

School of Physics and Astronomy



Towards Improving Detection Rates of
Gravitational Waves from Precessing Binary Black
Holes

Alistair William Muir

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

2018

Summary of thesis

According to Einstein’s theory of General Relativity, the acceleration of matter can cause ripples in the curvature of spacetime, given the name *gravitational waves*. Such ripples are negligible in magnitude for all but the most energetic astrophysical events, such as the coalescence of compact binary stars.

In 2015, gravitational waves were first directly detected from a binary black hole (BBH) coalescence [19]. This was achieved using two independent laser interferometers which each measured the fluctuations caused by the gravitational waves as they passed by. Matched filtering and other data analysis techniques were then employed to identify the properties of the source and measure the likelihood that the detection is a false alarm. The efficacy of these matched filtering techniques is pivotal to not only detecting gravitational waves, but drawing as much information about their sources as possible.

The methods for detecting a BBH involve the construction of a template bank; a group of synthesised *waveforms* which each represent a detectable series of gravitational waves that a BBH could produce. The characteristics of a BBH template are governed by the two masses and how they spin, the distance to the source, its orientation and its sky location. Current template banks do not include templates for sources where the spins are misaligned with the orbital momentum, which is the cause of precession in BBH. Thus, the algorithms are effective for detecting a non-precessing BBH, but much less sensitive towards precessing sources.

Creating a template bank which includes all possible precessing waveforms is computationally infeasible and would induce enough statistical noise to negate any extra sensitivity gained. However, many precessing signals would be undetectable or indistinguishable from non-precessing signals. Including such signals in a bank would result in no gain in its sensitivity. This thesis attempts to locate areas of precessing parameter space where waveforms are distinguishable from non-precessing sources, and begins work on forming a function which maps observable precession through parameter space.

Contents

1	Introduction	1
1.1	Gravitational Waves	1
1.2	Detection Methods	2
1.2.1	Weber Bars	2
1.2.2	Laser Interferometers	3
2	Gravitational Waves from Binary Black Hole Mergers	7
2.1	Inspiral	9
2.2	Merger	15
2.3	Ringdown	17
2.4	Stitching the Pieces Together	17
2.5	Precession	18
2.5.1	Orbital Precession	19
2.5.2	Transitional Precession	21
2.5.3	Eccentricity	21
3	Observing Precession in the Waveform	22
3.1	Modelling Precession	22
3.2	Waveform Parameters	23
3.2.1	Intrinsic Parameters	23
3.2.2	Extrinsic Parameters	23
3.3	Appropriate Parameters to Measure Observed Precession	25
3.3.1	Reduced Spin Parameters	25
3.3.2	Switching to $\hat{\mathbf{J}}$ as Reference for Orientation	27
3.4	Matched Filtering	29
3.5	Template Banks	31
3.5.1	Populating Template Banks	32
3.6	The Parameter Space of Interest	33
3.6.1	Relevant Parameter Space	33
3.6.2	Wave Frequency and Initial Conditions	34
3.7	Objectives	36
4	Results	37
4.1	Method	37
4.2	Case 1: Fixed Masses, Zero Aligned-Spin	38
4.2.1	Variation of Match with Precessing Spin and Orientation	39
4.2.2	Variation of Match with Orientation for Fixed Precessing Spin	42

4.2.3	Match and the Opening Angle	44
4.3	Case 2: Variation of Match with Aligned Spin	50
4.3.1	β Behaviour	54
4.4	Case 3: Variation of Match with Mass Ratio	56
4.5	Case 4: Variation of Match with Total Mass	62
4.6	Towards a Generalised Formula for Observable Precession	63
5	Conclusion	67

List of Figures

1.1	The effect of a GW on the plane of spacetime orthogonal to its propagation. Above, the red dashed lines indicate how a plus-polarised wave affects a ring in this plane: stretching the \hat{y} -axis whilst squeezing the \hat{x} -axis. The ring then returns to circular before stretching the \hat{x} -axis and squeezing the \hat{y} -axis, after which the ring returns to circular once more. Beneath, the blue dotted lines show the effect of a cross-polarised GW, where the effect is the same but with the axes rotated 45° . This cartoon shows an unrealistically high magnitude of peak strain to demonstrate the effect ($\Delta L/L \sim 0.25$), whilst a typical source would cause a peak strain of $\Delta L/L \sim 10^{-20}$	2
1.2	<i>Reproduction of the figure from [12].</i> Simplified diagram depicting the experimental set-up of a LIGO interferometer used to detect GW.	4
2.1	<i>Reproduction of FIG. 2 in [19].</i> The waveform of the first BBH GW detection with cartoons depicting the inspiral, merger and ringdown phases. Plots showing the evolution of the orbital separation and relative velocity are given underneath.	8
2.2	<i>Reproduction of FIG. 1 in [9].</i> Left: The frequency-dependent sensitivity curves for both LIGO detectors during the first observing run in the Advanced LIGO era, plotted with the frequency-dependent strain of the three candidate detections found in that period. The thickness of the bands represent the 90% confidence intervals. Right: Plots of <i>Strain vs Time</i> for each of the signals.	9
2.3	An example of the relative amplitude vs inclination (blue line) of instantaneous GW emission for an arbitrary polarisation, in the non-inertial reference frame of $\hat{\mathbf{L}}$. The black circle represents the unit sphere. In this co-rotational frame, we see $\hat{\mathbf{J}}$ precessing about $\hat{\mathbf{L}}$. To an inertial distant observer, $\hat{\mathbf{L}}$ would be seen to precess about the constant \mathbf{J} , describing a cone of identical shape. Line-of-sight $\hat{\mathbf{N}}$ is also plotted with its trajectory on the unit sphere describing a circle identical in size to that of the precession cone.	20
2.4	A precessing inspiral waveform overlaid with a non-precessing inspiral waveform from a source with the same masses and aligned spin. In both cases, the orbital frequency is seen to increase throughout. The precession effects upon the amplitude and phase are visible, as is the acceleration of the rate of precession towards merger.	20

3.1	Relative amplitude (to first-order) for instantaneous GW emission vs inclination, ι , for plus-polarised (+) and cross-polarised (\times) waves. The arrow represents $\hat{\mathbf{L}}$, the black circle represents the unit sphere.	24
3.2	The precession of $\hat{\mathbf{L}}$ in a co-rotating frame with the opening angle, β , labelled. A physical representation of the precession phase, ϕ_{JL} , and orbital phase, ϕ , are also labelled.	28
3.3	<i>Reproduction of Figure 1 from [70].</i> Noise-Frequency curve for the PSD used for all matched filtering calculations in this thesis (zero-detuned, high power), labelled here as "Total noise". It is calculated as a summation of theorised optimal contributions from all identifiable noise sources, which are also plotted.	30
4.1	Colour plots for \mathcal{M} vs χ_p vs θ_{JN} , with each row showing results with fixed $\psi = [0, 15, 30, 45^\circ]$ downwards and columns $\phi_{JL} = [0, 90^\circ]$ from left to right.	40
4.2	Colour plots for \mathcal{M} vs χ_p vs θ_{JN} , with each row showing results with fixed $\psi = [0, 15, 30, 45^\circ]$ downwards and columns $\phi_{JL} = [180, 270^\circ]$ from left to right.	41
4.3	Colour plots for \mathcal{M} vs ψ vs θ_{JN} , with $\chi_p = 0.8$ and fixing ϕ_{JL} at the specified values for each plane.	43
4.4	$\chi_p = 0.8$, $\phi_{JL} = 90^\circ$. A colour plot for \mathcal{M} vs ψ vs θ_{JN} , showing the same results as Figure 4.3c overlaid with the precession cone's tangent circle as calculated by the LAL formula: $\beta_{LAL} = 33.7^\circ$, and Lundgren's formula: $\beta_{Lund} = 35.9^\circ$	45
4.5	<i>Reproduction of Figure 3.1 in Section 3.2.</i> The relative amplitude of instantaneous emission for the two polarisations vs inclination, ι . The arrow represents $\hat{\mathbf{L}}$. A null line sits in the orbital plane ($\iota = 90^\circ$) for cross-polarised waves, where you can see that the emitted amplitude is zero.	46
4.6	$\chi_p = 0.8$ ($\beta_{LAL} = 30.7^\circ$), $\phi_{JL} = 90^\circ$. A colour plot for \mathcal{M} vs ψ vs θ_{JN} with longer waveforms, their initial conditions set at $f_{gw} = 20\text{Hz}$. The tangent circle to the β_{LAL} precession cone is also plotted.	46
4.7	$\chi_p = 0.8$, $\phi_{JL} = 45^\circ$. Colour plot for \mathcal{M} vs ψ vs θ_{JN} , using the same parameters as Figure 4.3b with longer waveforms (initial frequency $f_{gw} = 20\text{Hz}$).	48
4.8	Amplitude profiles for both polarisations, with fixed vectors $\hat{\mathbf{J}}$ and $\hat{\mathbf{N}}$. $\hat{\mathbf{L}}$ is shown to be precessing about $\hat{\mathbf{J}}$. The three plots show the profiles at different stages on the precession cone. The black circle represents the unit sphere. From left to right, ϕ_{JL} has passed through approximately 90° since the previous graph.	48
4.9	Colour plots for \mathcal{M} vs ψ vs θ_{JN} in planes of the given χ_{eff} . $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$, $\phi_{JL} = 90^\circ$ and initial frequency $f_{gw} = 30\text{Hz}$ for all waveforms.	51
4.10	Colour plots for \mathcal{M} vs ψ vs θ_{JN} where $\chi_{eff} < 0$. $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$ and initial frequency $f_{gw} = 30\text{Hz}$ for all waveforms.	52

4.11	Colour plots for \mathcal{M} vs ψ vs θ_{JN} where $\chi_{eff} < 0$. $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$, and initial frequency $f_{gw} = 20\text{Hz}$ for all waveforms.	53
4.12	A colour plot showing \mathcal{M} vs ψ vs θ_{JN} to demonstrate the tangent circle “wrapping around” on the ψ -axis as a result of the opening angle $\beta > 45^\circ$. In this case, $\beta_{LAL} = 53.0^\circ$	54
4.13	The relationship between β and χ_{eff} for different χ_p , where $m_1 = 20M_\odot$, $m_2 = 5M_\odot$ and $f_{gw} = 30\text{Hz}$. β was calculated using the LAL equation.	55
4.14	The frequency-evolution of β as calculated by the LAL algorithms for sources with given χ_p , $\chi_{eff} = -0.6$, and masses $m_1 = 20M_\odot$, $m_2 = 5M_\odot$	56
4.15	Colour plots for \mathcal{M} vs ψ vs θ_{JN} for the given mass ratios where total mass $M = 25M_\odot$, $\chi_p = 0.6$ and $\phi_{JL} = 90^\circ$	57
4.16	Top: The relationship between β and mass ratio, q , for various values of χ_p , where $\chi_{eff} = 0$, $M = 25M_\odot$ and frequency $f_{gw} = 30\text{Hz}$. β was calculated using Eq. 4.2. Bottom: The relationship between Q_p and β , for $M = 25M_\odot$ and $\chi_{eff} = 0$	58
4.17	Plots showing the variance in the Match for each $[\theta_{JN}, \psi]$ coordinate for combinations of $[\chi_p, q]$, spanning their full possible range, for four fixed values of Q_p . Descending: $Q_p = [1.2, 3.6, 6.0, 8.4]$. $\chi_{eff} = 0$ and $M = 25M_\odot$ for all waveforms. Note: the colour axes are not consistent.	60
4.18	Plots showing the variance in the Match for each $[\theta_{JN}, \psi]$ coordinate for $Q_p = 3.6$ with ϕ_{JL} fixed at four different values. Descending: $\phi_{JL} = [30, 60, 120, 150^\circ]$. Note: the colour axes are not consistent. . .	61
4.19	Colour plots for \mathcal{M} vs ψ vs θ_{JN} for different total masses where $q = 4$, $\chi_p = 0.6$, $\chi_{eff} = 0$ and $\phi_{JL} = 90^\circ$	62
4.20	Top: The relationship between β and Q_p for a range of sources where total mass is fixed at five given values and $\chi_{eff} = 0$. Bottom: The relationship between β and Q_p for a range of sources where $M = 25M_\odot$ and $\chi_{eff} = -0.6$	65

List of Tables

2.1	A description of the group of correction terms introduced at each PN order. See Section 2.5 for an explanation of precession.	15
4.1	Results from a number of variance calculations of the Matches over the $[\theta_{JN}, \psi]$ plane ($\phi_{JL} = 90^\circ$ for all), for the given fixed values of Q_p , aligned spin, χ_{eff} , and total mass, M . The mean of the variance over the whole plane ($\overline{\text{Var}(\mathcal{M} Q_p)}$) is given along with the percentage of points with variance of $\text{Var}(\mathcal{M} Q_p) > 0.01$. Every waveform was started at $f_{gw} = 30\text{Hz}$	64

To Mum and Dad, for being the parents I needed;
Bleddyn, for being the friend I needed;
Steve, for being the supervisor I needed;
And to failure, for being the success I needed.

Chapter 1

Introduction

1.1 Gravitational Waves

According to Einstein's theory of General Relativity (GR), gravity is the result of mass being affected by the curvature of spacetime. Spacetime is curved by the presence of any mass, and as such any change in the distribution of mass will cause the spacetime curvature to adapt [1]. The speed at which this change occurs encapsulates a significant difference between GR and Newtonian theory: Einstein's theory produces the result that gravity travels at light speed (more precisely, any change in the mass distribution causes a change in spacetime curvature which propagates from the source at light speed), whereas in a Newtonian universe any change of mass distribution would be felt everywhere instantaneously. In the case where a mass is accelerated in a manner which changes the quadrupole moment, GR predicts the creation of oscillations in spacetime curvature which propagate away from the mass, known as gravitational waves (GW).

Einstein published the GW solution to his Field Equations in 1916, a year after he first formalised and published his GR theory [2]. The equations reduce to the time-dependent wave equation when a weak non-stationary field is assumed very distant from a relativistic disturbance. When certain gauge conditions are chosen (e.g. the amplitude of curvature is only non-zero orthogonal to the propagation), the solution gives two polarisations of plane waves [3]. As a GW propagates it contorts spacetime in the plane perpendicular to the direction of propagation, squeezing one dimension whilst stretching the orthogonal dimension as demonstrated in Figure 1.1. The two polarisations are 45° apart, one is labelled as *plus-polarised* (+) whilst the other is *cross-polarised* (\times). A plus-polarised GW travelling along $\hat{\mathbf{z}}$ stretches the plane along the $\hat{\mathbf{y}}$ -axis while squeezing the $\hat{\mathbf{x}}$ -axis by the same factor, before doing the opposite and finally returning the plane to its original state. A cross-polarised GW

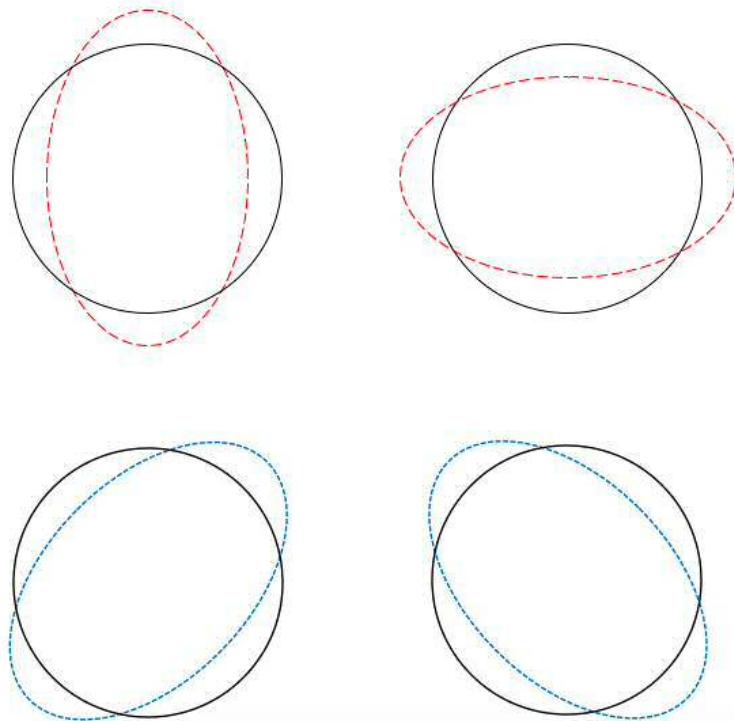


Figure 1.1: The effect of a GW on the plane of spacetime orthogonal to its propagation. Above, the red dashed lines indicate how a plus-polarised wave affects a ring in this plane: stretching the \hat{y} -axis whilst squeezing the \hat{x} -axis. The ring then returns to circular before stretching the \hat{x} -axis and squeezing the \hat{y} -axis, after which the ring returns to circular once more. Beneath, the blue dotted lines show the effect of a cross-polarised GW, where the effect is the same but with the axes rotated 45° . This cartoon shows an unrealistically high magnitude of peak strain to demonstrate the effect ($\Delta L/L \sim 0.25$), whilst a typical source would cause a peak strain of $\Delta L/L \sim 10^{-20}$.

would do the same, but with the axes rotated 45° about the direction of propagation. Any GW can be described as a superposition of these two polarisations.

The strongest astrophysical sources of GW are considered to be coalescences of the highest density stars with high mass, such as binary black hole (BBH) systems [4]. Such events produce a peak strain of $\Delta L/L \sim 10^{-20}$, dependent on their mass, distance and orientation to the observer. Such a strain is equivalent to a 100km long bungee rope stretching and contracting through the diameter of a proton.

1.2 Detection Methods

1.2.1 Weber Bars

The first experiments were carried out in the 1960s by Joseph Weber, who constructed resonant-mass bars. These were cylinders, roughly 3m in length and 30cm in diameter, with piezoelectric crystal attached around their waist [5]. As a gravi-

tational wave of a frequency close to the bar’s natural frequency travelled through, the bar would compress and depress accordingly. This would cause the crystals to produce a very weak electromagnetic signal. Weber attempted to isolate the crystals from seismic events and electromagnetic disturbances, purportedly reducing his noise threshold to roughly 10^{-16}m [6]. He reported several hundred detections over two years, but several attempts were made to reproduce Weber’s results without success. *Gibbons and Hawking* (1971) calculated that sources in the galactic centre or nearer would each have to convert $70M_{\odot}$ into gravitational radiation within 0.5 seconds, 10,000 times per day, to produce the signals at the frequency Weber observed [7]. For comparison, the most powerful source detected to date by more sensitive experiments, discussed in Section 1.2.2, was a BBH coalescence which lost an estimated $3M_{\odot}$ to GW emission over roughly the same timescale at a much greater distance (of order $\mathcal{O}(100)\text{Mpc}$) [8]. Such experiments have detected a total of six events in observation periods spanning several months [9, 10].

1.2.2 Laser Interferometers

An alternative method for detecting gravitational radiation was proposed by a number of scientists independently, including Weber and Rainer Weiss in 1967. Feasibility studies were carried out throughout the 1970s, and the first interferometer was built in 1978 [11]. With cited sensitivities only down to 10^{-16}m , it may not have been surprising that this experiment did not yield any detections, but its successful construction and operation acted as a proof of concept. Concept development continued and in 1990, a proposal to begin the construction of a pair of larger scale ground-based laser interferometers was approved [4]. *LIGO* - the Laser Interferometer Gravitational wave Observatory - was born.

The experiment involves directing a high-power laser at a beam splitter, causing two coherent beams to be sent down perpendicular arms. The tubes are each 4 kilometres long, and the photons in each beam are reflected a number of times between extremely smooth, extremely reflective test masses before returning to the beam splitter. At this point the beams reunite and the combined beam is directed into a photodetector.

If a gravitational wave was to propagate through the experiment, the path of one beam would be stretched while the other is squeezed, causing the two beams to travel marginally different distances before reuniting at the splitter. This slight difference in travel time causes the beams to be out of phase as they superpose and reflect to the photodetector. As the GW passes through, the phase-shift will vary with the strain it causes to the beam lengths. The photodetector records the

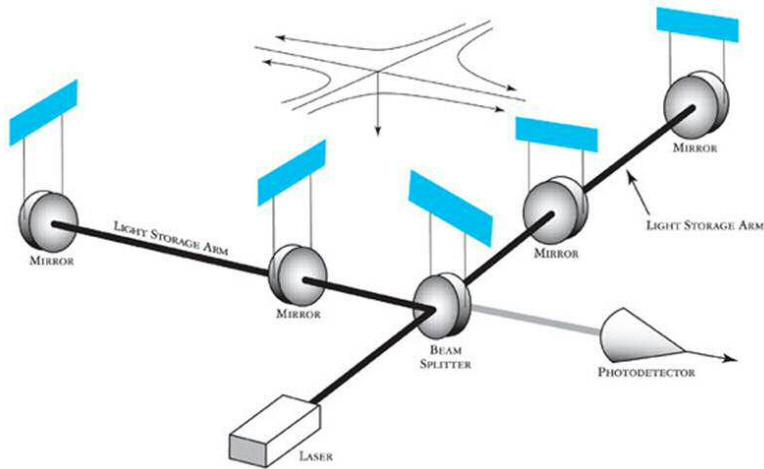


Figure 1.2: *Reproduction of the figure from [12].* Simplified diagram depicting the experimental set-up of a LIGO interferometer used to detect GW.

varying intensity of the superposed beam, and real-time computational analysis is used to identify specific patterns which have the hallmarks of a GW signal. If such a pattern is found, the detector response is analysed both around the time that the potential signal was registered and over the whole period that the detector was active. If the pattern is demonstrated to be of sufficient significance - most notably it's signal-to-noise ratio (SNR) - parameter estimation is carried out to identify the source's characteristics and from where the signal came.

The detectors must be sensitive to changes in arm length of $\Delta L \sim 10^{-18}\text{m}$ in order to observe GW from extragalactic sources [13]. This extreme sensitivity requires the isolation from numerous noise sources. The dominant noise sources include those of seismic, thermal and quantum origin [14]. Moreover, irregular transient noise caused by human activity, earthquakes and weather must also be mitigated. Noise is characterised in the frequency domain. Terrestrial sources of noise dominate in the lower frequency bands (0.1 - 10Hz), such as seismic activity, wind and human activity. Thermal noise is the chief contributor in the 10 - 100Hz range, but diminishes near 100Hz where quantum shot noise begins to dominate [14]. The test masses and beam-splitter are suspended on a sophisticated quadruple-pendulum system. These use both active-response and passive techniques to isolate the laser from low-frequency noise sources, damping vibrations and transients in this frequency region. Some contribution to the thermal noise is due to movements in the test masses due to their suspended position. Losses in the thermal fluctuations on the test mass surfaces also contribute. In order to minimise these effects, coatings on the test masses are optimised to maximise reflectivity whilst reducing losses. Quantum shot noise acts as a lower-bound for sensitivity in regions where it dominates, but

methods using quantum squeezing have been proposed to force that limit lower [15].

The perpendicular arms ensure that the strain caused by a GW can have a maximal effect on the laser paths. However, this set-up renders the interferometer totally insensitive to GW propagating from certain directions. Figure 1.2 shows the strain of a wave travelling through the detector at optimal orientation; one arm will be maximally squeezed whilst the other is maximally stretched. But if the GW was in the other polarisation (i.e. rotated 45°) both arms would be affected in an identical manner throughout the GW, causing the beam phase to be unaffected. Alternatively, if the GW was propagating parallel to one of the arms, only the other arm’s length would be affected. In this case the detector’s sensitivity is at most 50% of the overhead source in Figure 1.2 (when the GW is plus-polarised, inducing maximal strain in the arm orthogonal to propagation). This example demonstrates that a detector’s sensitivity to a given GW is heavily dependent on its direction of propagation as well as its polarisation [16].

The two original LIGO experiments in Hanford, Washington and Livingstone, Louisiana were built in alignment on Earth’s surface (their geographical positions mean that the planes on which the experiments lie are marginally different). This was done so that each detector could act as an effective check for “glitches” and noise transients in the other, and also to maximise the sensitivity for two detectors [17]. There is a caveat to arranging the two detectors in this manner: both detectors are insensitive to the same polarisation, reducing the capability to estimate the source’s distance, sky location and orientation. Multiple detectors with different orientations are able to place much more precise limits on these parameters. The fifth event detected by LIGO was simultaneously detected by another interferometer experiment based in Italy, VIRGO, which enabled the likely sky location (90% confidence region) to be constrained to 60 deg^2 [18]. In comparison, the sky location of previous events detected by LIGO alone were constrained to areas ranging from 520 deg^2 to 1200 deg^2 [19, 20, 21, 22]. Another detector based in Germany, GEO600, is used as a testbed for interferometer technology, whilst two more interferometers are in the process of being constructed in India (LIGO-India) and Japan (KAGRA) [23].

Binary black hole (BBH) coalescences are considered the most likely sources to produce detectable signals in LIGO. Since upgrades to the LIGO detectors were completed in 2014, beginning the “Advanced LIGO” era, two observational periods have been conducted in which five observations of GW from such sources have been made [13, 19, 20, 21, 22, 18]. Whilst coalescences involving neutron stars are not within the remit of the research in this thesis, it is very much worth mentioning that

LIGO also observed a GW signal from a binary neutron star coalescence. A gamma-ray burst was detected in coincidence by the *Fermi Gamma-ray Burst Monitor* [24], and a kilonova was also observed in data collected in days after the event [25]. Collaborative analysis strongly supports the hypothesis that the same event caused these phenomena [26].

Chapter 2

Gravitational Waves from Binary Black Hole Mergers

Binary black hole mergers emit the most powerful gravitational radiation of all known sources to which ground-based interferometers are sensitive [27]. As a BBH orbits, orbital energy is radiated away in the form of GW, causing the separation of the two black holes to reduce and the orbital frequency to increase. The black holes are now undergoing greater acceleration as they orbit, so they emit stronger GW which radiate more orbital energy, further reducing their separation and increasing their orbital frequency [28]. This process continues, the orbit accelerating rapidly as the black holes spiral inwards until they reach a separation too small to maintain stable orbit. At this point they plunge into each other, forming a single black hole which then experiences a series of oscillations as the spacetime reaches a new stable state [29]. The series of GW observed from such an event are collectively described as a *waveform*.

The waveform produced by a BBH coalescence can be described as a combination of three phases: inspiral, merger and ringdown. The *inspiral* phase is defined as the period where the objects are in quasi-circular orbit, where the orbital decay is dominated by the emission of gravitational waves [13]. Once their separation breaches the innermost stable orbit, the black holes plunge into each other. This is the *merger* phase [30]. Once the black holes have merged, the remnant body enters the *ringdown* phase, where its spacetime curvature undergoes oscillations as the stress-energy tends to a new stable state [31]. These oscillations in turn cause further GW emission, decaying in magnitude until the new black hole becomes axi-symmetric [29].

