SPIN-PRECESSING COMPACT BINARIES: GRAVITATIONAL WAVE MODELING AND INFORMATION EXTRACTION

by

Katerina Chatziioannou

A dissertation submitted in partial fulfillment
of the requirements for the degree

of

Doctor of Philosophy

in

Physics

MONTANA STATE UNIVERSITY
Bozeman, Montana

June, 2016

ACKNOWLEDGEMENTS

I would like to thank my collaborators Neil Cornish, Antoine Klein, Eric Poisson, Kent Yagi, and, in particular, my advisor Nicolás Yunes for teaching me how research is/should be done and how I can be a part of the research world.

Funding Acknowledgment

This work was supported by the Onassis Foundation, NSF Grant No. PHY-1114374 and NASA Grant No. NNX11AI49G, under subaward 00001944.

TABLE OF CONTENTS

TABLE OF CONTENTS - CONTINUED

TABLE OF CONTENTS - CONTINUED

LIST OF TABLES

LIST OF FIGURES

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

LIST OF FIGURES - CONTINUED

Figure Page

ABSTRACT

In this dissertation we study the effect of spin-precession on gravitational waves emitted by quasicircular compact binary systems. In their most generic configuration, compact objects in a binary system are subject to interactions between the spin and the orbital angular momenta. These interactions give rise to precessional effects that add rich structure to the emitted gravitational waveforms. We study this spin-induced structure with an emphasis on extracting the information it encodes. In particular, we construct gravitational wave models that accurately capture spin-precessional effects. We then use them to study how much information relevant to astrophysics and nuclear physics we can extract from future observations of gravitational waves from compact binary coalescences.

1

INTRODUCTION

The recent discovery of gravitational waves (GWs) has been heralded as one of the most important scientific achievements of the last decades. The study of the Universe in the GW spectrum, known as *gravitational wave astronomy*, promises to illuminate electromagnetically inaccessible phenomena as well as compliment conventional studies. Probably the most promising source of GWs –as well as the subject of the first detection– is the coalescence of compact objects such as black holes and neutron stars. When two compact objects inspiral under the influence of gravitational radiation, they exhibit an interesting feature: their spin angular momenta interact and cause the binary to change orientation in space, an effect known as *spin-precession*. The complexity of spin-precessing renders it difficult, yet important, to study owing to its influence on the GW emitted.

In this dissertation we study spin-precession, how it affects the GW emitted, and what we can learn by detecting such a wave. We first construct appropriate templates that can be used for the detection of GWs from spin-precessing binary systems as well as for extracting information about them. We, then, put these models to the test and study the quantity and quality of information we can extract about Nature by spin-precessing binary systems. We find that accurately accounting for spin-precession can improve the measurement of the system's masses and spins by up to an order of magnitude. Finally, we study what future GW observations can reveal about the internal composition of neutron stars, information pertaining to the yet-unknown equation of state of supranuclear matter.

The rest of this thesis is organized as follows. In Ch. 2, we describe spin-precession within the framework of General Relativity (GR). In Ch. 3, we discuss GW templates, their role in GW detection and characterization, as well as the difficulties

of constructing a template bank for spin-precessing compact binaries. In Ch. 4, we describe a first attempt to construct spin-precessing templates that apply to slowly spinning compact binaries. In Ch. 5, we derive a spin-precessing template bank that is valid for generic compact binaries. In Chs. 6 and 7, we describe how the use of spin-precessing templates enhances the relevant astrophysical information we can extract from GWs. In Ch. 8, we study the implications of spin-precessing templates on extracting information that is applicable to nuclear physics. Finally, in Ch. 9 we summarize and conclude.

Chapters 4, 5, 6, 7, and 8 are based on my original work; this work has been done in collaboration with a number of researchers. Chapters 4 [62], 5 [61], 6 [59], and 7 [60] have been carried out in collaboration with Dr. Antoine Klein, Prof. Nicolás Yunes, and Prof. Neil Cornish. Chapter 8 [64] has been carried out in collaboration with Dr. Kent Yagi, Dr. Antoine Klein, Prof. Nicolás Yunes, and Prof. Neil Cornish. All the work described in this dissertation has been either published or submitted for publication in peer-reviewed journals. My participation in all the above projects has been extensive. I am the first author in all the publications and I carried out the majority of the analytical or numerical calculations.

