930_133541	$0.7^{+1.3}_{-0.5}$	0.01

Table 6.1: Table showing the the precession SNR $\rho_{\rm p}$ and the Jensen Shannon Divergence $D_{\rm JS}$ between a precessing and non-precessing distribution (see text) for all events in the second gravitational wave catalogue [44]. For events where $D_{\rm JS}$ could not be calculated, due to a lack of publicly available non-precessing samples, we add a hyphen. Where applicable we report the median values along with the 90% symmetric credible intervals.

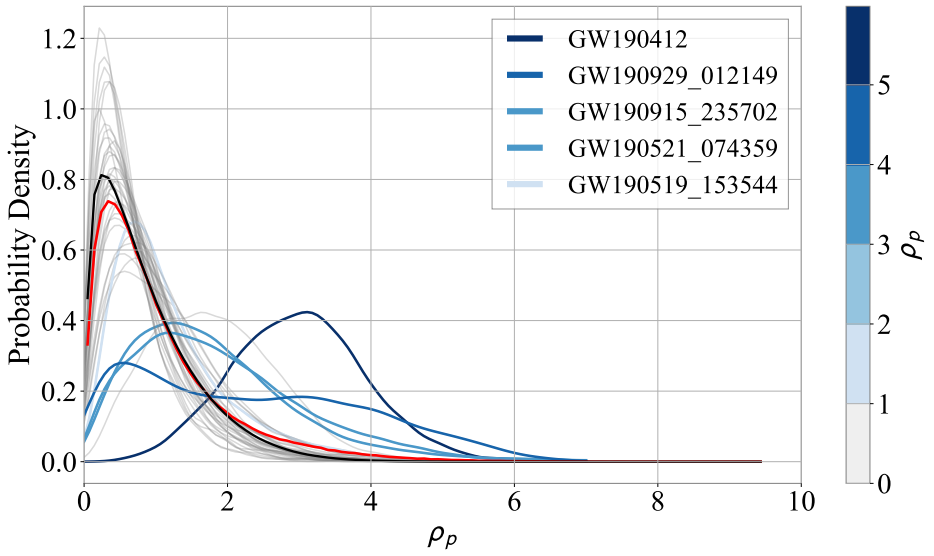

Figure 6.1: Plot showing the $\rho_{\rm p}$ distributions for all observations in the second gravitational wave catalogue (GWTC-2) colored by the median of their $\rho_{\rm p}$ distribution. The 5 events with the largest largest evidence for precession (based on their $D_{\rm JS}$ values, see text) are labelled. The average distribution is shown in red. The average distribution of $\rho_{\rm p}$ in a stretch of noisy data under the assumption that all sources are non-precessing is shown in black. The grey track which peaks close to $\rho_{\rm p} = 2$ corresponds to GW190814.

the data files re-weighed to a flat-in-comoving-volume distance prior [see Appendix C of 44] to match Ref. [44].

As part of the publicly available data files, numerous datasets containing posterior samples collected with different signal models are available. To calculate $\rho_{\rm p}$, we wish to use posterior samples obtained with signal models that include precession. Since the "PublicationSamples" dataset, as used in Ref. [44], always contains posterior samples obtained with precessing signal models (see Table VIII in Ref. [44]), we use these samples to calculate $\rho_{\rm p}$. These results may differ from previously published results [45, 47] since for the majority of candidates, the "PublicationSamples" dataset contains posterior samples obtained with precessing higher order multipole signal models. Our results are consistent with those reported in Ref. [44].

Table 6.1 presents a summary of the main results. All measurements are reported as symmetric 90% credible intervals around the median of the marginalized posterior distribution, unless otherwise specified. Figure 6.1 shows the inferred posteriors for $\rho_{\rm p}$ for all events in O3a.

6.3 Evidence for precession in O3a

Here, we calculate $\rho_{\rm p}$ for each event in O3a and compare it to the expected distribution in the absence of precession. This allows us to identify if any event in O3a

shows evidence for precession. In Chapter 2, we found that under the assumption of Gaussian noise, $\rho_{\rm p}$ is expected to follow a χ distribution with 2 degrees of freedom. This means that from noise alone, you would expect to measure $\rho_{\rm p} > 2.1$ 10% of the time. It was then shown in Chapter 3 that when parameter estimation is performed on a signal in zero noise, the inferred $\rho_{\rm p}$ is well described by the product of two terms, an *informed prior* – the $\rho_{\rm p}$ distribution calculated using the results from a non-precessing analysis and in-plane spins drawn from their respective priors – and a precessing likelihood, dependent on the *true* $\rho_{\rm p}$ of the signal. We demonstrated that this method works well for estimating the inferred $\rho_{\rm p}$ when parameter estimation is performed on a non-precessing signal (for this case the *true* $\rho_{\rm p} = 0$). In this chapter we take advantage of these calculations to estimate the expected distribution of $\rho_{\rm p}$ for each event in a stretch of noisy data under the assumption that the source is non-precessing (denoted by $\rho_{\rm p}^{\rm NP}$). We then compare the inferred $\rho_{\rm p}$ to $\rho_{\rm p}^{\rm NP}$ to identify if there is evidence for precession in the system.

Where available, we use the "AlignedSpinIMR" dataset for each gravitational wave candidate [184] to calculate the *informed prior* for each event. We then randomly draw a sample from the χ distribution with 2 degrees of freedom to represent the value of $\rho_{\rm p}$ consistent with a specific realisation of the noise. Since we assume the observed gravitational wave is in the presence of noisy data and its source is non-precessing, we use this value as the *true* $\rho_{\rm p}$ when calculating the precessing likelihood. We next combine the informed prior with the precessing likelihood to estimate the likely distribution of $\rho_{\rm p}$ for a specific realisation of the noise. We repeat this procedure 100 times to represent different realisations of the noise. We then take the median of these distributions as our estimate for $\rho_{\rm p}^{\rm NP}$. We show the average distribution of $\rho_{\rm p}^{\rm NP}$ across all events, $\bar{\rho}_{\rm p}^{\rm NP}$, in Figure 6.1.

Figure 6.1 shows that in general there is no strong evidence for precession in O3a as the average $\rho_{\rm p}$ distribution across all events $\bar{\rho}_{\rm p}$ is almost indistinguishable from $\bar{\rho}_{\rm p}^{\rm NP}$. Of all of the events in O3a, GW190412 [45], GW190929_012149 [44] and GW190814 [47] have the largest $\rho_{\rm p}$ with $\rho_{\rm p} = 2.99^{+1.58}_{-1.51}$, $\rho_{\rm p} = 2.13^{+3.04}_{-1.91}$ and $\rho_{\rm p} = 1.75^{+1.60}_{-1.23}$ respectively. To identify which events show evidence for precession we compute the Jensen-Shannon divergence $D_{\rm JS}$ [257] between $\rho_{\rm p}$ and $\rho_{\rm p}^{\rm NP}$ (as was done in Ref [44] between $\chi_{\rm p}$ and the conditioned $\chi_{\rm p}$ prior). This statistic is designed to quantify the difference between probability distributions. It is based on the Kullback–Leibler divergence ($D_{\rm KL}$) [258], which quantifies the information gain from distribution A to distribution B, but unlike $D_{\rm KL}$ it is symmetric and always has a finite value: $D_{\rm JS} \in [0, 1]$bits. When $D_{\rm JS} = 0\,(D_{\rm JS} = 1)$, the distributions A and B are identical (significantly different). Unlike other statistical tests (e.g. the Kolmogorov–Smirnov test [259, 260]), it is preferable to use $D_{\rm KL}$ or $D_{\rm JS}$ since they consider the entire distribution of A and B. In this work we use $D_{\rm JS}$ rather than $D_{\rm KL}$ since we found it to be more stable under random fluctuations and it also does not suffer from infinite values when distribution A has zero probability (for instance

when $\rho_{\rm p} = 0$).

We find that most events have $D_{\rm JS} < 0.05$bits which indicates the $\rho_{\rm p}$ and $\rho_{\rm p}^{\rm NP}$ agree well, see Table 6.1. This highlights that for most events, the observed signal is consistent with originating from a non-precessing binary. This is why $\bar{\rho}_{\rm p}$ is almost indistinguishable from $\bar{\rho}_{\rm p}^{\rm NP}$. However, there are a few notable exceptions: GW190412, GW190929_012149, GW190915_235702 have $D_{\rm JS} > 0.1$ bits and GW190521_074359, GW190519_153544 have $D_{\rm JS} > 0.05$ bits. We discuss why our analysis demonstrates GW190412, GW190929_012149 and GW190915_235702 show the largest evidence for precession in detail below. Since Ref. [44] identified GW190521 and GW190814 as having two of the most informative $\chi_{\rm p}$ distributions compared to the prior, we also describe why our analysis does not highlight these events.

GW190412 likely originated from a precessing system since the inferred $\rho_{\rm p}$ is significantly larger than $\rho_{\rm p}^{\rm NP}$ with $D_{\rm JS} = 0.36$ bits, see Figure 6.2. We understand this because a) GW190412 is consistent with an exceptionally large $\rho_{\rm p}$ (compared to the other events in O3a) as a result of the second largest mass ratio $q = 3.6^{+1.1}_{-1.1}$ and inclination angle constrained away from face-on $\theta_{JN} = 0.7^{+0.3}_{-0.2}$ rad (folded between $[0, \pi/2]$), see Chapter 3 for details, and b) GW190412's *informed prior* is broad $\rho_{\rm p} = 3.3^{+5.6}_{-2.7}$, meaning that $\rho_{\rm p}^{\rm NP}$ approximately peaks at the maximum of the precessing likelihood. Since the precessing likelihood peaks at $\rho_{\rm p} = 1.2$, with $\sim 1\%$ probability of random drawing a value consistent with GW190412's large $\rho_{\rm p} = 2.99$, $\rho_{p}^{\rm NP}$ peaks at much smaller values than the inferred $\rho_{\rm p}$. This results in a large Jensen-Shannon divergence. We find that GW190412 is inconsistent with a non-precessing system at $> 97\%$ probability. This result differs from the conclusions presented in Refs [45, 261] which are based upon the more commonly used Bayes factors between the precessing and non-precessing hypothesis. Figure 6.2 also shows that $\rho_{\rm p}^{\rm NP}$ is shifted to larger values than average. This follows from the fact that GW190412's aligned spin analysis also confidently identifies GW190412 as an unequal mass ratio system ($q = 3.2^{+1.1}_{-1.1}$). This means that for a given in-plane spin sample drawn from the prior, the calculated value of $\rho_{\rm p}$ will be larger for GW190412 than for an equal mass ratio binary.

GW190929_012149 shows the second largest evidence for precession with $D_{\rm JS} = 0.23$ bits. From Figure 6.2, we see that this is because GW190929_012149's $\rho_{\rm p}$ is distinctly different from $\rho_{\rm p}^{\rm NP}$. In fact, GW190929_012149's $\rho_{\rm p}$ distribution is bimodal. The reason for this is because GW190929_012149's mass ratio is bimodal: one peak at $q \sim 1$, which maps to the peak at low $\rho_{\rm p}$, and another at $q \sim 3$. If we restrict $q < 1.5$ (i.e. only consider samples from the equal mass peak), we find that $\rho_{\rm p}$ is indistinguishable from $\rho_{\rm p}^{\rm NP}$ with $D_{\rm JS} = 0.02$. If we restrict $q > 1.5$, we recover only the second peak of $\rho_{\rm p}$ and find $D_{\rm JS} = 0.28$. This means that GW190929_012149 has the second largest evidence for precession as a result of the high mass ratio peak. Unlike GW190412, GW190929_012149's $\rho_{\rm p}^{\rm NP}$ distribution is consistent with $\bar{\rho}_{\rm p}^{\rm NP}$ despite having support for large mass ratios. This is because the aligned spin analysis

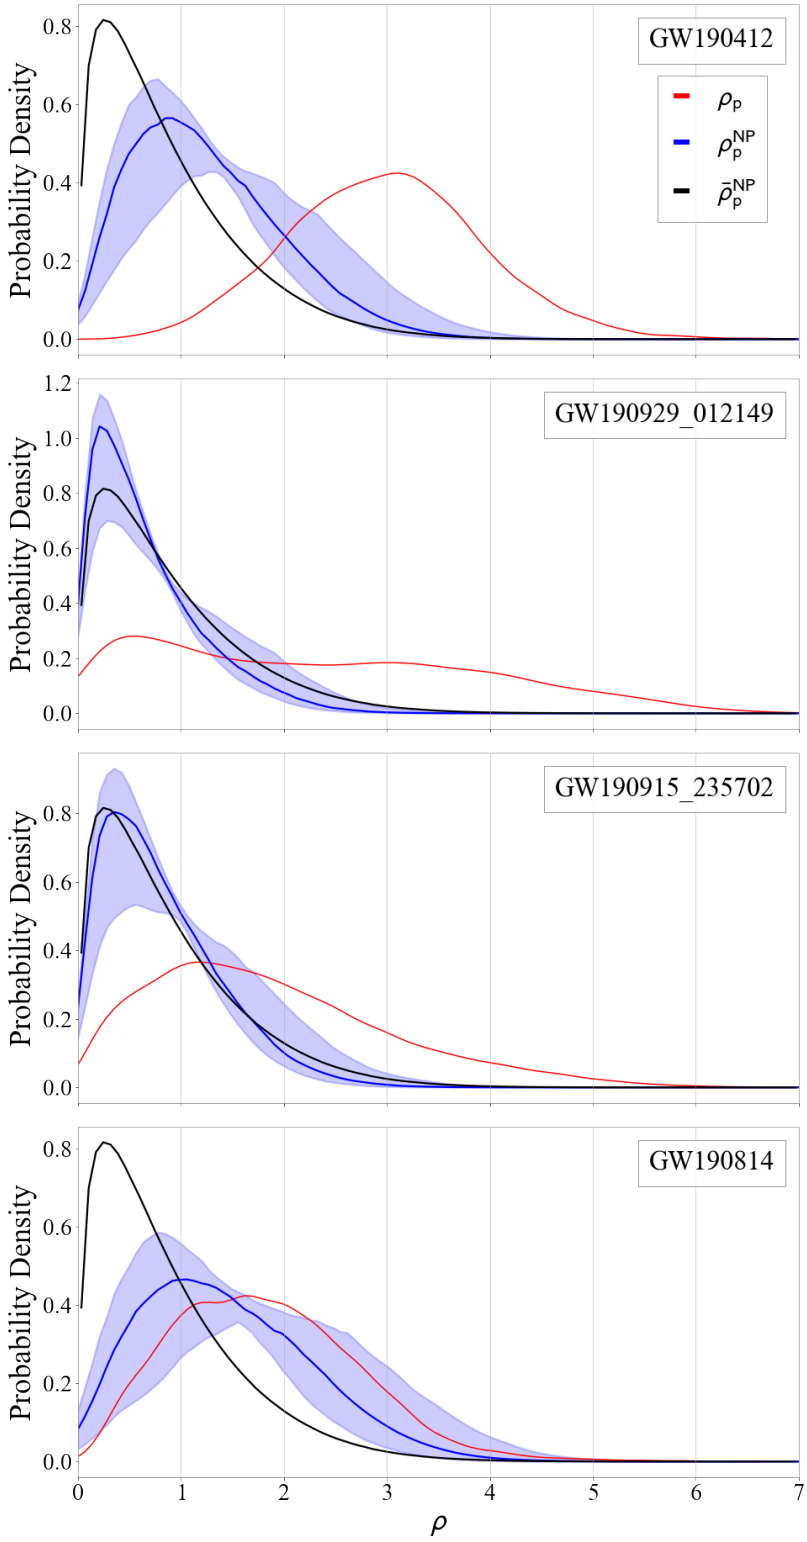

Figure 6.2: ρ_p distributions for *First row*: GW190412, *Second row*: GW190929_012149, *Third row*: GW190915_235702, *Bottom row*: GW190814. The blue line shows the expected distribution of ρ_p in a stretch of noisy data under the assumption that the source is non-precessing, ρ_p^{NP}. The blue shaded region shows the 1σ uncertainty of ρ_p^{NP}. The black line shows the average ρ_p^{NP} across all events.

identifies GW190929_012149 as a near equal mass ratio system, with no support for a second peak ($q = 2.0^{+2.8}_{-0.9}$).

GW190915_235702 has $D_{\rm JS} = 0.16$ bits. GW190915_235702 is consistent with one of the largest inclination angles ($\theta_{JN} = 0.9^{+0.5}_{-0.6}$ folded between $[0, \pi/2]$) of any event in O3a. This means that although GW190915_235702 is consistent with an equal mass binary ($q = 1.4^{+0.9}_{-0.4}$), $\rho_{\rm p}$ is the fourth largest of any event in O3a. The inclination dependence on $\rho_{\rm p}$ is discussed at length in Chapters 2, 3 and 4. As can be seen in Figure 6.2, the inferred $\rho_{\rm p}$ is different from $\rho_{\rm p}^{\rm NP}$ but not to the same extent as GW190412 and GW190929_012149. This explains why GW190412 and GW190929_012149 have larger Jensen-Shannon divergences. Interesting, $\rho_{\rm p}^{\rm NP}$ is consistent with $\bar{\rho}_{\rm p}^{\rm NP}$. This is because the aligned spin analysis is unable to break the distance – inclination degeneracy [99], and as such, simply recovers the prior for θ_{JN}. This reduces the inferred value of θ_{JN} considerably and as a result reduces $\rho_{\rm p}^{\rm NP}$. We understand that it is consistent with $\bar{\rho}_{\rm p}^{\rm NP}$ because the majority of aligned spin analyses in O3a recover the prior for the inclination angle and equal mass binaries.

We see that GW190814 is consistent with originating from a non-precessing system with $D_{JS} = 0.03$ bits. This is expected given the near-zero $\chi_{\rm p}$ measurement ($\chi_{\rm p} = 0.04^{+0.04}_{-0.03}$). Despite this, GW190814 has the third largest $\rho_{\rm p}$ in O3a. This apparent contradiction is a result of GW190814's large mass ratio ($q = 9.0^{+0.8}_{-0.6}$). As shown in Figure 6.2, we see that because of GW190814's extraordinary mass ratio, $\rho_{\rm p}^{\rm NP}$ peaks at significantly larger values than average, with $\rho_{\rm p}$ also entirely contained within the 1σ uncertainty. This means that although $\rho_{\rm p}$ is large for this system, it is still consistent with originating from a non-precessing binary. This explains why $D_{\rm JS} = 0.03$bits.

GW190521 had the largest inferred $\chi_{\rm p}$ in GWTC-2: $0.68^{+0.26}_{-0.44}$. Surprisingly, GW190521 has a small $\rho_{\rm p}$: $0.7^{+1.4}_{-0.6}$. We understand this because GW190521 is the largest mass event detected with LIGO/Virgo. This means that GW190521 is very short in duration: 4 cycles (2 orbits) within the detectors' sensitive frequency band. Consequently, GW190521 is decomposed into two near parallel "precession harmonics" (with overlap $|{\rm O}^{\rm prec}_{1,0}| = 0.97^{+0.01}_{-0.03}$). This means that any power orthogonal to the dominant harmonic is small and $\rho_{\rm p}$ is small by definition (we refer the reader to Chapter 2 for a detailed discussion). Several explanations for the large $\chi_{\rm p}$ have been suggested, including possible evidence for eccentricity [262] and head-on collisions [263]. Unfortunately we are unable to compare $\rho_{\rm p}$ to $\rho_{\rm p}^{\rm NP}$ since there are no publicly available non-precessing samples for GW190521.

6.4 Conclusion

In this chapter we have determined the significance of precession in the latest observations made by the LIGO Scientific, Virgo and KAGRA collaborations. We have demonstrated that in general there is no strong evidence for precession in O3a but

GW190412, GW190929_012149 and GW190915_235702 are most likely to have originated from precessing systems. We note that these results only compare the one dimensional distributions and we present a more detailed analysis in Chapter 7.

Chapter 7

Constraining black-hole spins with gravitational wave observations 2

This chapter reproduces the text of Hoy *et al.* (*in preparation*). The code used in this analysis was built upon work done by Tiwari *et al.* [119]. All extensions were written by Charlie Hoy. I wrote all Sections and produced all Figures.

7.1 Introduction

During the first and second gravitational wave (GW) observing runs (O1 and O2) [38] of the Advanced LIGO [5] and Advanced Virgo [6] GW observatories, the LIGO scientific and Virgo collaborations announced eleven GW candidates; ten from merging black-hole binaries (BBHs) [7, 33, 35, 36, 37, 38] and one from a binary neutron star coalescence [34]. Independent groups also reported additional GW candidates [39, 40, 41, 42]. By combining parameter estimates for these BBH observations, first attempts at deciphering the astrophysical distribution of black hole spins were conducted [118, 264, 119, 265, 120, 266]. However, the limited sample size meant that only weak constraints could be placed on the distribution of black hole spins.