Figure 2.1 shows a representation of LIGO’s first detected GW event with the three phases indicated. In the case of the first detected BBH events, the inspiral

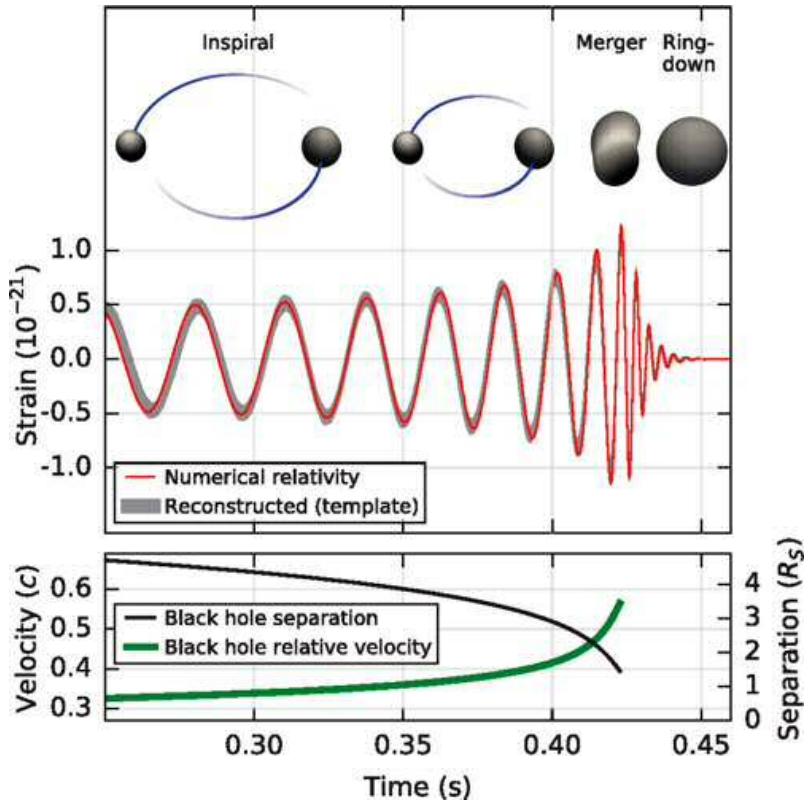


Figure 2.1: *Reproduction of FIG. 2 in [19].* The waveform of the first BBH GW detection with cartoons depicting the inspiral, merger and ringdown phases. Plots showing the evolution of the orbital separation and relative velocity are given underneath.

was the phase which contributed the most to the signal strength. This is because the majority of the signal duration was spent in the inspiral, and the GW frequencies in this phase sit within the detectors' most sensitive band as shown in Figure 2.2. Notice that, whilst the GW151226 signal is much longer than GW150914, the latter was more powerful and peaked at a frequency where the detectors were more sensitive. For this reason, GW150914 appeared as a much stronger signal in the LIGO data [9]. LVT151012 is a candidate signal which did not meet the statistical criteria to be accepted as a bona fide detection, but data analysis concluded it is more likely to have been a GW signal than a noise transient [9].

In order to detect GW from a BBH coalescence, it must be understood what gravitational waves would look like from each of the three phases. This requires the ability to simulate GW emission throughout a BBH coalescence. For the majority of BBH sources, the inspiral phase is longer in duration than the merger and ringdown. However, the merger phase produces the peak in GW strain, and much information about the source is theoretically calculable from the ringdown phase. Each of the three phases pose different challenges in modelling GW emission, which are described in Sections 2.1-2.3. To synthesise a complete waveform, the models for the three

phases must be somehow seamed together. This process is discussed in Section 2.4. The inspiral phase shall be described in particular detail, as it is during the inspiral where precession effects are observable [32]. Precession itself shall be described in Section 2.5.

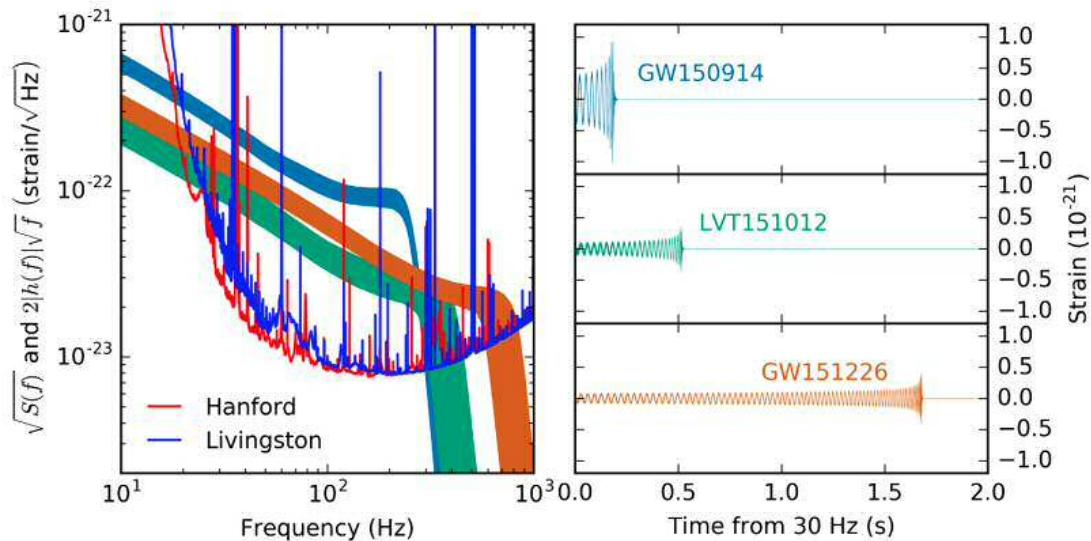


Figure 2.2: *Reproduction of FIG. 1 in [9].* Left: The frequency-dependent sensitivity curves for both LIGO detectors during the first observing run in the Advanced LIGO era, plotted with the frequency-dependent strain of the three candidate detections found in that period. The thickness of the bands represent the 90% confidence intervals. Right: Plots of *Strain* vs *Time* for each of the signals.

2.1 Inspiral

The inspiral phase is where the two black holes are orbiting each other and emanating a continuous flow of GW, thereby reducing the orbital energy, causing the black holes' separation to reduce. The closer they become, the more energy is released via GW, causing them to “spiral in” at an accelerating rate.

A non-trivial distribution of net-zero mass is only a hypothetical phenomenon, with no confirmed detections of such an occurrence and all commonly accepted theories of gravitation suggesting that mass dipoles do not exist [27]. In other words, “negative mass” has not been observed and we have little reason to believe it will be. As such, acceleration of matter can only cause gravitational radiation if there is a flux in the quadrupole moment. Therefore, to lowest order, GW are emitted in a quadrupole formation. The derivation of the quadrupole approximation is summarised below.

In linearised theory, the spacetime metric, $g_{\mu\nu}$, can be described as a Minkowski (i.e. flat) spacetime metric, $\eta_{\mu\nu}$, undergoing a slight perturbation, $h_{\mu\nu}$:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad ||h_{\mu\nu}|| \ll 1. \quad (2.1)$$

We follow the convention that Greek letters (μ, ν) refer to the four spacetime coordinates and Roman letters (i, j) refer to the three spacial coordinates. If the Lorentz gauge is chosen (Eq. 2.2), and two conditions are set: spacetimes are globally vacuum (Eq. 2.3) and asymptotically flat (Eq. 2.4), the transverse-traceless gauge is formed:

$$\partial^\mu \bar{h}_{\mu\nu} = 0 \quad (2.2)$$

$$T_{\mu\nu} = 0 \quad (2.3)$$

$$\lim_{r \rightarrow \infty} h_{\mu\nu} \rightarrow 0 \quad (2.4)$$

This is a convenient gauge choice for GW mathematics since all local gauge freedoms are entirely fixed; only physical information is contained in the metric perturbation $h_{\mu\nu}^{TT}$ [33]. The superscript “TT” denotes the metric is in the transverse-traceless gauge. This notation will be used henceforth.

In Lorentz gauges alone (without applying conditions 2.3 or 2.4), the linearised Field Equations are

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}, \quad (2.5)$$

where \square is the d’Alembertian operator and $T_{\mu\nu}$ is the energy-momentum tensor of matter [27]. The gravitational constant, G , and speed of light in vacuum, c , take their usual values. Solving Eq. 2.5 for radiation requires the *retarded Green’s function*:

$$G_{ret}(x - x') = -\frac{1}{4\pi|\mathbf{x} - \mathbf{x}'|} \delta(x_{ret}^0 - x'^0), \quad (2.6)$$

where $x'^0 = ct'$, $x_{ret}^0 = ct_{ret}$ and the *retarded time*, t_{ret} , accounts for the time it takes for information to propagate from \mathbf{x}' to \mathbf{x} :

$$t_{ret} = t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}. \quad (2.7)$$

The result,

$$\bar{h}_{\mu\nu} = \frac{4G}{c^4} \int d^3 x' \frac{1}{|\mathbf{x} - \mathbf{x}'|} T_{\mu\nu} \left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}, \mathbf{x}' \right), \quad (2.8)$$

can be placed in the TT gauge if the two restrictions Eqs. 2.3 and 2.4 are placed on the solution. These respectively place the observer distant to the source, and describe the background spacetime as flat (i.e. the source is conforming to Newto-

nian gravity, and not self-gravitating). The former choice implies $|\mathbf{x}| \gg |\mathbf{x} - \mathbf{x}'|$, permitting the substitution $r = |\mathbf{x}| = |\mathbf{x} - \mathbf{x}'|$ [33]. The next step is to perform an expansion about v/c , where v is the velocity of the internal motions of the source [33]. This requires the assumption $v \ll c$, which, along with the assumption of Newtonian gravity, is not strictly appropriate for CBC sources. We shall address these discrepancies later. The leading term in this expansion is the *Einstein quadrupole*:

$$\bar{h}_{ij}^{TT}(t, \mathbf{x}) = \frac{1}{r} \frac{4G}{c^4} \int d^3x T^{ij}(t - \frac{r}{c}, \mathbf{x}). \quad (2.9)$$

The higher-order terms form the multipole expansion, which will be discussed later in this section. As described earlier, GW carry energy away from the source; as the source emits GW it loses mass and momentum via back-reaction (confirmed by Hulse and Taylor’s discovery of a binary neutron star in decaying orbit [28, 34]). If the momenta of stress tensor T^{ij} is defined as

$$\frac{1}{2} \ddot{M}^{ij}(t) = \int d^3x T^{ij}(t, \mathbf{x}), \quad (2.10)$$

where M is the mass of the source, the power radiated can be expressed in terms of the rate of change of momenta. However, for this derivation the amplitude of emitted GW is desired, for which we can define the quadrupole moment:

$$Q^{ij} = M^{ij} - \frac{1}{3} \delta^{ij} M_{kk}. \quad (2.11)$$

Eq. 2.9 can now be written in the more typical form in terms of Q :

$$[h_{ij}^{TT}(t, \mathbf{x})]_{quad} = \frac{1}{r} \frac{2G}{c^4} \ddot{Q}_{ij}^{TT}(t - r/c), \quad (2.12)$$

with the subscript “quad” indicating that the metric is a quadrupole formation. Re-written in the more wieldy form of the two independent polarisations (aligning the direction of propagation $\hat{\mathbf{N}}$ with $\hat{\mathbf{z}}$ in a frame with spacial axes (x, y, z)), Eq. 2.12 becomes

$$h_+ = \frac{1}{r} \frac{G}{c^4} (\ddot{M}_{11} - \ddot{M}_{22}), \quad (2.13)$$

$$h_\times = \frac{2}{r} \frac{G}{c^4} \ddot{M}_{12}. \quad (2.14)$$

We are now prepared to consider quadrupole radiation for the specific case of compact binaries. In the centre-of-mass frame for a binary system in quasi-circular

orbit, the second mass moment $M^{ij} = \mu x_0^i(t)x_0^j(t)$ gives

$$M_{11} = \mu R^2 \frac{1 - \cos(2\omega_{orb}t)}{2} \quad (2.15)$$

$$M_{22} = \mu R^2 \frac{1 + \cos(2\omega_{orb}t)}{2} \quad (2.16)$$

$$M_{12} = -\frac{1}{2}\mu R^2 \sin(2\omega_{orb}t) \quad (2.17)$$

and all other components vanish. R is the radius of orbit with frequency ω_{orb} . $\mu = m_1 m_2 / (m_1 + m_2)$ is the reduced mass of the system. These equations lead to

$$\ddot{M}_{11} = -\ddot{M}_{22} = 2\mu R^2 \omega_{orb}^2 \cos(2\omega_{orb}t), \quad (2.18)$$

$$\ddot{M}_{12} = 2\mu R^2 \omega_{orb}^2 \sin(2\omega_{orb}t), \quad (2.19)$$

which, when plugged into Eqs. 2.13 and 2.14, give

$$h_+(t) = \frac{1}{r} \frac{4G\mu\omega_{orb}^2 R^2}{c^4} \left(\frac{1 + \cos^2 \iota}{2} \right) \cos(2\omega_{orb}t), \quad (2.20)$$

$$h_\times(t) = \frac{1}{r} \frac{4G\mu\omega_{orb}^2 R^2}{c^4} \cos \iota \sin(2\omega_{orb}t). \quad (2.21)$$

The inclination of $\hat{\mathbf{N}}$ from $\hat{\mathbf{L}}$ is denoted by ι . As this is a first-order linear approximation, Kepler's law holds:

$$\omega_{orb}^2 = \frac{Gm}{R^3}. \quad (2.22)$$

Substituting in Eq. 2.22 and the *chirp mass*,

$$M_c = \mu^{3/5} m^{2/5} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad (2.23)$$

we can express Eqs. 2.20 and 2.21 as

$$h_+(t) = \frac{4}{r} \left(\frac{GM_c}{c^2} \right)^{5/3} \left(\frac{\pi f_{gw}}{c} \right)^{2/3} \frac{1 + \cos^2 \iota}{2} \cos(2\pi f_{gw} t_{ret} + 2\phi), \quad (2.24)$$

$$h_\times(t) = \frac{4}{r} \left(\frac{GM_c}{c^2} \right)^{5/3} \left(\frac{\pi f_{gw}}{c} \right)^{2/3} \cos \iota \sin(2\pi f_{gw} t_{ret} + 2\phi), \quad (2.25)$$

where ϕ is the orbital phase at t_0 , $f_{gw} = \omega_{gw}/(2\pi)$ and $\omega_{gw} = 2\omega_{orb}$.

Four assumptions were made in this derivation: the source masses are in quasi-circular orbit; the observer of gravitational radiation is distant; the source is not self-gravitating and its internal velocities are slow ($v \ll c$). This is the multipolar

expansion. As the black holes approach each other, the spacetime perturbations within the source become very strong and the velocities reach the relativistic domain [35], so one would not expect the expansion to agree with simulations of such sources. Nonetheless, it was discovered that waveforms created with this expansion agreed remarkably well with numerical simulations for compact binary inspirals [36, 37, 38]. However, beyond the first few terms, the multipole expansion diverges as $r \rightarrow \infty$. In short, this is due to the non-linearity of GR: the assumption that a distant observer is at infinity (where we substituted $|\mathbf{x} - \mathbf{x}'|$ for r) cannot hold throughout the entire expansion [27]. Thus, accurate models must rely on a finite expansion, where the self-gravity of the source is considered. This expansion is named the *post-Newtonian* (PN) approximation, and creates additional correction terms for each multipole term [39].

The PN approximation expands about v/c , but for convenience the dimensionless parameter x is defined:

$$x \equiv \left(\frac{GM\omega_{orb}}{c^3} \right)^{2/3} = \mathcal{O}\left(\frac{v^2}{c^2}\right), \quad (2.26)$$

where $M = m_1 + m_2$ is the total mass of the binary system. The PN approximation can be expanded as a power series in $x^{1/2}$. By way of example, the equations of motion for a CBC take the form

$$\frac{d^2x^i}{dt^2} = -\frac{GM}{r^2}\hat{x}^i[1 + \mathcal{F}_1(\eta)x + \mathcal{F}_2(\eta)x^{3/2} + \mathcal{F}_3(\eta)x^2 + \mathcal{F}_4(\eta)x^{5/2} + \dots], \quad (2.27)$$

where \mathcal{F}_i are trivial (usually linear) functions of symmetric mass ratio $\eta = (m_1 + m_2)/M^2$ and remembering chirp mass, M_c , from Eq. 2.23. The number of orbital cycles in a waveform from a given frequency ($f_{gw} = \omega_{orb}/\pi$) is a calculation with the generic form of

$$\mathcal{N}_{cyc} = \frac{x^{-5/2}}{32\pi\eta}[1 + \mathcal{O}(x) + \mathcal{O}(x^{3/2}) + \mathcal{O}(x^2) + \mathcal{O}(x^{5/2}) + \dots]. \quad (2.28)$$

A specific PN correction term is henceforth labelled as “ n PN” where n is the power to which x is raised in that term (i.e. half the order to which v/c is taken in the expansion), e.g. $\mathcal{F}_2(\eta)x^{3/2}$ represents the 1.5PN term in Eq. 2.27. You may notice there are no 0.5PN terms. All such terms vanish in the expansion, for no known physical reason [27]. Alternative theories of gravitation do produce 0.5PN terms in their expansions, and so this disparity has been utilised to test the accuracy of GR [40]. The generic form of the PN expansion exemplified in Eq. 2.27 demonstrates how the contributions diminish in magnitude for higher expansion terms (since v/c

is, of course, always less than unity). However, we can see that the the number of orbital cycles diverges as $x \rightarrow 0$ for terms lower than 2.5PN. This is to be expected: the orbital frequency is a function of separation, and black holes at infinity to each other would take an infinite number of orbits to coalesce. As separation increases, the general relativistic effects diminish, and so beyond a certain PN order the terms contribute too weakly to diverge. In practice, waveforms begin at frequencies too high for \mathcal{N} to explode, but Eq. 2.28 indicates the importance of post-Newtonian corrections in forming accurate waveforms. Without terms to 2.5PN, the number of orbital cycles in the waveform would have errors of $\mathcal{O}(1)$ or higher [27].

The convergence properties of the PN expansion are not well understood, as it gets progressively more difficult to identify the next order of PN terms, but current models which incorporate all terms to 3.5PN are faithful enough to numerical models for the detection and parameter estimation of BBH sources [39, 19]. Every PN order contains a number of correction terms, with contributions from various GR effects on the source.

Different phenomena begin to contribute PN terms at different expansion orders, detailed in Table 2.1. Modelling a complete waveform numerically, known as *Numerical Relativity* (NR), costs a huge amount of computation time, so analytical methods have been developed to construct faithful waveforms quickly [44]. An analytical waveform model or “approximant” includes a truncation of the PN expansion, choosing PN terms from which it derives its own equations for the *polarisation amplitudes*, i.e. the amplitudes of each GW polarisation emitted at a given moment in a given direction, and *orbital phase* of gravitational radiation, and the equations of motion to evolve a BBH in time. To construct a waveform, the approximant is given parameter values pertaining to the source and its orientation, and the equations are used to describe the complete waveform. A given BBH approximant is calibrated within a specific range of values for the black hole masses and spins, outside of which it may give inaccurate waveforms - unfaithful to NR models [42].

Waveform models for BBH coalescences with generic spins include terms up to order 3.5PN [19, 45, 46, 20]. These models are not reliable for mass ratios above $q = m_1/m_2 = 9$ (here and henceforth, the convention $m_1 \geq m_2$ applies). Models for specific classes of BBH with higher mass ratios have incorporated 4PN terms and higher [47], and methods to go beyond 4PN for generic systems have been proposed [48, 49, 50]. However, there is currently no BBH model that incorporates terms beyond 3.5PN in waveform calculations which has been calibrated for the range of parameters values being investigated here [46].

It is worth noting that such terms would generally be dominated by the lower

PN Order	Significance
0 (Quad)	Chirp mass: The dominant term in the quadrupole formula is $\propto M_c^{5/3}$ (see Eqs. 2.24, 2.25). All PN terms are subject to this factor, thus M_c dominates the calculation.
0.5	No terms: All terms of this order vanish in the PN expansion. Alternative theories of gravity derive expansions which do include 0.5PN terms, and search algorithms have been developed to test GR using this difference [40].
1	Mass ratio: After leading order, terms are dependent on different combinations of M_c and η . Therefore, symmetric mass ratio is the second most dominant parameter.
1.5	Spin-orbit coupling: Leading order terms appear for interactions between orbital angular momentum and single-spin angular momentum. <i>Highest order included in polarisation amplitude calculations for current precessing BBH waveform models</i> [41].
2	Spin-spin coupling: Leading order terms appear for interactions between both black hole spin momenta and orbital angular momentum. <i>Only leading order PN terms known for spin-spin coupling</i> [39].
3	<i>Highest order included in polarisation amplitude calculations for current non-precessing BBH waveform models</i> [42].
3.5	<i>Highest known PN terms from spin-orbit coupling</i> [43]. <i>Highest PN order included in phase calculations for current precessing BBH waveform models</i> [41].

Table 2.1: A description of the group of correction terms introduced at each PN order. See Section 2.5 for an explanation of precession.

order contributions; models which include orders up to 3.5PN match very well ($> 99\%$) with NR simulations for the inspiral phase of most BBH [45, 39]. Disparity creeps up for a class of BBH with both high spins and high mass ratio ($m_1 \gg m_2$), where high-order *tail effects* (back-scattering of linear waves in the source's curved spacetime) begin to dominate [51, 50].

2.2 Merger

At some point during the late inspiral, the assumption of quasi-circular orbit underpinning the PN expansion may no longer be applied. As the BBH emits increasingly powerful radiation, its orbital decay accelerates. This leads to a huge expulsion of energy over a very short timescale. In GW150914 the remnant black hole was mea-

sured to be some $3M_\odot$ less massive than the total mass of its BBH ‘‘parents’’ [8]. Mergers of BBH with higher q will lose less mass [52, 20, 21], but of a similar order. For all cases the PN expansion becomes inappropriate somewhere shortly before merger, as the orbit loses energy too rapidly for the assumption of quasi-circularity to hold and spacetime curvature can no longer be accurately approximated as linear [30].

The velocity of the black holes peaks as they are simultaneously rapidly orbiting each other and rapidly losing orbital energy, causing them to accelerate. The GW amplitude peaks as they pass each other’s *innermost stable circular orbit* (ISCO), where they are undergoing the greatest acceleration whilst still independent bodies [29]. The first detection’s source produced its peak radiation at a frequency of $\omega_{orb} \approx 150\text{Hz}$ and relative velocities $v/c \approx 0.6$ [8]. In this period, both amplitude and phase of the waveform evolve rapidly, thus any attempt to model waveforms for this stage must be more finely-tuned compared to the inspiral to maintain the same level of conformance to NR simulations. Moreover, non-linear GR effects must be considered, dramatically inflating the computational time required to evolve a simulation with the needed precision.

Whilst performing NR simulations for an entire waveform is impractically costly, techniques have been developed where just the merger phase of a waveform is simulated in NR and ‘‘stitched’’ to both the preceding inspiral phase and succeeding ringdown phase, which are modelled independently. This is described in more detail in Section 2.4.

Modelling the merger numerically involves describing the system’s spacetime as a grid where the Field Equations are described in generalised harmonic coordinates and discretely applied to each unit [53]. Suppose any given unit is of length, l , and the BBH merger’s spacetime metric evolves on a timescale, \mathcal{T} , over a lengthscale \mathcal{L} . The resolution of the grid is increased in specific areas so that all units satisfy $l \ll \mathcal{L}$. The Field Equations are solved for each unit and advance the simulation in time increments, t , always satisfying $t \ll \mathcal{T}$.

Given that an NR simulation uses analytical solutions to the Field Equations for each segment in the spacetime mesh, errors are caused by insufficient grid resolution in the mesh or inaccuracies arising from deriving a solution for an observer fixed at a finite point. Both sources of error are well understood and their effects identifiable [53]. With appropriate treatment and sufficient computation time, NR simulations can incur arbitrarily small error. This allows them to be used for calibrating the less accurate, more computational efficient PN models in their applicable domain (the inspiral) and similarly for the models which approximate the ringdown phase

using perturbation theory, discussed in the next section.

2.3 Ringdown

Immediately following merger, the remnant black hole is in a dynamically unstable state. The angular momentum is not symmetrically distributed about the axis of rotation, thus the spacetime acts to bring itself towards equilibrium. The angular momentum redistributes itself in a series of oscillations, which can be modelled via perturbation theory [31]. The oscillations cause GW emission, radiating away some of the angular momentum and so acting to damp the oscillations. The overall effect is for post-merger GW emission to exponentially decay like the eponymous ringing down of a bell. Perturbation theory models the GW emission in the form of harmonic quasi-normal modes, with the $(2, 2, 0)$ mode as the fundamental. Models have been built which include modes beyond the fundamental, but they all suffer computational caveats, and models which only encompass the $(2, 2, 0)$ mode generally return errors of less than 1% [54].

2.4 Stitching the Pieces Together

Models which produce a complete waveform must incorporate all three phases. To do this they carry out the stitching process briefly described in Section 2.2 for both the transition from inspiral to merger, and the transition from merger to ringdown. The inspiral-merger stitching process involves matching the inspiral simulation with the NR simulation over a segment of the waveforms where the PN approximation is reliable (before merger), and choosing a frequency or time before merger where the two simulations cohere to transition from describing the waveform via the PN model to describing the waveform via the NR model. Similarly, the merger-ringdown stitching process matches a portion of the NR simulation immediately after the coalescence with the ringdown simulations and find an appropriately well-matched point in the waveform evolution to seam the two waveforms together [55].

Precession increases difficulty in the stitching process, as the waveform shape is dependent on more parameters and the orbital phase evolution is less linear (as we shall see in section 2.5) [56]. NR simulations have been carried out for a variety of different BBH sources, generally of mass ratio $q < 12$, such as the catalog described in *Mroué et al. (2016)* [56].

2.5 Precession

Post-Newtonian effects include spin-orbit and spin-spin coupling, as described in Table 2.1, which are interactions between the orbital angular momentum (Eq. 2.29) and spin angular momenta (Eq. 2.30) of the two black holes [57]. The angular momenta are defined in geometric coordinates ($G = c = 1$) as

$$\mathbf{L} = \mu M^{1/2} r^{1/2} \hat{\mathbf{L}}, \quad (2.29)$$

$$\mathbf{S}_i = \chi_i m_i^2 \hat{\mathbf{S}}_i, \quad (2.30)$$

where $i = 1, 2$, r is the orbital separation and χ is the dimensionless spin parameter $0 \leq \chi \leq 1$. A spin of $\chi = 1$ represents a spin at the extreme Kerr limit, where the mass has the maximum spin angular momentum possible.