3

SPIN PRECESSION IN GR

The stellar mass quasicircular compact binaries ground-based detectors advanced LIGO (aLIGO) and advanced Virgo (AdV) have already started detecting are highly relativistic systems; close to the final coalescence the binary components orbit each other with velocities exceeding half the speed of light and the gravitational field changes rapidly. These extreme conditions are ideal testbeds for a very interesting relativistic phenomenon with no Newtonian analogue: in GR every form of energy gravitates, implying that not only matter, but also the *flow* of matter affects the spacetime geometry. In the case of compact binaries, the evolution of the system and the GWs it emits depend not only on the masses of the bodies, but also on their spin angular momenta.

Spin is ubiquitous in Nature; under the influence of a generic perturbation any astrophysical system will rotate even if the initial state was perfectly spherically symmetric. When massive stars give birth to neutron stars (NSs) or black holes (BHs) through supernova explosions, the newly-born remnant typically spins rapidly even if the progenitor was spinning slowly. This is possibly due to asymmetries in the supernova explosion inducing a "kick" on the remnant causing it to rotate [203]. At a fundamental level, the physics at play is the same as that which causes a soccer ball or a pineapple to spin after kicked.

The spin angular momenta of the components of a compact binary system will not necessarily be aligned with the orbital angular momentum. For example, consider an isolated binary system of two stars with spins aligned with the orbital angular momentum. The most massive star will first fill its Roche lobe and transfer mass to the companion before going supernova. The compact remnant (BH or NS) receives a kick spinning it up [203] and tilting the orbital plane, since the kick's direction is typically

correlated with the spin of the exploding star [226,227]; the various angular momenta become misaligned [101,118]. Eventually, the second star also fills its Roche lobe and the binary enters a common envelope phase. In this phase, the angular momenta could partially align through tidal effects. But, again, the phase ends with the star going supernova and endowing the binary with a kick that spins the remnant up and typically tilts the orbital plane, misaligning the angular momenta yet again [101,118].

Few mechanism exist that could prevent misalignment or re-align the spins of compact binary components with the orbital angular momentum. One possibility is if the supernova kicks are in the orbital plane, i.e. *perpendicular* to the spin angular momentum, such that the orbital plane is not tilted [118]. This possibility is remote, with models and data suggesting that the kick is actually *aligned* with the spin angular momentum [119, 226, 227]. A mechanism for re-alignment is through torques exerted by a circumbinary accretion disk on BH binaries [45]. Such disks, however, are only expected in galaxy mergers of supermassive BHs, whose GWs would be outside the sensitivity band of aLIGO and AdV. Finally, we note that the binary might undergo spin-orbit resonances during its evolution [193]. These resonances, however, do not align the spins with the orbital angular moment, but rather maintain certain special ("resonant") configurations [99–101, 121, 193].

Compact binaries with significant spins that are also misaligned with the orbital angular momenta are qualitatively different from non-spinning binaries. The various spin-orbit and spin-spin interactions they exhibit have been the focus of extensive theoretical study in the past decades, leading to an ever improving framework to analyze and incorporate spins in GW science. Largely following the treatment of [43] and [174], we provide a brief description of that framework here.

Post-Newtonian Theory and Spin

To date, no exact analytic solutions exist to the two body problem in full GR. A number of perturbative approaches with generally distinct (though sometimes overlapping) regions of validity have been developed and tested. In the context of compact binaries, the *post-Newtonian* (PN) formalism has been shown to be extremely successful at describing systems that are not too close to coalescence. PN theory treats the gravitational field as Newtonian plus small corrections which for compact binaries means that the characteristic velocities are small compared to the speed of light and systematically solves the Einstein equations as an expansion in small velocities. Results obtained in this framework are Newtonian to leading order with corrections that are inversely proportional to the speed of light.

Single Body

Before we describe the complicated PN dynamics of compact binaries, we examine a single spinning body as a way to acquire some intuition on the effect of spin on the spacetime. Within the PN approximations, a straightforward –though lengthy– calculation gives the spacetime metric at $\boldsymbol{x}$ close to single spinning body of mass M and spin angular momentum 3-vector $\boldsymbol{S}$ at a coordinate system where the body is at rest and at the origin

$$g_{00} = -1 + \frac{2}{c^2}\frac{GM}{r} - \frac{2}{c^4}\left(\frac{GM}{r}\right)^2 + \mathcal{O}(c^{-6}), \tag{2.1}$$