During the first half of the third gravitational wave observing run (O3a), a further 39 GW candidates were announced [44]. Similar to those from O1 and O2 [85], most of these detection's remained largely uninformative about the presence of the General Relativistic phenomenon of spin-induced orbital precession[44] — the misalignment of the binary's orbital angular momentum and the spins of each compact object resulting in characteristic modulations to the GWs amplitude and phase [54]. A direct measurement of spin-induced orbital precession provides a unique insight into the astrophysical distribution of black hole spins [e.g. 71, 72].

Precession is often parameterised by a single effective parameter $\chi_{\rm p}$, ranging between 0 (no precession of the orbital plane) and 1 (maximal precesion) [64]. $\chi_{\rm p}$ is widely used for inferring the occurrence of precession in GW data [see e.g. 38, 44], although alternative metrics have also been proposed [e.g. 65, 66]. In Chapter 2, we introduced a different metric to quantify precession: the precession signal-to-noise ratio (SNR) $\rho_{\rm p}$, described as the contribution to the total SNR of the system that can be attributed to precession. $\rho_{\rm p}$ is calculated by decomposing a GW into two non-precessing harmonics and isolating the SNR contained in the harmonic orthogonal to the dominant one. By deconstructing a precessing gravitational wave in this form, the characteristic amplitude and phase modulations can be interpreted as the beating of these harmonics. If $\rho_{\rm p}$ is small ($\rho_{\rm p} \lesssim 2.1$), the amplitude of the second harmonic is insignificant and any beating of the harmonics is negligible. For this case, we would observe a GW which looks like the dominant non-precessing harmonic. Although $\rho_{\rm p}$ is dependent on the GW detector network and its sensitivity, if the harmonics are close to orthogonal, $\rho_{\rm p}$ can be scaled by the total SNR ρ to provide a detector invariant quantity $\rho_{\rm p}/\rho$ (see Eq. 2.34 noting that ρ is denoted as $\rho_{\rm 2harm}$). This is a useful quantity for population studies since it means results are independent of the detector network and the chosen detector sensitivity. This implies that the $\rho_{\rm p} \gtrsim 2.1$ criterion used in previous works for quantifying the measurability of precession [see e.g. 44, 45, 47, 86, 120, 145] becomes $\rho_{\rm p}/\rho \gtrsim 2.1/\rho$, which is bounded between $0 \leq \rho_{\rm p}/\rho \leq 1/\sqrt{2}$, where the upper bound implies equal power in both harmonics. Consequently, for systems with large $\rho_{\rm p}/\rho$, precession contributes significantly to

the total SNR of the system.

Ref. [170] recently showed, through a hierarchical Bayesian analysis [267, 268, 269], that the population of known BBHs provides clear evidence for misaligned spins despite no single event unambiguously exhibiting evidence for spin-induced orbital precession. By assuming a population model where the spin magnitude of each black hole is described by a beta function [140] and the orientation by a model allowing for both isotropic and aligned spins [138], it was shown that the underlying spin distribution peaks at spin magnitudes ~ 0.2 with preference for primarily aligned spins (although there is non-vanishing support for angles $> 90^\circ$ indicating the presence of misaligned component spins). This has significant implications, as measuring the spin distribution can distinguish between the two favoured binary formation mechanisms. These are field binaries which are formed from isolated stellar progenitors and expected to have spins distributed about the orbital angular momentum with some unknown misalignment angle [e.g. 69, 70, 71], and dynamic binaries which are expected to have randomly orientated spins and formed when two black holes become gravitationally bound in dense stellar environments [e.g. 72].

In this chapter we use publicly available data from the second gravitational wave catalogue (GWTC-2) [44] to draw inferences about the underlying spin distribution of black holes. We show that if we use the same set of BBHs, we obtain the same conclusions as Ref. [170] irrespective of whether we use $\chi_{\rm p}$ or $\rho_{\rm p}/\rho$ to parameterize precession, i.e. the current population of binary black holes prefers a spin distribution model with mild preference for aligned spins and spin magnitudes peaking at ~ 0.2. However, we show that the distribution in Ref. [170] is only marginally preferred to other distributions used in this chapter with odds ratios $\gtrsim 3:1$. By inspecting each BBH separately, we also show that preference for the distribution reconstructed in Ref. [170] is a result of a small subset of observations (13/44), with the majority of events preferring a distribution which has isotropic spins and significantly lower spin magnitudes. Finally, we hint at potential structure in the spin magnitude parameter space, with lower mass events preferring lower spin magnitudes and higher mass events preferring higher spin magnitudes.

7.2 Method

We use Bayesian model selection to calculate the odds ratio between 9 different spin distributions using the publicly released posterior samples from GWTC-2, made available through the Gravitational Wave Open Science Center [GWOSC; 184]. As in Ref. [170], we only consider BBHs with false alarm rates (FARs) $< 1{\rm yr}^{-1}$. This means we exclude 2 marginal events included in Ref. [44]: GW190719_215514 and GW190909_114149 (we consider GW190426_152155 to a be marginal neutron star black hole candidate). Unlike Ref. [170], we consider how our results change if GW190814 is included in our analysis. This is because GW190814 is more likely

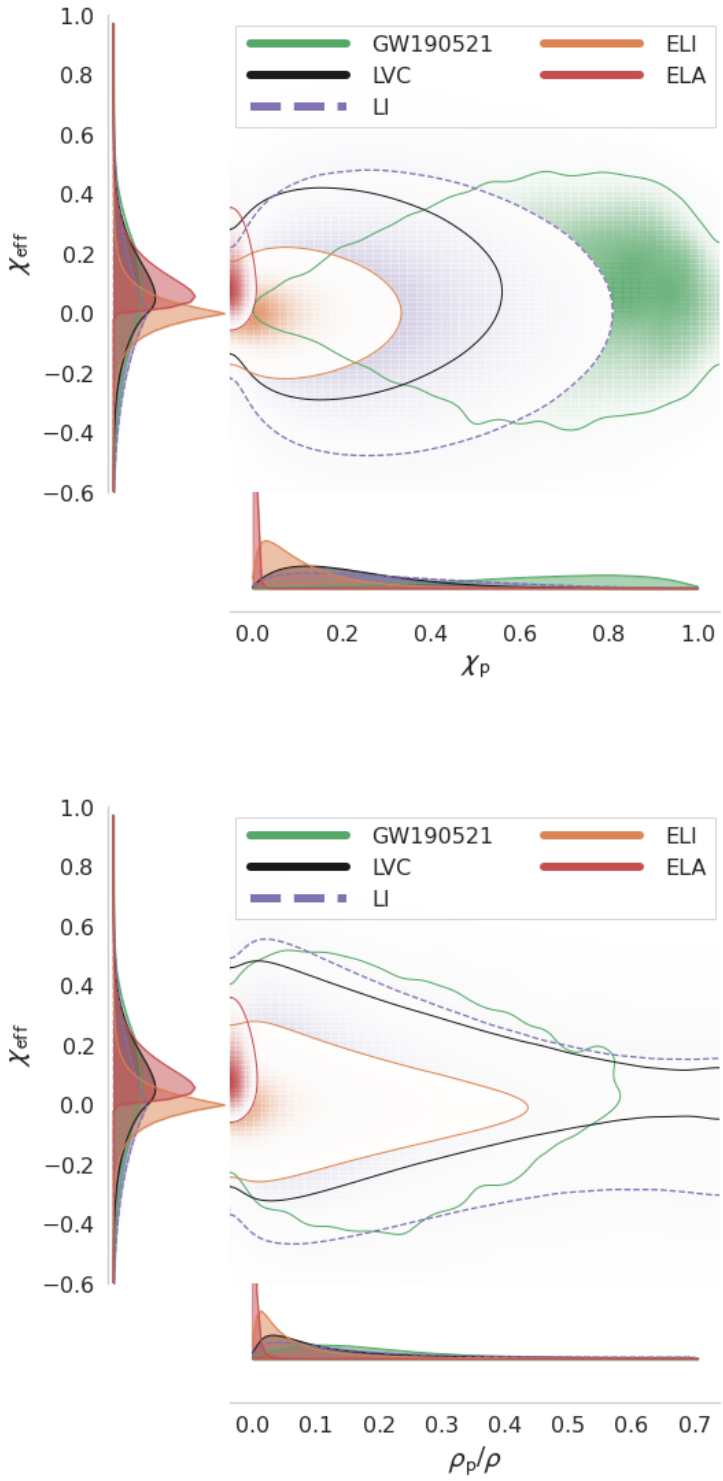

Figure 7.1: 2-dimensional probability density functions (PDFs) showing support across the *Top*: $\chi_{\rm eff}$–$\chi_{\rm p}$and *Bottom*: $\chi_{\rm eff}$–$\rho_{\rm p}/\rho$ parameter space for a selection of the different spin distributions used in our analysis (see text for definitions). Each contour shows the 90% credible interval. We show the preferred model from Ref. [170] in black. The 2-dimensional PDF for GW190521 [46] is shown in green. Gaussian kernel density estimates are used to estimate the probability density.

(71%) the result of a BBH merger [47]. For each GW candidate considered, we randomly draw 10^4 samples from the 'PublicationSamples' dataset to ensure a consistent number of samples across the population. Unless otherwise stated, the odds ratio is compared to the gentle aligned distribution displayed in Figure 10 of Ref. [170], denoted LVC. Given that population studies of BBHs from O1, O2 and O3a disfavour highly spinning black holes [118, 119, 120, 170], we only consider spin distributions with either low (L) [consistent with 118], very low (VL) [consistent with 85] or extremely low (EL) spin magnitudes,

$$\begin{aligned} p_{\rm EL}(a) &\propto e^{-8a} \\ p_{\rm VL}(a) &\propto e^{-5a} \\ p_{\rm L}(a) &\propto (1-a) \end{aligned} \tag{7.1}$$

where $a = |cS/(Gm^2)|$ is the spin magnitude of a black hole with mass m and spin S. We consider 3 distributions for the tilt angles: aligned (A), nearly aligned (NA) and isotropic (I). The A and NA distributions are both triangular in $\cos\theta$, with a peak at 1, taking values between $0.999 \leq \cos\theta \leq 1$ and $0.85 \leq \cos\theta \leq 1$ respectively. The A and NA distributions resemble field binaries with misalignment angles 3° and 30° respectively. The I distribution is uniform in $\cos\theta$ between -1 and 1 and resembles dynamic binaries. A *universe* modelling each spin distribution is generated by randomly drawing 10^7 binaries.

As it is difficult to constrain the individual black hole spins at typical SNRs [270], we use a mass weighted effective spin $\chi_{\rm eff}$ to describe the average projection of spins parallel to the orbital angular momentum [248]. Precession, arising from the projection of spins perpendicular to the orbital angular momentum, is described by $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$. Figure 7.1 shows how a subset of these spin distributions vary across the $\chi_{\rm p}$–$\chi_{\rm eff}$ and $\rho_{\rm p}/\rho$–$\chi_{\rm eff}$ parameter space. The A distributions can easily be distinguished from the I distributions as $\chi_{\rm eff} > 0$ and $\rho_{\rm p}/\rho$ is small by definition.

Following the methodology described in Ref. [119], we calculate the odds ratio between two spin distribution models λ_1 and λ_2 as,

$$O_{\lambda_1,\lambda_2} = \frac{p(\lambda_1|\{\mathbf{d}\})}{p(\lambda_2|\{\mathbf{d}\})} \approx \left[\frac{V_{\rm pop}(\lambda_1)}{V_{\rm pop}(\lambda_2)}\right]^{-N} \prod_{i=1}^{N} \left[\frac{\sum_j p(\theta_i^j|\lambda_1)/\pi(\theta_i^j)}{\sum_j p(\theta_i^j|\lambda_2)/\pi(\theta_i^j)}\right] \times \left[\frac{p(\lambda_1)}{p(\lambda_2)}\right], \tag{7.2}$$

where $p(\lambda|\{\boldsymbol{d}\})$ is the posterior distribution for the model λ given a set of BBH observations $\{\boldsymbol{d}\}$, $V_{\rm pop}(\lambda)$ is the sensitive volume for the model λ, θ_i^j is the jth posterior sample for observation i, $\sum_j p(\theta_i^j|\lambda)/\pi(\theta_i^j)$ is the sum over posterior samples re-weighted from the default *prior universe* used in the LVC analyses $\pi(\theta_i^j)$ (LAL prior), to the *universe* assuming a given model λ $p(\theta_i^j|\lambda)$, $p(\lambda)$ is the prior on the model and we restrict $\boldsymbol{\theta} = (\mathcal{M}, q, \chi_{\rm eff}, [\chi_{\rm p}, \rho_{\rm p}/\rho], \iota)$. As with Ref. [119], we assume all models are equally likely i.e. $p(\lambda_1)/p(\lambda_2) = 1$.

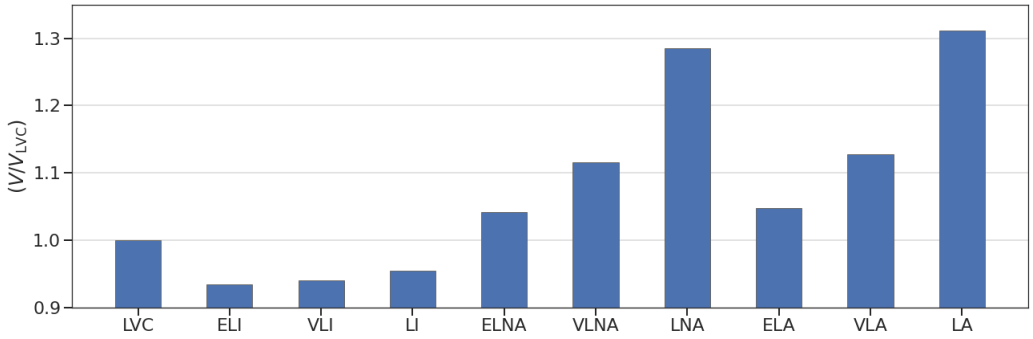

Figure 7.2: Sensitive volumes for 10 different spin distributions with respect to the preferred model in [170] (LVC), see text for model definitions. This ratio to the power of the number of observations accounts for the selection effects.

For LVC parameter estimation analyses, the *prior universe* is taken to be flat in m_1 and m_2 (with the condition that $m_1 \geq m_2$) and spin vectors are assumed to be uniform in spin magnitude and isotropic on the sphere [see Appendix B.1 of 38]. When generating a *universe* for a given model λ, we use the same mass distribution but vary the spin magnitude and orientation vectors. All other binary parameters are randomly drawn from the same distributions as used in Ref. [38].

The sensitive volume $V_{\rm pop}(\lambda)$ is essential for accounting for selection effects. It is estimated numerically by injecting GW signals drawn from model λ into GW strain data and searching for them assuming a given detection threshold [201]. Currently search pipelines employ non-precessing waveform approximants for matched filtering [127, 126]. This means that current techniques to estimate the sensitive volume omit precession (although see Ref. [271] which suggests an alternative method that includes precession). Since precessing signals will be recovered at lower probabilities than an equivalent precessing search pipeline, we can expect that the sensitive volume will be underestimated for systems where precession effects are observable [272]. However, in Chapter 2, we argued that for signals with low $\rho_{\rm p}$ this effect is minimal. Given that for most models used in this chapter $\rho_{\rm p}/\rho$ is small, we approximate $V_{\rm pop}(\lambda)$ by $V_{\rm pop}(\lambda^{\rm np})$: the sensitive volume for the non-precessing equivalent λ. This is the same assumption used in Ref. [170].

In Figure 7.2 we show $V_{\rm pop}(\lambda^{\rm np})$ for each spin distribution model compared to LVC. We see that as the spin magnitude increases and spin orientation becomes more aligned, the sensitive volume increases. We understand this because binaries with a larger aligned spin (larger $\chi_{\rm eff}$) can be observed at a greater distance. We therefore expect ELI to have the lowest (LA to have the largest) sensitive volume as it leads to a population with the smallest (largest) aligned spin. Given that the odds ratio (see Equation 7.2) involves dividing by the models sensitive volume, assuming all other parameters are equal, the model with the lower sensitive volume is preferred. Although for a single event this effect if minimal, for 44 observations the odds ratios increases by of ~ 10 in favour of ELI over LVC.

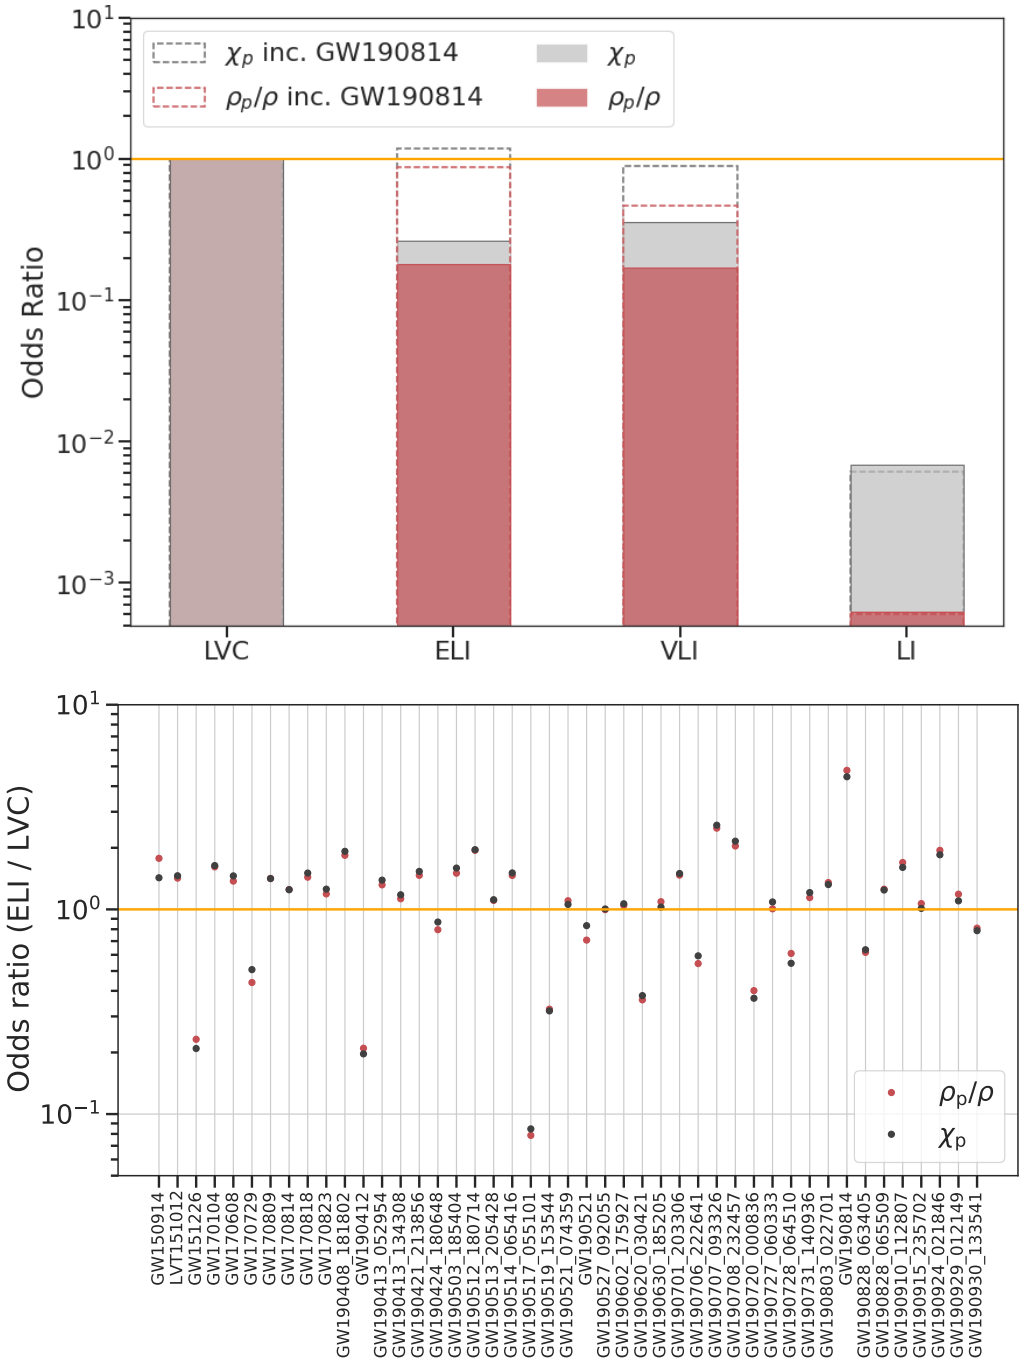

Figure 7.3: *Top*: Odds ratios for different spin distributions in reference to the preferred model in Ref. [170] (LVC), see text for model definitions. The grey and red bars indicate the inferred odds ratios when $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ is used to parameterise precession respectively. Dashed lines show the inferred odds ratio when GW190814 is included in the analysis. Only the largest four odds ratio are shown. Models that are not shown have odds ratios $< 10^{-5} : 1$. *Bottom*: Odds ratios for ELI against LVC for each binary black hole candidate considered in this analysis [44]. An orange line shows an odds ratio of 1 meaning that neither model is preferred. An odds ratio greater than 1 shows preference for ELI over LVC. In both cases odds ratios are calculated using two different paramerisations of precession: $\chi_{\rm p}$ (grey) as used in Ref. [44] and $\rho_{\rm p}/\rho$ (red). Odds ratios are calculated using the posterior samples released as part of GWTC-2 [44, 184].