The PN terms contributed by these effects include a correction to the orbital angular momentum vector, $\hat{\mathbf{L}}$ and both spin angular momentum vectors, $\hat{\mathbf{S}}_1, \hat{\mathbf{S}}_2$. The physical result is a rotation of $\hat{\mathbf{L}}$ about $\hat{\mathbf{J}}$, where $\hat{\mathbf{J}}$ is the total angular momentum vector:

$$\hat{\mathbf{J}} = \frac{\mathbf{J}}{|\mathbf{J}|} \equiv \frac{\mathbf{L} + \mathbf{S}}{|\mathbf{L} + \mathbf{S}|}, \quad (2.31)$$

where $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$. The precessional evolution of \mathbf{L} was first formulated and described in *Apostolatos et al* (1990)[57]). Eq. 11a of that paper gives the evolution of \mathbf{L} to 2PN order:

$$\begin{aligned} \dot{\mathbf{L}} = & \frac{1}{r^3} \left[\frac{4m_1 + 3m_2}{2m_1} \mathbf{S}_1 + \frac{4m_2 + 3m_1}{2m_2} \mathbf{S}_2 \right] \times \mathbf{L} \\ & - \frac{3}{2} \frac{1}{r^3} [(\mathbf{S}_2 \cdot \hat{\mathbf{L}}) \mathbf{S}_1 + (\mathbf{S}_1 \cdot \hat{\mathbf{L}}) \mathbf{S}_2] \times \hat{\mathbf{L}} \\ & - \frac{32}{5} \frac{\mu^2}{r} \left(\frac{M}{r} \right)^{5/2} \hat{\mathbf{L}}, \end{aligned} \quad (2.32)$$

where we recall the total mass and reduced mass: $M = m_1 + m_2$ and $\mu = m_1 m_2 / M$ respectively. The first term is the result of spin-orbit coupling whilst the second term due to spin-spin coupling. Notice that the final term is the only one which evolves the magnitude of \mathbf{L} . This is the *radiation reaction* term, which we can see is independent from \mathbf{S} , permitting the “twisting up” process `IMRPhenomPv2` employs on a non-precessing model to produce accurate precessing waveforms [46]. Equations for $\dot{\mathbf{S}}_1$ and $\dot{\mathbf{S}}_2$ are also given in [57], reproduced in Eq. 2.33 which contain terms

only which evolve their direction:

$$\begin{aligned}\dot{\mathbf{S}}_i &= \frac{1}{r^3} \frac{4m_i + 3m_{3-i}}{2m_i} \mathbf{L} \times \mathbf{S}_i \\ &+ \frac{1}{r^3} \left[\frac{1}{2} \mathbf{S}_{3-i} - \frac{3}{2} (\mathbf{S}_i \cdot \hat{\mathbf{L}}) \hat{\mathbf{L}} \right] \times \mathbf{S}_i\end{aligned}\quad (2.33)$$

where $i = 1, 2$. As in Eq. 2.32, the first terms are due to spin-orbit coupling and the second terms are due to spin-spin coupling. These evolutions are dominated by a $1/r^3$ coefficient on all terms, thus we can deduce that as a BBH inspirals, the rate of precession increases as exemplified in Figure 2.4. This is *orbital precession* - the subject of the next section.

2.5.1 Orbital Precession

Whilst we know angular momentum is lost to radiation, the assumption of quasi-circular orbit implies $\dot{\mathbf{J}} \approx 0$ over an orbital period. Therefore a rotation of \mathbf{L} necessitates an equal counter rotation of $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$ (see Eq. 11 of *Apostolatos et al. (1994)* [57]). When \mathbf{S} has a significant component perpendicular to \mathbf{L} , the rotation of \mathbf{L} about \mathbf{J} describes a cone, henceforth referred to as the *precession cone*. The opening angle of the precession cone is defined as

$$\beta = \arccos(\hat{\mathbf{J}} \cdot \hat{\mathbf{L}}). \quad (2.34)$$

Since instantaneous gravitational wave emission is fundamentally quadrupolar, a fixed distant observer will receive radiation modulating in both amplitude and phase, as the “line-of-sight”, $\hat{\mathbf{N}}$, varies in inclination to $\hat{\mathbf{L}}$. In Figure 2.3 the relative amplitude of a typical CBC source’s GW emission is plotted against inclination to $\hat{\mathbf{L}}$, with the vectors of $\hat{\mathbf{J}}$, $\hat{\mathbf{L}}$ and $\hat{\mathbf{N}}$. Observe that near *edge-on* to the orbital plane ($\iota \approx 90^\circ$), the relative amplitude decreases for both polarisations, but at different rates. The maximum amplitude is observed when *face-on* to the source ($\sin \iota = 0$). The graph is in the co-rotational frame of $\hat{\mathbf{L}}$, thus the fixed external observer is seen to follow a circular trajectory. In doing so, the observer witnesses a fluctuation in the received GW amplitude. Another effect of precession on the GW observed by $\hat{\mathbf{N}}$ is not obvious in Figure 2.3. The precession cone would also act to “speed up” and “slow down” the orbit from the observer’s perspective, causing modulations in the phase of the received waveform.

A comparison between a non-precessing and precessing waveform is given in

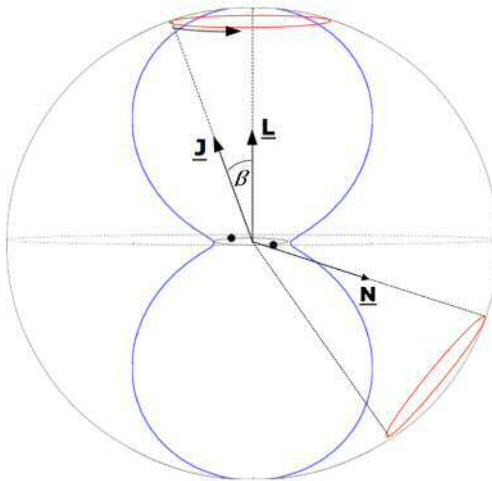


Figure 2.3: An example of the relative amplitude vs inclination (blue line) of instantaneous GW emission for an arbitrary polarisation, in the non-inertial reference frame of $\hat{\mathbf{L}}$. The black circle represents the unit sphere. In this co-rotational frame, we see $\hat{\mathbf{J}}$ precessing about $\hat{\mathbf{L}}$. To an inertial distant observer, $\hat{\mathbf{L}}$ would be seen to precess about the constant \mathbf{J} , describing a cone of identical shape. Line-of-sight $\hat{\mathbf{N}}$ is also plotted with its trajectory on the unit sphere describing a circle identical in size to that of the precession cone.

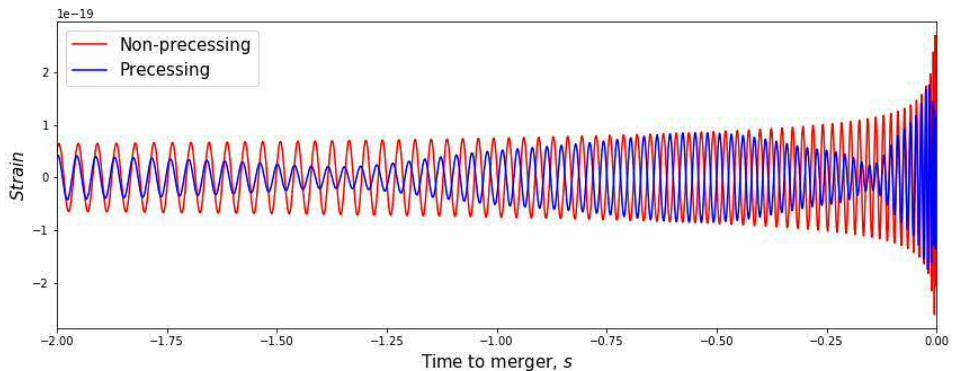


Figure 2.4: A precessing inspiral waveform overlaid with a non-precessing inspiral waveform from a source with the same masses and aligned spin. In both cases, the orbital frequency is seen to increase throughout. The precession effects upon the amplitude and phase are visible, as is the acceleration of the rate of precession towards merger.

Figure 2.4. An undulation of mean amplitude is clearly visible in the precessing waveform. The phase modulations in the former case is also apparent: notice that at $t = -2s$, the two waveforms are in phase with each other. By $t = -0.75s$ they are entirely out of phase, and in the last few discernible orbits before merger they return to being coherent. The amplitude of a waveform is difficult to precisely identify due to noise, whilst the orbital phase evolution can be tracked to a fractional accuracy of $\ll 10^{-3}$, [58]. For this reason, the amplitude oscillations of an observed precessing waveform in fact do not significantly contribute to the signal

power compared to that of the orbital phase modulations.

2.5.2 Transitional Precession

Apostolatos et al. in [57] describes another possible form of precession, which they named *transitional precession*. This phenomenon occurs when a BBH which first satisfies $|\mathbf{L}| > |\mathbf{S}|$ loses enough orbital angular momentum to satisfy $\mathbf{S} \approx -\mathbf{L}$. This requires \mathbf{S} to be high in magnitude and approximately anti-aligned with \mathbf{L} .

A BBH which experiences a transitional precession epoch “loses its gyroscopic bearings” and undergoes a rapid swing as $\hat{\mathbf{J}}$ reacquires equilibrium close to $\hat{\mathbf{S}}$. There are three criteria which must be satisfied for transitional precession to occur in the detectable region of a waveform:

- $\hat{\mathbf{S}} \approx -\hat{\mathbf{L}}$,
- the inspiral begins with $|\mathbf{L}| > |\mathbf{S}|$,
- $|\mathbf{L}| \approx |\mathbf{S}|$ becomes satisfied in the late inspiral.

Whilst the spin vector distribution in black hole populations is not well known, theories and detections to date suggest the first criterion is rarely satisfied [59]. The second criterion is true for most waveforms (simple arguments can confirm this using Eqs. 2.29 and 2.30), but combined with the final criterion it restricts the phenomenon to CBC with high mass ratios.

2.5.3 Eccentricity

The effects of eccentricity are disregarded in the derivations of the quadrupole formula and PN expansion as described in Section 2.1. Including eccentricity requires solving the EFE without the assumption of quasi-circular orbits [27], introducing considerably more PN terms. Waveform models have been constructed which include parameters for eccentricity [60, 61], with the effect of considerably increasing the number of theoretically distinguishable, observable systems.

Back-reaction from GW emission acts to circularise BBH orbits on timescales much shorter than the inspiral phase [27], implying that BBH coalescences of eccentric sources could only occur due to gravitational influences from external sources in the inspiral phase [62]. Possible causes for such events have been theorised and modelled [63, 64]. However, including such waveforms in search algorithms would act to reduce overall search sensitivity unless it is conducted with a targeted approach towards specific sources with eccentricity [65]. For this reason, eccentricity is not considered in this investigation.

Chapter 3

Observing Precession in the Waveform

3.1 Modelling Precession

There are two main classes of BBH waveform approximants: those which adopt phenomenological modelling techniques, and those which use the “effective one-body” approach, treating the BBH as a single body and tracking the motion of test particles on an external metric [66]. This investigation uses phenomenological waveforms, due to their quicker computation time [45, 46]. Phenomenological modelling involves empirically tailoring a model to describe observed phenomena consistent with theoretical predictions, without being derived from theory itself. In the case of BBH GW, phenomenological models are constructed in a manner which minimises the number of parameters required to fully describe the waveform [44]. This is done by combining several waveform parameters into one which models the effects of those parameters. In this way, degeneracies are removed and calculations become much less computationally expensive. Calibrations with NR waveforms are conducted over the parameter space in which the approximant is desired to be reliable.

The model used to synthesise all waveforms in this investigation is `IMRPhenomPv2` [32], which constructs waveforms in the fourier domain (frequency) rather time-evolving the waveform. `IMRPhenomPv2` constructs precessing waveforms by “twisting up” non-precessing waveforms whose inspiral rate is the same, treating precession effects as corrections to a basic waveform shape [46]. This relies on the fact that precession terms to first order do not affect the power radiated. Since the precession terms satisfy $|\dot{\mathbf{L}}| = 0$, instantaneous GW emission is of equal magnitude for a given polarisation irrespective of the amount of precession.

3.2 Waveform Parameters

The strain received by an observer of a BBH coalescence is described entirely by 15 parameters. Eight of these parameters are the physical attributes of the black holes themselves: their masses (m_1, m_2) and spin angular momenta ($\mathbf{S}_1, \mathbf{S}_2$). These shall collectively be referred to as the *intrinsic* parameters. The remaining seven relate to the observer: sky location (*latitude* and *longitude*), orientation angles (polarisation ψ ; phase ϕ ; inclination ι), time, t , and distance, d . These shall be referred to as the *extrinsic* parameters. A BBH approximant takes values for all parameters at a chosen point in the waveform; either when a given orbital frequency is reached, or at the instant of peak strain [44].

3.2.1 Intrinsic Parameters

The black hole masses, m_1, m_2 , will always be defined in this investigation such that $m_1 \geq m_2$ in order to omit degeneracy between waveforms where the masses are identical but reversed. Using total mass, $M = m_1 + m_2$, and mass ratio, $q = m_1/m_2$ with this criterion averts the degeneracy between waveforms where both masses are the same in each case, but our parameters label them the opposite way round. An alternative method is used in the PN equations (implied by the PN expansion (Section 2.1)), where chirp mass, M_c , and the symmetric mass ratio, $\eta = m_1 m_2 / M^2$, are the two mass parameters.

The spin vectors, $\mathbf{S}_1, \mathbf{S}_2$, of the two black holes include six independent parameters with numerous degeneracies. Waveform models reduce this number by combining the spin components aligned with $\hat{\mathbf{L}}$, and combining the remaining components, which all act in the orbital plane. This shall be described in detail in Chapter 3.

3.2.2 Extrinsic Parameters

Time, t , dictates at what point in time all the other parameters are given their prescribed values. Some parameters are constant throughout the waveform, but others oscillate or vary and so the point in the waveform at which they are given a certain value must be set. This is typically done via prescribing the parameters at a given orbital frequency and setting the time of coalescence, t_c , as a reference time for the rest of the waveform.

Distance, d , from the observer to the source scales the waveform amplitude as $1/d$. This is, of course, uniform over the full duration of the waveform and so irrelevant to the effects of precession. However it does contribute to the sensitivity of the detectors to a given signal; for a given waveform, there will be a horizon distance

beyond which the signal cannot be detected. This can be used for calculations of *effective volumes* of sensitivity towards waveforms in a certain class (see [67]).

Polarisation, ψ , dictates the contribution of the two independent GW polarisations to the amplitude of the observed strain at time, t . Figure 3.1 shows the instantaneous emission amplitude profile of plus-polarised and cross-polarised waves as a function of inclination. Notice that no cross-polarised radiation is emitted in the orbital plane.

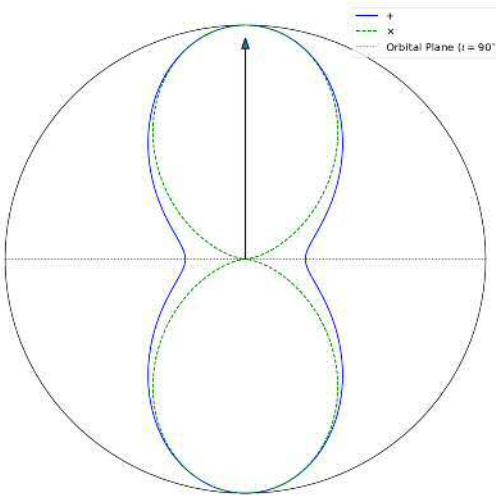


Figure 3.1: Relative amplitude (to first-order) for instantaneous GW emission vs inclination, ι , for plus-polarised (+) and cross-polarised (x) waves. The arrow represents $\hat{\mathbf{L}}$, the black circle represents the unit sphere.

Inclination, ι , between $\hat{\mathbf{L}}$ and the observer's line-of-sight, $\hat{\mathbf{N}}$, affects the amplitude of the received waveform. The amplitude profile over $\iota = [0, \pi]$ differs for the two independent polarisations, as can be seen in Figure 3.1, and so the amplitude of the polarisation being observed at any time is a function of both ι and ψ . From Figure 2.3 we know an observer of a precessing BBH sees an oscillation of polarisation and inclination over the course of a waveform. The two parameters ι and ψ together dictate where on the precession cone \mathbf{L} is pointed at time, t , relative to the observer. That is to say, in a precessing system ι, ψ do not remain constant but fluctuate in accordance with the precession cone and its orientation relative to the observer.

Sky location (right ascension, *R.A.*, and declination, *Dec.*) is degenerate with ϕ and ψ . Specifically, sky location dictates the orientation of the detector with respect to the source, which is equivalent to a transposition of the orientation of the

source. As a detector’s sensitivity is dependent upon the GW polarisation, it is also dependent on sky location. The detector response to a source is also dependent on sky location, since the direction of propagation affects the magnitude of arm-length oscillations as described in Section 1.2.2.

Orbital phase, ϕ , is a measurement of how far through one period the has BBH orbited, where the origin ($\phi = 0$) is defined in relation to the line-of-sight dependent upon the approximant being used. Some approximants set the orbital phase at the same time or frequency as the other parameters, whilst others set it as the coalescence phase.

3.3 Appropriate Parameters to Measure Observed Precession

In principle, estimates for all fifteen parameters described in Section 3.2 could be constrained for a detected signal with a high enough SNR. However, with current detector sensitivities signals are rarely going to be so loud [13], and there are a number of degeneracies between many of these parameters which cause many waveforms from different sources to look very similar. These degeneracies can be utilised to reduce the number of parameters needed to describe a complete waveform in order to increase the computational efficiency and sensitivity in GW search algorithms. This section describes the formulation of some such alternative parameters which are useful specifically for investigating precession.

3.3.1 Reduced Spin Parameters

The initial conditions of a CBC waveform are set in a coordinate system where the $\hat{\mathbf{z}}$ -axis is defined as $\hat{\mathbf{z}} \equiv \hat{\mathbf{L}}$. That is to say, all aligned spin is contained in the $\hat{\mathbf{z}}$ -component. If the black holes’ spin-vectors have non-zero S_x and S_y components, the precession cone must have a non-zero opening angle. The spin components aligned with the orbital momentum vector, S_{1z} and S_{2z} , will induce an effect of their own, decreasing or increasing the rate of inspiral, dependent upon whether the mass-weighted addition of the spins are aligned ($\uparrow\uparrow$) or anti-aligned ($\uparrow\downarrow$) with $\hat{\mathbf{L}}$ [68]. These effects and those of precession are entirely independent.

Regardless of spin magnitude or to which mass the spin belongs, the perpendicular components are exclusively significant for calculating the precession. Thus, it is possible to reduce the number of parameters defining the spins of the black holes by combining the two “parallel” spin components [S_{1z} , S_{2z}] and, independently, com-

bining the four “perpendicular” spin components (S_{1x}, S_{1y}, S_{2x} and S_{2y}). This has been done in previous work [44, 32], in which two dimensionless spin parameters were defined and have subsequently been used to accurately approximate the effects of the full 6-component spin-vector parameter space. Such a model (which only adopts first order spin-spin coupling terms) does forfeit some parameter estimation capabilities, but induces minimal inaccuracies when modelling inspiral waveforms.

The first of the two reduced spin parameters considers only the spin components aligned with the orbital momentum: S_1, S_2 . By first defining aligned dimensionless spin for black hole, i :

$$\chi_i = \frac{\mathbf{S}_i \cdot \hat{\mathbf{L}}}{m_i^2}, \quad (3.1)$$

the aligned spin parameter can be defined as

$$\chi_{eff} = \frac{m_1 \chi_1 + m_2 \chi_2}{M}. \quad (3.2)$$

where $M = m_1 + m_2$ is the total mass of the BBH system [42]. The other reduced spin parameter encompasses all four perpendicular spin components:

$$\chi_p = \frac{\max(B_1 S_{1\perp}, B_2 S_{2\perp})}{B_2}, \quad (3.3)$$

where the masses are chosen such that $m_1 > m_2$ and $S_{i\perp} = (S_{ix}^2 + S_{iy}^2)^{1/2}/m_i^2$ [32]. $B_i = 2 + (3m_{3-i})/(2m_i)$ is the PN-corrected mass-weighting for precessional spin-orbit coupling [57].

The aligned spin parameter, χ_{eff} , stores the magnitude of aligned-spin effects upon the orbit. The two spins are mass-weighted (with PN-corrections) and summed. The second new dimensionless parameter, χ_p , stores the magnitude of precession. The mass-weightings also act to normalise the parameters to the respective ranges of $-1 \leq \chi_{eff} \leq 1$ and $0 \leq \chi_p \leq 1$. Negative values of χ_{eff} represent an aligned spin anti-parallel to \mathbf{L} . After the precessional spins are mass-weighted the maximum value is adopted, ignoring the contribution from the other black hole entirely. This approximation is, in effect, an average over one orbital period. The less significant $S_{i\perp}$ component spends equal time in alignment with the greater $\mathbf{S}_{i\perp}$ as it does in opposition to it over one orbit, and so its time-averaged contribution is zero (to first order) [46].

Current waveform models use an aligned spin parameter with an additional term to χ_{eff} :

$$\chi_{PN} = \chi_{eff} - \frac{38}{113} \eta (\chi_1 + \chi_2), \quad (3.4)$$

recalling the symmetric mass ratio, $\eta = m_1 m_2 / M^2$, which adds the contributions from next-to-leading order PN terms (see Eq. 1.58 in [42]). The waveform approximant we use in this investigation adopts the additional term in its simulation, but it contributes little towards the parameter for the regions of parameter space this thesis investigates (less than 5% for all waveforms). For clarity, we refer to the value of effective aligned spin using χ_{eff} .

3.3.2 Switching to $\hat{\mathbf{J}}$ as Reference for Orientation

Precession is not observable in a single orbit, as its effects are only apparent when the waveform is long enough for the precession cycle to have progressed (introducing phase and amplitude modulation). The closer a parameter is to constant throughout a waveform, the more precisely its relationship with observed precession can be assessed. Due to the fact that ι is defined as the inclination from $\hat{\mathbf{L}}$, it oscillates throughout a precessing waveform. The two other source orientation parameters, ϕ and ψ , also track with $\hat{\mathbf{L}}$.

A description of the orientation of the source comprising of more inert coordinate parameters would help us understand their physical meaning. Such parameters can be facilitated by drawing them from $\hat{\mathbf{J}}$, rather than $\hat{\mathbf{L}}$. We define the angle between observer and $\hat{\mathbf{J}}$:

$$\theta_{JN} = \hat{\mathbf{J}} \cdot \hat{\mathbf{N}}, \quad (3.5)$$

which acts much like ι whilst having the benefit of being invariant through the precession cycle. Gravitational radiation draws energy from the orbit, thus $\dot{\mathbf{L}} < 0$, and as negligible energy is lost from \mathbf{S} , total angular momentum must also evolve: $\dot{\mathbf{J}} < 0$. The laws of conservation demand that the energy emitted via GW is equal to the orbital energy lost, and the energy lost does not affect $\hat{\mathbf{J}}$, from which we can infer $\dot{\mathbf{J}} \propto \mathbf{J}$. It is therefore appropriate to assume $\dot{\theta}_{JN} \approx 0$ holds true throughout the waveform for all CBC sources.

The precession cycle causes undulations in the evolution of ϕ , since the parameter is being affected by two rotations simultaneously - the orbit and the precession. An elegant way to separate these effects is to use a parameter describing the precession phase, i.e. how much of one precession cycle $\hat{\mathbf{L}}$ has completed. The precession phase is defined as the angle between a fixed unit vector, $\hat{\mathbf{x}}$, and the projection of $\hat{\mathbf{J}}$ into the orbital plane [32]:

$$\phi_{JL} = \hat{\mathbf{x}} \times (\hat{\mathbf{J}} \times \hat{\mathbf{L}}). \quad (3.6)$$

IMRPhenomPv2 defines $\hat{\mathbf{x}}$ as the positive vector (rightwards) in the orbital plane perpendicular to the projection of $\hat{\mathbf{N}}$ into the orbital plane. Orbital phase remains

a necessary independent parameter to describe a waveform, since ϕ_{JL} alone cannot define the position of the masses on their orbit, but we now have a phase parameter which is independent from orientation.

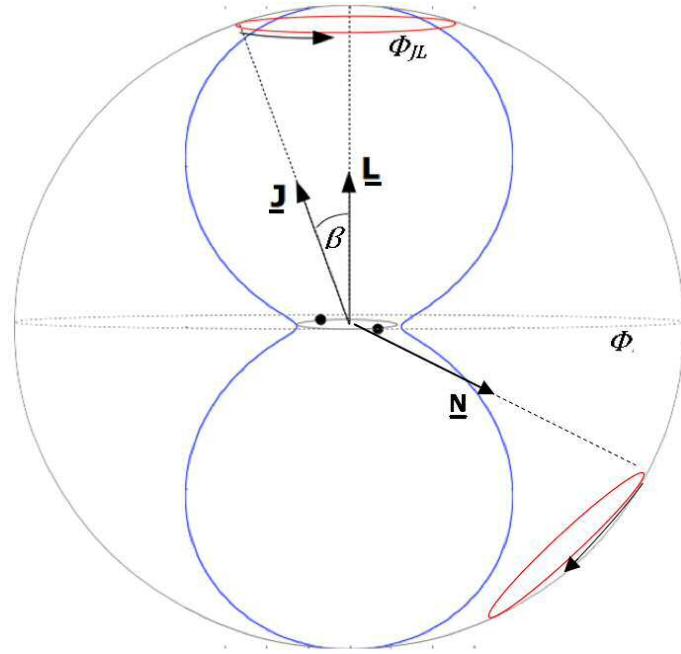


Figure 3.2: The precession of $\hat{\mathbf{L}}$ in a co-rotating frame with the opening angle, β , labelled. A physical representation of the precession phase, ϕ_{JL} , and orbital phase, ϕ , are also labelled.

Figure 3.2 is a similar diagram to Figure 2.3, with the two phase angles labelled along with β . It illustrates the advantages of using ϕ_{JL} : as $\hat{\mathbf{N}}$ fluctuates through ι , he will be rotating with the orbit at some points and against the orbit at others, causing the orbital frequency to appear to oscillate with an amplitude dependent on ι (the modulations described earlier and seen in Figure 2.4). In contrast, the precession rate, $\dot{\phi}_{JL}$, is dependent only on \mathbf{J} .

Polarisation oscillates in a precessing waveform in accordance with the precession cycle, much like the orbital phase. For any change in $\hat{\mathbf{S}}$ corresponding to a rotation about the precession cone, there is an equivalent change in $[\phi_{JL}, \psi]$. The amplitude profile of the two polarisations begin to differ as ι approaches the orbital plane (see Figure 3.1), and so the magnitude of the amplitude fluctuations observed is dependent on the inclinations through which \mathbf{L} is sweeping as it precesses. All of this is well-defined, as ψ is constant without precession, unlike orbital phase. Therefore, ψ need not be adapted or changed for this investigation.

3.4 Matched Filtering

The methods used to detect a GW signal from CBC in data collected by a LIGO detector heavily depend on the process called *matched filtering*. This is a technique generally used to find a specific signal embedded within a noisy data series. It requires a *template*: a synthesised exact copy of the signal intended to be found (created using either NR simulations or PN approximants). The template is treated as a vector, h , a data time-sample (of the same length as the template) is taken as another vector, s , and the inner product is calculated in a cross-correlation function for the two vectors [69]:

$$(h|s) \equiv 4\Re \int_0^\infty \frac{\tilde{h}^*(f)\tilde{s}(f)}{S_n(f)} df, \quad (3.7)$$

where an asterisk denotes complex conjugation, $S_n(f)$ is the one-sided power spectral density and $\tilde{h}(f)$ is the Fourier transform of $h(t)$:

$$\tilde{h}(f) \equiv \int_{-\infty}^\infty e^{2\pi ift} h(t) dt. \quad (3.8)$$

A power spectral density (PSD) describes the frequency-dependent power of the data within which the signal lies. If the detector is particularly noisy in a certain frequency range, the portion of the signal in that frequency range is not going to be as prominent, and so is not as valuable to the inner product. To this end, a PSD effectively weights the power of the frequencies in the signal according to their visibility in the detector data. If waves of a given frequency are loud compared to the PSD for that frequency, they will contribute more to the signal power than other frequencies. Similarly, if waves of a given frequency are no more powerful than others, but they are emitted over a longer duration, those waves will contribute more to the signal power. The PSD is “one-sided” as we restrict our search to positive frequencies.

PSDs are estimated via taking multiple “chunks” of data and analysing its behaviour through all frequencies. Here, we are not investigating a specific observing run, nor are we concerned with the noise characteristics of a detector, and so the PSD employed in this investigation shall be chosen to represent optimal data from a hypothetical Advanced LIGO observing run where the sensitivity has not been tuned for specific astrophysical sources, named the *aLIGO zero-detuned, high-power* (ZDHP) PSD. Figure 3.3 (taken from [70]) shows the noise-frequency curve of the ZDHP PSD along with the theoretically optimal curves for all major noise contributions (from which this PSD is calculated).