$$g_{0j} = \frac{2}{c^3}\frac{G(\boldsymbol{x}\times\boldsymbol{S})^j}{r^3} + \mathcal{O}(c^{-5}), \tag{2.2}$$

$$g_{jk} = \left(1 + \frac{2}{c^2}\frac{GM}{r}\right)\delta_{jk} + \mathcal{O}(c^{-4}), \tag{2.3}$$

where we employ the usual convention that latin indices run through $\{1, 2, 3\}$ denoting space coordinates and greek indices run through $\{0, 1, 2, 3\}$. In the above expression c is the speed of light, G is the gravitational constant, $r = |\boldsymbol{x}|$, and bold-faced symbols stand for vectors.

The metric is obtained by solving the Einstein equations perturbatively under the assumption of weak fields to second order. The metric is comprised of a part corresponding to the usual Minkowski metric of flat spacetime, a first correction corresponding to Newtonian gravity, and a second correction being the first pure relativistic effect we encounter in PN theory, usually referred to as 1PN. That means that it is of $\mathcal{O}(c^{-2})$ relative to its leading order controlling factor.

A test body of mass m and 4-velocity $u^\alpha = (c, \boldsymbol{u})$ will follow geodesics of this spacetime, described by the Lagrangian

$$\begin{aligned} L &= -mc\sqrt{-g_{\alpha\beta}u^\alpha u^\beta} \\ &= -mc^2 + \frac{1}{2}m(\dot{r}^2 + r^2\dot{\theta}^2 + r^2\sin^2\theta\dot{\phi}^2) + \frac{GMm}{r} - \frac{2GmS\sin^2\theta\dot{\phi}}{c^2 r} + \mathcal{O}(c^{-2}), \end{aligned} \tag{2.4}$$

where for simplicity we only keep Newtonian and spin terms and assume the spin is in the z direction and $S \equiv S_z$. The overdot indicates the usual derivative with respect to time.

In the above Lagrangian ϕ is a cyclic coordinate, implying the conservation of

$$h \equiv \frac{\partial L}{\partial \dot{\phi}} = mr^2\sin^2\theta\left(\dot{\phi} - \frac{2GS}{c^2 r^3}\right). \tag{2.5}$$

The conservation of h means that zero angular velocity and zero angular momentum are not equivalent. A test particle starting from rest at infinity and falling towards a spinning body will acquire an angular velocity even though it has

no angular momentum. This effect is known as *dragging of inertial frames* and gives an idea of how spacetime and its geodesics are "dragged along" by rotation.

Quasicircular Compact Binary

The equations of motion of a quasicircular compact binary can be computed in the PN framework through the energy-momentum conservation law and with knowledge of the appropriate PN metric from the Einstein equations. Carrying out this calculation we get the relative acceleration of the two binary components with masses (M_1, M_2), spins $(\boldsymbol{S}_1, \boldsymbol{S}_2)$, relative velocity $\boldsymbol{v}$, relative position $\boldsymbol{r} = r\boldsymbol{n}$, and total mass m in the barycentric frame

$$\boldsymbol{a} = a[\text{0PN}] + a[\text{1PN}] + a[\text{SO}] + a[\text{SS}] + \mathcal{O}(c^{-4}), \tag{2.6}$$

where

$$a[\text{0PN}] = -\frac{Gm}{r^2}\boldsymbol{n}, \tag{2.7}$$

is the well-known Newtonian gravitational acceleration,

$$a[\text{1PN}] = \frac{Gm}{c^2 r^2}\left(3 - \eta\right) v^2 \boldsymbol{n}, \tag{2.8}$$

is its first relativistic correction with $\eta \equiv M_1 M_2/m^2$, and $a[\text{SO}], a[\text{SS}]$ are the spin-orbit (linear in spin) and spin-spin (quadratic in spin) contributions respectively

$$a[\text{SO}] = \frac{G}{c^2 r^3}\left\{\frac{3}{2}(\boldsymbol{n}\times\boldsymbol{v})\cdot[3\boldsymbol{\sigma} + 4\boldsymbol{S}]\,\boldsymbol{n} - \boldsymbol{v}\times(3\boldsymbol{\sigma} + 4\boldsymbol{S})\right\}, \tag{2.9}$$

$$a[\text{SS}] = -\frac{3Gm}{c^2 r^4}\left[(\hat{\boldsymbol{S}}_1\cdot\hat{\boldsymbol{S}}_2)\boldsymbol{n} - 5(\hat{\boldsymbol{S}}_1\cdot\boldsymbol{n})(\hat{\boldsymbol{S}}_2\cdot\boldsymbol{n})\boldsymbol{n} + (\hat{\boldsymbol{S}}_1\cdot\boldsymbol{n})\hat{\boldsymbol{S}}_2 + (\hat{\boldsymbol{S}}_2\cdot\boldsymbol{n})\hat{\boldsymbol{S}}_1\right]. \tag{2.10}$$