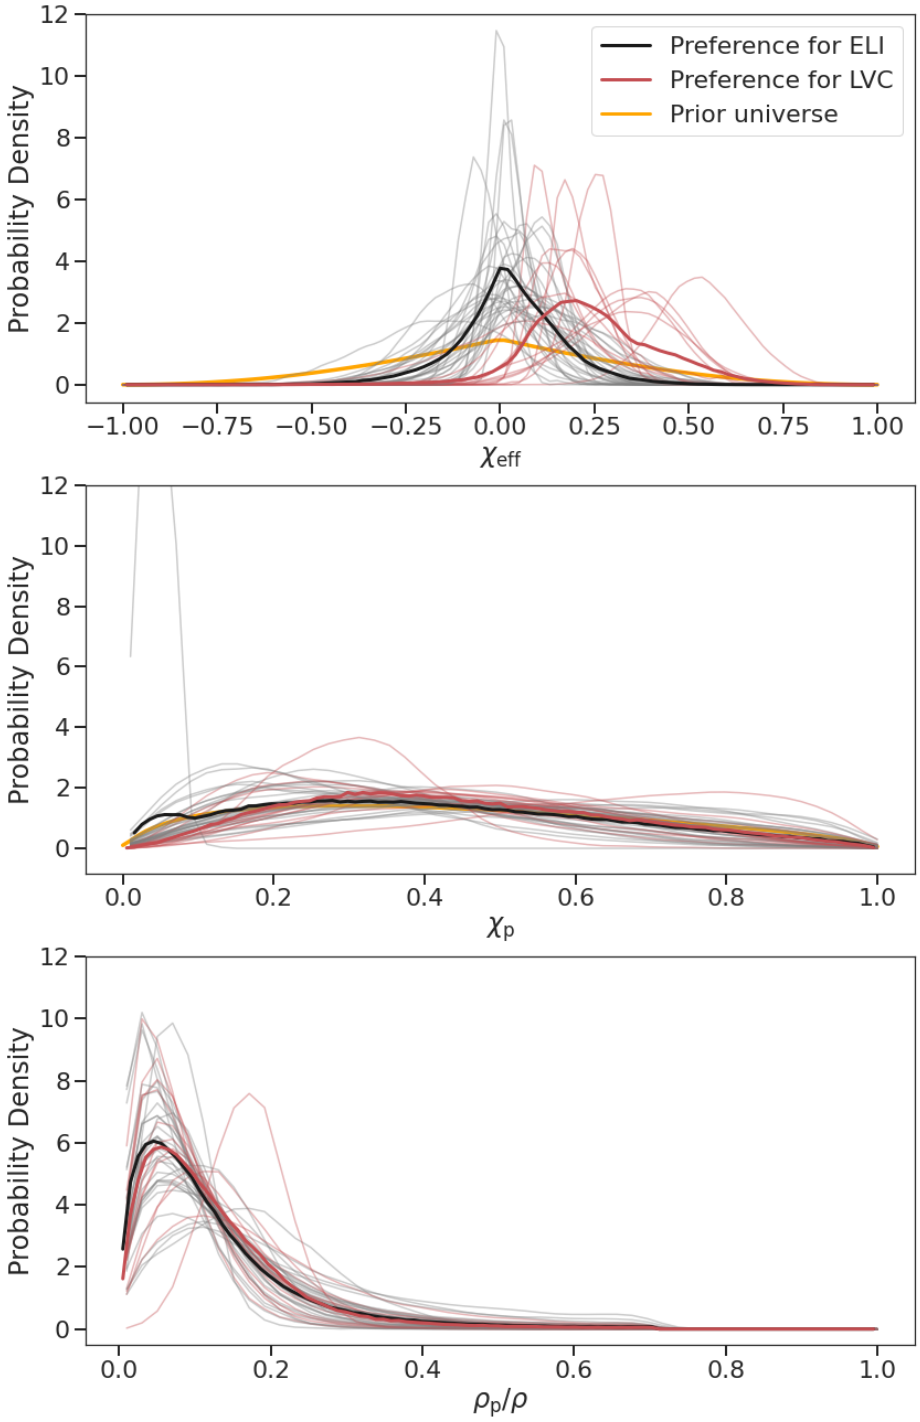

Figure 7.4: Posterior distributions for *Top*: $\chi_{\rm eff}$, *Middle*: $\chi_{\rm p}$, *Bottom*: $\rho_{\rm p}/\rho$ for binary black hole candidates in the second gravitational wave catalogue (GWTC-2) [44, 184]. Light grey and red traces show the posterior distributions for events which prefer ELI over LVC and LVC over ELI respectively (see bottom panel of Figure 7.3). Solid black and red curves shows the average of the light grey and red traces respectively. The orange curves show the default $\chi_{\rm eff}$ and $\chi_{\rm p}$ priors used in the LVC analyses. For the middle plot, we restrict the y axis to be less than 12 such that all posterior distributions can be distinguished. This means that for one of the posterior distributions, a portion of the distribution cannot be seen.

7.3 Results and Discussion

The top panel of Figure 7.3 plots odds ratios with respect to the LVC spin distribution for two different metrics for precession: $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$. Models that are not shown have odds ratios $< 10^{-5}:1$. For both metrics we see that when GW190814 is excluded from our analysis, the same conclusions as Ref. [170] are obtained: the LVC spin distribution is the preferred model for the selected BBH candidates. However, for the first time, we can quantify by how much this model is preferred. Our analysis infers that the LVC spin distribution is only marginally preferred with both VLI and ELI disfavoured by odds ratios $2.8:1$ and $3.8:1$ for the $\chi_{\rm p}$ analysis and $5.9:1$ and $5.5:1$ for the $\rho_{\rm p}/\rho$ analysis respectively. This equates to about 0.9σ and 1.1σ for the $\chi_{\rm p}$ analysis and 1.4σ and 1.3σ for the $\rho_{\rm p}/\rho$ analysis respectively. When GW190814 is included, ELI and LVC are equally likely with odds ratios for both metrics ~ 1 ($1.2:1$ and $0.9:1$ for the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses respectively).

In the bottom panel of Figure 7.3 we show how the odds ratio of ELI vs LVC changes as a function of GW candidate. This provides key insight into why the LVC spin distribution is marginally preferred and why this preference changes when GW190814 is included. In general, we see good agreement between the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses for each event. The largest disagreement is for GW190521 where the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses prefer LVC over ELI by $1.2:1$ and $1.4:1$ respectively. We explain why this event has the largest disagreement in Section 7.3.a. Given that the largest disagreement is $\sim 0.2\sigma$, we report the average odds ratio for a single event unless otherwise stated. Figure 7.3 shows that most BBHs detected through GWs (31/44) prefer ELI over LVC with GW190814 displaying the strongest preference: $5:1$. Given that the inferred $\chi_{\rm eff}$ and $\chi_{\rm p}$ distributions from GW190814 are strongly peaked near-zero, it is no surprise that ELI is preferred. In this particular region of parameter space ELI has $\sim 7\times$ more support than LVC. Since comparing support in a given region of the parameter space is effectively computing a simplified version of Eq. 7.2, we expect this calculation to be indicative of the odds ratio. It is therefore a good sanity check for our results.

As expected, all candidates that support $\chi_{\rm eff} > 0$ at more than 90% probability (GW151226, GW170729, GW190412, GW190517_055101, GW190519_153544, GW190620_030421, GW170706_222641, GW190720_000836, GW190728_064510, GW190828_063405, GW190930_133541) prefer LVC over ELI. This is because the majority of the LVC distribution supports $\chi_{\rm eff} > 0$, 70% compared to 50% for ELI, and has a longer tail up to larger $\chi_{\rm eff}$, ~ 0.32 compared to ~ 0.1 for ELI. GW190517_055101 [44] and GW190412 [45] show the largest preference for LVC with odds ratios $11:1$, $5:1$ respectively. GW190517_055101 has the largest $\chi_{\rm eff}$ observed so far, with $\chi_{\rm eff} = 0.52^{+0.19}_{-0.19}$. In this particular region of the $\chi_{\rm eff}$ parameter space, LVC has $\sim 15\times$ more support than ELI, suggesting that the odds ratio is heavily influenced by GW190517_055101's $\chi_{\rm eff}$ measurement. We find that if this

single event is removed from our analysis, ELI is preferred over LVC with odds ratios $3:1$ and $2:1$ for the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses respectively. Meanwhile, GW190412 supports $\chi_{\rm eff} > 0.15$ at 90% confidence and is consistent with a mildly precessing system, $\chi_{\rm p} = 0.31^{+0.19}_{-0.16}$. For this region of parameter space, LVC has $\sim 5\times$ more support than ELI.

In Figure 7.4 we plot the posterior distributions for $\chi_{\rm eff}$, $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ for all events used in this analysis. We see that of those events which prefer ELI over LVC, the $\chi_{\rm eff}$ distribution is on average strongly peaked at zero with width comparable to ELI (see Figure 7.1). We also see that on average there is no information from precession, with $\chi_{\rm p}$ resembling the LAL prior and near zero $\rho_{\rm p}/\rho$. The strong peak at $\chi_{\rm eff}= 0$ and lack of information from precession, means that isotropic spin distributions are preferred. As the width of the $\chi_{\rm eff}$ distribution is narrow, extremely low spin magnitudes are preferred. Of those events which prefer LVC over ELI, the average $\chi_{\rm eff}$ distribution is positive, peaking at ~ 0.2, with little support for $\chi_{\rm eff} \leq 0$. We also see that the average $\chi_{\rm p}$ distribution peaks at slightly larger values than the LAL prior while the average $\rho_{\rm p}/\rho$ resembles the average distribution for those events which prefer ELI over LVC. Because of this, any model which allows for large in-plane spins will be favoured more by the $\chi_{\rm p}$ analysis. For this case, the combination of the positive $\chi_{\rm eff}$ measurement disfavouring isotropic spin distributions, and the non-zero $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ disfavours aligned spin models. Given that LVC has support for large positive $\chi_{\rm eff}$ while still allowing for mild precession, this model is preferred for this subset of binaries. In general, we note that the $\chi_{\rm eff}$ measurement is the prevailing quantity for why the population prefers LVC over ELI.

Next we comment briefly on the difference between the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses on a population level. From Figure 7.3 we see a 0.4σ, 0.7σ and 1.7σ difference between the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses for the ELI, VLI and LI spin distributions respectively. We see that the difference becomes larger for larger spin magnitudes. We understand this trend because when calculating $\rho_{\rm p}/\rho$ for all BBHs considered, we see that the distribution has a narrow peak near zero (implying a lack of measurable precession in the majority of signals) while the $\chi_{\rm p}$ distribution is broad and resembles the prior, see Figure 7.4. This implies that any model which has $\rho_{\rm p}/\rho$ peaking in a different location to the majority of signals, will be heavily disfavoured. However, because the $\chi_{\rm p}$ distribution is broad, any model will still receive reasonable support compared to the $\rho_{\rm p}/\rho$ case. Therefore, we argue that this growing difference is expected, given that larger spin magnitudes predict larger degrees of precession and consequently larger $\rho_{\rm p}/\rho$. We stress here that the purpose of this chapter is not to determine which metric is optimal, only to provide evidence that neither model should be rejected or favoured.

7.3.a GW190521

GW190521 is an exceptional event, it is consistent with a merger of two black holes with masses $85^{+21}_{-14}M_\odot$ and $66^{+17}_{-18}M_\odot$, and $\chi_{\rm p} = 0.68^{+0.26}_{-0.44}$. Owing to the large total mass, GW190521 is very short in duration, with only 4 cycles (2 orbits) within the sensitive frequency band of the GW observatories [46]. In contrast to the $\chi_{\rm p}$ measurement, the inferred $\rho_{\rm p}/\rho$ demonstrates a lack of measurable precession within the gravitational wave signal: $\rho_{\rm p}/\rho = 0.16^{+0.33}_{-0.13}$. This apparent contradiction between the inferred $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ measurements is due to GW190521's short duration. With only 4 visible cycles, GW190521 is decomposed into two near parallel "precession" harmonics (with overlap $|{\rm O}^{\rm prec}_{1,0}| = 0.97^{+0.01}_{-0.03}$). This means that any SNR orthogonal to the dominant harmonic is small and $\rho_{\rm p}/\rho$ is near-zero as a result. This contradiction between precession measurements highlights the difficulty in extracting source properties from a GW with only 4 visible cycles. It has been suggested that the large $\chi_{\rm p}$ could be a consequence of, for example, eccentricity [262] or head-on collisions [263]. However, neither affect $\rho_{\rm p}$ since the short duration of this signal means that the overlap between harmonics is always large. Precession in the system is not the only observed contradiction. Recently, Nitz *et al.* [273] challenged the inferred source properties of GW190521. They demonstrated that it is possible to obtain component masses which straddle the pair instability mass gap [274, 275, 276, 46, 146] with $q \sim 10$. A region of parameter space which was not explored in Ref. [46] owing to prior constraints.

As discussed above, when calculating the odds ratios for GW190521, the $\rho_{\rm p}/\rho$ and $\chi_{\rm p}$ analyses disagree more than any other event. Both demonstrate preference for LVC over ELI, but the $\rho_{\rm p}/\rho$ analysis calculates a greater preference for LVC than the $\chi_{\rm p}$ analysis. The reason for this difference is because the LVC model predicts $\rho_{\rm p}/\rho = 0.10^{+0.45}_{-0.09}$, which is consistent with GW190521's inferred distribution, and $\chi_{\rm p} = 0.19^{+0.30}_{-0.15}$, which is inconsistent, see Figure 7.1. Given that the LVC model has more support in the $\rho_{\rm p}/\rho$ parameter space than the $\chi_{\rm p}$ parameter space, a larger odds ratio is expected.

7.3.b Possible structure in the preferred spin distribution

By analysing BBHs in GWTC-2, it has been suggested that more massive BBHs could have formed from hierarchical mergers [277, 278, 279, 146], where the remnant of a previous binary becomes part of a new one [e.g. 280]. Given that the remnant of a previous equal mass binary is expected to have dimensionless spin $\chi \sim 0.7$ (inherited from the angular momentum of the previous binary) [281], we expect more massive BBHs to favour spin distributions with larger spin magnitudes if they were formed from hierarchical mergers. We therefore investigate whether the preferred spin distribution changes as a function of mass. For the purpose of this study we simply use 2 bins: one for *low mass* BBHs where both component masses

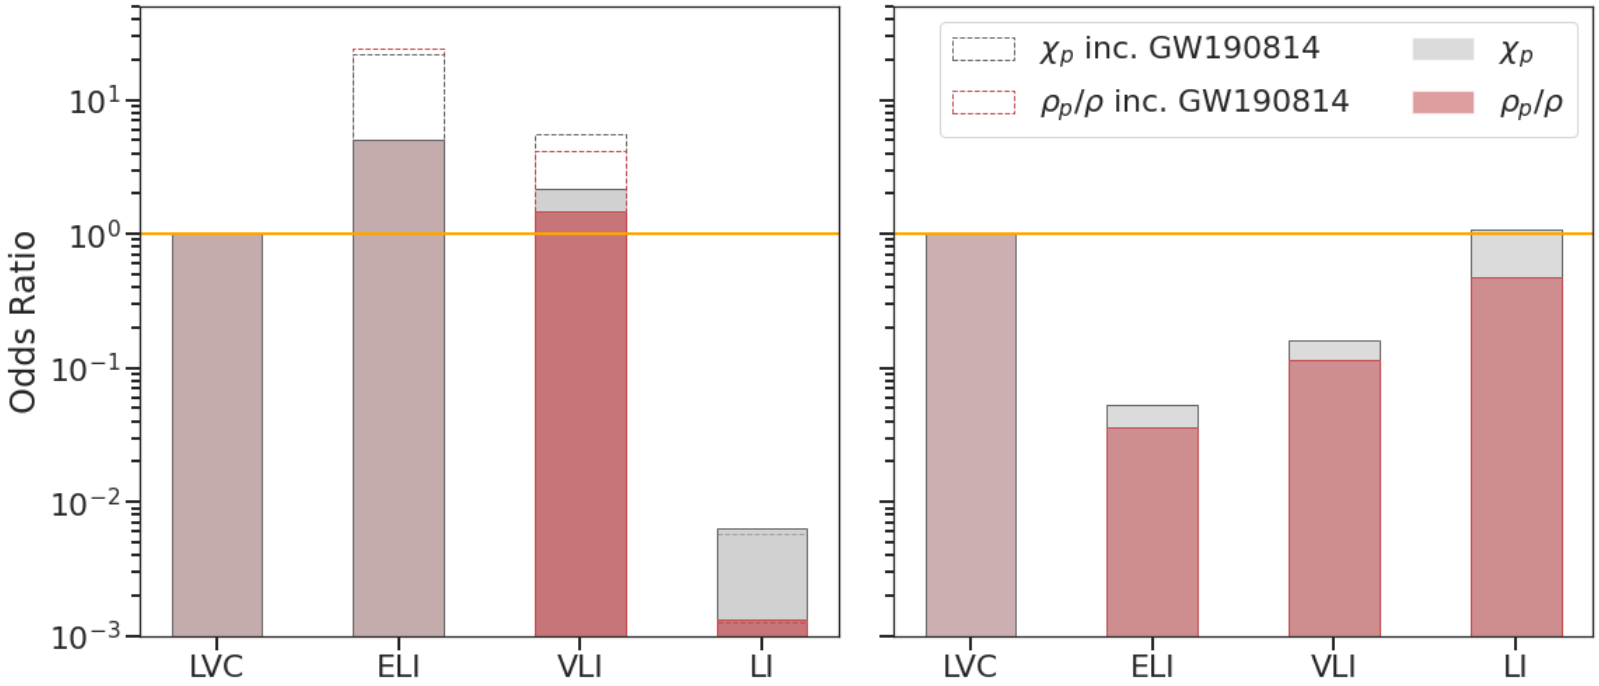

Figure 7.5: Odds ratios for the ELI, VLI and LI models in reference to the preferred model in Ref. [170] (LVC), see text for model definitions. The grey and red bars indicate the inferred odds ratios when $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ is used to parameterise precession respectively. Dashed lines show the inferred odds ratio when GW190814 is included in the analysis. The *Left* panel considers BBHs with both component masses less than $50M_{\odot}$ and the *Right* panel examins BBHs where at least one component mass is $50M_{\odot}$.

are less than $50M_\odot$ and another for *high mass* BBHs where at least one component mass is greater than $50M_\odot$. These bins were chosen since it is predicted that there exists a mass gap within the range $\sim 50 - 120M_\odot$ as a result of the pair-instability supernova theory [see e.g. 274, 275, 276, 46, 146]. This means that if an observed BBH has one or more component masses greater than $\gtrsim 50M_\odot$, it is likely that it was formed through a hierarchical merger. The following candidates have at least one component mass greater than $50M_\odot$ (based on the median of their mass distribution): GW170729, GW190519_153544, GW190602_175927, GW190620_030421, GW190521, GW190701_203306, GW190706_222641, GW190929_012149.