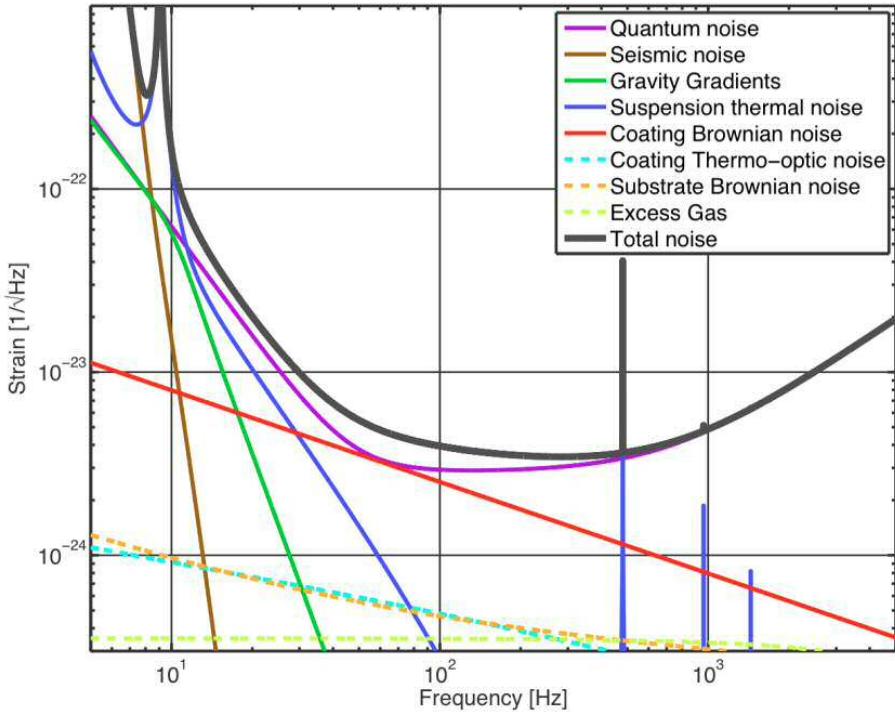


Figure 3.3: Reproduction of Figure 1 from [70]. Noise-Frequency curve for the PSD used for all matched filtering calculations in this thesis (zero-detuned, high power), labelled here as "Total noise". It is calculated as a summation of theorised optimal contributions from all identifiable noise sources, which are also plotted.

For our purposes, a *template* is a waveform produced by an approximant; the properties of the BBH source are known, and the signal is cut off at a specified frequency. In a matched filtering search for a signal, the inner product between data and template is calculated repeatedly, incrementally shifting the coalescence time of the template along the detector data time-series. During this process, the inner product is also maximised over the coalescence phase of the waveform.

The maximum possible signal power retrieved for any signal in a given PSD is equal to the inner product of the signal's perfectly matching template with itself: $(h|h)$ (in effect, this is the magnitude of the template vector). To first approximation, the SNR of a waveform is $\rho(h) \approx \sqrt{(h|h)}$.

Replacing the detector output with another template in Eq. 3.7, the inner product between two waveforms can be written as [71]

$$(h_1|h_2) \equiv 4 \int_0^\infty \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)} df. \quad (3.9)$$

The correlation is maximised over the waveforms' time and phase, and normalised

over the product of the template magnitudes to give the Match [72]:

$$\mathcal{M} = \frac{\max_{t_c, \phi_c}(h_1|h_2)}{\sqrt{(h_1|h_1)}\sqrt{(h_2|h_2)}}, \quad (3.10)$$

where t_c denotes coalescence time and ϕ_c denotes the coalescence phase: $\phi_c = \phi(t_c)$. A match of $\mathcal{M} = 1.0$ indicates the two waveforms are identical. A low Match indicates they are very dissimilar. What constitutes a “low” Match is dependent upon the purposes for which it is being employed.

3.5 Template Banks

A comprehensive search for CBC gravitational waves requires matched filtering to be carried out for all possible signals any CBC can produce. This requires a large number of prepared templates, where a matched filtering algorithm is carried out between the data and those templates to find the best-matching template. If that Match is high enough, the signal is assumed to have been produced by a source with similar parameters to that template. Subsequently, various analyses take place such as false-alarm checks, parameter estimation and error calculation. To summarise, having found the best-matching template, the significance of the detection must be evaluated - if the data Matches well with the template, but also looks characteristic of detector noise around that time, then the likelihood of the signal being of astrophysical origin is significantly reduced. The SNR of a signal is defined as the cross-correlation of the signal with the template, divided by the SNR of the template with itself (using the same PSD for each calculation) [73]:

$$\rho(s) = \frac{(s|h)}{\sqrt{(h|h)}} \quad (3.11)$$

The catalogue of templates used in searching the data for signals is called a *template bank*. Each template in a bank is unique - specific to a source of given parameters.

Template banks employed for BBH searches in the first two Advanced LIGO observing runs cover the whole of non-precessing parameter space for $M \leq 100M_\odot$ and $\chi_{eff} \leq 0.99$ [74]. There is no reason to doubt the existence of higher mass binaries, but searches for such events are often treated separately, both due to the natural trend towards extreme mass ratios in such populations and to restrict the size of the template bank to a manageable level. Additionally, **IMRPhenomPv2** is not calibrated for BBH systems with mass ratios above $q = 12$. For these reasons we

focus only on systems with total masses of $M \sim \mathcal{O}(10M_{\odot})$.

No template bank including precessing BBH waveforms has yet been implemented in search algorithms. Previous studies have concluded that such a search would be too inefficient to carry out a matched filtering process in a reasonable time frame, and would also increase the *false-alarm rate* (the statistical background) so much as to counteract any extra sensitivity gained towards precessing BBH systems [75, 76]. However, it has been shown that a significant proportion of parameter space remains uncovered by non-precessing template banks [67].

The drawbacks of a fully comprehensive search may be overcome if a selective template bank is formed, where precessing templates are included only if they are significantly dissimilar to all templates in the non-precessing region of parameter space (as well as already-accepted precessing templates). This would ensure precessing templates are only added to the bank in areas of parameter space where they will have the greatest impact on the bank's effectiveness. Areas of parameter space where precessing templates are not significantly dissimilar to those of non-precessing waveforms would act to inflate the size of a template bank without contributing much to the overall sensitivity of the search, thus omitting such areas from the bank would help reduce the aforementioned problems concerning precessing template banks. As previously stipulated, this thesis works towards identifying areas in parameter space where observed precession is most prevalent. Further work could incorporate the results found here into analysis on source-specific detector sensitivity in order to create a method of choosing precessing waveforms to include into an optimally filled BBH template bank.

3.5.1 Populating Template Banks

The process of constructing a template bank involves testing a candidate template, and if it does not give a Match value above a predetermined level with any already-accepted templates, it is added to the bank. A template bank is populated by one of two types of algorithm: the *geometric method* and the *stochastic method*. The stochastic method adds templates with randomly chosen parameter values, only rejecting those that match too closely with already-selected templates [77]. The geometric method involves creating a series of two-dimensional hexagonal grids of waveforms in the parameter space (the planes must be in the space of two parameters which exhibit no degeneracy with any other parameters), upon which the rest of the parameter space is filled using the stochastic method [78]. No suitable two-dimensional plane has yet been found to form a geometric grid of templates which includes precession, nor is there any reason to assume such a plane exists - the

number of degenerate parameters is too large for the geometric approach to be applicable [75]. Therefore, any attempt to construct a “precessing template bank” would adopt the stochastic method.

The sensitivity of a matched filtering search is limited by the fitting factor of the template bank being used. A bank uses the criterion $\mathcal{M} \leq \mathcal{M}_{min}$ to assess whether a candidate template is accepted. If the Match satisfies this criterion, the candidate template is accepted into the bank. For a perfect template bank, $\mathcal{M}_{min} = 1$. However, this would require an infinite number of templates - ensuring a perfect Match for any waveform within the designated parameter space. This is, of course, computationally impossible. Lowering the chosen value of \mathcal{M}_{min} reduces the number of templates needed to cover the desired parameter space. However, the lower \mathcal{M}_{min} is set, the more likely it is that false detections are made and real signals are missed. Typically, $\mathcal{M}_{min} = 0.97$ is chosen, since this corresponds to an effective volume of $V_{eff} = V\mathcal{M}_{min}^3 \approx 0.9V$, where V is the total volume of space to which a perfect search is sensitive [79].

In an ideal stochastic template bank, the templates would lie close enough in parameter space so that the Match of any potential signal gives a value $\mathcal{M} \geq \mathcal{M}_{min}$ with at least one adjacent template. In practice, it would be extremely computationally costly to fill a bank in such a way, thus the construction of a template bank stops after an appropriate *fitting factor* is reached (the fraction of total parameter space which gives a Match with at least one template satisfying $\mathcal{M} \geq \mathcal{M}_{min}$). Generally, a bank is taken to be full when its fitting factor, FF , satisfies $FF \geq \mathcal{M}_{min}$ [77, 80]. Naturally, the number of templates required to fill a bank is dependent on \mathcal{M}_{min} and the size of parameter space being covered.

3.6 The Parameter Space of Interest

3.6.1 Relevant Parameter Space

The investigation will cover a number of different combinations of values for m_1, m_2 to identify where in that two-dimensional parameter space precession presents. Mass ratios will not exceed $q = 12$ since `IMRPhenomPv2` is not considered reliable beyond that limit for high masses [32]. Total mass will be kept within order $\mathcal{O}(10M_\odot)$, for reasons given in Section 3.5.

The ranges of $\chi_p \in [0, 0.99]$ and $\chi_{eff} \in [-0.8, 0.8]$ shall be explored, given that $\chi_p = 0.99$ is the calibration limit of `IMRPhenomPv2`. It should be noted that for cases where $\chi_{eff} \neq 0$, a physical limit is placed on χ_p , since the spin magnitude of a black hole cannot exceed the Kerr limit: $|\mathbf{S}| < 1$.

We shall include the full ranges of $[\theta_{JN}, \psi, \phi_{JL}]$ in our investigation, but geometry dictates that for a given source, half the coordinate sphere of $[\theta_{JN}, \psi]$ is a reflection of the other half, so we only take values of $0 \leq \theta_{JN} \leq \pi$.

Sky location affects the polarisation observed and the overall sensitivity (using arguments discussed in Section 1.2.2), and so the difference in sensitivity between two sky locations is dependent upon the orientation of the source. We do not account for detector sensitivity in this investigation - we are only concerned with the observability of precession in waveforms. We also do not consider multiple detector analysis (which would both act to improve sensitivity and dramatically complicate things), and so the effects of sky location upon observed precession are left for future research. The sky location for all waveforms used in this investigation will be directly above the detector.

Time is, of course, needed to track the waveform evolution. But a reference time (e.g. coalescence time t_c , at which the waveform strain peaks) is only relevant if the data and, thusly, noise is time-dependent. This investigation is using a theoretical PSD to Match templates - we are not using real data. Therefore, time has no influence on our results.

Distance, d , to the source would only be relevant if noise and the detectable threshold of a given signal are being assessed. Since the calculation of \mathcal{M} maximises over amplitude (given that the templates are treated as normalised vectors), distance makes no difference in Match calculations. All waveforms constructed in our research are given a distance of $d = 410\text{Mpc}$.

3.6.2 Wave Frequency and Initial Conditions

There is another variable which we must care to consider when matching waveforms. The original set of 15 waveform parameters alone describe all possible CBC waveforms, but they do not specify the “beginning” of a waveform. In order to create a template, an approximant must be given the initial conditions of the source from which which the equations governing its evolution and GW emission can be calculated. The initial conditions are simply the values assigned to the 15 parameters when the waveform begins. Some orientation parameters, such as the orbital phase, are not constant, so the point at which they are assigned may affect the waveform shape and the Match. The gravitational wave frequency, f_{gw} , is a convenient parameter to use to dictate where the waveform begins, as it is hermetically set by the intrinsic source parameters (GW frequency is double the orbital frequency) and has a well-defined albeit non-linear relationship with waveform length. Using f_{gw} to set the initial conditions is also useful for analysing the signal power contribution from

different frequencies. The end of a waveform is cut off during the ringdown once the mean amplitude drops below a given threshold [44].

Carrying out matched filtering searches on template banks whose waveforms are very long would be computationally costly, since more templates would be needed (longer templates naturally lead to smaller deviations in the parameter values having a greater effect on the Match) and the Match calculations taking longer to compute. This would have very little benefit if most of the waveform length is in a region of poor detector sensitivity (see Figure 3.3). Conversely, a short template may not recover the whole signal, leading to artificially inflating the Match if the missed portion of the signal would have significantly deviated from the template in sensitive frequencies. If a template is longer than a visible signal, a portion of the template is redundant in the Match calculation, since the signal power will be dominated by the portion of template which overlaps. To optimise the inner product calculation's faithfulness to true parity between a template and signal whilst minimising computational cost in constructing the template bank, the template must be long enough to cover the whole of the visible signal, but not significantly longer.

The signal power of all GW observations to date have received negligible contributions from the portions of waveforms with wave frequencies lower than 30Hz (e.g. [19, 21]). This is due to the low signal power contribution from these frequencies even though the sources spend a comparatively long time in this frequency region, since the background noise rises rapidly as frequency decreases below $f_{gw} = 30\text{Hz}$. However, the ZDHP PSD used in this investigation represents a much lower noise curve than the current detectors, and so noticeable contributions to the signal power are possible from frequencies lower than 30Hz. Our investigation will predominately focus on waveforms beginning with wave frequencies of 30Hz to approximately cohere with current detector capabilities. The reason we use the ZDHP PSD is so that we may observe differences in observable precession for longer waveforms (where the wave frequency is started at 20Hz) in order to provide further insight. Less sensitive PSDs (such as those derived from real LIGO data) would not permit enough signal power to be recovered in that frequency band to observe a significant difference in the results. Furthermore, we shall see that $f_{gw} = 30\text{Hz}$ sits in the frequency range which contributes the most signal power for many of the sources we investigate, since the comparatively low signal strength is overcompensated by the longer duration spent in these frequencies. Thus choosing this frequency to set the parameter values means we know the parameters at the detector's peak sensitivity. This does not apply for all waveforms, and we shall explore the reasons for the exceptions when we find them.

3.7 Objectives

The Match equation (Eq. 3.10) enables us to quantify the difference between waveforms. We can use it to investigate how much a waveform changes if a parameter value is varied. The main objective of this thesis is to identify relationships between waveform parameters and observable precession in order to find areas of parameter space where precession has the greatest effect on the observed waveform. Using Eq. 3.10 to compare non-precessing waveforms with waveforms of precessing systems allows us to quantify how much precession we would observe if we were to detect such a signal. If a precessing waveform Matches well with a non-precessing waveform, we may deduce that our current template banks are already sensitive to such a signal, and so we need not attempt to improve our methods to increase our sensitivity to such sources. Conversely, if the Match between a precessing waveform and a non-precessing waveform is poor, we may deduce that adding such a waveform to our template banks could improve our overall sensitivity, and crucially, sensitivity towards precessing BBH sources.

Calculating a waveform's contribution to a search algorithm's sensitivity requires more than analysis of Match calculations - sensitivity to a source is also highly dependent on the power of the signal. A template is only worth including in a bank if the volume of space for which a detector would be sensitive to such a source is above a certain threshold, assuming the population density of such a source is also sufficiently high. In this thesis, we aim only to investigate where in parameter space we could find precessing BBH given that the SNR is high enough and attempt to identify trends, leaving the application of specific detector sensitivities for future research.

Using only the “nearest” non-precessing template to the precessing signal also allows us to identify relationships between parameters and observable precession without conflating them with trends caused by the effectualness of the aligned-spin bank recovering precessing signals as it varies across parameter space. Understanding both of these effects independently would be beneficial on the path to creating an effective precessing template bank, but here we choose to focus solely on trying to understand how parameters affect observable precession.

Chapter 4

Results

4.1 Method

We calculate the Match between precessing and non-precessing waveforms, changing the parameter values to cover a section of parameter space. We make the assumption that the non-precessing waveform which produces the best Match with a given precessing waveform has identical intrinsic parameter values besides the non-aligned spin components. In reality, this is not the case. A more rigorous assessment of observable precession where each precessing template is matched filtered against a complete non-precessing template bank would produce higher Matches in some areas of parameter space (for example, the centre graph of FIG. 1 in [67] shows an aligned-spin bank for NSBH coalescences can recover certain highly precessing signals). Such an investigation would produce a more precise set of results indicating where in parameter space precession is undetectable. However, our objectives are to identify trends between waveform parameters and observable precession, whilst maximising over a full template bank would induce other patterns in the results caused by the varying efficacy of the aligned spin bank recovering precessing signals across parameter space, as explained in Section 3.7. The more basic approach used here allows future research to identify those patterns caused by maximising the Match over a template bank, which would otherwise be very difficult to distinguish from the trends caused by the influence on observable precession by individual parameters.

The choice to use the theoretically optimal PSD for LIGO will maximise the signal power and thus the disparity between waveforms. Thus, one may consider the results presented in this thesis as the most optimistic portrayal of the disparity between precessing and non-precessing signals. We may consider the areas of good Matches we find as the minimum proportion of parameter space (in the given do-

main) in which precession cannot be observed by LIGO, as those Matches represent waveforms of negligible observable precession, which our choices have maximised.

For non-precessing waveforms, the orientation parameters (and distance) only affect the observed amplitude. Since we are using an approximant with reduced spin parameters, only three parameters relevant to our investigation affect the Match of a non-precessing waveform: m_1, m_2, χ_{eff} . (More generally, there are four non-precessing parameters which contribute to the Match: $m_1, m_2, \mathbf{S}_1 z, \mathbf{S}_2 z$.) When referring to orientation parameters henceforth, we exclusively mean those of the precessing waveform, where they could affect the Match since the orientation of a precessing source is not constant. All waveforms are generated from an initial frequency of $f_{gw} = 30\text{Hz}$, with their parameter values set at that frequency, unless otherwise specified. We perform Match calculations for precessing waveforms in 2D planes of parameter space where the non-precessing waveforms are given the same values for those intrinsic parameters as the precessing waveforms with which they are being Matched. Each 2D plane will cover ranges of two parameters within the group $[\chi_p, \theta_{JN}, \psi, \phi_{JL}]$, whilst the remaining two parameters are fixed. From these results we will be able to identify the influence those parameters have on observed precession, and also how they influence each other's effects.

Henceforth, all angles will be expressed in degrees for the sake of clarity. The range of \mathcal{M} will be constant for all plots: $[0.42 \leq \mathcal{M} \leq 1]$, to facilitate easy by-eye comparisons. In the few cases where Matches gave values $\mathcal{M} < 0.42$, they are coloured the same as $\mathcal{M} = 0.42$. The results are arranged into four cases. Case 1 investigates Matches for waveforms of zero aligned spin and a single combination of masses: $[m_1 = 20M_\odot, m_2 = 5M_\odot, \chi_{eff} = 0]$. This is our focal set. Cases 2, 3 and 4 investigate the influence of aligned spin, mass ratio and total mass respectively. In order to keep the scope of this investigation manageable, we choose not to change the other intrinsic parameters from its value in our focal set in each case.

4.2 Case 1: Fixed Masses, Zero Aligned-Spin

We begin our investigation by considering BBH waveforms from our focal set: sources with masses of $m_1 = 20M_\odot$, $m_2 = 5M_\odot$ and zero aligned spin, $\chi_{eff} = 0$. These masses were chosen to lie approximately within the 90% confidence interval of a detected BBH coalescence, GW151226 [81], thus it is safe to assume a waveform from such a source is typically observable. All the parameter values are set at the initial frequency of each waveform, $f_{gw} = 30\text{Hz}$.

4.2.1 Variation of Match with Precessing Spin and Orientation

We choose to first investigation 2D planes in $[\chi_p, \theta_{JN}]$, fixing $[\phi_{JL}, \psi]$. An array of Matches were calculated between a template of a non-spinning waveform spanning the full range of θ_{JN} and χ_p , repeating this whilst setting ψ and ϕ_{JL} at different values for each plane. The planes are shown in Figures 4.1 and 4.2.

Immediately, we see that there is a large region where Matches are good on every plot. For $\chi_p \lesssim 0.3$, $\mathcal{M} > 0.97$ except where $\psi \gtrsim 30^\circ$ and $\theta_{JN} \approx 90 \pm 20^\circ$. Very few Matches of $\mathcal{M} < 0.9$ occur below $\chi_p = 0.3$, almost all of which are within $\theta_{JN} = 90 \pm 10^\circ$.

Match falls rapidly as χ_p increases beyond $\chi_p > 0.3$, and the worst Matches are generally centred upon $\theta_{JN} = 90^\circ$. Recall that waveforms beginning edge-on have an inclination of $\iota = 90^\circ$. Depending on the orientation parameters, such cases have a value of θ_{JN} in the range: $90^\circ - \beta \leq \theta_{JN} \leq 90^\circ + \beta$, where β is the opening angle of the precession cone. Therefore, waveforms beginning edge-on would not always be at $\theta_{JN} = 90^\circ$, but the range they cover would be symmetric about that line in these plots. We see the worst Matches following that same pattern, generally in the realm of very high precessional spin. We can also see patches of marginally better Matches within the poorest areas at very high χ_p . However, the plots are not perfectly symmetric about $\theta_{JN} = 90^\circ$, largely due to the non-central positioning of these patches.

Let us now focus on the influence of ϕ_{JL} and ψ . In cases where the waveform starts entirely plus-polarised or entirely cross-polarised (top and bottom rows, respectively), a shift in ϕ_{JL} of 180° acts to reflected the plot through $\theta_{JN} = 90^\circ$.

Comparing plots with polarisations of $\psi = [15^\circ, 30^\circ]$ (the second and third rows, respectively), we see that the symmetry somewhat holds between plots of $\phi_{JL} = 90^\circ$ and $\phi_{JL} = 270^\circ$. Yet there is no such symmetry in the planes for $\phi_{JL} = 0^\circ, 180^\circ$. This suggests a relationship between ϕ_{JL} and ψ which cannot be well understood using these plots. As ψ approaches total cross-polarisation, we see the poorest Matches fall to lower values of χ_p and become more tightly constrained towards $\theta_{JN} = 90^\circ$. We see in the planes where $\psi = 45^\circ$ even a marginal quantity of precessional spin produces a very poor Match, but only when $\theta_{JN} = 90^\circ$. Deviating from $\theta_{JN} = 90^\circ$ only slightly launches the Match to $\mathcal{M} > 0.9$. We see no noticeable influence of ψ upon waveforms near $\sin \theta_{JN} = 0$ (where the observer is inside the precession cone), and this remains the case for any value of ϕ_{JL} . In effect, Matches at $\sin \theta_{JN} = 0$ represent the maximum Match (i.e. minimum observable precession) for a source with a given χ_p .

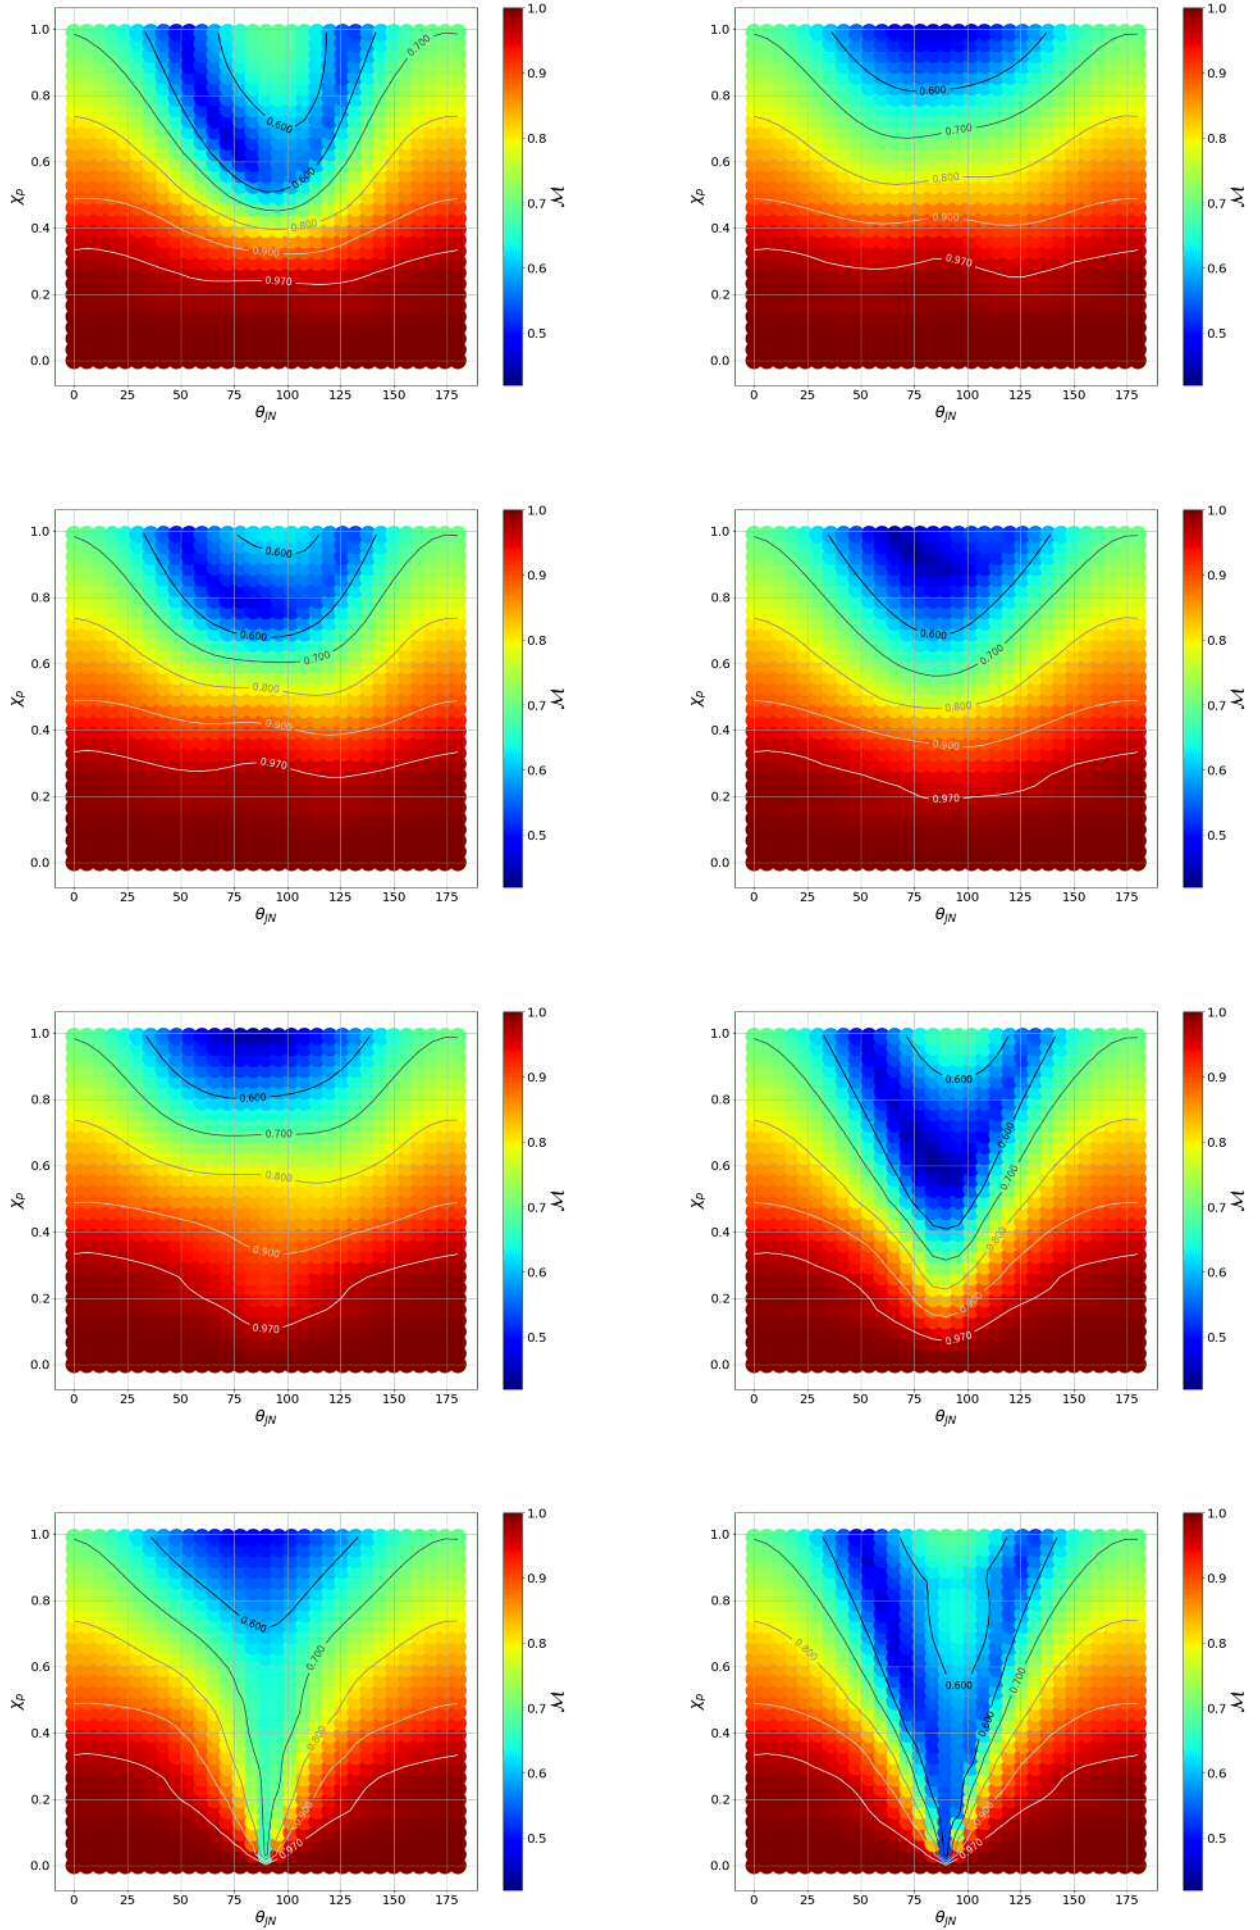


Figure 4.1: Colour plots for \mathcal{M} vs χ_p vs θ_{JN} , with each row showing results with fixed $\psi = [0, 15, 30, 45^\circ]$ downwards and columns $\phi_{JL} = [0, 90^\circ]$ from left to right.