In the above expressions we have defined $\boldsymbol{\sigma} \equiv (M_2/M_1)\boldsymbol{S}_1+(M_1/M_2)\boldsymbol{S}_2$, $\boldsymbol{S} \equiv \boldsymbol{S}_1+\boldsymbol{S}_2$, $\hat{\boldsymbol{S}}_1 \equiv \boldsymbol{S}_1/M_1$, and $\hat{\boldsymbol{S}}_2 \equiv \boldsymbol{S}_2/M_2$.

Spin Evolution

With these results at hand, we can derive an evolution equation for the binary components' spins. In GR, the spin of body A in the center-of-mass frame is

$$\boldsymbol{S}_A = \int_A \rho^* \bar{\boldsymbol{x}} \times \bar{\boldsymbol{v}}\, d^3\bar{x}, \tag{2.11}$$

where the integration is carried over a volume enclosing only body A, $\bar{\boldsymbol{x}}$ and $\bar{\boldsymbol{v}}$ are position and velocity in the center of mass frame and $\rho^* \equiv \rho\sqrt{-g}$, where ρ is the mass density, and g the metric determinant.

The evolution of spin can then be obtained through

$$\frac{dS_A^j}{dt} = \epsilon^{jpq} \int_A \rho^* \bar{x}^p \frac{d\bar{v}^q}{dt} d^3\bar{x}, \tag{2.12}$$

where ϵ^{jpq} is the Levi-Civita symbol. The quantity $d\bar{v}^q/dt$ is similar to the acceleration derived in Eq. (2.6) and the integral of Eq. (2.12) can be solved by a calculation similar to the one that led to Eq. (2.6). A final lengthy calculation reveals that in the case of quasicircular compact binaries, the evolution of the spin of body A when interacting with body B is given by

$$\frac{d\bar{\boldsymbol{S}}_A}{dt} = (\boldsymbol{\Omega}_A[\mathrm{SO}] + \boldsymbol{\Omega}_A[\mathrm{SS}]) \times \bar{\boldsymbol{S}}_A \tag{2.13}$$

with

$$\boldsymbol{\Omega}_A[\mathrm{SO}] \equiv \frac{2\eta G m}{c^2 r^2}\left(1+\frac{3M_B}{4M_A}\right)\boldsymbol{n}\times\boldsymbol{v}, \tag{2.14}$$

$$\boldsymbol{\Omega}_A[\mathrm{SS}] \equiv \frac{Gm}{c^2 r^3}\left[3(\bar{\boldsymbol{S}}_B\cdot\boldsymbol{n})\boldsymbol{n}-\bar{\boldsymbol{S}}_B\right], \tag{2.15}$$

describing spin-orbit and spin-spin interactions respectively.

The spin-precession equation (2.13) has been written in terms of $\bar{\boldsymbol{S}}_A = \boldsymbol{S}_A + \Delta\boldsymbol{S}_A/c^2$, where $\Delta\boldsymbol{S}_A$ are PN corrections that arise from the bodies internal interactions, as well as from external interactions with its environment. The internal interactions take the form of total time derivatives that are moved from the right-hand side of Eq. (2.12) to the left-hand side. The external interaction corrections are chosen such that $\bar{\boldsymbol{S}}_A$ corresponds to the spin vector measured in a non-inertial frame that is moving with the body (*proper* or *comoving* spin). In the following chapters we will study GWs from spinning compact binaries, in a framework that has been developed in terms of $\bar{\boldsymbol{S}}_A$. For simplicity, we will drop the bar and call $\bar{\boldsymbol{S}}_A$ "the spin of body A".[1]

Equation (2.13) describes the evolution of the spins in a compact binary system. Its form implies that the magnitudes of the spins remain constant while they precess around the vector $\boldsymbol{\Omega}_A[\mathrm{SO}] + \boldsymbol{\Omega}_A[\mathrm{SS}]$.