Figure 7.5 shows that for BBHs with both component masses less than $50M_\odot$, ELI is the preferred spin distribution model even when GW190814 is excluded from the analysis. We find that that when GW190814 is excluded, ELI is preferred to LVC with an odds ratio of 5 : 1 for both $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses. This is compared to 22 : 1 and 24 : 1 when GW190814 is included for the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses respectively. VLI, which has greater support for lower spin magnitudes than LVC, is the second preferred model with odds ratios 2.2 : 1 and 1.5 : 1 for the $\chi_{\rm p}$ and $\rho_{\rm p}/\rho$ analyses. Isotropic spin models are preferred to aligned models by an odds ratio $> 10^6 : 1$. This suggests that lower mass BBHs have lower spin magnitudes. On the other hand, we see the opposite trend for BBHs with at least one component mass greater than $50M_\odot$. For this case models with higher spin magnitudes (LVC, LI) are favoured against models with lower spin magnitudes (ELI, VLI). We have therefore shown that for lower mass BBHs, extremely low isotropic spins are preferred over low spins by $\sim 100 : 1$. Whereas for higher mass BBHs, this is reversed and low spins are preferred over extremely low by $\sim 20 : 1$. This suggests that using a single model to describe all BHs is not appropriate, especially with a growing catalog of BBH observations.

Ref. [170] also investigated whether there is evidence for a mass dependence in the BH spin distribution through a hierarchical Bayesian analysis of the population of known BBHs. Similar to the work presented here, Ref. [170] also found a weak preference for higher spin magnitudes in higher mass events (see e.g. Figure 13 of Ref. [170]). However, since the uncertainty on their measurement was broad, a mass dependence could not be confidently claimed. We anticipate that by using a discrete set of 9 different spin distributions, rather than a continuous model, we are able to reduce the possibility of over-fitting the noise and more clearly identify a possible mass dependence in the spin distribution. However, further work needs to be done to confirm the results presented in this section.

7.4 Conclusion

We have shown that, although we find preference for the same model as Ref. [170] when using the same sample of BBH mergers in GWTC-2, this distribution is only

marginally preferred to other models used in this chapter. We also show that the preference can change depending on whether GW190814 is included in the analysis or not. We also show that the majority of BBH mergers in GWTC-2 prefer a spin distribution model which has lower spin magnitudes and more isotropic spins than that predicted in Ref. [170]. Finally, we demonstrate that there is potential structure in the preferred spin distribution with BBHs containing at least one component mass $> 50M_{\odot}$ preferring higher spin magnitudes than BBHs with no component mass $> 50M_{\odot}$. We emphasise that these conclusions depends on our choices of possible spin distributions.

Chapter 8

PESummary: the code agnostic Parameter Estimation Summary page builder

This chapter is based upon the text of Hoy and Raymond [282] submitted to SoftwareX as arXiv:2006.06639. Charlie Hoy wrote all sections and created all Figures. Charlie Hoy is the primary author of the PESummary sofware. All authors edited the text.

Required Metadata

C1	Current code version	0.6.0
C2	Permanent link to code/repository used for this code version	https://github.com/pesummary/pesummary/tree/v0.6.0
C3	Legal Code License	MIT
C4	Code versioning system used	git
C5	Software code languages, tools, and services used	Python, Javascript, HTML, CSS, Bootstrap, Unix/MacOS
C6	Compilation requirements, operating environments & dependencies	numpy≥ 1.15.4, h5py, matplotlib, seaborn, statsmodels, corner, tables, deepdish, pandas, pygments, astropy≥ 3.2.3, lalsuite≥ 6.70.0, ligo-gracedb, configparser, gwpy, plotly, tqdm≥ 4.44.0
C7	If available Link to developer documentation/manual	See https://lscsoft.docs.ligo.org/pesummary
C8	Support email for questions	charlie.hoy@ligo.org

Table 8.1: Code metadata (mandatory)

8.1 Motivation and significance

Bayesian inference is central for many areas of science [e.g. 283, 284, 285, 286] as it allows for the identification of model parameters which best describes the collected data [e.g. 287, 288], see Section 8.2.a for more details. Since the first detection of gravitational-waves (GWs) [7], cosmic ripples predicted by Einstein's theory of General Relativity [289, 290], Bayesian inference has been used to infer the astrophysical source properties from GWs measurements [see e.g. 8, 243, 38], understanding the properties of noise within the GW detectors [291] and understanding the astrophysical population of GW sources [292]. For a detailed review of how Bayesian inference is used within the GW community see [267]. In addition, with more areas of science entering the 'open data era', there is a need for a robust, easy to use and code agnostic software to interpret and distribute the output from all Bayesian inference codes.

PESummary, the Parameter Estimation Summary page builder, is a Python package that provides an intuitive, object-oriented user interface for displaying, interpreting, comparing and distributing posterior samples from any parameter estimation code. PESummary's unified input/output (I/O) system enables the comparison of samples from codes that previously stored data in incompatible formats.

Since its first use in 2019, PESummary has grown to be a key component in the workflow of the LIGO Scientific [293], Virgo [294] and KAGRA [295] collaborations (LVK), as shown in Figure 8.1. It is used to analyse and compare the

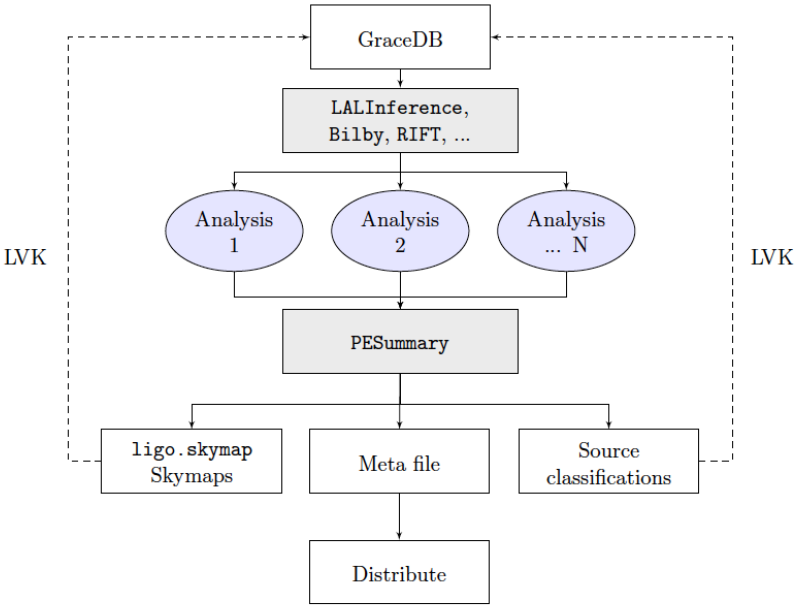

Figure 8.1: Flow chart showing the role of PESummary within the gravitational wave community. Dashed paths are specific for the LIGO Scientific, Virgo and KAGRA collaborations (LVK). Once a gravitational-wave is observed and uploaded to GraceDB, numerous Bayesian inference analyses are launched to extract the astrophysical source properties from the gravitational-wave measurement. The output data from all analyses is then passed to `PESummary` for post-processing. Alongside webpages for displaying, interpreting and comparing the Bayesian inference data, `PESummary` produces skymaps and source classifications, and a single universal 'metafile' containing information about each Bayesian inference analyses undertaken. This file can then distributed to the wider community. For the LVK workflow, skymaps and source classification probabilities are uploaded to GraceDB for circulation.

outputs from the popular GW Bayesian inference software packages: `LALInference` [296], `RIFT` [297, 298, 299], `PyCBC` [92] and `Bilby` [300, 301, 215] (both `PyCBC` and `Bilby` provide an interface for a number of popular sampling packages [e.g. 302, 303]), as well as circulating both skymaps through the GW candidate event database GraceDB [252] and Bayesian inference data through the LIGO Document Control Center. PESummary was critical in the parameter estimation analysis of GW190412 [45], GW190425 [43] and the re-analysis of the first gravitational-wave transient catalog (GWTC-1) [38] using BILBY. The released data were also in the PESummary data format [304, 305].

This chapter does not present a complete record of all the capabilities of PESummary, for more information, see https://lscsoft.docs.ligo.org/pesummary/stable_docs/index.html.

8.2 Software description

The main PESummary interface is implemented in pure Python [306] with the core library relying on a number of established scientific programming packages [307, 308, 309, 310, 311, 312, 313, 314, 315] and the GW library requiring custom GW data analysis software: `LALSuite` [175] and `GWpy` [316].

PESummary can be used for more than simply post-processing parameter estimation results. This library includes, but not limited to, new Kernel Density Estimators [317] to estimate PDFs of random variables within specified domains (one and two dimensional), a unified infrastructure for reading data in different formats, a release ready `HDF5` [308] file for storing one or more Bayesian inference analyses, a comprehensive plotting suite and webpage module for analysing and comparing Bayesian inference data, and checkpointing, a requirement for cloud based computing.

PESummary is designed to be as modular and adaptable as possible, ensuring the code will age gracefully with advances in Bayesian inference software. Where possible, the code is kept general meaning additional post-processing techniques can easily be implemented into the PESummary framework.

8.2.a Software Architecture

PESummary is structured around a small number of class objects. Each object provides class and instance methods to provide the user with a complete interface for all operations.

Tabular Data

Bayesian inference returns the best fit model parameters through a *posterior* probability distribution function (PDF) calculated through Bayes' theorem [318],

$$p(\theta|d) = \frac{p(\theta)p(d|\theta)}{p(d)} \tag{8.1}$$

where $p(\theta|d)$ is the posterior probability of the model having parameters θ given the data d, $p(\theta)$ is the prior probability of the model having parameters θ, $p(d|\theta)$ is the likelihood of observing the data given a model with parameters θ and $p(d)$ is the evidence of the data, a model independent quantity. Often the likelihood and the posterior probability distribution function are unknown analytically. Consequently, most Bayesian inference packages are designed to draw samples from the unknown posterior PDF [319, 320]. Samples approximating a posterior PDF for a single parameter, otherwise known as a marginalized posterior PDF, are then identified by integrating the posterior PDF over all other parameters. Often, these posterior samples are output in the form of a table saved in a 'result file'.

PESummary's stores posterior samples in a custom high level table structure, `pesummary.utils.samples_dict.SamplesDict`, a subclass of the Python builtin `dict` object. This class offers numerous methods for analysing the stored data, including the `.plot` method. Each marginalized posterior distribution is stored as a `pesummary.utils.samples_dict.Array` object, an inherited class of `numpy.ndarray` [321]. This structure provides direct access to the optimised array functions from NumPy [307] and the usability and familiarity of dictionaries.

A popular algorithm for sampling from the posterior distribution is Markov-Chain Monte-Carlo [319]. For this case, multiple Markov chains are often run in parallel to test convergence through the Gelman-Rubin statistic [322]. PESummary stores multiple Markov chains in the `pesummary.utils.samples_dict.MCMCSamplesDict`) class, a dictionary where each chains posterior samples are represented by a `pesummary.utils.samples_dict.SamplesDict` object. Algorithms for removing burnin are available via the `.burnin()` method. Convergence between chains can be measured via the `.gelman_rubin()` method.

Often multiple Bayesian inference analyses are performed to identify how the PDFs vary for different settings. Consequently, PESummary provides the `pesummary.utils.samples_dict.MultiAnalysisSamplesDict` class, inherited from the the Python builtin `dict` object, for storing multiple `pesummary.utils.samples_dict.SamplesDict` tables, each with an assigned label.

8.2.b Packages

PESummary has followed the same methodology as BILBY by separating the top level code into 2 packages: `core` and `gw`. The `core` package provides all of the necessary code for analysing, displaying and comparing data files from general inference problems. The `gw` specific package contains GW functionality, including converting posterior distributions, deriving event classifications and GW specific plots, see Sec.8.2.b.

The `core` package

The `pesummary.core` package provides all of the code required for post-processing general inference result files. It is designed to be generic and therefore work with the output from any Bayesian inference software. It provides a unified interface for producing plots, calculating useful statistics and generating webpages.

Plots are generated via the `pesummary.core.plots.plot` module. One dimensional marginalized posterior distributions are visualised as either a histogram, a kernel density estimate (`scipy.stats.gaussian_kde` or `pesummary.core.plots.plots.bounded_1d_kde.Bounded_1d_kde`) or cumulative distributions. Diagnostic plots displaying the marginalized trace or autocor-

relation are also available. Comparison plots between multiple result files can be achieved by producing comparison one dimensional marginalized posterior distributions, cumulative distributions, box plots, violin plots, two dimensional contour and scatter plots or Jensen–Shannon [257] and Kolmogorov–Smirnov [323, 324] statistics.

The `pesummary.core.webpage.webpage` module is a Python wrapper for writing HTML. The `pesummary.core.webpage.webpage.page` class provides functionality for generating multi-level navigation bars, tables of images and/or data, modal carousels and more. The design and functionality of the webpages are controlled through the `pesummary.core.css` and `pesummary.core.js` modules; each containing custom style sheets (CSS) or JavaScript files (JS) respectively.

The `combine_corner.js` script contains functionality to generate custom corner plots by manipulating a pre-made figure [made with e.g. 310]. By providing an interface for the user to specify parameters they wish to compare, the PESummary webpages allow for the PDF, and its correlations, to be interactively explored.

The `gw` package

The `gw` package provides the functionality for parameter estimation specific to GW astronomy. Building on the `core` package, the `gw` module provides additional methods and classes for handling GW specific data files. Although the `gw` package is tailored for compact binary coalescence data files, we provide GW specific methods which can be applied to the Bayesian inference data from any transient GW source.

The `gw` package provides a comprehensive conversion module (`pesummary.gw.file.conversions`) for deriving alternative posterior distributions from the input, e.g. the primary mass and secondary mass from chirp mass and mass ratio. Assuming a binary black hole merger, the conversion class provides multiple methods for estimating the properties of the final black hole: final mass (or equivalently the energy radiated in gravitational-waves), final spin, peak luminosity and final kick velocity [79, 224, 75, 80, 81, 325, 326, 327, 328, 329, 330, 331, 332]. All fits are calibrated to numerical relativity and are interchangeable via keyword arguments provided to the conversion class.

On the 17th August 2017, alerts [333] were sent out to more than 60 international groups of astronomers notifying them that a GW had just been detected: GW170817 [334]. Each group began observing the night sky to independently observe the source of the GW. GW170817 opened the window of a long-awaited multi-messenger astronomy [335]. Through interfacing with the `ligo.skymap` [336] package, PESummary provides an intuitive method for generating skymaps: data files showing the most likely source location of the GWs. These skymaps are distributed to astronomers through GraceDB, see Figure 8.1. In the last GW observing period, 26 alerts were sent out regarding possible GW candidates, each containing skymaps [337]. PESummary also interfaces with the `PEPredicates` [338] and `P-`

Package	Format	Class	Description
`core`	.dat	`.default.Default`	'.dat' format with parameter names as header
`core`	.txt	`.default.Default`	'.txt' format with parameter names as header
`core`	.h5/.hdf5	`.default.Default`	'.hdf5' format with data stored in a group called 'posterior' or 'posterior_samples'
`core`	.json	`.default.Default`	'.json' format with posterior samples stored in a group called 'posterior' or 'posterior_samples'
`core`	BILBY	`.bilby.Bilby`	Result file produced by BILBY
`gw`	BILBY	`.bilby.Bilby`	Result file produced by BILBY
`gw`	LALINFERENCE	`.lalinference.LALInference`	Result file produced by LALINFERENCE

Table 8.2: A selection of formats that can be read in with PESummary and accessible through the unified I/O interface for the listed class object(s). The core classes are all prepended by `pesummary.core.file.formats`, the gw classes are all prepended by `pesummary.gw.file.formats`.

`Astro` [339] packages to provide source classification probabilities for the type of compact binary merger observed; such as the probability that the system is a binary black hole or a binary neutron star. Both the skymap and source classification probabilities are vital for electromagnetic follow-up campaigns.

PESummary takes advantage of the `plotly` [315] interactive plotting package to produce interactive corner plots for extrinsic and intrinsic parameters. The `pesummary.gw.plots.publication` module also allows for 'publication' quality plots to be produced, for instance those in [38].

8.2.c Software Functionalities

Unified input/output

Bayesian inference packages output their posterior samples in varying formats. In just the GW community alone, `LALInference`, `Bilby` and `RIFT` output their data in 3 different formats: `HDF5` [308], `JSON` [340] and `dat`. Although `Bilby` is able to output it's data in `HDF5`, it's posterior samples are stored differently to `LALInference`: `LALInference` stores posterior samples as a `numpy.recarray`

while Bilby writes each marginalized posterior distribution as a separate dataset, and an alternative software package for reading HDF5 files is required. Each data file also stores different metadata in different locations. This incompatibility makes it difficult to compare the contents from different Bayesian inference samplers both in and out of the GW community.

PESummary provides a unified infrastructure for reading data in different formats via the pesummary.io module. The universal pesummary.io.read() function reads data stored in all common file formats via the pesummary.core.file.formats module. The GW specific package extends the allowed file formats via the pesummary.gw.file.formats module. See Table 8.2 for a reduced list. Once read, the posterior samples from different data files can be compared through the common '.samples_dict' property. All read result files can be written to a specified file format via the common .write method. This method calls the universal pesummary.io.write() function.

PESummary offers the pesummary.core.file.meta_file module for producing a universal file containing posterior samples and metadata for one or more analyses. This single 'metafile' aims to store all information regarding the Baysian inference analysis. This includes prior samples, configuration files, injection information, environment information for reproducibility etc. For GW analyses, the pesummary.gw.file.meta_file module allows for the PSD, skymap data and calibration information to also be saved.

Executables

PESummary provides multiple executable Python scripts to act as an intermediary between the command line and the core functionality. See Table 8.3. summarypages is the main executable and combines summaryclassification, summaryclean, summarycombine, summarydetchar, summaryplots and summarypublication. summarypages is the most general executable provided by PESummary. Its GW workflow is described in Fig. 8.2. The core workflow is similar, except GW specific plots and webpages are not generated and the core classes are used.

The summarypages executable takes one or more result files as input via the --samples command line argument. A plethora of optional command line arguments are also available to allow for complete customisation[1]. All options are then passed to the pesummary.gw.inputs.GWInput class (inherited from the pesummary.core.inputs.Input class) which checks the inputs. Amongst these checks, the samples are extracted from all result files (see Section 8.2.c), converted to generate additional PDFs, and stored in a pesummary.utils.samples_dict.MultiAnalysisSamplesDict object with an assigned label.

[1]For details run summarypages --help or vist the online documentation available https://lscsoft.docs.ligo.org/pesummary/stable_docs/core/cli/summarypages.html.

Executable	Description
`summaryclassification`	Generate GW event classification probabilities
`summaryclean`	Clean and convert an input result file
`summarycombine`	Combine multiple result files into a single format
`summarycompare`	Compare the contents of multiple files
`summarydetchar`	Generate GW Detector characterisation plots
`summarygracedb`	Retrieve data from the online GW Candidate Event Database [252]
`summarymodify`	Modify a PESummary meta file
`summarypages`	Generate webpages, plots and a meta file for N result files
`summarypageslw`	Generate webpages, plots and meta files for a select number of parameters
`summaryplots`	Generate specific plots for a given result file
`summarypipe`	Generate a `summarypages` executable given a GW run directory
`summarypublication`	Generate publication plots for N result files
`summaryrecreate`	Launch the GW analysis that was used to generate the PESummary meta file
`summaryversion`	Version of PESummary installed

Table 8.3: A selection of executable python scripts provided by PESummary

All properties of the `GWInput` class are then used to initiate the `pesummary.cli.summaryplots.PlotGeneration` class. Here, all plots are generated on multiple CPUs, if specified, with the `.generate_plots` method. This includes source classification plots, marginalized posteriors, interactive plots and more. Initial skymaps are generated based on a two-dimensional histogram of the right ascension and declination samples. A subprocess is then launched to generate the more accurate `ligo.skymap` skymap. This implementation allows for the rest of the pipeline to continue without having to wait for the often time-consuming `ligo.skymap` skymap to be produced. Once complete, the initial skymap is overwritten.

All webpages are then produced by running the `.generate_webpages` method once the `pesummary.cli.summarypages.WebpageGeneration` class is initialised with all `GWInput` properties. This includes single webpages for each parameter in each result file, interactive corner pages, comparison pages for all common parameters, a version page providing version and environment information and a downloads page. Finally a single metafile containing all information from the run is produced. This includes environment information, posterior samples, command line arguments and more. This file is available for download via the downloads page.