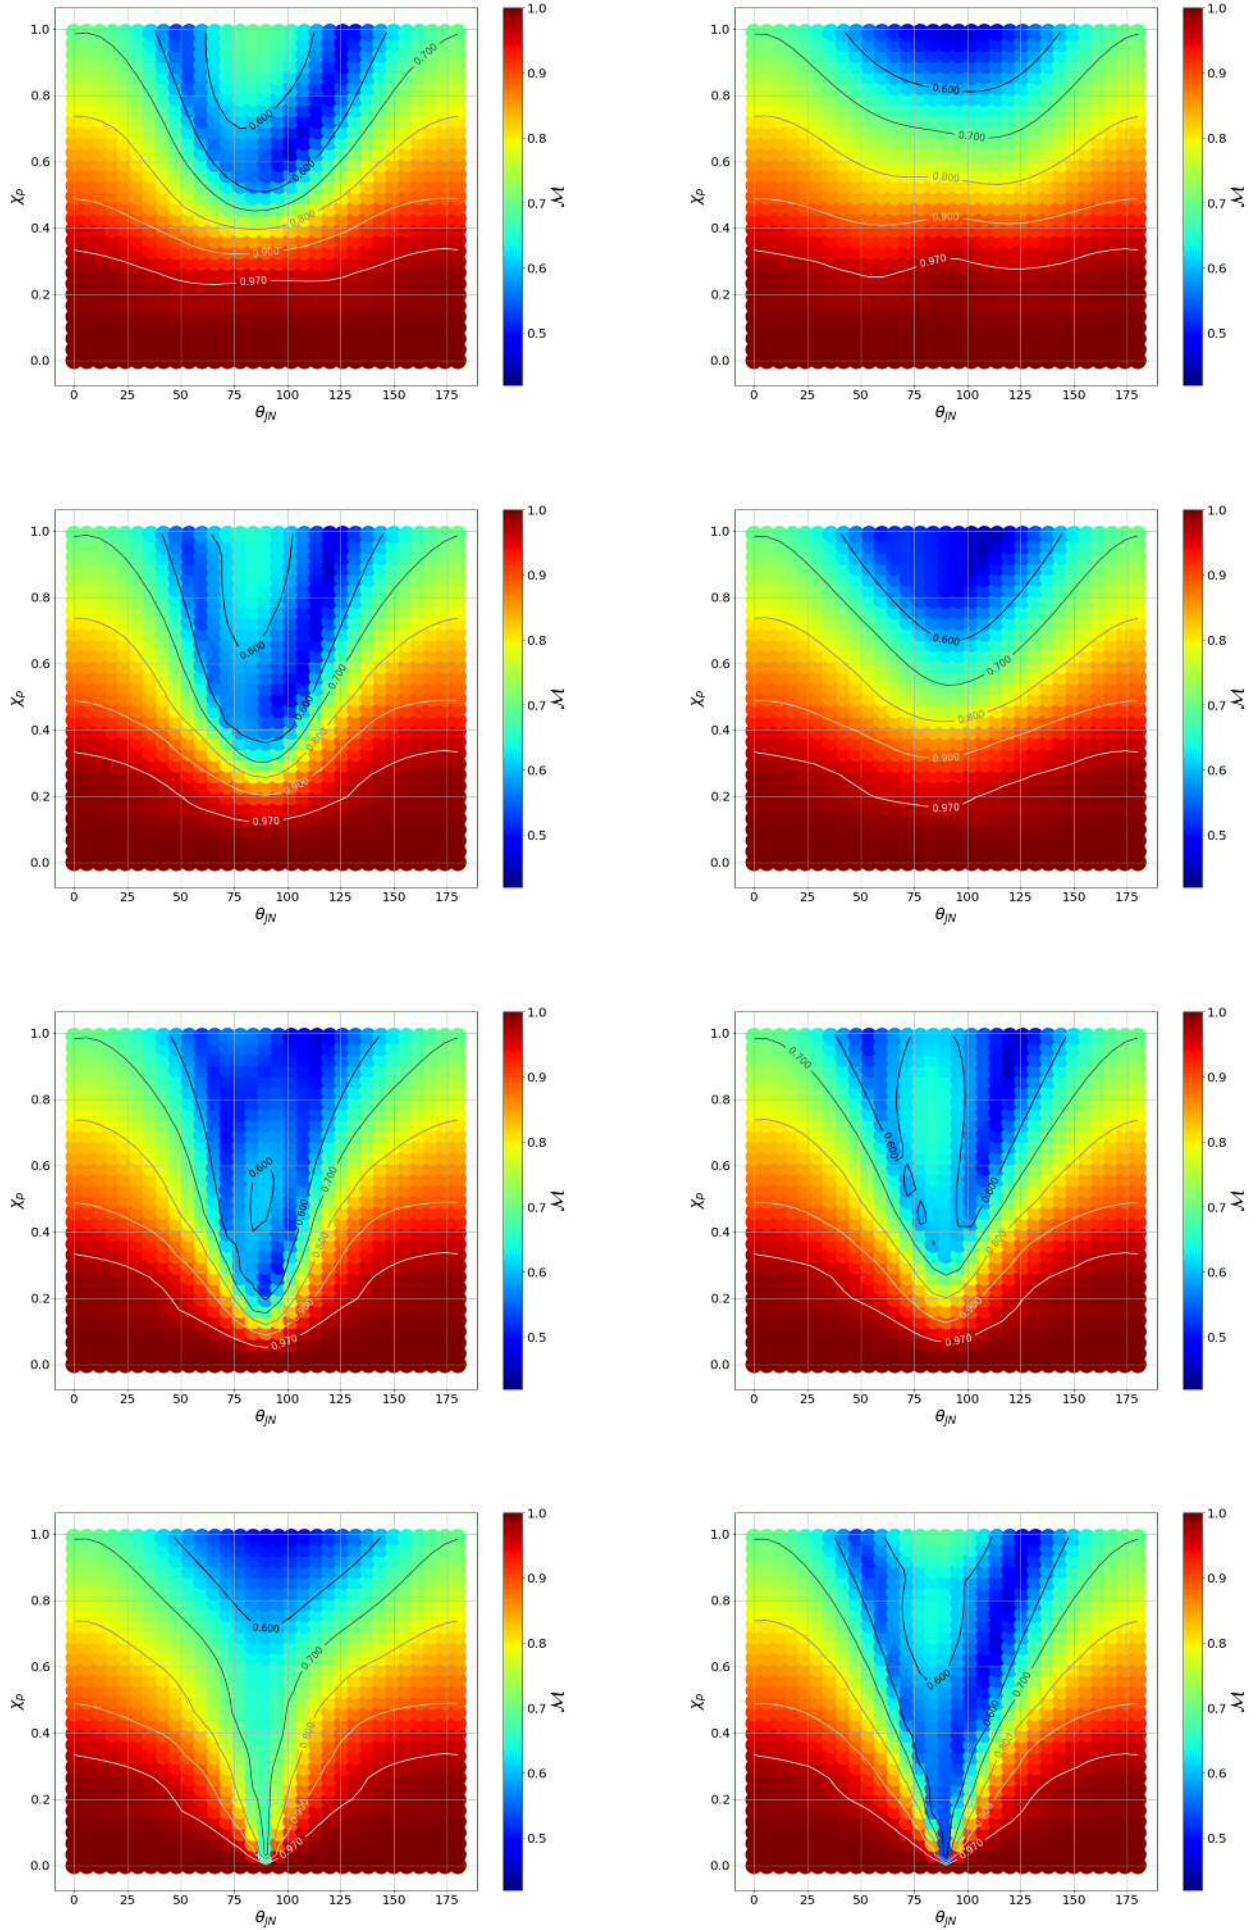


Figure 4.2: Colour plots for \mathcal{M} vs χ_p vs θ_{JN} , with each row showing results with fixed $\psi = [0, 15, 30, 45^\circ]$ downwards and columns $\phi_{JL} = [180, 270^\circ]$ from left to right.

4.2.2 Variation of Match with Orientation for Fixed Precessing Spin

Now we present results keeping χ_p constant in order to better understand the influence on the orientation parameters. Figure 4.3 (overleaf) shows results for four planes of constant ϕ_{JL} in $[\psi, \theta_{JN}]$, where each waveform was given $\chi_p = 0.6$.

The poorest Matches form a ring, where the centre is positioned somewhere on the line $\theta_{JN} = 90^\circ$, at a value of ψ dependent on ϕ_{JL} . The patches of better Matches seen within the group of worst Matches in those Figures now appear within the regions of worst Matches in these plots. These regions of worst Matches form rings, all centred upon $\theta_{JN} = 90^\circ$. Plots for $\phi_{JL} = [180, 225, 270, 315^\circ]$ (not shown) produce the same formation of Matches seen in (a), (b), (c) and (d) respectively, but rotated 180° . This is the same symmetry as seen in Figures 4.1 and 4.2: the reflection in those plots is equivalent to a rotation in a perpendicular plane. Given that we have now fixed χ_p and ψ is being varied, we are viewing such a plane. Planes of $[\theta_{JN}, \psi]$ for values of ϕ_{JL} which differ by 180° represent waveforms which began on opposite sides of the precession cone.

Notice that Matches are $\mathcal{M} < 0.8$ across the entire plane, irrespective of ϕ_{JL} . This coheres with Figures 4.1 and 4.2, where all Matches at $\chi_p = 0.8$ were similarly low. Outside the range $30^\circ \lesssim \theta_{JN} \lesssim 150^\circ$ the Matches become near independent from ψ and never drop below $\mathcal{M} = 0.7$. The lowest Matches forming the ring have a Match of $\mathcal{M} < 0.6$.

The ring of poorest Matches can be understood as a result of the degeneracy between θ_{JN} and ψ in defining the initial direction of $\hat{\mathbf{L}}$ with respect to the observer. Consider Figure 4.3c, where the line $\psi = 45^\circ$ diametrically intersects the circle of poorest Match. At $[\theta_{JN} = 90^\circ, \psi = 45^\circ]$, the observer is orthogonal to $\hat{\mathbf{J}}$. An edge-on observer must be orthogonal to $\hat{\mathbf{L}}$. Therefore, if we define a vector tangent to $\hat{\mathbf{L}}$ in the $[\hat{\mathbf{L}}, \hat{\mathbf{N}}]$ plane, as the source precesses this tangent describes a circle the same size as the precession cone (i.e. a radius of β) centred upon $[\theta_{JN} = 90^\circ, \psi = 45^\circ]$.

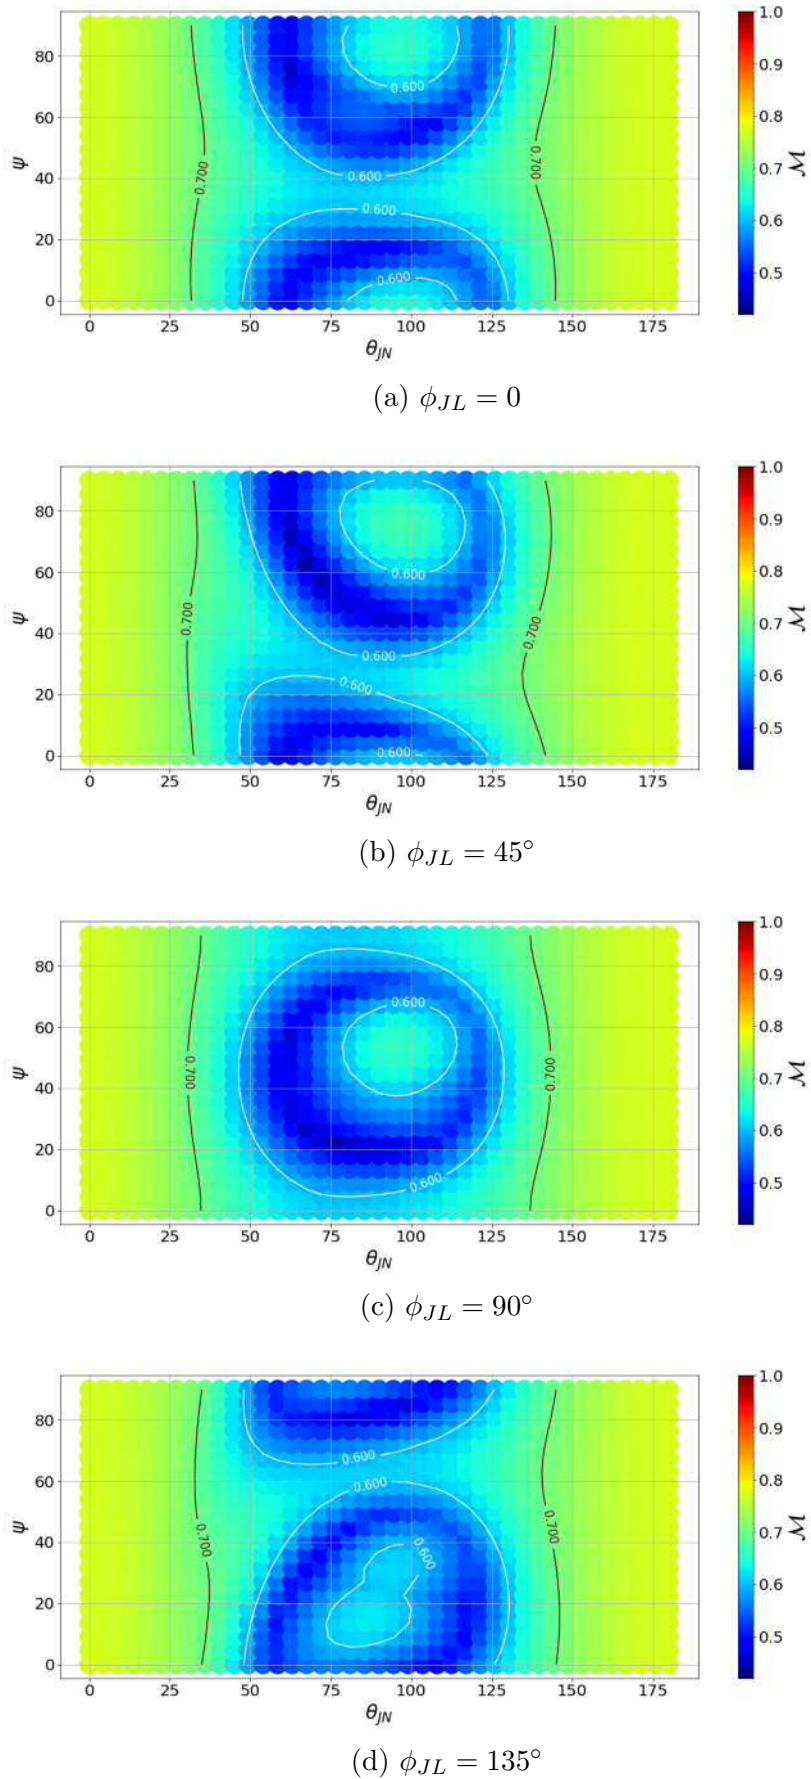


Figure 4.3: Colour plots for \mathcal{M} vs ψ vs θ_{JN} , with $\chi_p = 0.8$ and fixing ϕ_{JL} at the specified values for each plane.

4.2.3 Match and the Opening Angle

Figure 4.4 is a replication of Figure 4.3c with a precession cone *tangent circle* drawn over the top, centred on $[\theta_{JN} = 90^\circ, \psi = 45^\circ]$. This tangent circle represents the group of $[\theta_{JN}, \psi]$ combinations which place the observer initially perpendicular to $\hat{\mathbf{L}}$, and in the course of a full precession cycle, the change in orientation describes the same circle. Two separate circles are plotted: the white circle from a first-order approximation by *Lundgren and O'Shaughnessy* (2014) [82] (an elegant arrangement of Eq. 2.34 using only the magnitude of the spin and orbital angular momenta); the magenta circle for the calculation extracted from the algorithms used by the approximant itself. Lundgren defines the opening angle as

$$\cos(\beta_{Lund}) = \hat{\mathbf{L}}_{\mathbf{N}} \cdot \hat{\mathbf{J}} = \frac{1 + \kappa\gamma}{\Gamma_J}, \quad (4.1)$$

where $\Gamma_J = |S|/|J|$, $\gamma = |\mathbf{S}_1|/|\mathbf{L}|$ and $\kappa = L_N \cdot S$. The algorithms upon which `IMRPhenomPv2` depends to simulate waveforms, the *LIGO Algorithm Library* (LAL) [83], does not explicitly calculate the opening angle, but one can retrieve all the equations necessary to give a 2PN solution:

$$\cos\left(\frac{\beta_{LAL}}{2}\right) = \left(\frac{1 + \frac{1}{\sqrt{1 + s^2}}}{2}\right)^{\frac{1}{2}}, \quad (4.2)$$

where $s = \frac{m_1^2 \chi_p}{L_{2PN} + m_1^2 \chi_{eff}}$ and recalling the symmetric mass ratio $\eta = m_1 m_2 / M^2$ for

$$L_{2PN} = \frac{\eta \left(1 + \left(1.5 + \frac{\eta}{6}\right) v^2 + \left(3.375 - \frac{19.0\eta}{8} - \frac{\eta^2}{24}\right) v^4\right)}{v}.$$

We can see a relatively small difference between the two opening angle calculations, caused by the PN terms for β_{LAL} . The purpose of this comparison is to demonstrate the accuracy of first-order precessing terms.

The lowest Matches lie slightly within the tangent circle, with precession becoming increasingly independent from ψ as the distance from the circle increases. Given that all the points on the circle represent waveforms from sources with the same precession cone, one expects to see very similar Matches for all waveforms which lie on the circle, as is the case.

A *null line* [84] is an orientation of a CBC source which renders a single detector entirely insensitive to the GW it emits. Looking at Figure ??, we can see that one null line lies within the orbital plane, where $\iota = 90^\circ, \psi = 45^\circ$. This orientation would render the detector sensitive to purely cross-polarised waves, which are not emitted in the orbital plane, and so no GW would be received from a source at this orientation. A precessing waveform oscillates in both ι and ψ as it precesses, thus

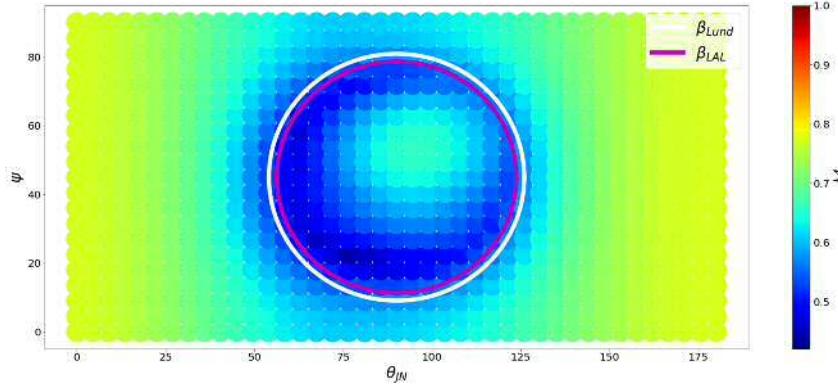


Figure 4.4: $\chi_p = 0.8$, $\phi_{JL} = 90^\circ$. A colour plot for \mathcal{M} vs ψ vs θ_{JN} , showing the same results as Figure 4.3c overlaid with the precession cone's tangent circle as calculated by the LAL formula: $\beta_{LAL} = 33.7^\circ$, and Lundgren's formula: $\beta_{Lund} = 35.9^\circ$.

a null line would appear as a near instantaneous period of zero amplitude in the waveform. Any waveform which lies on the tangent circle sweeps through a null line in each precession cycle. For the plane we are investigating, we deduce that a null line sits on $[\psi = 45^\circ, \theta_{JN} = 56^\circ]$, the left-most point on the tangent circle, using this formula for inclination:

$$\iota = \theta_{JN} + \beta \cos \phi_{JL} = 90^\circ, \quad (4.3)$$

and the LAL equation for the opening angle, which gives $\beta_{LAL} = 33.7^\circ$ for these sources. Eq. 4.3 is dependent upon the vector from which ϕ_{JL} is measured, described in Section 3.3.2 (Eq. 3.6), reproduced here:

$$\phi_{JL} = \hat{\mathbf{x}} \times (\hat{\mathbf{J}} \times \hat{\mathbf{L}}) \quad (4.4)$$

where $\hat{\mathbf{x}}$ is the vector in the orbital plane, perpendicular to both $\hat{\mathbf{J}}$ and the projection of $\hat{\mathbf{N}}$ in the orbital plane, pointed rightward with respect to the observer.

Null lines may be the reason for the non-concentric patch of better Matches within the circle. We carried out the same plot of Matches but fixing $\phi_{JL} = 270^\circ$, thus placing the null line in the plane on the opposite side of the tangent circle, at $\theta_{JN} = 124^\circ$. The results were effectively the same but rotated 180° (as described before), placing the patch of better Matches within the circle now slightly lower and left of the circle's centre.

Waveforms which sweep through null lines exhibit the greatest magnitude of fluctuations in polarisation and amplitude (see Figure 4.5). Thus, waveforms which sweep through null lines, or near them, exhibit the most observable precession and so return the lowest Matches. Any waveform on the tangent circle passes through the null line which sits at $\iota = 90^\circ$, $\psi = 45^\circ$ on its precession cycle. If the number of precession cycles in a waveform is sufficiently small, the position of a null line “crossing” in a waveform may have an impact on its Match, since it will reduce the contribution to the signal power from the frequencies emitted during that period.

To test this hypothesis, longer waveforms can be Matched and plotted. The

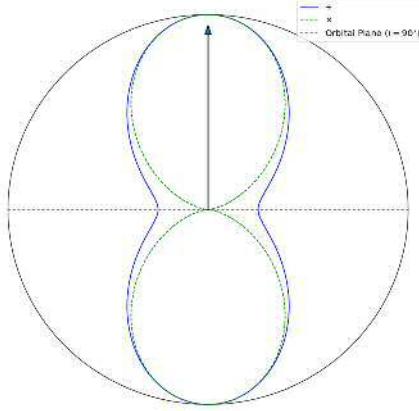


Figure 4.5: *Reproduction of Figure 3.1 in Section 3.2.* The relative amplitude of instantaneous emission for the two polarisations vs inclination, ι . The arrow represents $\hat{\mathbf{L}}$. A null line sits in the orbital plane ($\iota = 90^\circ$) for cross-polarised waves, where you can see that the emitted amplitude is zero.

waveforms are giving the same parameter values, but the waveform begins from lower initial orbital frequencies (i.e. the black holes are initially set further apart). This ensures there is a longer waveform to recover with the inner product, reducing the impact of null lines and any difference between the waveforms being Matched has an increased effect on the inner product. If the position of null lines do not affect the signal power of frequencies which contain high observable precession for waveforms of specific orientations, we expect the patch to remain regardless of waveform length.

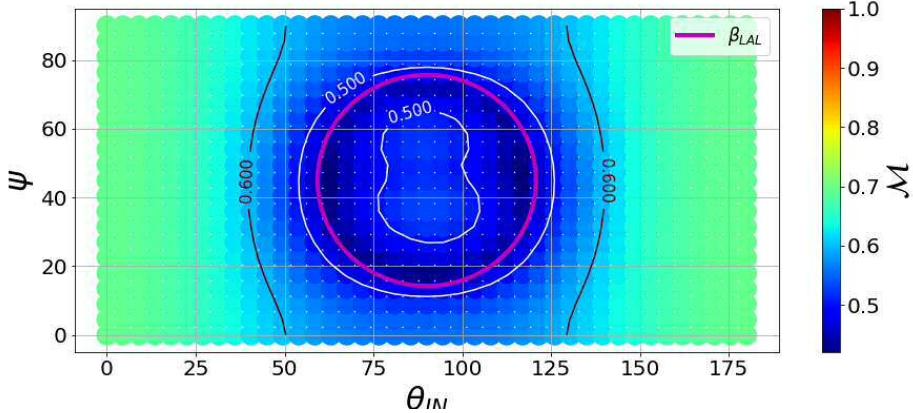


Figure 4.6: $\chi_p = 0.8$ ($\beta_{LAL} = 30.7^\circ$), $\phi_{JL} = 90^\circ$. A colour plot for \mathcal{M} vs ψ vs θ_{JN} with longer waveforms, their initial conditions set at $f_{gw} = 20\text{Hz}$. The tangent circle to the β_{LAL} precession cone is also plotted.