[1]Besides the definition of the comoving spins, additional freedom remains in the selection of the frame in which we define the world line of the body's center of mass. The choice of this frame is known in the literature as *supplemental spin condition* and it affects the equations of motion, but leaves the comoving spin and the spin-precession equations of this section unaltered. For the equations of motion of the previous section, and specifically Eq. (2.9), we employ a representative world line that tracks the body's center of mass in a global inertial frame.

Orbit-Averaged Spin Evolution

Neglecting the effects of radiation reaction which cause the orbit to shrink (in a PN sense, these effects are of higher order) the motion of the binary can be divided in two parts: orbital motion, evolving by Kepler's law on a timescale of

$$T_{\rm orb} \sim r^{3/2} \sim v^{-3}, \tag{2.16}$$

and precessional motion, evolving on

$$T_{\rm pr} \sim |\mathbf{\Omega}_A[\mathrm{SO}]|^{-1} \sim \frac{r^2}{v} \sim v^{-5}. \tag{2.17}$$

Since $T_{\rm orb}/T_{\rm pr} \sim v^2 \ll 1$ in a PN sense, we can average over the fast orbital motion and obtain the *orbit-averaged spin-precession equations*

$$\frac{d\bar{\boldsymbol{S}}_A}{dt} = \frac{2\eta G m v}{c^2 r^2}\left(1+\frac{3M_B}{4M_A}\right)\hat{\boldsymbol{L}}\times\bar{\boldsymbol{S}}_A - \frac{3Gm}{2c^2r^3}(\bar{\boldsymbol{S}}_B\cdot\hat{\boldsymbol{L}})\hat{\boldsymbol{L}}\times\bar{\boldsymbol{S}}_A + \frac{Gm}{c^2r^3}\bar{\boldsymbol{S}}_B\times\bar{\boldsymbol{S}}_A. \tag{2.18}$$

where $\boldsymbol{L}$ is the orbital angular momentum of the system, and $\hat{\boldsymbol{L}}$ its direction. Notice that when the spins are aligned with the angular momentum the above differential equation becomes trivial and there is no precession. The misalignment of the bodies' spins with the orbital angular momentum is essential for the system to be precessing and indeed this is the astrophysically expected scenario, as argued at the beginning of this chapter.

Finally, we can derive an equation for the evolution of $\boldsymbol{L}$ by noting that in the absence of radiation reaction the total angular momentum needs to be conserved

$$\frac{d\boldsymbol{L}}{dt} = -\frac{d\bar{\boldsymbol{S}}_A}{dt} - \frac{d\bar{\boldsymbol{S}}_B}{dt}. \tag{2.19}$$

Equations (2.18), (2.18) with $A \leftrightarrow B$, and (2.19) form a system of 9 coupled differential equations for 9 unknowns and prescribe the evolution of all angular momenta of a binary system in the absence of radiation reaction. In Chapters 4 and 5 we will include radiation reaction and solve the spin-precession equations in order to construct waveforms for spin-precessing, quasicircular compact binaries. In some cases we will include PN corrections, or express the equations in alternative, yet equivalent, forms. We will, therefore, state which precise form of the spin-precession equations we are studying in each chapter and why.

12

SPIN-PRECESSING GW TEMPLATES

Quasicircular compact binaries emitting GWs in the sensitivity band of second-generation ground based detectors such as aLIGO [68,108] and AdV [5] will have spins with arbitrary magnitudes and directions, though the spin of NSs will typically be smaller than that of BHs[1]. Such GWs are very different from those emitted by non-spinning binaries or binaries with spins aligned with the orbital angular momentum for the reason described in Chapter 2: *spin-precession*.

When spins are misaligned with the orbital angular momentum, all angular momenta precess about the total angular momentum, causing the orbital plane to precess too. This induces amplitude and phase modulations in the GWs, for example through the changing inclination angle of the system relative to the line of sight. As an example, in Fig. 3.1 we plot the scaled leading PN order amplitude for spin-precessing, spin-aligned, and nonspinning GWs as a function of the GW frequency. The precessing waveform amplitude is qualitatively different from the nonprecessing ones, something that is also encountered in the GW phase. The latter is more important from a data analysis point of view, but more difficult to visualize, so we have plotted the amplitude only. Accurately modeling precessional modulations can be important for detection and parameter estimation of GW sources [59, 60].