The output from `summarypages` is completely self-contained; allowing for it to distributed if required. An example of the `summarypages` output can be seen here: https://pesummary.github.io/GW190412/home.html. This output was pro-

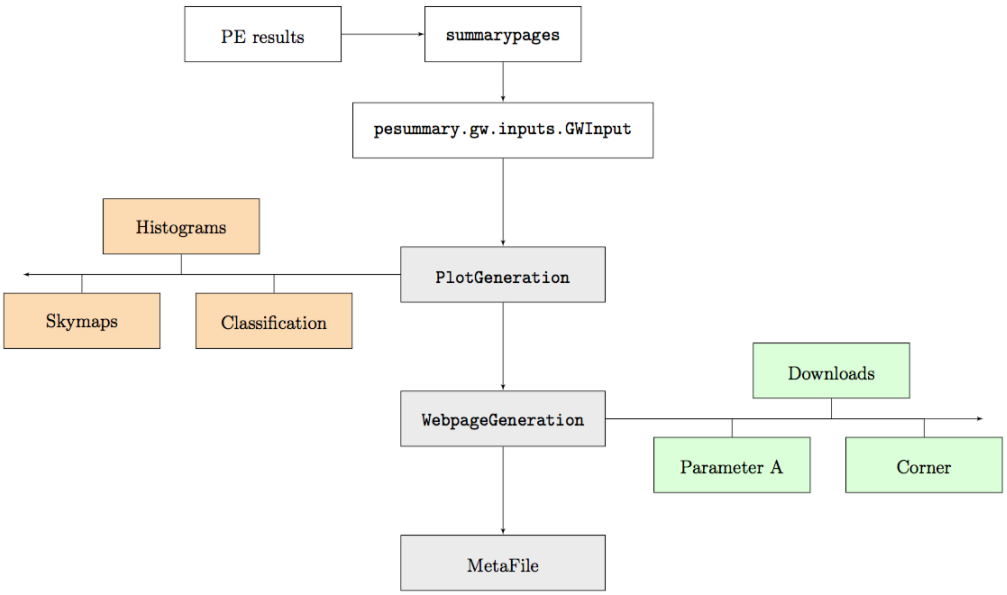

Figure 8.2: Workflow for the most general executable provided by PESummary: `summarypages`. Here we show the workflow when the `gw` package is used.

duced from Listing 8.4.

Amidst the collection of optional command line arguments, the `pesummary.gw` module also provides a dynamic argument parser. This allows for dependent arguments to be passed from the command line (see section 8.3.e).

8.3 Illustrative Examples

Below we provide a limited set of examples to demonstrate some of the features of PESummary. All data and scripts that are used as part of this section can be downloaded from https://github.com/pesummary/pesummary-paper-data. More examples can be found in the PESummary repository. The following examples assume that you have cloned the pesummary-paper-data repository and you are in the base directory.

8.3.a Example 1: Running with `emcee`

This example shows the flexibility of the PESummary post-processing package. We run the *Fitting a model to data* tutorial provided as part of the `emcee`[302] sampling package. We save the posterior samples to a '.dat' file and run with 8 chains. All post-processing is then handled with PESummary, see Figure 8.3 for example output. The `emcee` code snippet, as well as the posterior samples can be downloaded from the pesummary-paper-data repository.

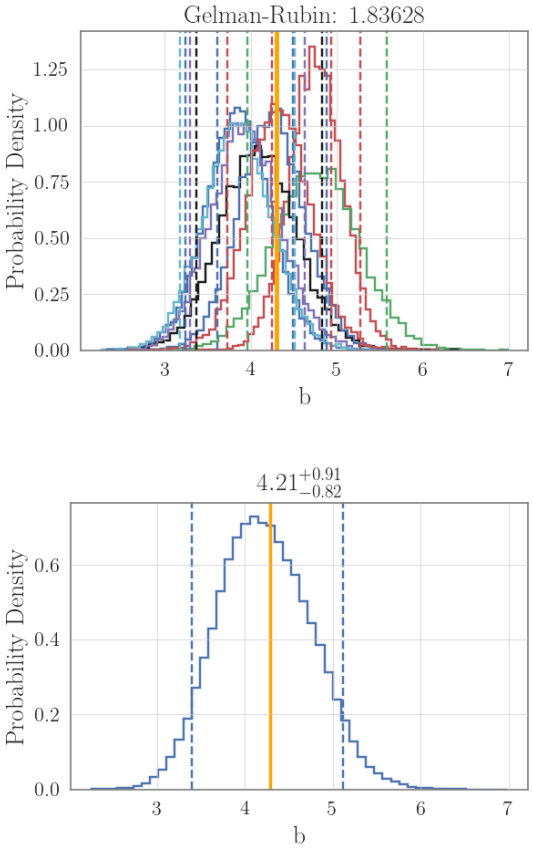

Figure 8.3: The output marginalized posterior for each chain (Top) and the posterior from combining all chains (Bottom) resulting from Listing 8.1. The injected value is shown in orange.

Listing 8.1: Running PESummary on the `emcee` output.

```
python emcee_tutorial.py
chains=($(ls emcee_output/chain_*.dat))
summarypages --webdir ./webpage --samples ${chains[@]}
             --labels tutorial --mcmc_samples --inj_file
                 emcee_output/injected.txt
```

8.3.b Example 2: Reading a result file

This example demonstrates how to read in a PESummary 'metafile' using the `pesummary.io` module. We use the `posterior_samples.h5` file produced as part of Listing 8.1 (running Listing 8.1 is not necessary for this example, the metafile is available in the pesummary-paper-data repository) and specify that we would like to use the `core` package meaning only the core file formats are allowed. We show how to extract the posterior samples for a specified analysis, how to access the marginalized posterior samples, and how to extract the prior samples and any

stored configuration files. Some of the attributes of each object are also shown. For full documentation, run the inbuilt Python help function `help` [341].

Listing 8.2: Example code to read in a PESummary result file with PESummary.

```
from pesummary.io import read

f = read("posterior_samples.h5", package="core")
multi_analysis_samples = f.samples_dict
print(type(multi_analysis_samples))
analysis = multi_analysis_samples.labels[0]
posterior_samples = multi_analysis_samples[analysis]
print(type(posterior_samples))
print(posterior_samples.keys())
marginalized = posterior_samples["m"]
print(type(marginalized))
print(marginalized.average(type="median"))
priors = f.priors["samples"]
config = f.config
```

8.3.c (GW) Example 3: analysing public LIGO and Virgo posterior samples

This example demonstrates how to extract and analyse public LIGO and Virgo posterior samples. It includes demonstrations of how to produce 'standard' plots for the 'combined' analysis stored in the PESummary 'metafile' through the `.plot` method, see Figure 8.4. We use the publicly available posterior_samples.h5 file, which has been copied to the pesummary-paper-data repository.

Listing 8.3: Extracting data from public LIGO and Virgo posterior samples

```
from pesummary.io import read

f = read("GW190412_posterior_samples.h5", package="gw")
analysis = "combined"
posterior_samples = f.samples_dict[analysis]
psds = f.psd[analysis]
calibration_envelope = f.priors["calibration"][analysis]
prior_samples = f.priors["samples"][analysis]

hist = posterior_samples.plot("mass_1", type="hist")
hist.show()
skymap = posterior_samples.plot(type="skymap")
skymap.show()
plot = psds.plot(fmin=20)
plot.show()
```

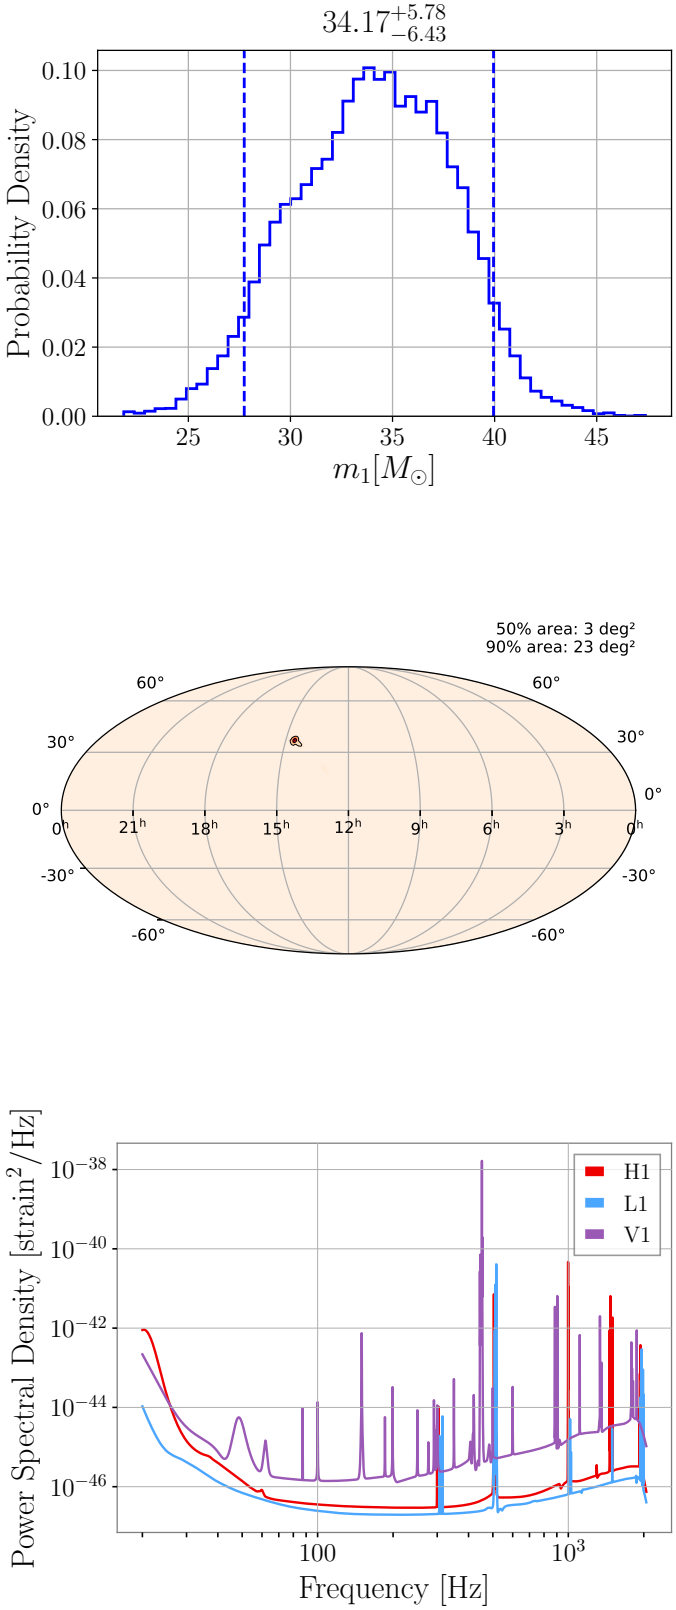

Figure 8.4: The output marginalized posterior (top), skymap (middle) and PSD plot (bottom) from Listing 8.3.

8.3.d (GW) Example 4: Producing a summary page for public LIGO and Virgo posterior samples

This example demonstrates how a summary page can be generated from publicly available LIGO and Virgo data files. We use the same publicly available data file as Listing 8.3. This 'metafile' contains a total of 10 Bayesian inference analyses, each using different standard models used by the LVK. In this example, we choose to compare only a select subset: an analysis conducted with the IMRPhenomPv3HM model [81] and the SEOBNRv4PHM model [224, 162, 79] (see [45] for details). We specify that we would like to use the `gw` package (`--gw`), run on 15 CPUs (`--multi_process 15`), and generate 'publication' quality plots (`--publication`). The output page can be found at: https://pesummary.github.io.

Listing 8.4: Generating a html page to compare and analyse the public LIGO and Virgo posterior samples. This may take some time.

```
summarypages --webdir ./GW190412 \
             --samples GW190412_posterior_samples.h5 \
             --gw --compare_results IMRPhenomPv3HM SEOBNRv4PHM \
             --publication \
             --multi_process 15
```

8.3.e (GW) Example 5: PESummary's dynamic argument parser

`argparse`, the Python module for handling command-line arguments (CLA) [342], requires a known list of arguments. This means that they cannot depend on another variable. Through the `pesummary.gw.command_line` module, PESummary allows the user to specify CLAs which change depending on the provided label. Below we show how PSDs can be provided to `summarypages` dynamically through the `--{}_psd` CLA. Result files and psd data was created using the 'make_data_for_listing5.py' script made available in the pesummary-paper-data repository.

Listing 8.5: Example usage of PESummary's dynamic argument parser

```
summarypages --webdir ./webpage \
             --samples test.hdf5 test.json \
             --labels hdf5_example json_example \
             --hdf5_example_psd H1:psd_H1.dat \
             --json_example_psd V1:psd_V1.dat L1:psd_L1.dat \
             --gw
```

8.3.f (GW) Example 6: Reproducing LIGO and Virgo plots

This example demonstrates how to reproduce a subset of plots in the first GW transient catalog (GWTC-1) [Figures 4 and 5 in Ref. 38]. We use publicly available posterior samples released as part of GWTC-1 (ignoring the GW170809 [38]

prior choices file): https://dcc.ligo.org/public/0157/P1800370/005/GWTC-1_sample_release.tar.gz, which have been copied to the pesummary-paper-data repository. Figure 8.5 shows an example of the output. When running this Listing, PESummary will print multiple warnings to `stdout`. These warnings are expected and shows that PESummary is robust to potential failures and will continue to produce an output while still warning the reader appropriately. One such example is a message warning the user that the './GWTC-1_sample_release/GW170817_GWTC-1.hdf5' cannot be read in and data will not be added to the plot as this file contains 'multiple posterior sample tables: IMRPhenomPv2NRT_highSpin_posterior, IMRPhenomPv2NRT_lowSpin_posterior' and we have not specified which we wish to load from the command line.

Listing 8.6: Example code to generate GW plots with PESummary. This may take some time.

```
FILES=($(ls ./GWTC-1_sample_release/*_GWTC-1.hdf5))

LABELS=()
for i in ${FILES[@]}; do
    label=`python -c "print('${i}'.split('_GWTC-1.hdf5')[0])"`
    LABELS+=($(python -c "print('${label}'.split('./GWTC-1
        _sample_release/')[1])"))
done

COLORS=('#00a7f0' '#9da500' '#c59700' '#55b300' '#f77b00' '#ea65ff
    ' '#00b1a4' '#00b47d' '#00aec2' '#9f8eff')
LINESTYLES=(solid dashed solid solid dashed solid solid dashed
    dashed dashed solid)

summarypublication --plot 2d_contour \
                   --webdir ./GWTC-1_sample_release \
                   --samples ${FILES[@]} \
                   --labels ${LABELS[@]} \
                   --parameters mass_1_source mass_2_source \
                   --colors ${COLORS[@]} \
                   --linestyles ${LINESTYLES[@]} \
                   --publication_kwargs xlow:0 xhigh:80 ylow:0
                       yhigh:50

summarypublication --plot violin \
                   --webdir ./GWTC-1_sample_release \
                   --samples ${FILES[@]} \
                   --labels ${LABELS[@]} \
                   --parameters mass_ratio
```

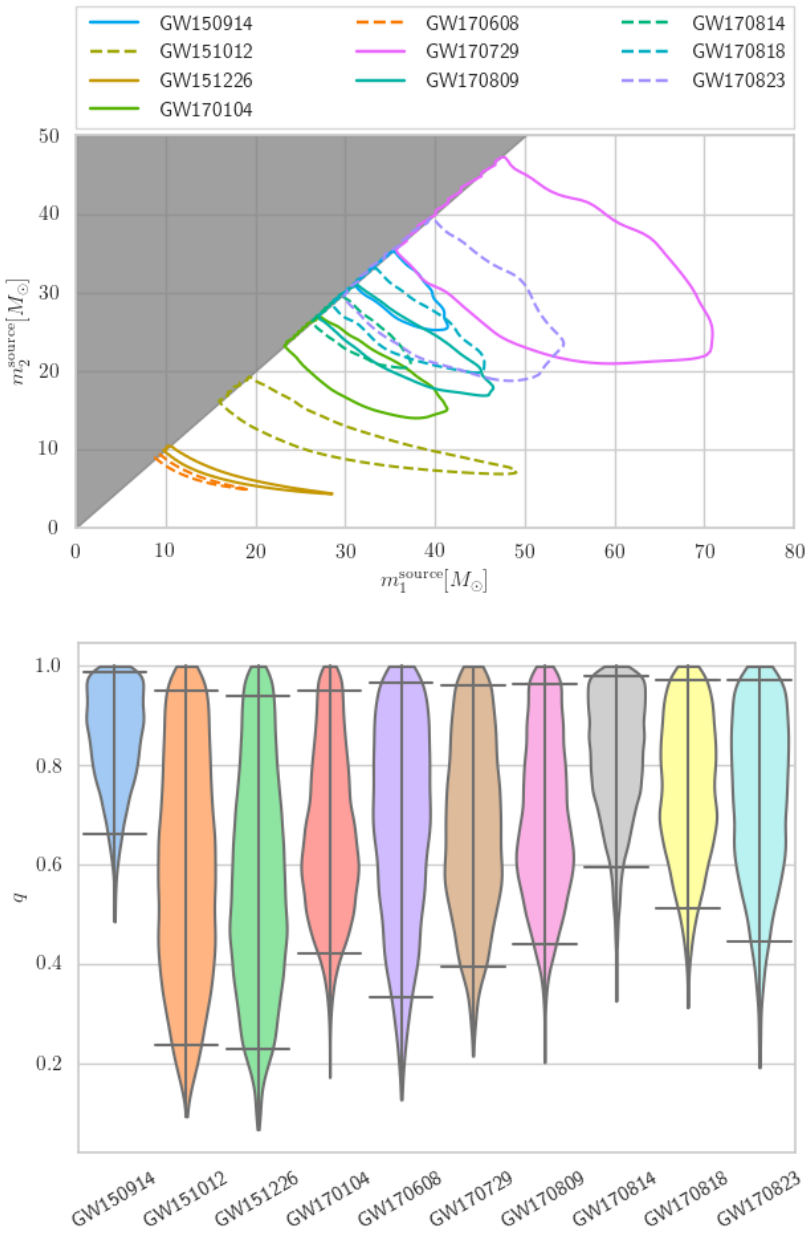

Figure 8.5: Two plots output from Listing 8.6. Top: Recreation of Figure 4 in Ref. [38], Bottom: Recreation of Figure 5 in Ref. [38].

8.4 Impact

PESummary is now a widely used library in the LVK. In just over one year, PESummary has become the post-processing software for the main LVK GW parameter estimation codes [300, 301, 296, 343, 344, 345, 250] and is relied upon for the open data release of the LVK parameter estimation data products [346]. Rather than releasing multiple data products in different formats spread across numerous URLs, PESummary has greatly simplified this to simply releasing a single file [304, 305].

This can propel future research in the field, as researches no longer have to create custom scripts for combining, retrieving, and recreating the initial analysis.

The flexibility and intuitive Python executables provided by PESummary, has significantly improved the efficiency of researchers (often early-career graduates or graduate scientists), by reducing the need for repetitive tasks. In particular, PESummary's interactive corner plots has led to the increased knowledge of degeneracies between parameters for researchers within the LVK.

PESummary is also fully incorporated into the automated low-latency alert workflow [347]. Posterior based classifications are therefore automatically produced and posted to the online GW Candidate Event Database [252] from the automatic parameter estimation follow-up. This information will greatly aid electromagnetic followup campaigns.

8.5 Conclusions

In this chapter we have described PESummary, the Parameter Estimation Summary page builder. This Python package provides a modern interface for displaying interpreting, comparing and distributing Bayesian inference data from any parameter estimation code. PESummary has been used extensively by the international GW community, and has been crucial for Advanced LIGO's [5] and Advanced Virgo's [6] third gravitational wave observing run.

Although PESummary is primarily used for post-processing GW Bayesian inference data from compact binary coalescenses, looking forward, we plan to incorporate other GW fields, e.g. Tests of General Relativity [see e.g. 153]. This will enable PESummary to be the driving force behind the post-processing and distribution of all Bayesian inference data output from the LVK. We will also continue to expand our already comprehensive plotting suite to improve its versatility for any Bayesian inference analysis.