Figure 4.6 shows the Match results of the same waveform parameters as in Figure

4.4 except for the initial waveform frequency, which was set to $f_{gw} = 20\text{Hz}$. The waveform length is roughly five times longer when started from 30Hz for these intrinsic parameter values. The Match has dropped for all waveforms, due to the differences between the non-precessing template being amplified over the longer inner product. We no longer see an off-centre patch of significantly better Matches within the circle, but a shallow concentric slope towards the circle's origin. This is to be expected: as the observer moves within the tangent circle, the range of inclinations they observe through the precession cycle becomes shallower and the observed phase oscillations reduce in amplitude, just as in the case where the observer lies just outside the tangent circle. We can therefore deduce that the patches of higher Matches seen in Figure 4.3 are caused by the location of null lines in waveforms of few precession cycles.

The value of β_{LAL} given is the opening angle at the initial frequency $f_{gw} = 20\text{Hz}$. It is smaller than the value calculated for the waveforms in Figure 4.4, since at lower frequency the black holes have greater orbital momentum. We can see that the tangent circle produced by this value of β accurately traces the lowest Matches. From this, we can deduce that most of the signal power is collected from frequencies near $f_{gw} = 20\text{Hz}$, since if higher frequencies contributed more power, the larger opening angle at those frequencies would cause the worst Matches to occur in a larger circle than the tangent circle calculated from the initial value of β . In effect, we can use the location of the worst Matches in relation to the tangent circle for the initial value of β to assess whether lower frequencies or higher frequencies contribute more to the signal power. This will be useful in Cases 2,3 and 4.

It is interesting to note that extending these waveforms to much lower initial frequencies would not affect the Matches a great deal more than seen in Figure 4.6, since the ZDHP PSD increases sharply below $f_{gw} = 20\text{Hz}$ which significantly reduces the contribution of this region to the signal power.

Figure 4.7 shows Match results for waveforms with the same parameters as those in Figure 4.3b, but with an initial frequency of $f_{gw} = 20\text{Hz}$. These results demonstrate that, whilst the groups of better Matches around $\theta_{JN} = 90^\circ$ only occur for shorter waveforms, the effect of ϕ_{JL} on the shape of the ring of poorest Matches is not dependent on waveform length.

For values of ϕ_{JL} where $|\sin \phi_{JL}| \neq 1$, the ring is distorted in the same manner as we see for the shorter waveforms used in Figure 4.3. The ψ coordinate for the centre of the ring also varies with ϕ_{JL} , just as we saw in Figure 4.3. We may deduce that the distortion of the ring is the result of the degeneracy between ψ and ϕ_{JL} .

The degeneracy between ψ and ϕ_{JL} can be understood if we consider the diagrams in Figure 4.8, which show the line-of-sight, $\hat{\mathbf{N}}$, and momentum vectors of a BBH at three stages during a precession cycle in the inertial frame of the observer. In the left and right diagrams, $\hat{\mathbf{J}}$, $\hat{\mathbf{L}}$ and $\hat{\mathbf{N}}$ share the same plane, thus $\hat{\mathbf{N}}$ is observing only plus-polarised GW. The middle diagram shows the inclination of $\hat{\mathbf{N}}$ to $\hat{\mathbf{L}}$ to be at roughly the same as the inclination of $\hat{\mathbf{N}}$ to $\hat{\mathbf{J}}$, but naturally $\hat{\mathbf{L}}$ and $\hat{\mathbf{J}}$ are separated by the opening angle, β . From this, we can deduce that the observer in the middle diagram receives waves of polarisation $\psi = \beta$. Since neither the precession cone nor $\hat{\mathbf{N}}$ changes throughout the three diagrams, a change in ϕ_{JL} must induce a sinusoidal change in polarisation where the peak is $\Delta\psi = \beta$ above the mean.

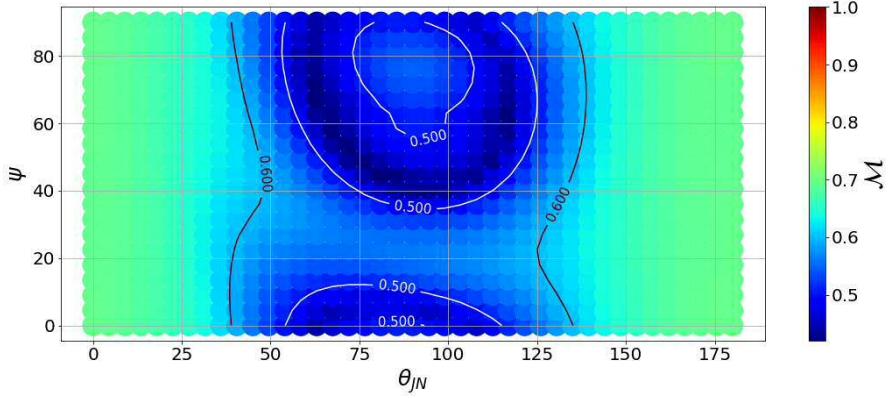


Figure 4.7: $\chi_p = 0.8$, $\phi_{JL} = 45^\circ$. Colour plot for \mathcal{M} vs ψ vs θ_{JN} , using the same parameters as Figure 4.3b with longer waveforms (initial frequency $f_{gw} = 20\text{Hz}$).

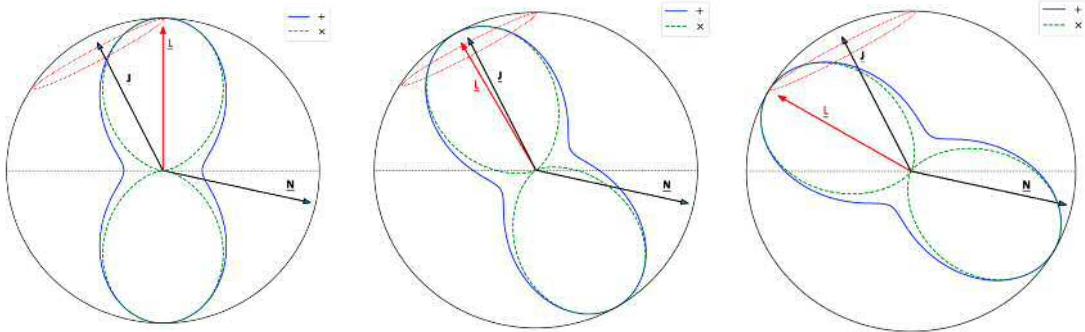


Figure 4.8: Amplitude profiles for both polarisations, with fixed vectors $\hat{\mathbf{J}}$ and $\hat{\mathbf{N}}$. $\hat{\mathbf{L}}$ is shown to be precessing about $\hat{\mathbf{J}}$. The three plots show the profiles at different stages on the precession cone. The black circle represents the unit sphere. From left to right, ϕ_{JL} has passed through approximately 90° since the previous graph.

Now let us understand why this relationship causes the circle of poorest Matches to distort. From Eq. 4.4 we can deduce $\cos \phi_{JL} = 0$ sets the three vectors $[\hat{\mathbf{L}}, \hat{\mathbf{J}}, \hat{\mathbf{N}}]$ coplanar. Using this information and defining ψ_{co} as the polarisation observed when all three vectors are coplanar, we can loosely describe the initial polarisation as a function of ψ_{co} , ϕ_{JL} and θ_{JN} :

$$\psi = \psi_{co} + \beta \cos \phi_{JL} \mathcal{F}(\sin(\theta_{JN}), \phi_{JL}), \quad (4.5)$$

where \mathcal{F} is some function dependent upon $[\sin \theta_{JN}, \phi_{JL}]$.

We recall that θ_{JN} is constant throughout a waveform, thus ψ oscillates with ϕ_{JL} maximally when $\sin \theta_{JN} = 1$ and does not change at all when $\sin \theta_{JN} = 0$. Note that observers at $\sin \theta_{JN} = 0$ witness no fluctuations in amplitude either, since $\hat{\mathbf{J}}$ is directed at the observer, maintaining the same inclination to $\hat{\mathbf{L}}$. We have not considered here the time-dependence of the opening angle, which we shall discuss later in this chapter.

We can see that ψ has an average value of $\psi = \psi_{co}$ (given the limit of many pre-

cession cycles). Therefore, the centre of the ring of poor Matches will be positioned at $[\theta_{JN} = 90^\circ, \psi_{co}]$, since this represents the orientation about which the tangent to $\hat{\mathbf{L}}$ oscillates.

For plots where $\cos \phi_{JL} = 0$ (Figure 4.3c), $\psi = \psi_{co}$, and so the three vectors $[\hat{\mathbf{J}}, \hat{\mathbf{L}}, \hat{\mathbf{N}}]$ begin coplanar for all waveforms. In these cases, the waveforms which begin cross-polarised are also the waveforms which are maximally cross-polarised throughout their duration. Hence, they produce the poorest Matches, and sit on the tangent circle.

For plots where $|\cos \phi_{JL}| = 1$ (Figure 4.3a), we see the Matches *wrap around* the ψ axis, where the circle is centred on some polarisation $\psi \neq 45^\circ$ and overlaps $\psi = 0 = 90^\circ$ (i.e. $\psi \equiv \psi + 90^\circ$). In these plots, the waveforms begin where $\psi = \psi_{co}$, and the centre of the ring will be positioned at $\psi = 45^\circ + \beta$, since the $[\hat{\mathbf{J}}, \hat{\mathbf{L}}]$ plane is perpendicular to the $[\hat{\mathbf{J}}, \hat{\mathbf{N}}]$ plane, so the polarisation at which the three vectors are coplanar is maximally distant from ψ . Thus, the tangent circle simply translates along the ψ axis by β degrees (upwards for $\phi_{JL} = 0^\circ$) and downwards for $\phi_{JL} = 180^\circ$.

In cases where the $[\mathbf{J}, \mathbf{L}]$ plane is neither parallel nor perpendicular to the $[\mathbf{J}, \hat{\mathbf{N}}]$ plane, $\Delta\psi = \psi - \psi_{co}$ is dependent upon θ_{JN} such that the precession cone tangent appears elliptical. Whilst we fall short of being able to fully describe this shape as a function of ϕ_{JL} , we know the tangent circle is truly circular at $\phi_{JL} = [0^\circ, 90^\circ, 180^\circ, 270^\circ]$ and its centre can be approximated by the coordinates:

$$(\theta_{JN}, \psi)_{cen} = (90^\circ, \beta \cos \phi_{JL} + 45^\circ). \quad (4.6)$$

The ellipticity of the circle is not very high, and reduces with lower χ_p . We may approximate the ring of poorest Matches as a circle for all planes in ϕ_{JL} , giving:

$$\sqrt{(\theta_{JN} - \theta_{JNcen})^2 + (\psi - \psi_{cen})^2} - \beta = 0 \quad (4.7)$$

To conclude, we have constructed a formula which approximates where the poorest Matches can be found in the parameter volume $[\theta_{JN}, \psi, \phi_{JL}]$ for the focal set's values of $[\chi_{eff}, q, M]$. Matches improve in a roughly radial form away from the circle, thus Eq. 4.7 could be used as a basis for further research to define a function mapping to all Matches in parameter space.

4.3 Case 2: Variation of Match with Aligned Spin

Now we have an understanding of the effects of orientation on observed precession, we can assess how the intrinsic parameters affect these patterns. First, we study how χ_{eff} affects the Matches in the planes we've investigated - keeping the masses the same.

We choose the $[\theta_{JN}, \psi]$ plane where $\phi_{JL} = 90^\circ$ to assess how the circle of poorest Matches is affected by aligned spin ($\chi_{eff} > 0$). The precessional spin is set to $\chi_p = 0.6$ for all waveforms. This value allows for aligned spin magnitudes up to $|\chi_{eff}| = 0.8$ whilst the source still permitting significant inherent precession. Four plots are shown in Figure 4.9, where in each plane χ_{eff} was fixed at a positive value.

It is clear that the circle remains well defined, slightly shrinking in size with increasing χ_{eff} . This is a result of the increase in \mathbf{J} . Aligned spin increases the inertia of the system, acting to reduce the size of the precession cone for a given χ_p . Moreover, more aligned spin requires more radiation to decay the orbit to the point of merger, which takes more time and so the waveform length increases as χ_{eff} increases. As we saw in Section 4.2, longer waveforms include more precession cycles which acts to reduce the Match for cases where precession is observable whilst not significantly affecting better Matches. Hence, we see Matches above $\mathcal{M} > 0.8$ occupy a very similar area for all $\chi_{eff} \geq 0.2$, whilst Matches within the circle decrease. The patch of better Matches within the circle entirely vanishes between $\chi_{eff} = 0$ and $\chi_{eff} = 0.4$. This, too, is a result of the waveforms being longer, as we saw in Figure 4.6.

Matches near $\sin \theta_{JN} \approx 0$ are not significantly affected by $\chi_{eff} > 0$, but there is a marginal improvement in this area between $\chi_{eff} = 0$ and $\chi_{eff} = 0.4$. This may be due to the increased inertia of the system arresting the evolution of β , reducing the influence of precession on the waveforms. As \mathbf{J} loses energy to GW, the opening angle would increase. Such an effect is only significant near merger, and sources where \mathbf{L} is larger would not experience as much change in β . One would expect this effect to increase its influence as $\chi_{eff} \rightarrow 1$, which we do not see. Instead we see a small drop in Matches at the extremes of θ_{JN} as χ_{eff} increases from 0.4 to 0.8. This is, perhaps, a result of the longer waveforms presenting enough precession to have a greater effect than that caused by the small reduction in β and its arrested evolution.

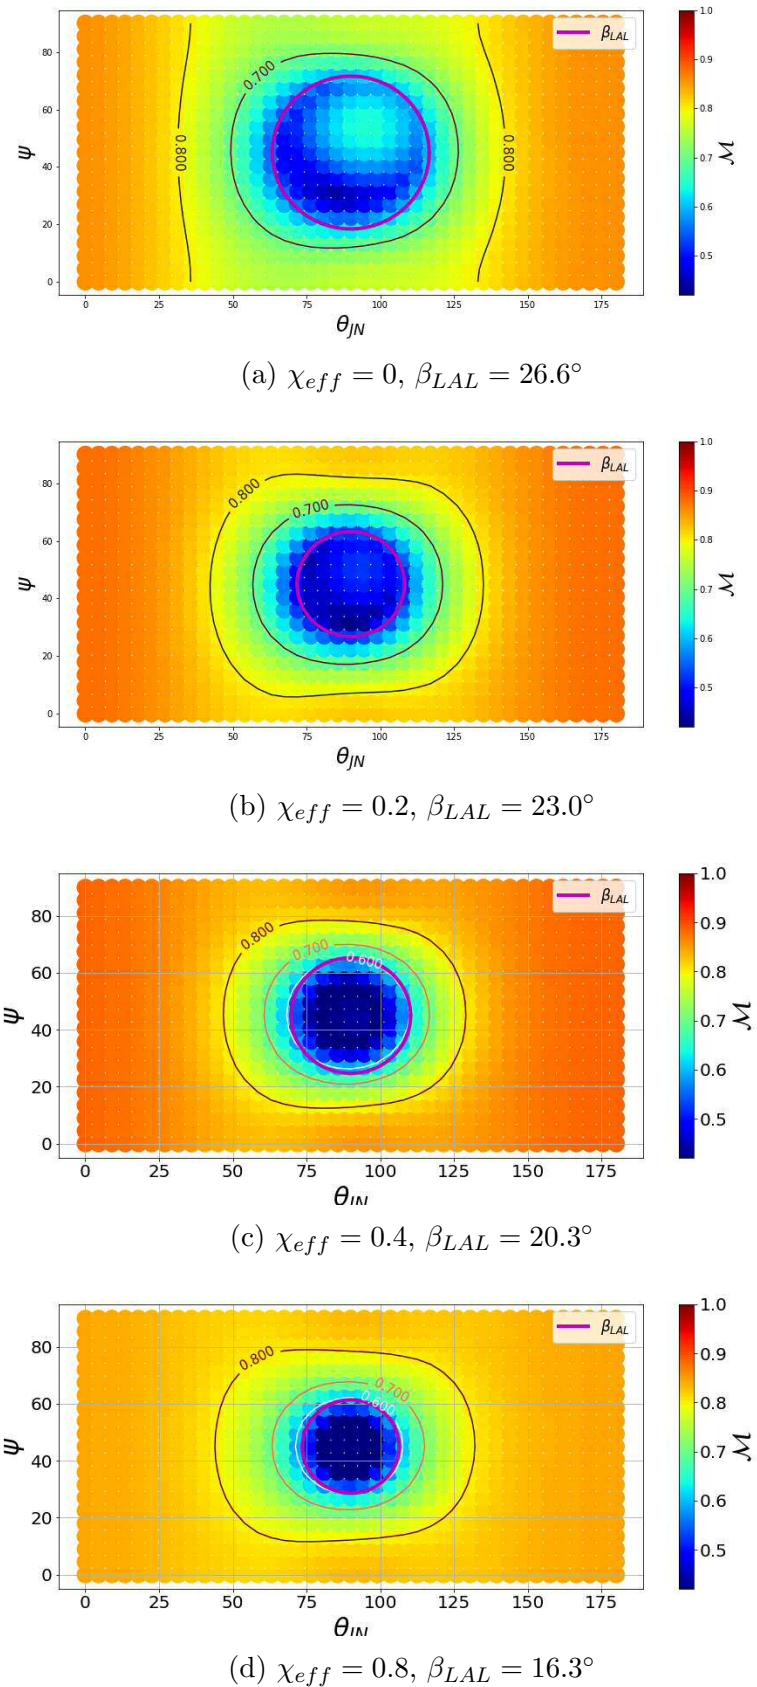


Figure 4.9: Colour plots for \mathcal{M} vs ψ vs θ_{JN} in planes of the given χ_{eff} . $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$, $\phi_{JL} = 90^\circ$ and initial frequency $f_{gw} = 30\text{Hz}$ for all waveforms.

Now let us investigate the effects of anti-aligned spin. The relationship between Match and the tangent circle begins to deteriorate, as seen in Figure 4.10. Outside the range $[30^\circ \lesssim \theta_{JN} \lesssim 150^\circ]$, the Matches are independent from ψ just as with all previous results, and the worst Matches still occur somewhere on the tangent circle, but not all the way around it.

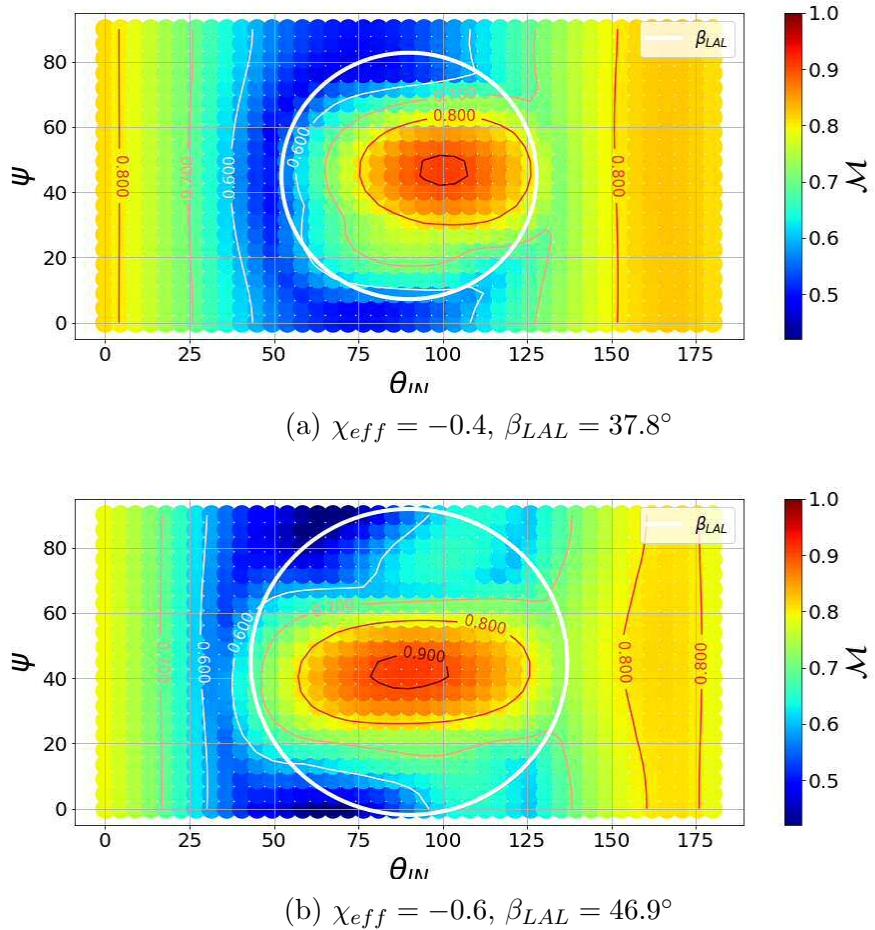
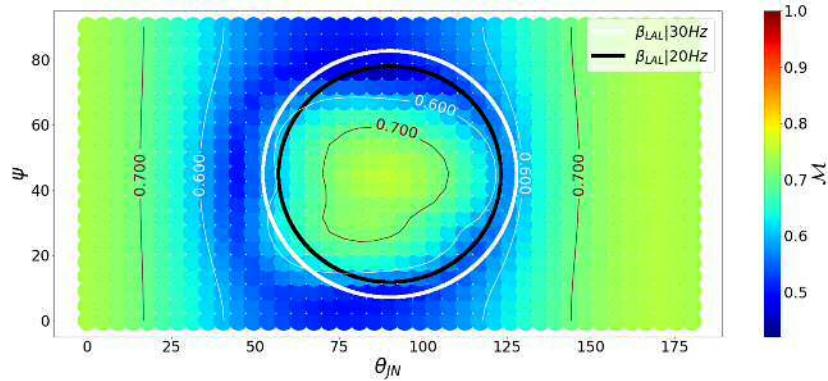


Figure 4.10: Colour plots for \mathcal{M} vs ψ vs θ_{JN} where $\chi_{eff} < 0$. $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$ and initial frequency $f_{gw} = 30\text{Hz}$ for all waveforms.

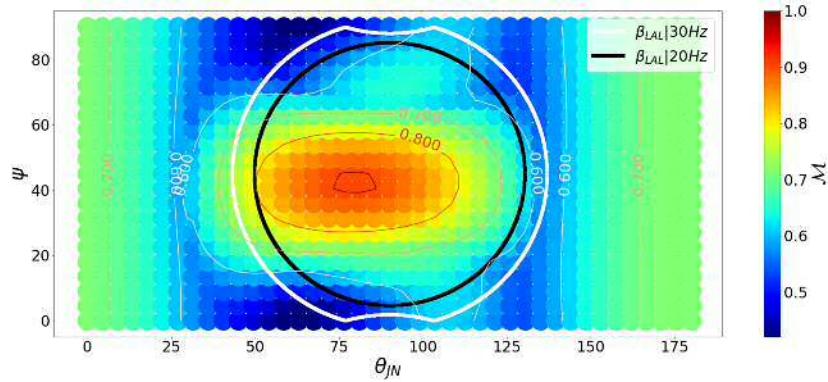
Whilst Matches are still much higher towards $\sin \theta_{JN} = 0$, the best Matches appear within the circle - the exact opposite to what we find when $\chi_{eff} > 0$. This behaviour is peculiar - it appears that for anti-aligned sources, predominantly cross-polarised waveforms exhibit less precession when the source is near edge-on to the observer. Using the same argument as for the patches in Figure 4.3, null lines may be a factor. We recreate the plots in Figure 4.10 using longer waveforms, starting at $f_{gw} = 20\text{Hz}$, but setting the parameter values at $f_{gw} = 30\text{Hz}$ as before (to ensure the only possible influence is additional waveform length, and not the parameters varying between 20Hz and 30Hz). The results are shown in Figure 4.11.

Most of the best Matches at $\chi_{eff} = -0.4$ are now near $\sin \theta_{JN} \sim 0$, whilst the poorest Matches are more radially symmetric and proximate to the tangent circle. However, the patch is still found within the tangent circle, and there are

still asymmetries (especially in the $\chi_{eff} = -0.6$ plot) which, if were caused by the waveform length and null lines, would have been smoothed out. If we look at Figure 4.11b we see negligible reduction in the Matches within the circle, whilst Matches near $\sin \theta_{JN} \sim 0$ are noticeably lower. The cause of these results is unknown. It may be in relation to the evolution of β through the waveform, especially near merger as \mathbf{L} becomes dominated by \mathbf{S} and the orbital frequency reaches its maximum (along with the precession cycle's frequency). In this process, the opening angle may grow to $\beta > 45^\circ$, where waveforms near $[\theta_{JN} = 90^\circ, \psi = 45^\circ]$ would in fact spend the majority of a precession cycle nearer face-on than edge-on. This must occur early enough before merger to have a significant impact on the Match in order to explain the results we see. Since our objective is to find regions of poor Match in parameter space, we leave it to future research to investigate this phenomenon further.



(a) $\chi_{eff} = -0.4$, $\beta_{LAL}(20\text{Hz}) = 33.1^\circ$, $\beta_{LAL}(30\text{Hz}) = 37.8^\circ$



(b) $\chi_{eff} = -0.6$, $\beta_{LAL}(20\text{Hz}) = 40.3^\circ$, $\beta_{LAL}(30\text{Hz}) = 46.9^\circ$

Figure 4.11: Colour plots for \mathcal{M} vs ψ vs θ_{JN} where $\chi_{eff} < 0$. $\chi_p = 0.6$, $m_1 = 20M_\odot$, $m_2 = 5M_\odot$, and initial frequency $f_{gw} = 20\text{Hz}$ for all waveforms.

Two tangent circles are plotting on each graph, one for the initial waveform frequency ($f_{gw} = 20\text{Hz}$) and one for the frequency at which the parameter values are set ($f_{gw} = 30\text{Hz}$). Since β depends on $|\mathbf{J}|/|\mathbf{L}|$, and the same system at a lower frequency means orbital momentum $|\mathbf{L}|$ is greater, the reason for the difference between the two circles is clear. As discussed, the initial opening angle may not

reflect its value in the frequency bandwidth of the waveform where most of its signal power is recovered. If β is high enough in that bandwidth, it causes the tangent circle to wrap around the ψ axis, as can be seen in Figure 4.11b. This effect is discussed further in Section 4.3.1.

4.3.1 β Behaviour

Figure 4.13 (overleaf) shows the relationship between β and χ_{eff} at the start of a waveform for a selection of values of χ_p . At some point as χ_{eff} decreases, the opening angle (i.e. the radius of the circle) exceeds 45° , and so the tangent circle begins to wrap around the ψ -axis, the top emerging at the bottom and vice versa. An example is given in Figure 4.12.

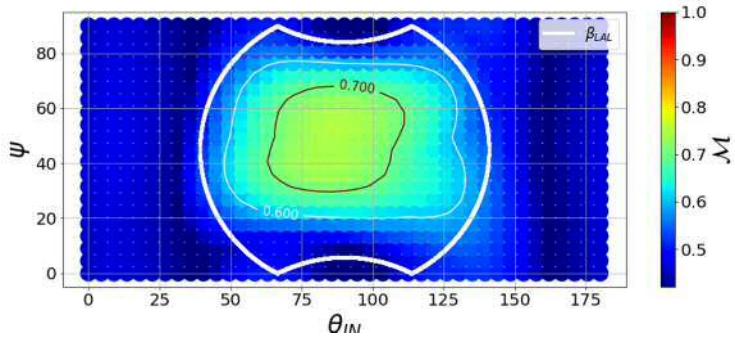


Figure 4.12: A colour plot showing \mathcal{M} vs ψ vs θ_{JN} to demonstrate the tangent circle “wrapping around” on the ψ -axis as a result of the opening angle $\beta > 45^\circ$. In this case, $\beta_{LAL} = 53.0^\circ$.