GWs with such rich and complex structure are double-edged swords: on the one hand, this structure can encode new information about the source and break degeneracies in parameter estimation (see Chapter 7); on the other hand, this intrinsic complexity comes at the cost of an increased difficulty to model these waves. First, the temporal evolution of the orbital phase depends on the angular momenta, which

[1] In this thesis, we only consider quasicircular binaries since radiation reaction tends to circularize systems formed with eccentricity before they enter the detectors' sensitivity band [171].

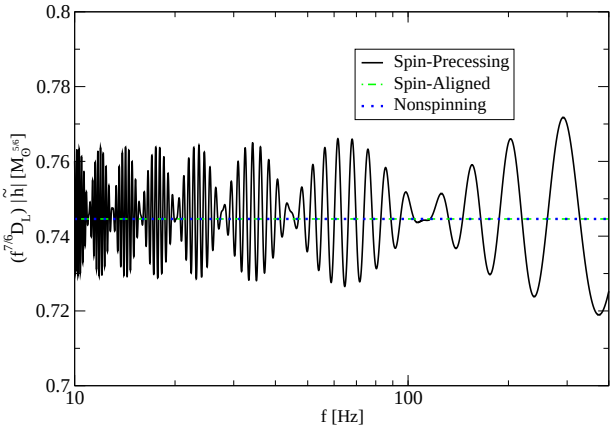

Figure 3.1: Leading PN order amplitude of a spin-precessing (black solid), a spin-aligned (green dot-dashed), and a nonspinning (blue dotted) GW as a function of the GW frequency f (twice the orbital frequency) scaled by $f^{7/6}$ and the luminosity distance of the source D_L.

themselves satisfy certain precession equations, increasing the overall complexity of the differential system. Second, precession introduces *mathematical catastrophes* when computing the Fourier transform of the GWs: degenerate critical points in the orbital phase, where the first and second time derivatives vanish, violating the assumptions of the standard *stationary phase approximation*[2] (SPA) and rendering it non-applicable [125]. This second point is important since aLIGO and AdV work with frequency-domain models for which data analysis is simplified and sped-up significantly[3].

Accurate modeling of GWs is important for detection and crucial for parameter estimation with ground-based detectors, since the expected signals could be deeply buried in the detector noise. The optimal strategy for extracting known signals from

[2]The SPA is the leading-order term in the asymptotic expansion of a Fourier integral through the method of steepest descent [37, 84].

[3]This is because for stationary noise the noise covariance matrix is diagonal in the frequency domain, but not in the time domain.

noise is by fitting *waveform models* to the interferometric data and minimizing the residual. Obtaining some parametrized waveform model and using it to filter the data is therefore essential for recovering the optimal signal-to-noise ratio (SNR)

$$\mathrm{SNR}^2 = 4 \int \frac{|\tilde{h}(f)|^2}{S_n(f)} df \tag{3.1}$$

from the signal. In this expression $S_n(f)$ is the detector's spectral noise density and $\tilde{h}(f)$ is the GW expressed in the Fourier-domain. The efficiency of this method relies on the accuracy of the waveform model, with parameter estimation placing more stringent requirements on the accuracy of the models used. Inaccuracies in the models can lead to missed signals or systematic errors in the extracted parameters.

This has motivated the construction of waveform models for coalescing compact binaries. During the inspiral, the binary can be modeled with the post-Newtonian formalism, an expansion in small characteristic velocities and weak gravitational fields [43]. When the binary components are not spinning or when their spin is aligned/anti-aligned with the orbital angular momentum, the equations of motion have been derived and solved up to 3.5PN order[4] including radiation reaction due to GW energy loss. When the spins are misaligned with the orbital angular momentum, the orbital equations of motion and the precession equations have been derived to 2.5PN order. In this case, a closed form solution has not been obtained due to the complexity of the differential system.

To this day, four main representations of GWs from spin-precessing compact binaries exist. The first representation is based on the fact that the precession equations admit a closed-form analytic solution when only one object is spinning [18–

[4]A term is of APN order is proportional to $(u/c)^{2A}$ relative to its controlling factor, where u is some characteristic velocity and c the speed of light.