Chapter 9

Conclusion

Almost a century after Einstein first proposed his revolutionary theory of General Relativity [1], a Nobel prize-winning team [348] witnessed one of the most violent processes in the universe: the cataclysmic collision of two black holes, releasing tiny cosmic ripples known as gravitational waves [7, 8]. Excitingly, over the past five years, around fifty gravitational waves have been observed, each contributing to the growing census of known black holes [7, 32, 33, 35, 36, 37, 38, 39, 40, 41, 42, 44, 45, 46]. As existing ground based gravitational wave detectors [4, 5, 6] continue to increase their sensitivity, and future detectors join the network [11, 349], the number of gravitational wave observations is expected to increase. For instance, the fourth gravitational wave observing run (O4), scheduled to start in mid 2022 [350] with a network of four gravitational wave detectors [4, 5, 6, 11], is predicted to observe $\sim 5\times$ more binary black hole coalescences than the third gravitational wave observing run (O3) [4]. With this ever increasing catalog of binary mergers, we hope to gain greater understanding of the underlying properties of binary black holes in the universe. One such property which is pivotal to our understanding is spin-induced orbital precession [54] as it is one of the most promising tracers for understanding the formation mechanism of binary black holes. We therefore require a method to easily quantify if spin-induced orbital precession has been observed in gravitational wave data.

To summarize, this thesis has provided details of an elegant new representation of orbital precession in binary systems. I demonstrated that this novel approach made it possible to easily quantify if there is any measurable precession in gravitational wave data, allowed us to identify which binaries are most likely to emit gravitational waves with measurable precession, and used the observed gravitational wave detections from the first, second and first half of the third gravitational wave observing runs to a) calculate the probability of observing precession for a multitude of possible spin distributions, b) predict that precession is likely to be observed during O3, c) demonstrate that several events in O3a are likely to have originated from precessing systems and d) make inferences about the underlying spin distribution of

black holes in the universe. This thesis has described techniques and tools that are necessary to make crucial steps towards understanding the formation mechanism of binary black holes. The first portion of this thesis introduced an innovative tool for quantifying orbital precession in gravitational wave data and demonstrates its efficacy. This tool formed the basis for subsequent chapters.

Chapter 2 demonstrated that a gravitational wave emitted from a precessing system can be decomposed into a power series of 5 non-precessing gravitational waves, where the characteristic amplitude and phase modulations are a consequence of their constructive and destructive interference. We showed that in the majority of the parameter space, the leading two terms are enough to reconstruct a precessing gravitational wave (the "two-harmonic approximation"). This allowed us to introduce the notion of a "precession signal-to-noise ratio" which was used to determine, for the first time, whether precession effects are observable in a given system.

Chapter 3 presents an extensive parameter-estimation study to identify the binary parameters where we are able to accurately identify precession. Through performing a series of one-dimensional investigations of the parameter space, and comparing the inferred precession signal-to-noise ratio to the often used scalar quantity $\chi_{\rm p}$, we verified that the precession signal-to-noise ratio provides a suitable and intuitive metric for determining whether or not precession effects are observable in a given system. We also demonstrated that there is an approximate mapping between the precession signal-to-noise ratio and the Bayes factor, which suggests that future large scale studies of precession can be made considerably cheaper through the use of the precession signal-to-noise ratio.

The second portion of this thesis demonstrated that the precession signal-to-noise ratio can be used to make inferences about the properties of black holes based on data collected from the first, second and third gravitational wave observing runs. It also presented a method for analysing systems without performing expensive Bayesian analyses through taking advantage of the precession signal-to-noise ratio.

Chapter 4 exploited the precession signal-to-noise ratio to calculate the probability of observing a black hole binary with measurable precession for nine astrophysically-motivated populations. We then identified which of the models considered is most likely given the data from the first and second gravitational wave observing runs. We made the prediction that for the most likely spin distribution model, there is an 83% chance of observing precession after ~ 40 observations. We also presented the surprising new result that most observations of precession will be in comparable-mass binaries. A study of this nature was previously not possible owing to the difficulty in classifying observable precession.

Chapter 5 discussed the properties of potentially the first neutron star black hole binary ever detected. Not only does this observation potentially point to a brand new subclass of compact binary objects, but also demonstrated that compact objects exist in the hypothesised lower mass gap between the populations of known neutron

stars and black holes; the secondary object in this binary is either the lightest black hole or the heaviest neutron star ever observed in a double compact-object system. This observation challenges our current understanding of astrophysical models. In the final section of this chapter, I demonstrated that the in-plane spin and inclination angle of this system can be estimated to within a good accuracy without performing expensive Bayesian analyses. This was done by utilizing the precession signal-to-noise ratio. This method indicates that it may be possible to estimate all source properties to within a good estimate by simply using information from the gravitational wave search pipelines. The hope is that this method can be optimized and extended to provide an algorithm that can estimate the source properties within several minutes after the gravitational wave was first observed.

Chapter 6 calculated the precession signal-to-noise ratio for all of the events in the first half of the third gravitational wave observing run (O3a). By comparing the average precession signal-to-noise ratio across all events to the expected distribution from noise, we demonstrated that in general there is no strong evidence for precession in the population of BBHs in O3a. However, we highlighted that three gravitational wave candidates could have originated from precessing systems, with GW190412 disfavouring the non-precessing scenario by $> 97\%$ probability.

Chapter 7 took advantage of the precession signal-to-noise ratio and presented a detailed model selection analysis to make inferences about the most likely spin distribution of black holes given the gravitational wave data from the first, second and third gravitational wave observing runs. We first demonstrated that the results presented in Ref. [170] are robust. However, we then highlighted that the majority of events prefer a different distribution: one with lower spin magnitudes and isotropic spins. We then pointed to potential structure in the black hole spin distribution, with high mass events preferring larger spin magnitudes than low mass events. This chapter combines the work from Chapters 2, 3, 4 and 6.

Finally, in the third portion of this thesis, Chapter 8 presented a new and innovative software package to analyse, display and combine posterior samples. This package has become one of the major workhorses of the LIGO, Virgo and KAGRA collaborations and is widely distributed through the gravitational wave data analysis computing environment. This package is central to making gravitational wave data open and easily reproducible.

Through the use of the precession signal-to-noise ratio, this thesis has demonstrated that there is no strong evidence for precession in the population of binary black holes. It then used this information to constrain the underlying spin distribution of black holes in the universe. However, these conclusions are dependent on the small number of gravitational wave candidates observed to date. Since O4 is predicted to detect $\sim 5\times$ more binary black hole events than O3, we expect to observe more events in extreme regions of the parameter space (asymmetric mass ratios and high spin magnitudes) for which spin-induced orbital precession, if present,

will be easier to measure (see Chapter 3 for details). O3 has already hinted at events in these extreme regions of the parameter space [45, 46, 47]. Consequently, as demonstrated in this thesis, any future study which investigates the observability of precession in the growing population of binary black holes would benefit greatly from the precession signal-to-noise ratio. This is because of it's ability to provide a more intuitive understanding and a reduction in the computational cost of multiple analyses. As Chapter 2 eluded to, future work which incorporates the two harmonic approximation into existing search pipelines would be extremely beneficial and pertinent as it increases our sensitivity to possible precessing binaries. It also means that estimates for the precessing sensitive volume can be calculated. Leading on from this, any future study which builds upon the work conducted in Chapter 5 and constructs a pipeline which bounds the source properties from the output of the search pipelines would be extremely valuable. This is because it has the potential to greatly reduce the computational cost of Bayesian analyses through incorporating custom jump proposals [see e.g. 301]. Finally, as Chapter 7 identified, any future studies which estimate the spin distribution of black holes should consider how their inference changes as a function of mass. The is because high mass candidates, which are not expected to form via the same method as low mass candidates, could influence results and cause incorrect conclusions to be drawn about the formation mechanism of binary black holes.

Bibliography

[1] Albert Einstein. Fundamental ideas of the general theory of relativity and the application of this theory in astronomy. *Preussische Akademie der Wissenschaften, Sitzungsberichte*, 315, 1915.

[2] Isaac Newton. *The Principia: mathematical principles of natural philosophy*. Univ of California Press, 1999.

[3] Albert Einstein. Näherungsweise integration der feldgleichungen der gravitation. *Albert Einstein: Akademie-Vorträge: Sitzungsberichte der Preußischen Akademie der Wissenschaften 1914–1932*, pages 99–108, 2005.

[4] B.P. Abbott et al. Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Rev. Rel.*, 21(1):3, 2018.

[5] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.

[6] F Acernese, et al. Advanced virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, 2014.

[7] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.

[8] LIGO Scientific Collaboration and Virgo Collaboration, et al. Properties of the Binary Black Hole Merger GW150914. *Phys. Rev. Lett.*, 116(24):241102, June 2016.

[9] B. S. Sathyaprakash and Bernard F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Relativ.*, 12(1):2, December 2009.

[10] KL Dooley, et al. Geo 600 and the geo-hf upgrade program: successes and challenges. *Classical and Quantum Gravity*, 33(7):075009, 2016.

[11] Yoichi Aso, et al. Interferometer design of the kagra gravitational wave detector. *Physical Review D*, 88(4):043007, 2013.

[12] Benjamin P Abbott, et al. First targeted search for gravitational-wave bursts from core-collapse supernovae in data of first-generation laser interferometer detectors. *Physical Review D*, 94(10):102001, 2016.

[13] Benjamin P Abbott, et al. All-sky search for continuous gravitational waves from isolated neutron stars using advanced ligo o2 data. *Physical Review D*, 100(2):024004, 2019.

[14] Keith Riles. Recent searches for continuous gravitational waves. *Modern Physics Letters A*, 32(39):1730035, 2017.

[15] Pau Amaro-Seoane, et al. Laser interferometer space antenna. *arXiv preprint arXiv:1702.00786*, 2017.

[16] P. Bender et al. LISA Pre-Phase A Report Second Edition, 1998.

[17] George Hobbs, et al. The international pulsar timing array project: using pulsars as a gravitational wave detector. *Classical and Quantum Gravity*, 27(8):084013, 2010.

[18] Robert Crittenden, J Richard Bond, Richard L Davis, George Efstathiou, and Paul J Steinhardt. Imprint of gravitational waves on the cosmic microwave background. *Physical Review Letters*, 71(3):324, 1993.

[19] Albert A Michelson and Edward W Morley. On the relative motion of the earth and of the luminiferous ether. *Sidereal Messenger, vol. 6, pp. 306-310*, 6:306–310, 1887.

[20] Joseph Weber. Detection and generation of gravitational waves. *Physical Review*, 117(1):306, 1960.

[21] RWP Drever, FJ Raab, KS Thorne, R Vogt, and R Weiss. A laser interferometer gravitational-wave observatory (ligo), 1989.

[22] Michele Maggiore. *Gravitational waves: Volume 1: Theory and experiments*, volume 1. Oxford university press, 2008.

[23] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015.

[24] RWP Drever. The detection of gravitational waves, edited by dg blair, 1991.

[25] P. B. Covas et al. Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97(8):082002, 2018.

[26] A. Effler et al. Environmental Influences on the LIGO Gravitational Wave Detectors during the 6th Science Run. *Class. Quantum Grav.*, 32(3):035017, 2015.

[27] SM Aston, et al. Update on quadruple suspension design for advanced ligo. *Classical and Quantum Gravity*, 29(23):235004, 2012.

[28] Fabrice Matichard, et al. Seismic isolation of advanced ligo: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, 2015.

[29] A Effler, et al. Environmental influences on the ligo gravitational wave detectors during the 6th science run. *Classical and Quantum Gravity*, 32(3):035017, 2015.

[30] Christopher J Moore, Robert H Cole, and Christopher PL Berry. Gravitational-wave sensitivity curves. *Classical and Quantum Gravity*, 32(1):015014, 2014.

[31] Tyson B. Littenberg and Neil J. Cornish. Bayesian inference for spectral estimation of gravitational wave detector noise. *PhRvD*, 91(8):084034, 2015.

[32] B. P. Abbott et al. Binary Black Hole Mergers in the first Advanced LIGO Observing Run. *Phys. Rev.*, X6(4):041015, 2016. [erratum: Phys. Rev.X8,no.3,039903(2018)].

[33] B.P. Abbott, et al. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Physical Review Letters*, 116(24), June 2016.

[34] B. P. Abbott, et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters*, 119(16), October 2017.

[35] LIGO Scientific and Virgo Collaboration, et al. GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2. *Phys. Rev. Lett.*, 118(22):221101, June 2017.

[36] B. P. Abbott, et al. GW170608: Observation of a 19 Solar-mass Binary Black Hole Coalescence. *The Astrophysical Journal*, 851(2):L35, December 2017.

[37] Benjamin P Abbott, et al. Gw170814: a three-detector observation of gravitational waves from a binary black hole coalescence. *Physical review letters*, 119(14):141101, 2017.

[38] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev.*, X9(3):031040, 2019.

[39] Alexander H Nitz, et al. 2-ogc: Open gravitational-wave catalog of binary mergers from analysis of public advanced ligo and virgo data. *The Astrophysical Journal*, 891(2):123, 2020.

[40] Tejaswi Venumadhav, Barak Zackay, Javier Roulet, Liang Dai, and Matias Zaldarriaga. New binary black hole mergers in the second observing run of advanced ligo and advanced virgo. *Physical Review D*, 101(8):083030, 2020.

[41] Barak Zackay, Tejaswi Venumadhav, Liang Dai, Javier Roulet, and Matias Zaldarriaga. Highly spinning and aligned binary black hole merger in the advanced ligo first observing run. *Physical Review D*, 100(2):023007, 2019.

[42] Barak Zackay, Liang Dai, Tejaswi Venumadhav, Javier Roulet, and Matias Zaldarriaga. Detecting gravitational waves with disparate detector responses: two new binary black hole mergers. *arXiv preprint arXiv:1910.09528*, 2019.

[43] B.P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$. *Astrophys. J. Lett.*, 892(1):L3, 2020.

[44] R Abbott, et al. Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *arXiv preprint arXiv:2010.14527*, 2020.

[45] R. Abbott et al. GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses. *Phys. Rev. D*, 102(4):043015, 2020.

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

[47] R. Abbott et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. *Astrophys. J.*, 896(2):L44, 2020.

[48] Clifford E Rhoades Jr and Remo Ruffini. Maximum mass of a neutron star. *Physical Review Letters*, 32(6):324, 1974.

[49] Vassiliki Kalogera and Gordon Baym. The maximum mass of a neutron star. *The Astrophysical Journal Letters*, 470(1):L61, 1996.

[50] Tanja Hinderer. Tidal Love numbers of neutron stars. *Astrophys. J.*, 677:1216–1220, 2008.

[51] Katerina Chatziioannou. Neutron star tidal deformability and equation of state constraints. 6 2020.

[52] S. Chandrasekhar. The Mathematical Theory of Black Holes. *Fundam. Theor. Phys.*, 9:5–26, 1984.

[53] Kostas D. Kokkotas and Bernd G. Schmidt. Quasi-normal modes of stars and black holes. *Living Reviews in Relativity*, 2(1):2, Sep 1999.

[54] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev.*, D49:6274–6297, 1994.

[55] Bernd Brügmann, et al. Calibration of moving puncture simulations. *Phys. Rev. D*, 77:024027, Jan 2008.

[56] Sascha Husa, Jose A. Gonzalez, Mark Hannam, Bernd Bruegmann, and Ulrich Sperhake. Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing. *Class. Quant. Grav.*, 25:105006, 2008.

[57] E. Fauchon-Jones et al. BAM catalogue of binary black hole simulations. 2021. in preparation.

[58] Lawrence E. Kidder. Coalescing binary systems of compact objects to post-Newtonian 5/2 order. 5. Spin effects. *Phys. Rev.*, D52:821–847, 1995.

[59] Duncan A. Brown, Andrew Lundgren, and R. O'Shaughnessy. Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation. *Phys. Rev.*, D86:064020, 2012.

[60] Curt Cutler and Eanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral wave form? *Phys. Rev.*, D49:2658–2697, 1994.

[61] Eric Poisson and Clifford M. Will. Gravitational waves from inspiraling compact binaries: Parameter estimation using second postNewtonian wave forms. *Phys. Rev.*, D52:848–855, 1995.

[62] Emily Baird, Stephen Fairhurst, Mark Hannam, and Patricia Murphy. Degeneracy between mass and spin in black-hole-binary waveforms. *Phys. Rev.*, D87(2):024035, 2013.

[63] R. O'Shaughnessy, et al. Parameter estimation of gravitational waves from precessing black hole-neutron star inspirals with higher harmonics. *Phys. Rev.*, D89(10):102005, 2014.

[64] Patricia Schmidt, Frank Ohme, and Mark Hannam. Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter. *Phys. Rev.*, D91(2):024043, 2015.

[65] Davide Gerosa, et al. A generalized precession parameter $\chi_{\rm p}$ to interpret gravitational-wave data. 11 2020.

[66] Lucy M. Thomas, Patricia Schmidt, and Geraint Pratten. A new effective precession spin for modelling multi-modal gravitational waveforms in the strong-field regime. 12 2020.

[67] Ilya Mandel and Richard O'Shaughnessy. Compact binary coalescences in the band of ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(11):114007, 2010.

[68] Ilya Mandel and Alison Farmer. Merging stellar-mass binary black holes. *arXiv preprint arXiv:1806.05820*, 2018.

[69] Vassiliki Kalogera. Spin orbit misalignment in close binaries with two compact objects. *Astrophys. J.*, 541:319–328, 2000.

[70] Ilya Mandel and Richard O'Shaughnessy. Compact Binary Coalescences in the Band of Ground-based Gravitational-Wave Detectors. *Class. Quant. Grav.*, 27:114007, 2010.

[71] Davide Gerosa, et al. Spin orientations of merging black holes formed from the evolution of stellar binaries. *Phys. Rev. D*, 98(8):084036, 2018.

[72] Carl L Rodriguez, Michael Zevin, Chris Pankow, Vasilliki Kalogera, and Frederic A Rasio. Illuminating black hole binary formation channels with spins in advanced ligo. *The Astrophysical Journal Letters*, 832(1):L2, 2016.

[73] Mark Hannam, et al. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014.

[74] Patricia Schmidt, Mark Hannam, and Sascha Husa. Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals. *Phys. Rev.*, D86:104063, 2012.

[75] Sascha Husa, et al. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev.*, D93(4):044006, 2016.

[76] Sebastian Khan, et al. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev.*, D93(4):044007, 2016.

[77] Prayush Kumar, et al. Accuracy of binary black hole waveform models for aligned-spin binaries. *Phys. Rev.*, D93(10):104050, 2016.

[78] Andrea Taracchini et al. Effective-one-body model for black-hole binaries with generic mass ratios and spins. *Phys. Rev.*, D89(6):061502, 2014.

[79] Stanislav Babak, Andrea Taracchini, and Alessandra Buonanno. Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity. *Phys. Rev.*, D95(2):024010, 2017.

[80] Sebastian Khan, Katerina Chatziioannou, Mark Hannam, and Frank Ohme. Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects. *Phys. Rev.*, D100(2):024059, 2019.

[81] Sebastian Khan, Frank Ohme, Katerina Chatziioannou, and Mark Hannam. Including higher order multipoles in gravitational-wave models for precessing binary black holes. *"PhRvD"*, 101(2):024056, January 2020.

[82] R. O'Shaughnessy, L. London, J. Healy, and D. Shoemaker. Precession during merger: Strong polarization changes are observationally accessible features of strong-field gravity during binary black hole merger. *Phys. Rev.*, D87(4):044038, 2013.

[83] Davide Gerosa, Richard O'Shaughnessy, Michael Kesden, Emanuele Berti, and Ulrich Sperhake. Distinguishing black-hole spin-orbit resonances by their gravitational-wave signatures. *Phys. Rev.*, D89(12):124025, 2014.

[84] Davide Gerosa, Michael Kesden, Ulrich Sperhake, Emanuele Berti, and Richard O'Shaughnessy. Multi-timescale analysis of phase transitions in precessing black-hole binaries. *Phys. Rev.*, D92:064016, 2015.

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

[86] Stephen Fairhurst, Rhys Green, Charlie Hoy, Mark Hannam, and Alistair Muir. Two-harmonic approximation for gravitational waveforms from precessing binaries. *Phys. Rev. D*, 102:024055, 2020.

[87] T. A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev.*, D52:605–620, 1995.

[88] Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Phys. Rev.*, D67:104025, 2003. [Erratum: Phys. Rev.D74,029904(2006)].

[89] Marc van der Sluys, et al. Parameter estimation of spinning binary inspirals using Markov-chain Monte Carlo. *Class. Quant. Grav.*, 25:184011, 2008.

[90] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys. Rev.*, D81:062003, 2010.

[91] J. Veitch et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev.*, D91(4):042003, 2015.

[92] C. M. Biwer, et al. PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals. *Publ. Astron. Soc. Pac.*, 131(996):024503, 2019.

[93] Rory Smith, et al. Fast and accurate inference on gravitational waves from precessing compact binaries. *Phys. Rev.*, D94(4):044031, 2016.