We can see that χ_{eff} dramatically changes the opening angle, even when there is very little precession. The gradient is near constant and only slightly increases with χ_p when $\chi_{eff} > 0$. Below $\chi_{eff} = 0$, the relationship changes heavily dependent upon χ_p . This Figure can be used to track the differences we see in the $[\theta_{JN}, \psi]$ plots for various combinations of χ_p and χ_{eff} . The difference between the plots in Figure 4.9 can be assessed using the purple curve (for $\chi_p = 0.6$). When we do this, we opening angle changes by $\sim 10^\circ$ through all positive aligned spin, and $\sim 34^\circ$ for anti-aligned spin.

Consider a waveform where there is little precessing spin, $\chi_p = 0.2$, but high anti-aligned spin, $\chi_{eff} \approx -0.9$. Figure 4.13 shows that the opening angle is $\beta \approx 50^\circ$, which is larger than most configurations of spin with higher values of χ_p and smaller anti-aligned spin. Waveforms for these values of χ_p showed negligible precession for almost all orientations in Figure 4.1 and 4.2, yet anti-aligned spin reduces \mathbf{J} so much (for these masses) that sources with low precessing spins have larger opening angles than most higher-precessing sources with smaller anti-aligned spins. This indicates that the richest area for observably precessing waveforms in the parameter space we have explored so far is in the domain of anti-aligned spin.

Whilst the tangent circle does not reach $\psi = 0^\circ$ in these plots, the fact that the opening angle increases throughout the waveform means that our calculation may

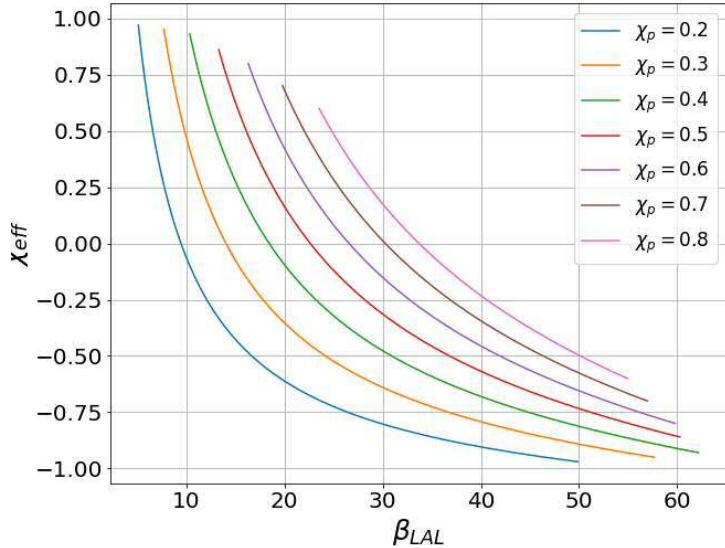


Figure 4.13: The relationship between β and χ_{eff} for different χ_p , where $m_1 = 20M_\odot$, $m_2 = 5M_\odot$ and $f_{gw} = 30\text{Hz}$. β was calculated using the LAL equation.

not truly reflect its influence on the Match. If the contribution to the signal power is significant from the portion of the waveform where β has become noticeably larger, this wrap-around effect could be observable as is the case in Figure 4.11b.

Curves for the evolution of the opening angle as a function of frequency for different values of χ_p are shown in Figure 4.14. To compare, GW151226 reached a peak frequency of approximately 450Hz [20], thus we know it is possible for sources similar to those we are investigation to reach such frequencies. The rate of change of β is greatest at lower frequencies, and steeper for higher χ_p . Yet even for cases of low precessing spin, we see β increase by a factor of two between 30Hz and 100Hz.

How much difference this evolution of β makes to the Matches outside the tangent circle we've draw for the initial waveform frequency is dependent upon which frequencies dominate the signal power. If the length of time β spends at higher angles is comparatively short, then the lower frequencies will contribute more signal power and the Matches outside the tangent circle will be less affected. For short waveforms (e.g. high χ_{eff}), the lower frequencies will contribute less to the signal power, and so the influence of the opening angle's evolution on the Matches will increase.

We conclude that positive aligned spin has a traceable but small impact on the observable precession in a waveform. The orientation of the most clearly precessing sources is unchanged, and higher magnitudes of χ_{eff} produce fewer of them. Adapting Eq. 4.7 to encompass positive aligned spin for these masses appears possible by multiplying β with a linear function of χ_{eff} .

Anti-aligned spin increases the observability of precession through all orientations outside the tangent circle, whilst reducing observable precession within it. It has a significant and non-linear effect on the orientation where the most observably precessing waveforms lie, due to its relationship with the opening angle. Given

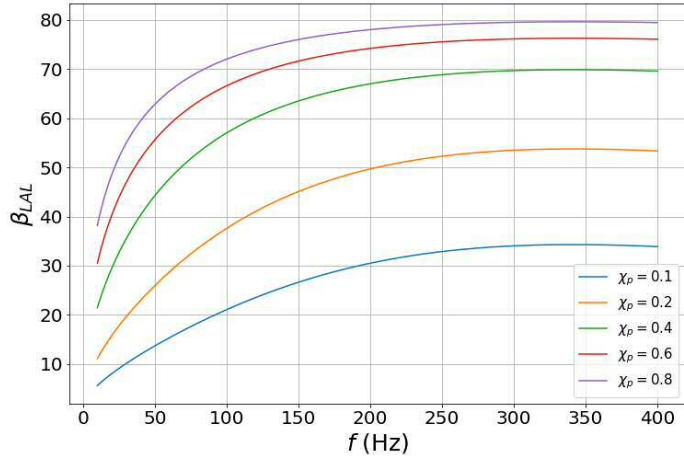


Figure 4.14: The frequency-evolution of β as calculated by the LAL algorithms for sources with given χ_p , $\chi_{eff} = -0.6$, and masses $m_1 = 20M_\odot, m_2 = 5M_\odot$.

that the opening angle's evolution after $f_{gw} = 30\text{Hz}$ in these waveforms noticeably contributes to the signal power, defining an appropriate function to capture these effects within a formula to approximate the location of the most observably precessing waveforms (such as Eq. 4.7 for zero aligned spin) is more challenging than for positive aligned spin.

4.4 Case 3: Variation of Match with Mass Ratio

In this Case we maintain the same total mass as in Cases 1 and 2: $M = 25M_\odot$, and remove aligned spin ($\chi_{eff} = 0$). We choose to explore two mass ratios: $q = 2$ and $q = 9$, to explore the parameter below and above the focal set.

Figure 4.15 shows colour plots for Matches in the $[\theta_{JN}, \psi]$ plane for $\phi_{JL} = 90^\circ$, with $\chi_p = 0.6$. There is an immediately obvious difference between the two plots: one is very red, the other is very blue. Nearly 80% of the waveforms in Figure 4.15a produced Matches of $\mathcal{M} > 0.9$, compared to zero such waveforms in Figure 4.15b.

As $q \rightarrow 1$, the effects of precession are significantly reduced in a similar manner to what we saw by increasing χ_{eff} . The ratio $|\mathbf{J}|/|\mathbf{L}|$ is being decreased in both cases, thus decreasing β and reducing the circle of poor Matches. We can see in Figure 4.15a that the Matches outside the circle are much better than we saw for positive aligned spin, demonstrating that q more effective at “polishing out” precession in this region. This is caused by a difference in signal power. Lower mass ratio sources radiate more powerful GW within the most sensitive frequency band in the PSD, thus improving the contribution to the Match in frequencies earlier on in the inspiral, where there is little observable precession in waveforms far from the tangent circle.

As mass ratio increases, we see all Matches across the plane drop, as shown in Figure 4.15b. The opening angle is larger for higher mass ratios since the orbital angular momentum is smaller, increasing $|\mathbf{J}|/|\mathbf{L}|$. As discussed in Case 2, we see wrap-around effects as the tangent circle encroaches on $\psi = 0 = 90^\circ$. The Matches

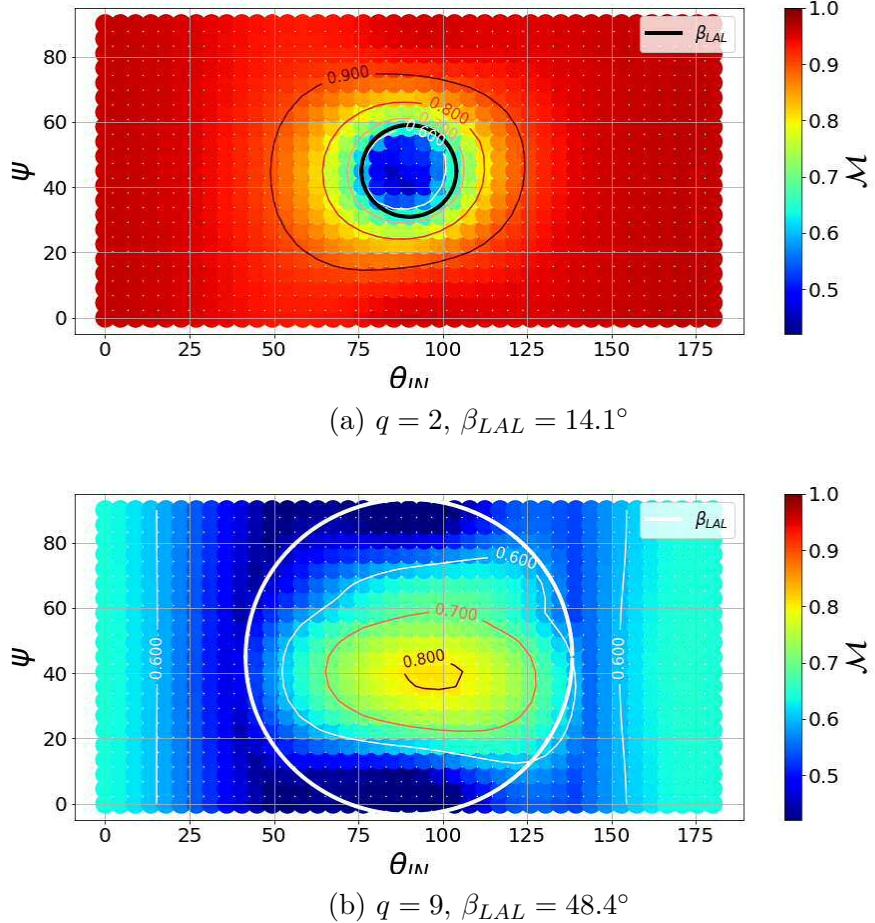


Figure 4.15: Colour plots for \mathcal{M} vs ψ vs θ_{JN} for the given mass ratios where total mass $M = 25M_\odot$, $\chi_p = 0.6$ and $\phi_{JL} = 90^\circ$.

are similarly asymmetric to those seen in Figure 4.10, which were largely reduced using longer waveforms. When the same technique was applied to these results we found the same outcome, where the Matches within the circle had reduced to a peak of $\mathcal{M} \approx 0.7$, and every waveform from the circle outwards produced Matches $\mathcal{M} < 0.6$.

We found that reducing χ_p when $q = 2$ decreased the size of the circle of poor Matches (all Matches of $\mathcal{M} < 0.6$ lay within $\beta_{LAL} \approx 7^\circ$) and improved the Matches outside it (95% of all Matches were $\mathcal{M} > 0.90$, and 80% were $\mathcal{M} > 0.97$). In the case of $q = 9$, we found that setting $\chi_p = 0.3$ gave very similar results to those to the top plane in Figure 4.9, where $\chi_p = 0.6, q = 4$ ($\chi_{eff} = 0$ in both cases). Further exploration revealed a very strong correlation between Matches from different combinations of χ_p and q . If we define a new parameter:

$$Q_p = \chi_p q, \quad (4.8)$$

we hypothesise that $\mathcal{M}(Q_p) \approx const.$ for any point in orientation parameter space $[\theta_{JN}, \psi, \phi_{JL}]$ (fixing $M = 25M_\odot, \chi_{eff} = 0$). Figure 4.16 lends credence to this

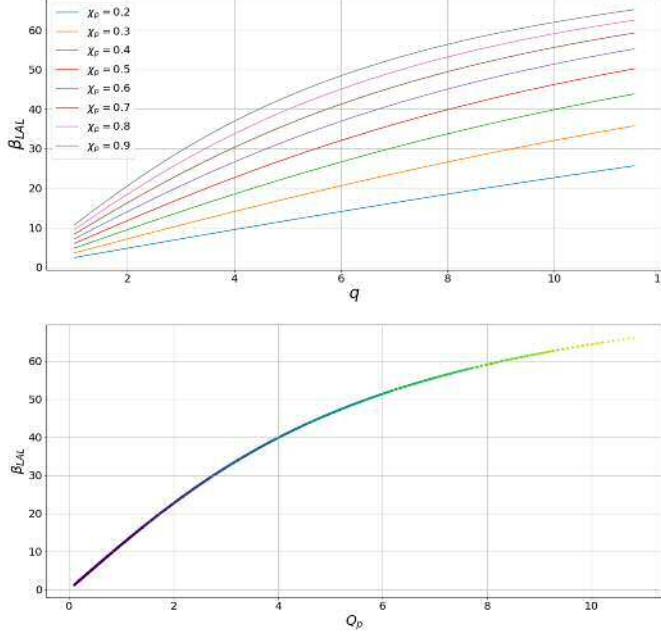


Figure 4.16: Top: The relationship between β and mass ratio, q , for various values of χ_p , where $\chi_{eff} = 0$, $M = 25M_\odot$ and frequency $f_{gw} = 30\text{Hz}$. β was calculated using Eq. 4.2. Bottom: The relationship between Q_p and β , for $M = 25M_\odot$ and $\chi_{eff} = 0$.

idea. We see that, whilst the relationship between β and q changes depending upon χ_p , our new parameter behaves in a very well-defined, one-to-one relation with β . Thus, the tangent circle will be the same size for all waveforms of equal Q_p , and the location of the ring in the $[\theta_{JN}, \psi]$ plane will also be the same (as it only depends on ϕ_{JL} and β).

To test this theory, we calculate the variance of the Match for constant Q_p (Eq. 4.9 at every point on the $[\theta_{JN}, \psi]$ plane (using the full range of possible combinations of χ_p, q , being careful not to breach the mass ratio calibration limit of **IMRPhenomPv2** for precessing BBH: $1 \leq q \leq 12$):

$$\text{Var}(\mathcal{M}|Q_p) = \frac{1}{n_{\chi_p} - 1} \sum_{i=1}^{n_{\chi_p}} (\mathcal{M}_i - \bar{\mathcal{M}})^2 \quad (4.9)$$

where n_{χ_p} is the number of different combinations of $[\chi_p, q]$ for which a Match is calculated and $\bar{\mathcal{M}}$ is the mean Match for that group.

The variance for points in the $[\theta_{JN}, \psi]$ plane for four different values of Q_p are given in Figure 4.17. Values of Q_p were chosen to be distributed evenly across the possible range. Each point represents the variance in 16 Match calculations. Figure 4.18 shows four more planes for $\text{Var}(\mathcal{M}|Q_p = 3.6)$ where ϕ_{JL} was set at different values. These plots demonstrate what Figure 4.16 implies: the coherence between

the Matches from different combinations of $[\chi_p, q]$ is not significantly affected by orientation. It should be noted that for Q_p to be high, both χ_p and q must be high in their range. Most possible combinations of $[\chi_p, q]$ will give values of Q_p clustered in the middle of the curve in the bottom graph of Figure 4.16 (around the dark blue to green segment).

We conclude that the effect of mass ratio on Matches can be approximated by considering q an amplification factor on χ_p , given $\chi_{eff} = 0$ and $M = 25M_\odot$. Combinations of $[\chi_p, q]$ which give the same value of Q_p produce approximately the same amount of observable precession, for any given orientation. As such, we can reliably identify areas of observable precession across five parameters: $[\chi_p, q, \theta_{JN}, \psi, \phi_{JL}]$, for $\chi_{eff} = 0$ and $M = 25M_\odot$. The influence of aligned spin and total mass on this relation is discussed in Section 4.6.

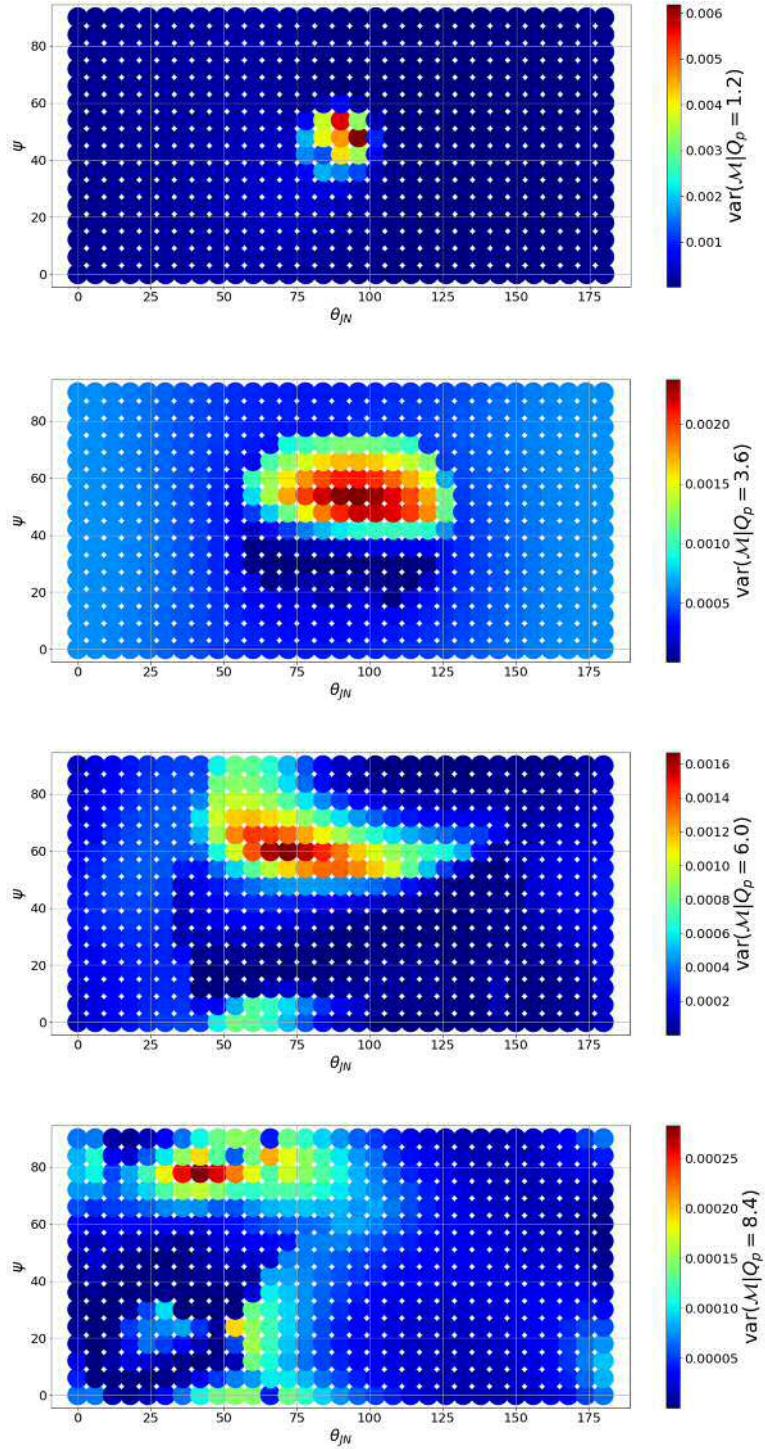


Figure 4.17: Plots showing the variance in the Match for each $[\theta_{JN}, \psi]$ coordinate for combinations of $[\chi_p, q]$, spanning their full possible range, for four fixed values of Q_p . Descending: $Q_p = [1.2, 3.6, 6.0, 8.4]$. $\chi_{eff} = 0$ and $M = 25M_\odot$ for all waveforms. Note: the colour axes are not consistent.

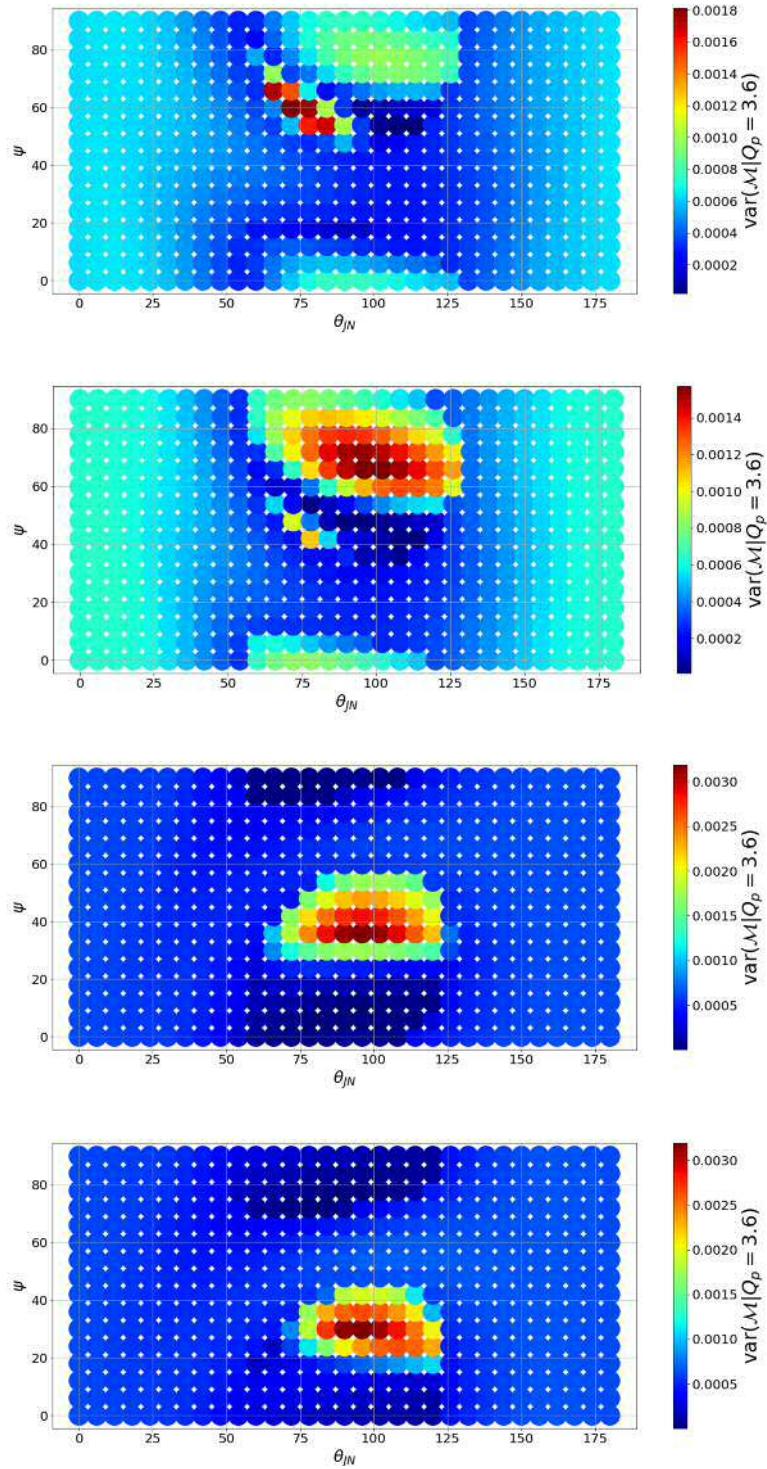


Figure 4.18: Plots showing the variance in the Match for each $[\theta_{JN}, \psi]$ coordinate for $Q_p = 3.6$ with ϕ_{JL} fixed at four different values. Descending: $\phi_{JL} = [30, 60, 120, 150^\circ]$. Note: the colour axes are not consistent.

4.5 Case 4: Variation of Match with Total Mass

We maintain zero aligned spin and return the mass ratio to $q = 4$. We choose two total masses to explore: $M = 12.5M_{\odot}$ and $M = 50M_{\odot}$, a factor of 2 lower and a factor of two higher than our focal set's total mass, respectively. Figure 4.19 shows plots for Matches in the same plane and for the same orientation as Figure 4.15.

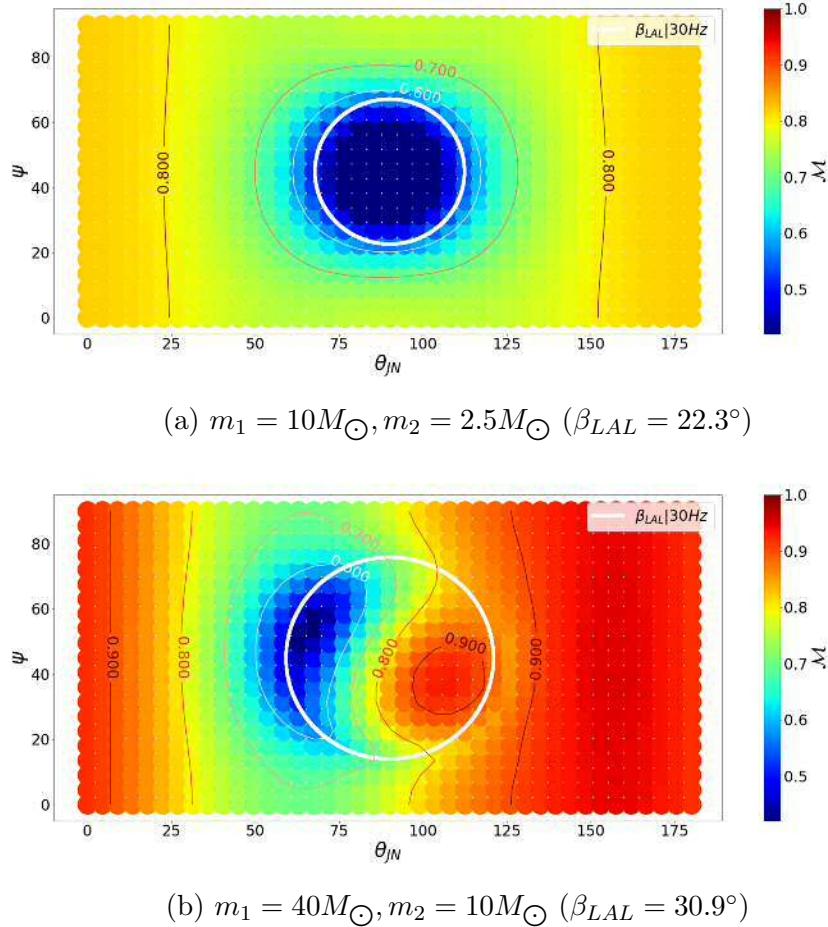


Figure 4.19: Colour plots for \mathcal{M} vs ψ vs θ_{JN} for different total masses where $q = 4$, $\chi_p = 0.6$, $\chi_{eff} = 0$ and $\phi_{JL} = 90^\circ$.