[94] Ryan N. Lang and Scott A. Hughes. Measuring coalescing massive binary black holes with gravitational waves: The Impact of spin-induced precession. *Phys. Rev.*, D74:122001, 2006. [Erratum: Phys. Rev.D77,109901(2008)].

[95] Stephen Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New J. Phys.*, 11:123006, 2009. [Erratum: New J. Phys.13,069602(2011)].

[96] Mark Hannam, Duncan A. Brown, Stephen Fairhurst, Chris L. Fryer, and Ian W. Harry. When can gravitational-wave observations distinguish between black holes and neutron stars? *Astrophys. J.*, 766:L14, 2013.

[97] Leo P. Singer and Larry R. Price. Rapid Bayesian position reconstruction for gravitational-wave transients. *Phys. Rev.*, D93(2):024013, 2016.

[98] Michele Vallisneri. Testing general relativity with gravitational waves: a reality check. *Phys. Rev.*, D86:082001, 2012.

[99] Samantha A. Usman, Joseph C. Mills, and Stephen Fairhurst. Constraining the Inclinations of Binary Mergers from Gravitational-wave Observations. *Astrophys. J.*, 877(2):82, 2019.

[100] Vijay Varma, et al. Surrogate models for precessing binary black hole simulations with unequal masses. *Phys. Rev. Research.*, 1:033015, 2019.

[101] Geraint Pratten et al. Let's twist again: computationally efficient models for the dominant and sub-dominant harmonic modes of precessing binary black holes. 4 2020.

[102] Salvatore Vitale, Ryan Lynch, John Veitch, Vivien Raymond, and Riccardo Sturani. Measuring the spin of black holes in binary systems using gravitational waves. *Phys. Rev. Lett.*, 112(25):251101, 2014.

[103] Daniele Trifir, et al. Distinguishing black-hole spin-orbit resonances by their gravitational wave signatures. II: Full parameter estimation. *Phys. Rev.*, D93(4):044071, 2016.

[104] Salvatore Vitale, et al. Parameter estimation for heavy binary-black holes with networks of second-generation gravitational-wave detectors. *Phys. Rev. D*, 95(6):064053, 2017.

[105] Benjamin Farr, Evan Ochsner, Will M. Farr, and Richard O'Shaughnessy. A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves. *Phys. Rev.*, D90(2):024018, 2014.

[106] Tyson B. Littenberg, Ben Farr, Scott Coughlin, and Vicky Kalogera. Systematic errors in low latency gravitational wave parameter estimation impact electromagnetic follow-up observations. *Astrophys. J.*, 820(1):7, 2016.

[107] Ben Farr et al. Parameter estimation on gravitational waves from neutron-star binaries with spinning components. *Astrophys. J.*, 825(2):116, 2016.

[108] Salvatore Vitale and Matthew Evans. Parameter estimation for binary black holes with networks of third generation gravitational-wave detectors. *Phys. Rev.*, D95(6):064052, 2017.

[109] R. O'Shaughnessy, P. Nepal, and A. Lundgren. A semianalytic Fisher matrix for precessing binaries with a single significant spin. *Class. Quant. Grav.*, 37(11):115006, 2020.

[110] A. Lundgren and R. O'Shaughnessy. Single-spin precessing gravitational waveform in closed form. *Phys. Rev.*, D89(4):044021, 2014.

[111] Ian W. Harry, et al. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. *Phys. Rev.*, D89(2):024010, 2014.

[112] Diego Fazi. *Development of a Physical-Template Search for Gravitational Waves from Spinning Compact-Object Binaries with LIGO.* PhD thesis, Bologna U., 2009.

[113] I. W. Harry and S. Fairhurst. A coherent triggered search for single spin compact binary coalescences in gravitational wave data. *Class. Quant. Grav.*, 28:134008, 2011.

[114] Ian Harry, Stephen Privitera, Alejandro Boh, and Alessandra Buonanno. Searching for Gravitational Waves from Compact Binaries with Precessing Spins. *Phys. Rev.*, D94(2):024012, 2016.

[115] Michael Boyle, Robert Owen, and Harald P. Pfeiffer. A geometric approach to the precession of compact binaries. *Phys. Rev.*, D84:124011, 2011.

[116] Ian W. Harry and Stephen Fairhurst. A targeted coherent search for gravitational waves from compact binary coalescences. *Phys. Rev.*, D83:084002, 2011.

[117] Serge Droz, Daniel J. Knapp, Eric Poisson, and Benjamin J. Owen. Gravitational waves from inspiraling compact binaries: Validity of the stationary phase approximation to the Fourier transform. *Phys. Rev.*, D59:124016, 1999.

[118] Will M. Farr, et al. Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves. *Nature*, 548:426, 2017.

[119] Vaibhav Tiwari, Stephen Fairhurst, and Mark Hannam. Constraining black-hole spins with gravitational wave observations. *Astrophys. J.*, 868(2):140, 2018.

[120] Stephen Fairhurst, Rhys Green, Mark Hannam, and Charlie Hoy. When will we observe binary black holes precessing? *Phys. Rev. D*, 102(4):041302, 2020.

[121] Fabio Antonini, Carl L. Rodriguez, Cristobal Petrovich, and Caitlin L. Fischer. Precessional dynamics of black hole triples: binary mergers with near-zero effective spin. *Mon. Not. Roy. Astron. Soc.*, 480(1):L58–L62, 2018.

[122] Carl L. Rodriguez and Fabio Antonini. A Triple Origin for the Heavy and Low-Spin Binary Black Holes Detected by LIGO/Virgo. *Astrophys. J.*, 863(1):7, 2018.

[123] Eanna E. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences: 2. The Waves' information and its extraction, with and without templates. *Phys. Rev.*, D57:4566–4587, 1998.

[124] Mark A. Miller. Accuracy requirements for the calculation of gravitational waveforms from coalescing compact binaries in numerical relativity. *Phys. Rev.*, D71:104016, 2005.

[125] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A.J. Weinstein. Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Phys. Rev. D*, 89(8):084041, 2014.

[126] Cody Messick et al. Analysis Framework for the Prompt Discovery of Compact Binary Mergers in Gravitational-wave Data. *Phys. Rev.*, D95(4):042001, 2017.

[127] Samantha A. Usman et al. The PyCBC search for gravitational waves from compact binary coalescence. *Class. Quant. Grav.*, 33(21):215004, 2016.

[128] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev.*, D53:6749–6761, 1996.

[129] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev.*, D60:022002, 1999.

[130] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash. A Template bank to search for gravitational waves from inspiralling compact binaries. I. Physical models. *Class. Quant. Grav.*, 23:5477–5504, 2006.

[131] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev.*, D85:122006, 2012.

[132] Bruce Allen. χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev.*, D71:062001, 2005.

[133] S. Babak et al. Searching for gravitational waves from binary coalescence. *Phys. Rev.*, D87(2):024033, 2013.

[134] Tito Dal Canton and Ian W. Harry. Designing a template bank to observe compact binary coalescences in Advanced LIGO's second observing run. *arXiv e-prints*, page arXiv:1705.01845, May 2017.

[135] Debnandini Mukherjee, et al. The GstLAL template bank for spinning compact binary mergers in the second observation run of Advanced LIGO and Virgo. *arXiv e-prints*, page arXiv:1812.05121, December 2018.

[136] Yi Pan, Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. A Physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single spin binaries. *Phys. Rev.*, D69:104017, 2004. [Erratum: Phys. Rev.D74,029905(2006)].

[137] Alexander H. Nitz, Thomas Dent, Tito Dal Canton, Stephen Fairhurst, and Duncan A. Brown. Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search. *Astrophys. J.*, 849(2):118, 2017.

[138] Colm Talbot and Eric Thrane. Determining the population properties of spinning black holes. *Phys. Rev.*, D96(2):023012, 2017.

[139] Colm Talbot, Rory Smith, Eric Thrane, and Gregory B. Poole. Parallelized Inference for Gravitational-Wave Astronomy. *Phys. Rev. D*, 100(4):043030, 2019.

[140] Daniel Wysocki, Jacob Lange, and Richard O'Shaughnessy. Reconstructing phenomenological distributions of compact binaries via gravitational wave observations. *Phys. Rev. D*, 100(4):043012, 2019.

[141] Gregory Ashton, et al. BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy. *Astrophysical Journal, Supplement*, 241(2):27, Apr 2019.

[142] Ben Farr, Daniel E. Holz, and Will M. Farr. Using Spin to Understand the Formation of LIGO and Virgos Black Holes. *Astrophys. J.*, 854(1):L9, 2018.

[143] Maya Fishbach and Daniel E. Holz. Where Are LIGO's Big Black Holes? *Astrophys. J. Lett.*, 851(2):L25, 2017.

[144] Maya Fishbach, Daniel E. Holz, and Will M. Farr. Does the Black Hole Merger Rate Evolve with Redshift? *Astrophys. J. Lett.*, 863(2):L41, 2018.

[145] Rhys Green, et al. Identifying when Precession can be Measured in Gravitational Waveforms. 10 2020.

[146] R. Abbott et al. Properties and astrophysical implications of the 150 Msun binary black hole merger GW190521. *Astrophys. J. Lett.*, 900:L13, 2020.

[147] Bernard F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323(6086):310–311, September 1986.

[148] M. Soares-Santos et al. First measurement of the Hubble constant from a dark standard siren using the Dark Energy Survey galaxies and the LIGO/Virgo binary-black-hole merger GW170814. *Submitted to: Astrophys. J.*, 2019.

[149] B. P. Abbott, et al. Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *The Astrophysical Journal*, 848(2):L13, October 2017.

[150] Benjamin P Abbott, et al. Gw170817: Measurements of neutron star radii and equation of state. *Physical review letters*, 121(16):161101, 2018.

[151] Benjamin P Abbott, et al. Gw170817: implications for the stochastic gravitational-wave background from compact binary coalescences. *Physical review letters*, 120(9):091101, 2018.

[152] LIGO Scientific Collaboration, et al. A gravitational-wave standard siren measurement of the hubble constant. *Nature*, 551(7678):85–88, 2017.

[153] BP Abbott, et al. Tests of general relativity with the binary black hole signals from the ligo-virgo catalog gwtc-1. *Physical Review D*, 100(10):104036, 2019.

[154] B. S. Sathyaprakash et al. Cosmology and the Early Universe. 2019.

[155] Eugenio Bianchi, Anuradha Gupta, Hal M. Haggard, and B. S. Sathyaprakash. Quantum gravity and black hole spin in gravitational wave observations: a test of the Bekenstein-Hawking entropy. 2018.

[156] K. E. Saavik Ford, et al. Multi-Messenger Astrophysics Opportunities with Stellar-Mass Binary Black Hole Mergers. 2019.

[157] Ben Farr, et al. Parameter estimation on gravitational waves from neutron-star binaries with spinning components. *The Astrophysical Journal*, 825(2):116, 2016.

[158] Ken KY Ng, et al. Gravitational-wave astrophysics with effective-spin measurements: Asymmetries and selection biases. *Physical Review D*, 98(8):083007, 2018.

[159] Philip B. Graff, Alessandra Buonanno, and B. S. Sathyaprakash. Missing Link: Bayesian detection and measurement of intermediate-mass black-hole binaries. *Phys. Rev. D*, 92(2):022002, 2015.

[160] Carl-Johan Haster, et al. Inference on gravitational waves from coalescences of stellar-mass compact objects and intermediate-mass black holes. *Mon. Not. Roy. Astron. Soc.*, 457(4):4499–4506, 2016.

[161] Hang Yu et al. Prospects for detecting gravitational waves at 5 Hz with ground-based detectors. *Phys. Rev. Lett.*, 120(14):141102, 2018.

[162] Yi Pan, et al. Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism. *Phys. Rev. D*, 89(8):084006, 2014.

[163] R Abbott, et al. Open data from the first and second observing runs of advanced ligo and advanced virgo. *arXiv preprint arXiv:1912.11716*, 2019.

[164] Alberto Vecchio. Lisa observations of rapidly spinning massive black hole binary systems. *Physical Review D*, 70(4):042001, 2004.

[165] Katerina Chatziioannou, et al. Measuring the properties of nearly extremal black holes with gravitational waves. *Phys. Rev. D*, 98(4):044028, 2018.

[166] Salvatore Vitale, Ryan Lynch, John Veitch, Vivien Raymond, and Riccardo Sturani. Measuring the spin of black holes in binary systems using gravitational waves. *Physical Review Letters*, 112(25):251101, 2014.

[167] Benjamin P Abbott, et al. Effects of waveform model systematics on the interpretation of gw150914. *Classical and Quantum Gravity*, 34(10):104002, 2017.

[168] Emanuele Berti, Alessandra Buonanno, and Clifford M Will. Estimating spinning binary parameters and testing alternative theories of gravity with lisa. *Physical Review D*, 71(8):084025, 2005.

[169] Geraint Pratten, Patricia Schmidt, Riccardo Buscicchio, and Lucy M. Thomas. On measuring precession in GW190814-like asymmetric compact binaries. 6 2020.

[170] R. Abbott et al. Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog. 10 2020.

[171] Alessandra Buonanno, Yan-bei Chen, Yi Pan, and Michele Vallisneri. A Quasi-physical family of gravity-wave templates for precessing binaries of spinning compact objects. 2. Application to double-spin precessing binaries. *Phys. Rev. D*, 70:104003, 2004. [Erratum: Phys.Rev.D 74, 029902 (2006)].

[172] Bohé, Alejandro and Hannam, Mark and Husa, Sascha and Ohme, Frank and Puerrer, Michael and Schmidt, Patricia. Phenompv2 - technical notes for lal implementation. Technical Report LIGO-T1500602, LIGO Project, 2016.

[173] A Buikema, et al. Sensitivity and performance of the advanced ligo detectors in the third observing run. *Physical Review D*, 102(6):062003, 2020.

[174] Lee S Finn. Detection, measurement, and gravitational radiation. *Physical Review D*, 46(12):5236, 1992.

[175] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. free software (GPL), 2018.

[176] Chinmay Kalaghatgi, Mark Hannam, and Vivien Raymond. Parameter estimation with a spinning multimode waveform model. *Phys. Rev. D*, 101(10):103004, 2020.

[177] Salvatore Vitale and Hsin-Yu Chen. Measuring the hubble constant with neutron star black hole mergers. *Physical review letters*, 121(2):021303, 2018.

[178] Michele Vallisneri. Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Phys. Rev. D*, 77:042001, 2008.

[179] Manuela Campanelli, C.O. Lousto, and Y. Zlochower. Spinning-black-hole binaries: The orbital hang up. *Phys. Rev. D*, 74:041501, 2006.

[180] Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge university press, 2003.

[181] Neil Cornish, Laura Sampson, Nicolas Yunes, and Frans Pretorius. Gravitational wave tests of general relativity with the parameterized post-einsteinian framework. *Physical Review D*, 84(6):062003, 2011.

[182] Cameron Mills and Stephen Fairhurst. Measuring gravitational-wave higher-order multipoles. *Phys. Rev. D*, 103(2):024042, 2021.

[183] Chinmay Kalaghatgi and Mark Hannam. Investigating the effect of in-plane spin directions for Precessing BBH systems. 8 2020.

[184] Michele Vallisneri, Jonah Kanner, Roy Williams, Alan Weinstein, and Branson Stephens. The LIGO Open Science Center. *J. Phys. Conf. Ser.*, 610(1):012021, 2015.

[185] Meredith L Rawls, et al. Refined neutron star mass determinations for six eclipsing x-ray pulsar binaries. *The Astrophysical Journal*, 730(1):25, 2011.

[186] H.T. Cromartie et al. Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar. *Nature Astron.*, 4(1):72–76, 2019.

[187] Emmanuel Fonseca et al. Refined Mass and Geometric Measurements of the High-Mass PSR J0740+6620. 4 2021.

[188] Charles D Bailyn, Raj K Jain, Paolo Coppi, and Jerome A Orosz. The mass distribution of stellar black holes. *The Astrophysical Journal*, 499(1):367, 1998.

[189] Feryal Özel, Dimitrios Psaltis, Ramesh Narayan, and Jeffrey E McClintock. The black hole mass distribution in the galaxy. *The Astrophysical Journal*, 725(2):1918, 2010.

[190] Feryal Özel, Dimitrios Psaltis, Ramesh Narayan, and Antonio Santos Villarreal. On the mass distribution and birth masses of neutron stars. *The Astrophysical Journal*, 757(1):55, 2012.

[191] Will M Farr, et al. The mass distribution of stellar-mass black holes. *The Astrophysical Journal*, 741(2):103, 2011.

[192] Jea Abadie, et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(17):173001, 2010.

[193] Masaru Shibata and Keisuke Taniguchi. Coalescence of black hole-neutron star binaries. *Living Reviews in Relativity*, 14(1):6, 2011.

[194] Francesco Pannarale, Luciano Rezzolla, Frank Ohme, and Jocelyn S Read. Will black hole-neutron star binary inspirals tell us about the neutron star equation of state? *Physical Review D*, 84(10):104017, 2011.

[195] Krzysztof Belczynski, Ronald E Taam, Emmanouela Rantsiou, and Marc Van Der Sluys. Black hole spin evolution: Implications for short-hard gamma-ray bursts and gravitational wave detection. *The Astrophysical Journal*, 682(1):474, 2008.

[196] Michela Mapelli and Nicola Giacobbo. The cosmic merger rate of neutron stars and black holes. *Mon. Not. Roy. Astron. Soc.*, 479(4):4391–4398, 2018.

[197] Ajit Kumar Mehta, Chandra Kant Mishra, Vijay Varma, and Parameswaran Ajith. Accurate inspiral-merger-ringdown gravitational waveforms for non-spinning black-hole binaries including the effect of subdominant modes. *Phys. Rev. D*, 96(12):124010, 2017.

[198] BP Abbott, et al. Model comparison from ligo–virgo data on gw170817s binary components and consequences for the merger remnant. *Classical and Quantum Gravity*, 37(4):045006, 2020.

[199] Michael W Coughlin and Tim Dietrich. Can a black hole–neutron star merger explain gw170817, at2017gfo, and grb170817a? *Physical Review D*, 100(4):043011, 2019.

[200] Tanja Hinderer, et al. Distinguishing the nature of comparable-mass neutron star binary systems with multimessenger observations: Gw170817 case study. *Physical Review D*, 100(6):063021, 2019.

[201] Vaibhav Tiwari. Estimation of the Sensitive Volume for Gravitational-wave Source Populations Using Weighted Monte Carlo Integration. *Class. Quant. Grav.*, 35(14):145009, 2018.

[202] Alex Nitz, et al. gwastro/pycbc: Pycbc release v1.15.2, December 2019.

[203] Jonathan E. Thompson, et al. Modeling the gravitational wave signature of neutron star black hole coalescences: PhenomNSBH. page arXiv:2002.08383, 2020.

[204] Andrew Matas et al. An aligned-spin neutron-star–black-hole waveform model based on the effective-one-body approach and numerical-relativity simulations. page arXiv:2004.10001, 2020.

[205] Benjamin P Abbott, et al. Upper limits on the rates of binary neutron star and neutron star–black hole mergers from advanced ligos first observing run. *The Astrophysical Journal Letters*, 832(2):L21, 2016.

[206] Will M. Farr, Jonathan R. Gair, Ilya Mandel, and Curt Cutler. Counting And Confusion: Bayesian Rate Estimation With Multiple Populations. *Phys. Rev. D*, 91(2):023005, 2015.

[207] K. Cannon, R. Cariou, A. Chapman, et al. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. *The Astrophysical Journal*, 748:136, 2012.

[208] S. Privitera, S. R. P. Mohapatra, P. Ajith, et al. Improving the sensitivity of a search for coalescing binary black holes with nonprecessing spins in gravitational wave data. *Physical Review D*, 89(2):024003, 2014.

[209] Chad Hanna, et al. Fast evaluation of multidetector consistency for real-time gravitational wave searches. *Physical Review D*, 101(2):022003, January 2020.

[210] S. Sachdev, et al. The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs. *arXiv:1901.08580*, 2019.

[211] T. Adams, et al. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *CQGra*, 33(17):175012, Sep 2016.

[212] Alexander H. Nitz, Tito Dal Canton, Derek Davis, and Steven Reyes. Rapid detection of gravitational waves from compact binary mergers with PyCBC Live. *Physical Review D*, 98(2):024050, 2018.

[213] M. Kasprzack and H. Yu. Beam position from angle to length minimization. Technical Report LIGO-T1600397, LIGO Project, 2017. `https://dcc.ligo.org/LIGO-T1600397/public/main`.

[214] LIGO Scientific Collaboration, Virgo Collaboration. *The Gamma-ray Coordinates Network*, 25324, 2019.