In comparison with the top plot in Figure 4.9, where $M = 25M_{\odot}$ and all other intrinsic parameters are the same, we see that waveforms of lower M produce similar Matches outside the circle, whilst Matches within the tangent circle are poorer ($\mathcal{M} < 0.5$ for all waveforms inside). This can be explained by the fact that the ISCO radius scales with mass, and so lower mass systems reach smaller radii and faster orbital velocities, causing them to precess more quickly. As such, just as with longer waveforms, the effects of precession are exacerbated. The lowest Matches are found inside the tangent circle rather than on it, just as we saw with positive aligned spin (Figure 4.9). Since these waveforms were started and parameters set at $f_{gw} = 30\text{Hz}$, lower frequencies where β would have been smaller have no influence,

and so these results indicate total mass affects observable precession in a manner not related to the precession cone. Again, we leave such an investigation for future research.

Figure 4.19b shows the results for sources with total mass $M = 50M_{\odot}$. Significant increase in the Matches is seen outside the tangent circle. Moreover, we see higher Matches for most of the waveforms within the circle. The number of Matches below $\mathcal{M} < 0.7$ is roughly equal for both plots, but there are many more Matches below $\mathcal{M} < 0.6$ in the $M = 12.5M_{\odot}$ plane. To test for the influence of null lines, we carried out the same plane of Matches using longer waveforms, and found that initial frequency dictates where on the circle the patch of worst Matches lies (whilst ubiquitously reducing the Match for all waveforms). As the length of the waveforms was reduced, the location of the lowest Matches rotated on the tangent circle. When the initial frequency was set to $f_{gw} = 20\text{Hz}$, the lowest Matches were found near the bottom of the circle, rotating clockwise as the initial frequency was increased. They reached the top of the circle when the waveforms were started at $f \approx 40\text{Hz}$.

From these results, we can discern that a waveform with initial frequency $f_{gw} = 20\text{Hz}$ completes only 180° of the precession cycle by the time it reaches $f \approx 40\text{Hz}$. Larger masses do not reach orbital frequencies as high as lower masses before merger (owing to their greater ISCO radius), and so their precession cycles during the late inspiral are slower. Hence, where the worst Matches appear on the tangent circle is dependent upon where the masses are on the precession cone as they enter the most sensitive frequency band in the PSD. The source completes very few precession cycles in the detectable frequency bandwidth, thus inducing a bias in which orientations present the most precession, dependent upon the orientation of the waveform as it enters those frequencies. Where on the precession cycle the source enters the PSD's most sensitive region is a function of where the waveform begins on the precession cone, which is dependent upon the precession phase (i.e. ϕ_{JL}) and how long it takes for the source to reach those frequencies (i.e. f_{gw}).

We conclude that lower total masses have only a significant effect on the Matches of the most observably precessing sources at the mass ratio, $q = 4$, investigated. Higher M increases the influence of the initial orientation and reduces the amount of observable precession in all but the most pronounced cases ($\mathcal{M} < 0.7$).

4.6 Towards a Generalised Formula for Observable Precession

Here we analyse the compatibility of the functions we've formed to approximate the trends found in the Match plots. Eq. 4.6 can be used to trace the approximate location of the poorest Matches (i.e. the most observably precessing sources) on the $[\theta_{JN}, \psi]$ plane. We've seen that the parameter $Q_p = \chi_p q$ tracks lines of constant Match to a good approximation through $[\chi_p, q, \theta_{JN}, \psi, \phi_{JL}]$ where $[\chi_{eff} = 0, M = 25M_{\odot}]$. Therefore, for any given solution to Eq. 4.7, Q_p can identify other solutions which produce the same Matches throughout all orientations. We found in Sections 4.3 and 4.5 that that Eq. 4.7 does not reliably trace the location of poorest Matches

Q_p	χ_{eff}	$M (M_\odot)$	$\overline{\text{Var}(\mathcal{M} Q_p)}$	$\text{Var}(\mathcal{M} Q_p) > 0.01$
1.2	-0.6	25	0.0060	21%
1.2	0.0	25	0.0002	0%
1.2	0.6	25	0.0003	0%
2.4	-0.6	50	0.0050	13%
2.4	0.0	50	0.0011	0%
2.4	0.6	50	0.0022	0%
3.6	-0.4	25	0.0074	24%
3.6	0.0	25	0.0006	0%
3.6	0.4	25	0.0003	0%
4.8	-0.4	12.5	0.0072	26%
4.8	0.0	12.5	0.0001	0%
4.8	0.8	12.5	0.0001	0%
6.0	-0.6	50	0.0004	0%
6.0	0.0	50	0.0002	0%
6.0	0.8	50	0.0011	0%

Table 4.1: Results from a number of variance calculations of the Matches over the $[\theta_{JN}, \psi]$ plane ($\phi_{JL} = 90^\circ$ for all), for the given fixed values of Q_p , aligned spin, χ_{eff} , and total mass, M . The mean of the variance over the whole plane ($\overline{\text{Var}(\mathcal{M}|Q_p)}$) is given along with the percentage of points with variance of $\text{Var}(\mathcal{M}|Q_p) > 0.01$. Every waveform was started at $f_{gw} = 30\text{Hz}$.

for cases where $\chi_{eff} \neq 0$ or $M \neq 25$, but the relation between Q_p and lines of constant Match shows promise to be more general than purely applicable to our focal set. Table 4.1 shows results for the variance in Match for a selection of different combinations of fixed $[Q_p, \chi_{eff}, M]$ spanning their full range within the calibrated space of `IMRPhenomPv2`. We define the “mean variance” as

$$\overline{\text{Var}(\mathcal{M}|Q_p)} = \frac{1}{n_\theta} \sum_{\theta_{JN}, \psi} \text{Var}(\mathcal{M}|Q_p), \quad (4.10)$$

where n_θ is the number of points in the $[\theta_{JN}, \psi]$ plane for which $\text{Var}(\mathcal{M}|Q_p)$ was evaluated. The mean variance summarises how well Q_p tracks constant Match for the whole plane. The mean variance for each combination of $[Q_p, \chi_{eff}, M]$ is given in Table 4.1 along with the percentage of points in that plane which gave a variance of $\text{Var}(\mathcal{M}|Q_p) > 0.01$ (i.e. where the standard deviation in Match is greater than 0.1).

Whilst these results alone are far from comprehensive, they indicate that χ_{eff} has a greater impact than M on the reliability of Q_p to track constant Match. Anti-aligned spin significantly increases the number of orientations which give $\text{Var}(\mathcal{M}|Q_p) > 0.01$, whilst positive aligned spin tends to have a modest effect on the mean variance and no effect on the number of highly-variant points. The influence of changing M

is small, with a slight increase in variance for increased total mass, but the trend would be difficult to identify without more comprehensive results. However, given that even in the least well-behaved plane, no more than 26% of the points had a variance of $\text{Var}(\mathcal{M}|Q_p) > 0.01$ for all combinations of $[\chi_p, q]$, these results indicate that Q_p can approximately trace constant Match throughout the vast majority of parameter space represented in Table 4.1. Further investigation is required to test this relation more rigorously.

Adapting Q_p to create a formula which traces lines of constant Match through varying $[\chi_p, q, \chi_{eff}, M]$ appears quite possible. Figure 4.20 shows the relationship between β and Q_p for various sources: the top graph is shows the relationship where total mass was kept constant at five different values (with zero aligned spin), the bottom graph shows the relationship for sources with $\chi_{eff} = 0.6$ and $M = 25M_\odot$.

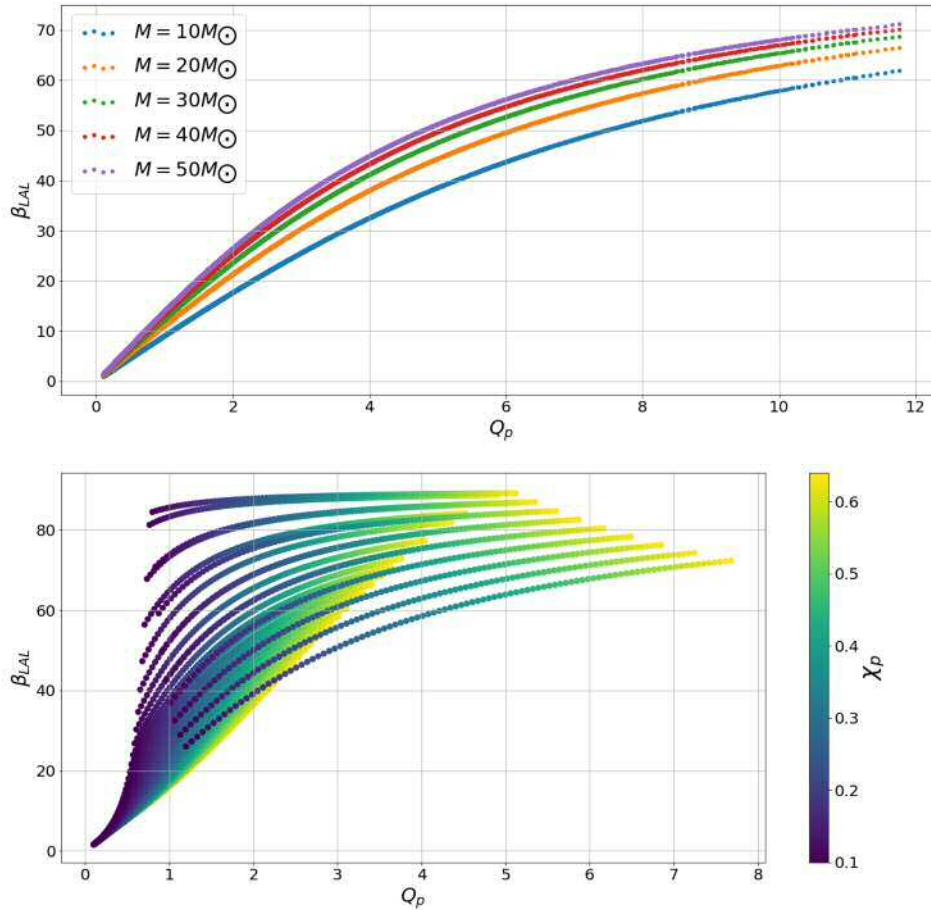


Figure 4.20: Top: The relationship between β and Q_p for a range of sources where total mass is fixed at five given values and $\chi_{eff} = 0$. Bottom: The relationship between β and Q_p for a range of sources where $M = 25M_\odot$ and $\chi_{eff} = -0.6$.

In the top graph we see that when there is no aligned spin, the relation remains one-to-one for any fixed total mass, which has only has a relatively small effect. This graph suggests Q_p can be augmented with some simple function of M to enable it to track lines of constant Match through $[\chi_p, q, M]$.

The bottom graph shows that non-zero aligned spin breaks the one-to-one relationship, indicating that there is a degeneracy between the three parameters $[\chi_p, \chi_{eff}, q]$. If this relationship is identified, it may be applied to Q_p in order to generalise its relation with constant Match to cover aligned spin as well. Since anti-aligned spin significantly increases the observability of precession across all orientations, our results suggest this step is most important (assuming equal populations of sources and detector sensitivities, which will be discussed in the following Chapter). However, the relation will be affected by M as well. Creating a general formula to track lines of constant Match through $[\chi_p, q, M, \chi_{eff}]$ will require this inter-dependent relationship between all four parameters and β to be understood. This is far from trivial, but if such an undertaking was successful, the resultant formula would be a powerful tool in the pursuit of creating an effective precessing template bank.

Chapter 5

Conclusion

In order to improve current BBH search sensitivities for precessing sources we must identify which precessing waveforms cannot be found using current non-precessing template banks, and which of those would reach high enough SNR for LIGO to detect. It is computationally infeasible to incorporate all possible precessing waveforms in a template bank, and so any attempt to construct one must use a selective process to find precessing waveforms which contribute to the overall sensitivity to BBH sources [75].

We have made numerous assumptions in order to simplify this initial investigation. We computed the Match only between precessing waveforms and non-precessing waveforms of identical masses and aligned spins, ignoring the possibility that non-precessing waveforms of differing masses and aligned spin may return higher Matches for certain precessing waveforms. We did not consider detector sensitivity, and so the parameter space we have identified as containing “observable precession” may include waveforms which would not likely be powerful enough for LIGO to detect. We have also not considered the population densities of sources: if few possible sources of a given waveform are expected to lie within the volume of space to which a detector is sensitive for such a signal (effective volume, defined in Section 3.5.1), then the waveform contributes less to the search sensitivity than another waveform of more common sources. However, there is as yet little knowledge of precessing BBH population densities [85], whilst we know the loss of coverage to precessing signals from using non-precessing template banks could be as high as 50% [67]. Thus, we forgive ourselves for omitting population densities in our investigation and assuming all signals are equally likely to be detected. However, we also only consider a single-detector observer. Matched-filtering signals across multiple detectors significantly increases the complexity of this problem, as the amplitude of signals and sky location cannot be considered independent. Including such analytic techniques would vastly broaden the scope of the problem, but would be necessary for an “observable precession” formula to be applicable, since a signal which multiple detectors is able to observe must be considered more beneficial to a template bank than a signal only observable by one detector.

However, we have made some first steps towards the ultimate goal of a precessing template bank. We identified where, on the $[\theta_{JN}, \psi]$ plane of constant χ_p , the most

observably precessing waveforms lie in relation to the opening angle for all parameter space investigated. We identified trends in these plots, and analysed why we observe such trends.

We found that waveforms with zero aligned spin and low precessing spin ($\chi_p < 0.3$) do not differ significantly from non-precessing waveform unless the polarisation is near $\psi = 45^\circ$. Anti-aligned spin increases the opening angle and amplifies the effects of precession regardless of orientation, such that even sources with minimal precessing spin could produce observably precessing waveforms if the anti-aligned spin magnitude is sufficiently large. We found that mass ratio acts as an amplification factor on χ_p , and this factor can be used to trace lines of constant Match through the orientation parameter space. Total mass demonstrated less influence on the opening angle compared to χ_{eff} or q when it was reduced, but had a significant impact on the location of observable precession when it was increased. We found that waveform length affects the region of most observable precession some areas of parameter space, whilst the length of time a waveform spends within the detectable frequency bandwidth affects which orientations may see observe the effects of precession for a high mass source.

We have formed a function (Eq. 4.6) which gives the location of the tangent to the precession cone in the $[\theta_{JN}, \psi]$ plane, for any $\phi_{JL}, \chi_p, \chi_{eff}, q, M$. We approximated this tangent to be a circle to form a function for the region of lowest Matches in that plane (Eq. 4.7), applicable to cases of zero aligned spin and masses of our focal set. We propose this function could be augmented with relatively simple dependencies upon χ_{eff} and q to be applicable for much larger areas of parameter space.

We have formed a simple combined parameter, $Q_p = \chi_p q$, which traces lines of approximately constant Match at fixed $[\chi_{eff}, M, \theta_{JN}, \psi, \phi_{JL}]$. We've shown reason to believe it is possible to augment Q_p so that it maps to areas of constant Match through the three-dimensional parameter space $[\chi_p, q, M]$, and perhaps through χ_{eff} as well, although adding the latter would be more challenging.

We conclude that the relationships between parameters and observable precession found in this research suggest that a formula can be constructed to map observable precession over the majority the parameter space in $[\chi_p, \chi_{eff}, q, M, \theta_{JN}, \psi, \phi_{JL}]$. Such a formula would be a powerful tool in the pursuit of identifying which precessing templates would help improve search sensitivities towards precessing BBH.

Our virtues and our failings are inseparable, like force and matter.

Bibliography

- [1] Albert Einstein. *Relativity*. Methuen and Co. Ltd., 1920.
- [2] Albert Einstein. Die grundlage der allgemeinen relativitätstheorie. *Annalen der Physik* 49, 769-822, 1916.
- [3] Bernard F. Schutz. *A first course in general relativity*. Cambridge University Press, 1990.
- [4] A. Abramovici et al. LIGO - The Laser Interferometer Gravitational-Wave Observatory. *Science* 256, 325-333, 1992.
- [5] David Lindley. Focus: A fleeting detection of gravitational waves. *Phys. Rev. Focus*: <https://physics.aps.org/story/v16/st19>, 2005.
- [6] Joseph Weber. Gravitational-wave-detector events. *Phys. Rev. Lett.* 20, 1307, 1968.
- [7] G.W. Gibbons and S.W. Hawking. Theory of the detection of short bursts of gravitational radiation. *Phys. Rev. D* 4, 8, 1971.
- [8] The LIGO Scientific Collaboration. Properties of the binary black hole merger gw150914. *Phys. Rev. Lett.* 116, 241102, 2016.
- [9] The LIGO Scientific Collaboration. Binary black hole mergers in the first advanced ligo observing run. *Phys. Rev. X* 6, 041015, 2016.
- [10] The LIGO Scientific Collaboration. All-sky search for periodic gravitational waves in the o1 ligo data. *Phys. Rev. D* 96, 062002, 2017.
- [11] R.L. Forward. Wideband laser-interferometer gravitational-radiation experiment. *Phys. Rev. D* 17, 379–390, 1978.
- [12] <https://www.ligo.caltech.edu/>. Ligo experiment diagram [public use].
- [13] The LIGO Scientific Collaboration. Advanced ligo. *Class. Quant. Grav.* 32:074001, 2014.
- [14] D.V. Martynov et al. The sensitivity of the advanced ligo detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93, 2016.
- [15] J. Aasi et al. Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light. *Nature Photonics* 7, 613–619, 2013.

- [16] The LIGO Scientific Collaboration. Detector description and performance for the first coincidence observations between ligo and geo. *Nucl. Instrum. Meth. A* **517**, 154-179, 2004.
- [17] The LIGO Scientific Collaboration. Ligo: The laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.* **72**, 076901, 2007.
- [18] The LIGO Scientific Collaboration. Gw170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Phys. Rev. Lett.* **119**, 141101, 2017.
- [19] The LIGO Scientific Collaboration. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.* **116**, 061102, 2016.
- [20] The LIGO Scientific Collaboration. Gw151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.* **116**, 241103, 2016.
- [21] The LIGO Scientific Collaboration. Gw170104: Observation of a 50-solar-mass binary black hole coalescence at redshift 0.2. *Phys. Rev. Lett.* **118**, 221101, 2017.
- [22] The LIGO Scientific Collaboration. Gw170608: Observation of a 19-solar-mass binary black hole coalescence. *Astrophys. J.* **851** no.2, 35, 2017.
- [23] Stephen Fairhurst. Source localization with an advanced gravitational wave detector network. *Class. Quant. Grav.* **28**, 105021, 2011.
- [24] The LIGO Scientific Collaboration. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.* **119**, 161101, 2017.
- [25] Iair Arcavi et al. Optical emission from a kilonova following a gravitational wave-detected neutron-star merger. *Nature* **551**, 64-66, 2017.
- [26] A. Goldstein et al. An ordinary short gamma-ray burst with extraordinary implications: Fermi-gbm detection of grb 170817a. *Astrophys. J.* **848** no.2, 14, 2017.
- [27] Michele Maggiore. *Gravitational Waves Volume 1: Theory and Experiments*. Oxford University Press, 2008.
- [28] R.A. Hulse and J.H. Taylor. Discovery of a pulsar in a binary system. *Astrophys. J.* **195**, 51-53, 1975.
- [29] Eanna E. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences: I. signal to noise for inspiral, merger, and ring-down. *Phys. Rev. D* **57**:4535-4565, 1997.
- [30] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D* **62**, 064015, 2000.

- [31] C.V. Vishveshwara. Stability of the schwartzchild metric. *Phys. Rev. D* **1**, 2870, 1970.
- [32] 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. D* **91**, 024043, 2014.
- [33] Éanna É. Flanagan and Scott A. Hughes. The basics of gravitational wave theory. *New J.Phys.*, 2005.
- [34] J.M. Weisberg, J.H. Taylor, and L.A. Fowler. Gravitational waves from an orbiting pulsar. *Sci. Am.* **245**, 1981.
- [35] The LIGO Scientific Collaboration. The basic physics of the binary black hole merger gw150914. *Ann. Phys. (Berlin)*, **529**, No.1-2, 2016.
- [36] J.G. Baker et al. Consistency of post-newtonian waveforms with numerical relativity. *Phys. Rev. Lett.* **99**, 181101, 2007.
- [37] M. Boyle et al. High-accuracy comparison of numerical relativity simulations with post-newtonian expansions. *Phys. Rev. D* **76**, 124038, 2007.
- [38] M. Hannam et al. Where post-newtonian and numerical-relativity waveforms meet. *Phys. Rev. D* **77**, 044020, 2008.
- [39] Mark the god Hannam. Modelling gravitational waves from precessing black-hole binaries: progress, challenges and prospects. *www.springerlink.com (Gen. Relativ. Grav.)*, 2014.
- [40] The LIGO Scientific Collaboration. Tests of general relativity with gw150914. *Phys. Rev. Lett.* **116**, 221101, 2016.
- [41] K.G. Arun et al. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev. D* **79**, 104023, 2009.
- [42] Sebastian Khan. *Numerical modelling of black-hole-binary mergers*. <http://orca.cf.ac.uk/93841/>, 2016.
- [43] Alejandro Bohé, Sylvian Marsat, and Luc Blanchet. Next-to-next-to-leading order spin-orbit effects in the gravitational wave flux and orbital phasing of compact binaries. *Class. Quant. Grav.* **30**, 135009, 2013.
- [44] Sebastian Khan et al. Frequency-domain gravitational waves from non-precessing black-hole binaries. ii. a phenomenological model for the advanced detector era. *Phys. Rev. D* **93**, 044007, 2015.
- [45] Stanislav Babak, Andrea Taracchini, and Alessandra Buonanno. Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity. *Phys. Rev. D.*, 2016.

- [46] Mark Hannam et al. Simple model of complete precessing black-hole-binary gravitational waveforms. *Phys. Rev. Lett.* **113**, 151101, 2014.
- [47] Alejandro Bohé et al. An improved effective-one-body model of spinning, non-precessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Phys. Rev. D* **95**, 024010, 2017.
- [48] Thibault Damour, Piotr Jaranowski, and Gerhard Schäfer. Fourth post-newtonian effective one-body dynamics. *Phys. Rev. D* **91**, 084024, 2015.
- [49] Francesco Messina, Alberto Meldarella, and Alessandro Naga. Factorization and resummation: A new paradigm to improve gravitational wave amplitudes. ii: the higher multipolar modes. *Phys. Rev. D* **97**, 084016, 2018.
- [50] Sylvain Marsat et al. Next-to-leading tail-induced spin-orbit effects in the gravitational radiation flux of compact binaries. *Class. Quant. Grav.* **31**, 025023, 2013.
- [51] Tanguy Marchand, Luc Blanchet, and Guillaume Faye. Gravitational-wave tail effects to quartic non-linear order. *Class. Quant. Grav.* **33**, 244003, 2016.
- [52] James Healy, Carlos Lousto, and Yosef Zlochower. Remnant mass, spin, and recoil from spin aligned black-hole binaries. *Phys. Rev. D* **90**, 104004, 2014.
- [53] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.* **95**:121101, 2005.
- [54] Lionel London, James Healy, and Deirdre Shoemaker. Modeling ringdown: Beyond the fundamental quasinormal modes. *Phys. Rev. D* **90**, 124032, 2017.
- [55] Mark Hannam. Status of black-hole-binary simulations for gravitational-wave detection. *Class. Quant. Grav.* **26**:114001, 2009.
- [56] Abdul H. Mouré et al. Catalog of 174 binary black hole simulations for gravitational wave astronomy. *Phys. Rev. Lett.* **111**:241104, 2013.
- [57] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries. *Phys. Rev. D* **49**, 6274, 1994.
- [58] Curt Cutler et al. The last three minutes: Issues in gravitational-wave measurements of coalescing compact binaries. *Phys. Rev. D* **70**, 2984, 1993.
- [59] Will M. Farr et al. Distinguishing spin-aligned and isotropic black hole populations with gravitational waves. *Nature*, 2017.
- [60] E.A. Huerta et al. Complete waveform model for compact binaries on eccentric orbits. *Phys. Rev. D* **95**, 024038, 2017.
- [61] E.A. Huerta et al. Accurate and efficient waveforms for compact binaries on eccentric orbits. *Phys. Rev. D* **90**, 084016, 2014.

- [62] F. Antonini and B.P. Hagai. Secular evolution of compact binaries near massive black holes: Gravitational wave sources and other exotica. *Astrophys. J.* **757**, 1, 2012.
- [63] Ryan O’Leary et al. Gravitational waves from scattering of stellar-mass black holes in galactic nuclei. *Mon.Not.Roy.Astron.Soc.* **395**, 4, 2009.
- [64] Johan Samsing et al. The formation of compact binary inspirals and the role of gravitational wave emission in binary-single stellar encounters. *Astrophys. J.* **784**, 1, 2014.
- [65] Duncan A. Brown and Peter J. Zimmerman. The effect of eccentricity on searches for gravitational-waves from coalescing compact binaries in ground-based detectors. *Phys. Rev. D* **81**, 024007, 2009.
- [66] Alessandra Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D.*, 1999.
- [67] Tito Dal Canton, Andrew Lundgren, and Alex Nielson. Impact of precession on aligned-spin searches for neutron-star - black-hole binaries. *Phys. Rev. D* **91**, 062010, 2015.
- [68] 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. D* **86**, 104063, 2012.
- [69] Benjamin Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D* **60**, 022002, 1998.
- [70] The LIGO Scientific Collaboration. Instrument science white paper. *LIGO Public Documents: LIGO-T1400316-v4*, 2014.
- [71] Curt Cutler and Éanna Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Phys. Rev. D* **49**, 2658, 1994.
- [72] Kumar Prayush et al. Template banks for binary black hole searches with numerical relativity waveforms. *Phys. Rev. D* **89**, 042002, 2013.
- [73] The LIGO Scientific Collaboration. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Phys. Rev. D* **72**, 082001, 2005.
- [74] The LIGO Scientific Collaboration. Gw150914: First results from the search for binary black hole coalescence with advanced ligo. *Phys. Rev. D* **93**, 122003, 2016.
- [75] Ian Harry, Stephen Privitera, Alejandro Bohé, and Alessandra Buonanno. Searching for Gravitational Waves from Compact Binaries with Precessing Spins. *Phys. Rev. D* **94**, 024012, 2016.

- [76] Chris Van Den Broeck et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Phys. Rev. D* **80**, 024009, 2009.
- [77] Ian Harry, Bruce Allen, and B. S. Sathyaprakash. A stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D* **80**, 104014, 2009.
- [78] T Cokelaer. Gravitational waves from inspiralling compact binaries: Hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D* **76**, 102004, 2007.
- [79] Emily Baird et al. Degeneracy between mass and spin in black-hole-binary waveforms. *Phys. Rev. D* **87**, 024035, 2012.
- [80] S. Babek et al. A template bank to search for gravitational waves from inspiralling compact binaries: I. physical models. *Class. Quant. Grav.* **23** 5477-5504, 2006.
- [81] The LIGO Scientific Collaboration. Gw151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.* **116**, 241103, 2016.
- [82] Andrew Lundgren and R. O’Shaughnessy. A single-spin precessing gravitational wave in closed form. *Phys. Rev. D* **89**, 044021, 2014.
- [83] LIGO Algorithm Library [LALsuite]. <https://git.ligo.org/lscsoft/lalsuite>.
- [84] Duncan Brown, Andrew Lundgren, and Richard O’Shaughnessy. Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation. *Phys. Rev. D* **86**, 064020, 2012.
- [85] The LIGO Scientific Collaboration and The Virgo Collaboration. Binary black hole population properties inferred from the first and second observing runs of advanced ligo and advanced virgo. *arXiv:1811.12940 [not yet published]*, 2018.