[215] Rory Smith and Gregory Ashton. Expediting Astrophysical Discovery with Gravitational-Wave Transients Through Massively Parallel Nested Sampling. page arXiv:1909.11873, September 2019.

[216] L. K. Nuttall. Characterizing transient noise in the LIGO detectors. *RSPTA*, 376(2120):20170286, May 2018.

[217] David J. Kaup. Klein-Gordon Geon. *PhRv*, 172(5):1331–1342, August 1968.

[218] Pawel O. Mazur and Emil Mottola. Gravitational vacuum condensate stars. *PNAS*, 101:9545–9550, 2004.

[219] Alejandro Bohé, et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *PhRvD*, 95(4):044028, 2017.

[220] F Foucart, et al. Gravitational waveforms from spectral einstein code simulations: Neutron star-neutron star and low-mass black hole-neutron star binaries. *Physical Review D*, 99(4):044008, 2019.

[221] Francois Foucart, et al. First direct comparison of nondisrupting neutron star-black hole and binary black hole merger simulations. *Physical Review D*, 88(6):064017, 2013.

[222] Lionel London, et al. First higher-multipole model of gravitational waves from spinning and coalescing black-hole binaries. *Phys. Rev. Lett.*, 120(16):161102, 2018.

[223] Roberto Cotesta, et al. Enriching the Symphony of Gravitational Waves from Binary Black Holes by Tuning Higher Harmonics. *Phys. Rev.*, D98(8):084028, 2018.

[224] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. *Phys. Rev. D*, 102(4):044055, 2020.

[225] Francois Foucart. Black-hole–neutron-star mergers: Disk mass predictions. *Physical Review D*, 86(12):124007, 2012.

[226] Francesco Pannarale, Emanuele Berti, Koutarou Kyutoku, Benjamin D. Lackey, and Masaru Shibata. Aligned spin neutron star-black hole mergers: a gravitational waveform amplitude model. *Phys. Rev. D*, 92(8):084050, 2015.

[227] Zachariah B Etienne, Yuk Tung Liu, Stuart L Shapiro, and Thomas W Baumgarte. General relativistic simulations of black-hole–neutron-star mergers: effects of black-hole spin. *Physical Review D*, 79(4):044024, 2009.

[228] Francois Foucart, et al. Black-hole–neutron-star mergers at realistic mass ratios: Equation of state and spin orientation effects. *Physical Review D*, 87(8):084006, 2013.

[229] Francesco Zappa, Sebastiano Bernuzzi, Francesco Pannarale, Michela Mapelli, and Nicola Giacobbo. Black-hole remnants from black-hole–neutron-star mergers. *Physical review letters*, 123(4):041102, 2019.

[230] Francesco Pannarale. Black hole remnant of black hole-neutron star coalescing binaries. *Physical Review D*, 88(10):104025, 2013.

[231] Francesco Pannarale. Black hole remnant of black hole-neutron star coalescing binaries with arbitrary black hole spin. *Physical Review D*, 89(4):044045, 2014.

[232] Kent Yagi and Nicolás Yunes. Approximate universal relations for neutron stars and quark stars. *Physics Reports*, 681:1–72, 2017.

[233] Thibault Damour and Alessandro Nagar. Relativistic tidal properties of neutron stars. *Phys. Rev. D*, 80:084035, 2009.

[234] Taylor Binnington and Eric Poisson. Relativistic theory of tidal Love numbers. *Phys. Rev. D*, 80:084018, 2009.

[235] Philippe Landry and Eric Poisson. Tidal deformation of a slowly rotating material body. External metric. *Phys. Rev. D*, 91:104018, 2015.

[236] Francois Foucart, et al. First direct comparison of nondisrupting neutron star-black hole and binary black hole merger simulations. *Phys. Rev. D*, 88(6):064017, 2013.

[237] Prayush Kumar, Michael Prrer, and Harald P. Pfeiffer. Measuring neutron star tidal deformability with Advanced LIGO: a Bayesian analysis of neutron star - black hole binary observations. *PhRvD*, 95(4):044039, 2017.

[238] Francois Foucart, Matthew D. Duez, Lawrence E. Kidder, and Saul A. Teukolsky. Black hole-neutron star mergers: effects of the orientation of the black hole spin. *Phys. Rev. D*, 83:024005, 2011.

[239] Francois Foucart, et al. Black hole-neutron star mergers at realistic mass ratios: Equation of state and spin orientation effects. *Phys. Rev. D*, 87:084006, 2013.

[240] Eanna E. Flanagan and Tanja Hinderer. Constraining neutron star tidal Love numbers with gravitational wave detectors. *Phys. Rev. D*, 77:021502, 2008.

[241] Will M. Farr and Katerina Chatziioannou. A Population-Informed Mass Estimate for Pulsar J0740+6620. *Research Notes of the American Astronomical Society*, 4(5):65, May 2020.

[242] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW170817: Measurements of Neutron Star Radii and Equation of State. *Physical Review Letters*, 121(16):161101, Oct 2018.

[243] BP Abbott, et al. Properties of the binary neutron star merger gw170817. *Physical Review X*, 9(1):011001, 2019.

[244] Justin Alsing, Hector O. Silva, and Emanuele Berti. Evidence for a maximum mass cut-off in the neutron star mass distribution and constraints on the equation of state. *Monthly Notices of the Royal Astronomical Society*, 478(1):1377–1391, Jul 2018.

[245] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Prospects for observing and localizing gravitational-wave transients with advanced ligo and advanced virgo. *arXiv:1304.0670v10*, 2020.

[246] Thibault Damour. Coalescence of two spinning black holes: an effective one-body approach. *Phys. Rev.*, D64:124013, 2001.

[247] Etienne Racine. Analysis of spin precession in binary black hole systems including quadrupole-monopole interaction. *Physical Review D*, 78:044021, 2008.

[248] P. Ajith, et al. Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins. *Physical Review Letters*, 106(24), June 2011.

[249] L. Santamaría, et al. Matching post-Newtonian and numerical relativity waveforms: Systematic errors and a new phenomenological model for nonprecessing black hole binaries. *Physical Review D*, 82(6), September 2010.

[250] Jacob Lange, Richard O'Shaughnessy, and Monica Rizzo. Rapid and accurate parameter inference for coalescing, precessing compact binaries. *arXiv:1805.10457*, 2018.

[251] Katerina Chatziioannou, Neil Cornish, Antoine Klein, and Nicolas Yunes. Spin-Precession: Breaking the Black Hole–Neutron Star Degeneracy. *Astrophys. J. Lett.*, 798(1):L17, 2015.

[252] Tanner Prestegard, Alexander Pace, Brian Moe, Branson Stephens, and Patrick Brady. `https://gracedb.ligo.org`, 2019.

[253] Larne Pekowsky, James Healy, Deirdre Shoemaker, and Pablo Laguna. Impact of higher-order modes on the detection of binary black hole coalescences. *Phys. Rev. D*, 87(8):084008, 2013.

[254] Vijay Varma, et al. Gravitational-wave observations of binary black holes: Effect of nonquadrupole modes. *Phys. Rev. D*, 90(12):124004, 2014.

[255] Juan Calderón Bustillo, Sascha Husa, Alicia M. Sintes, and Michael Pürrer. Impact of gravitational radiation higher order modes on single aligned-spin gravitational wave searches for binary black holes. *Phys. Rev. D*, 93(8):084019, 2016.

[256] Collin Capano, Yi Pan, and Alessandra Buonanno. Impact of higher harmonics in searching for gravitational waves from nonspinning binary black holes. *Phys. Rev. D*, 89(10):102003, 2014.

[257] J. Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151, Jan 1991.

[258] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[259] Andrey Kolmogorov. Sulla determinazione empirica di una lgge di distribuzione. *Inst. Ital. Attuari, Giorn.*, 4:83–91, 1933.

[260] Nickolay Smirnov. Table for estimating the goodness of fit of empirical distributions. *The annals of mathematical statistics*, 19(2):279–281, 1948.

[261] Marta Colleoni, et al. Towards the routine use of subdominant harmonics in gravitational-wave inference: Reanalysis of GW190412 with generation X waveform models. *Phys. Rev. D*, 103(2):024029, 2021.

[262] Isobel Romero-Shaw, Paul D Lasky, Eric Thrane, and Juan Calderón Bustillo. Gw190521: orbital eccentricity and signatures of dynamical formation in a binary black hole merger signal. *The Astrophysical Journal Letters*, 903(1):L5, 2020.

[263] Juan Calderón Bustillo, Nicolas Sanchis-Gual, Alejandro Torres-Forné, and José A. Font. Confusing head-on and precessing intermediate-mass binary black hole mergers. 9 2020.

[264] Ben Farr, Daniel E. Holz, and Will M. Farr. Using Spin to Understand the Formation of LIGO and Virgo's Black Holes. *ApJL*, 854:L9, 2018.

[265] BP Abbott, et al. 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):L24, 2019.

[266] Sylvia Biscoveanu, Maximiliano Isi, Salvatore Vitale, and Vijay Varma. A new spin on ligo-virgo binary black holes. *arXiv preprint arXiv:2007.09156*, 2020.

[267] Eric Thrane and Colm Talbot. An introduction to Bayesian inference in gravitational-wave astronomy: parameter estimation, model selection, and hierarchical models. *Publ. Astron. Soc. Austral.*, 36:e010, 2019. [Erratum: Publ.Astron.Soc.Austral. 37, e036 (2020)].

[268] Ilya Mandel, Will M Farr, and Jonathan R Gair. Extracting distribution parameters from multiple uncertain observations with selection biases. *MNRAS*, 486:1086, 2019.

[269] Salvatore Vitale. One, No One, and One Hundred Thousand – Inferring the properties of a population in presence of selection effects. 7 2020.

[270] Michael Pürrer, Mark Hannam, and Frank Ohme. Can we measure individual black-hole spins from gravitational-wave observations? *Physical Review D*, 93(8):084042, 2016.

[271] Davide Gerosa, Geraint Pratten, and Alberto Vecchio. Gravitational-wave selection effects using neural-network classifiers. *Phys. Rev. D*, 102(10):103020, 2020.

[272] Juan Caldern Bustillo, Pablo Laguna, and Deirdre Shoemaker. Detectability of gravitational waves from binary black holes: Impact of precession and higher modes. *Phys. Rev.*, D95(10):104038, 2017.

[273] Alexander H. Nitz and Collin D. Capano. GW190521 may be an intermediate mass ratio inspiral. *Astrophys. J. Lett.*, 907(1):L9, 2021.

[274] S. E. Woosley. Pulsational Pair-Instability Supernovae. *Astrophys. J.*, 836(2):244, 2017.

[275] K. Belczynski et al. The Effect of Pair-Instability Mass Loss on Black Hole Mergers. *Astron. Astrophys.*, 594:A97, 2016.

[276] Davide Gerosa and Emanuele Berti. Are merging black holes born from stellar collapse or previous mergers? *Phys. Rev. D*, 95(12):124046, 2017.

[277] Vaibhav Tiwari and Stephen Fairhurst. The emergence of structure in the binary black hole mass distribution: Potential evidence for hierarchical mergers. 2020.

[278] Chase Kimball et al. Evidence for hierarchical black hole mergers in the second LIGO–Virgo gravitational-wave catalog. 11 2020.

[279] Davide Gerosa, Salvatore Vitale, and Emanuele Berti. Astrophysical implications of GW190412 as a remnant of a previous black-hole merger. *Phys. Rev. Lett.*, 125(10):101103, 2020.

[280] Fabio Antonini and Frederic A. Rasio. Merging black hole binaries in galactic nuclei: implications for advanced-LIGO detections. *Astrophys. J.*, 831(2):187, 2016.

[281] Alessandra Buonanno, Lawrence E. Kidder, and Luis Lehner. Estimating the final spin of a binary black hole coalescence. *Phys. Rev. D*, 77:026004, 2008.

[282] Charlie Hoy and Vivien Raymond. PESummary: the code agnostic Parameter Estimation Summary page builder. 6 2020.

[283] Simon Jackman. *Bayesian analysis for the social sciences*, volume 846. John Wiley & Sons, 2009.

[284] Aaron M Ellison. Bayesian inference in ecology. *Ecology letters*, 7(6):509–520, 2004.

[285] Mark A Beaumont and Bruce Rannala. The bayesian revolution in genetics. *Nature Reviews Genetics*, 5(4):251–261, 2004.

[286] Roberto Trotta. Bayes in the sky: Bayesian inference and model selection in cosmology. *Contemporary Physics*, 49(2):71–104, 2008.

[287] Andrew Gelman, et al. *Bayesian data analysis*. CRC press, 2013.

[288] Hanns Ludwig Harney. *Bayesian Inference*. Springer, 2016.

[289] Albert Einstein. Approximative Integration of the Field Equations of Gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1916:688–696, 1916.

[290] Albert Einstein. Über Gravitationswellen. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1918:154–167, 1918.

[291] Neil J Cornish and Tyson B Littenberg. Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches. *Classical and Quantum Gravity*, 32(13):135012, 2015.

[292] Benjamin P Abbott, et al. Binary black hole mergers in the first advanced ligo observing run. *Physical Review X*, 6(4):041015, 2016.

[293] LIGO Scientific Collaboration. `https://www.ligo.org`, 2021.

[294] Virgo Collaboration. `https://www.virgo-gw.eu`, 2021.

[295] KAGRA Collaboration. `https://gwcenter.icrr.u-tokyo.ac.jp`, 2021.

[296] John Veitch, et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library. *Physical Review D*, 91(4):042003, 2015.

[297] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy. Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences. *Phys. Rev. D*, 92(2):023002, 2015.

[298] J. Lange et al. Parameter estimation method that directly compares gravitational wave observations to numerical relativity. *Phys. Rev. D*, 96(10):104041, 2017.

[299] D. Wysocki, R. O'Shaughnessy, Jacob Lange, and Yao-Lung L. Fang. Accelerating parameter inference with graphics processing units. *Phys. Rev. D*, 99(8):084026, 2019.

[300] Gregory Ashton, et al. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, 2019.

[301] I. M. Romero-Shaw et al. Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue. *Mon. Not. Roy. Astron. Soc.*, 499(3):3295–3319, 2020.

[302] Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306, March 2013.

[303] Joshua S Speagle. dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, 493(3):3132–3158, 2020.

[304] B.P. Abbott et al. `https://dcc.ligo.org/public/0163/P190412/008/posterior_samples.h5`, 2019.

[305] B.P. Abbott et al. `https://dcc.ligo.org/public/0165/P2000026/001/GW190425_posterior_samples.h5`, 2019.

[306] Python Language Reference The Python Software Foundation. `http://www.python.org`.

[307] Oliphant Travis E. a guide to numpy.

[308] Andrew Collette. *Python and HDF5*. O'Reilly, 2013.

[309] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

[310] Daniel Foreman-Mackey. corner.py: Scatterplot matrices in python. *The Journal of Open Source Software*, 1(2):24, jun 2016.

[311] Wes McKinney. Data structures for statistical computing in python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 51 – 56, 2010.

[312] Astropy Collaboration, et al. Astropy: A community Python package for astronomy. *Astronomy and Astrophysics*, 558:A33, October 2013.

[313] A. M. Price-Whelan, et al. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *The Astronomical Journal*, 156:123, September 2018.

[314] Casper da Costa-Luis. tqdm: A fast, extensible progress meter for python and cli. *Journal of Open Source Software*, 4(37):1277, 2019.

[315] Plotly Technologies Inc. Collaborative data science, 2015.

[316] D M Macleod et al. GWpy: a python package for gravitational-wave astrophysics. In preparation, 2020.

[317] Bernard W Silverman. *Density estimation for statistics and data analysis*, volume 26. CRC press, 1986.

[318] James Joyce. Bayes' theorem. 2003.

[319] Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. *Journal of the American statistical association*, 44(247):335–341, 1949.

[320] John Skilling et al. Nested sampling for general bayesian computation. *Bayesian analysis*, 1(4):833–859, 2006.

[321] Stéfan van der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.

[322] Stephen P Brooks and Andrew Gelman. General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, 7(4):434–455, 1998.

[323] A. N. Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *G. Ist. Ital. Attuari.*, 4:8391, 1933.

[324] N. Smirnov. Table for estimating the goodness of fit of empirical distributions. *Ann. Math. Statist.*, 19(2):279–281, 06 1948.

[325] Vijay Varma, et al. Surrogate models for precessing binary black hole simulations with unequal masses. *Physical Review Research*, 1(3):033015, 2019.

[326] Vijay Varma, et al. Surrogate model of hybridized numerical relativity binary black hole waveforms. *Physical Review D*, 99(6):064045, 2019.

[327] Fabian Hofmann, Enrico Barausse, and Luciano Rezzolla. The final spin from binary black holes in quasi-circular orbits. *ApJL*, 825(2):L19, 2016.

[328] Nathan K. Johnson-McDaniel et al. Determining the final spin of a binary black hole system including in-plane spins: Method and checks of accuracy. Technical Report LIGO-T1600168, LIGO Project, 2016. `https://dcc.ligo.org/LIGO-T1600168/public/main`.

[329] James Healy and Carlos O. Lousto. Remnant of binary black-hole mergers: New simulations and peak luminosity studies. *PhRvD*, 95(2):024037, 2017.

[330] Xisco Jiménez-Forteza, et al. Hierarchical data-driven approach to fitting numerical relativity data for nonprecessing binary black holes with an application to final spin and radiated energy. *PhRvD*, 95:064024, Mar 2017.

[331] David Keitel, et al. The most powerful astrophysical events: Gravitational-wave peak luminosity of binary black holes as predicted by numerical relativity. *Physical Review D*, 96:024006, Jul 2017.

[332] Vijay Varma, Maximiliano Isi, and Sylvia Biscoveanu. Extracting the Gravitational Recoil from Black Hole Merger Signals. *Physical Review Letters*, 124(10):101104, 2020.

[333] LIGO Scientific Collaboration and Virgo Collaboration. GCN 21505, 21509, 21513, 2017.

[334] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.

[335] B. P. Abbott et al. Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophys. J. Lett.*, 848(2):L12, 2017.

[336] L. Singer. ligo.skymap `https://lscsoft.docs.ligo.org/ligo.skymap/`, 2019.

[337] LIGO Scientific Collaboration and Virgo Collaboration. GCN 25871, 25829, 25753, 25503, 25497, 25324, 25187, 25164, 25115, 25012, 24998, 24950, 24922, 24717, 24632, 24621, 24598, 24570, 24522, 24503, 24377, 24237, 24168, 24141, 24098, 24069 `https://gcn.gsfc.nasa.gov/lvc.html`, 2019.

[338] W. Farr et al. `https://pypi.org/project/pepredicates/`, 2019.

[339] D. Chatterjee et al. `https://lscsoft.docs.ligo.org/p-astro/`, 2020.

[340] Bob Ippolito et al. `https://github.com/python/cpython/blob/3.8/Lib/json/__init__.py`, 2019.

[341] Python Software Foundation. Built-in functions `https://docs.python.org/3/library/functions.html`, 2021.

[342] The Python Software Foundation. argparse `https://docs.python.org/3/library/argparse.html`, 2021.

[343] Rory Smith and Gregory Ashton. Expediting astrophysical discovery with gravitational-wave transients through massively parallel nested sampling. *arXiv preprint arXiv:1909.11873*, 2019.

[344] Gregorio Carullo, Walter Del Pozzo, and John Veitch. Observational Black Hole Spectroscopy: A time-domain multimode analysis of GW150914. *Phys. Rev. D*, 99(12):123029, 2019. [Erratum: Phys.Rev.D 100, 089903 (2019)].

[345] Maximiliano Isi, Matthew Giesler, Will M. Farr, Mark A. Scheel, and Saul A. Teukolsky. Testing the no-hair theorem with GW150914. *Phys. Rev. Lett.*, 123(11):111102, 2019.

[346] LIGO Scientific Collaboration. https://dcc.ligo.org/LIGO-M1000066/public, 2017.

[347] L. P. Singer et al. GWCelery. https://git.ligo.org/emfollow/gwcelery, 2020.

[348] Nobel Prize Outreach AB 2021. The nobel prize in physics 2017 https://www.nobelprize.org/prizes/physics/2017/summary/, 2017.

[349] C. S. Unnikrishnan. IndIGO and LIGO-India: Scope and plans for gravitational wave research and precision metrology in India. *Int. J. Mod. Phys. D*, 22:1341010, 2013.

[350] Virgo Collaboration LIGO Scientific Collaboration and KAGRA Collaboration. Ligo, virgo and kagra observing run plans https://www.ligo.org/scientists/GWEMalerts.php, 2020.