21]. The ensuing motion is *simple precession* and the resulting waveform is ideal for BHNS systems [144]. The second representation is based on the *effective-one-body* formulation of the general relativistic two-body problem [53, 75]. The resulting waveform is ever improving through fits to numerical relativity simulations [166, 210] and describes the full coalescence, albeit at the expense of prohibitive computational cost. The third representation utilizes a coordinate frame in which precessional effects are minimized [49, 191] to compute a simpler waveform and map it back to the source frame [190, 192]. This approach applies to the full binary coalescence and the waveform was found to be sufficiently good for detection, but could introduce parameter biases in parameter estimation [107].

The final representation of GWs from inspiraling spin-precessing systems is through a multiple scale analysis (MSA) [37], a well-known mathematical technique to solve differential systems that have distinct characteristic scales by expanding in the ratio of these scales. For the problem at hand, the orbital timescale is much shorter than the precession timescale, which in turn is much shorter than the radiation reaction timescale. These waveforms are constructed so that they are purely analytical, and thus computationally inexpensive, but also so that they can be systematically improved by carrying out calculations at higher order in perturbation theory. Reference [124] implemented such a perturbative framework for compact binaries with spins almost aligned with the orbital angular momentum, but of arbitrary magnitude. In this dissertation, we first apply MSA to slowly spinning [62] systems and show the resulting waveforms to be accurate representations of NSNS inspirals, both for detection and parameter estimation [59, 60, 64]. We then solve the same problem but for generic precession [61] and construct a waveform that is an accurate representation of generic compact binary inspirals of arbitrary masses and spins.

In the next two chapters we construct first a GW template valid for slowly spinning binary components and then one valid for arbitrary binary parameters. Throughout, we use geometric units where $G = c = 1$ and employ the following conventions:

- Vectors are written in boldface, with components $\boldsymbol{A} = [A_x, A_y, A_z]$ and magnitude A. Unit vectors are denoted with a hat, e.g. $\hat{\boldsymbol{A}}$.

- The masses of the two binary components are m_A, with $\{A, B\} \in \{1, 2\}$, the total mass $M \equiv m_1 + m_2$ is set equal to 1, the mass ratio is $q \equiv m_2/m_1 < 1$, the symmetric mass ratio is $\eta \equiv m_1 m_2$ and the mass difference is $\delta m = m_1 - m_2$.

- The Newtonian orbital angular momentum of the system is $\boldsymbol{L}$, the spin angular momentum of each body is $\boldsymbol{S}_A$, and the total angular momentum is $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}_1 + \boldsymbol{S}_2$. The dimensionless spin parameter of each object is $\chi_A \equiv S_A/m_A^2$ with $A \in \{1, 2\}$.

- The orbital angular frequency in a frame fixed to the orbital plane is ω, while the PN expansion parameter we use is $v \equiv \omega^{1/3} = \eta L^{-1}$.

17

GW TEMPLATES FOR SLOWLY SPINNING BINARIES

NSs will take a long time to evolve into the sensitive band of ground-based detectors; by then, they are expected to be old, cold, slowly spinning, and with orbits that have circularized. Accurate predictions of the range of eccentricities and spins expected is currently lacking, but estimates suggest values smaller than 0.1 in dimensionless units [106]. This is why the modeling of NSNS inspirals has so far been restricted to mostly quasicircular systems that are nonspinning or have aligned/antialigned spins are are thus not precessing.

In this chapter we propose a new template family that fits in the analytical class for quasicircular inspiraling compact binaries, whose spin angular momentum is arbitrarily oriented, but small relative to their mass. This family is ideal to efficiently model precessing NSNS inspirals for detection and parameter estimation purposes. The waveforms are obtained through MSA, which allows us to accurately solve the precession equations in the time domain, and the SPA method, which allows us to model the Fourier transform of the waveform. Such a procedure is a direct application of the general framework of [124] but for a different physical system.

The time-domain waveforms are obtained by expanding the differential equations describing precession in the dimensionless spin parameters, which then leads to a separation of time scales that is amenable to MSA: the differential system is reexpanded in the ratio of the precession to the radiation-reaction time scales and solved order by order. We thus obtain an analytic solution to the precession equations for the orbital and spin angular momenta as a bivariate expansion in the dimensionless spin and the ratio of the precession to the radiation-reaction time scales. We work here to first order in both expansion parameters, but the formalism can easily be extended to higher order. An analytic approximation for the time evolution of the

momenta then leads to an analytic time-domain waveform that accurately models the precession of the spin and orbital angular momenta.

The resulting, analytic, time-domain waveform contains pieces of different PN and spin order. Clearly, the zeroth order in spin terms are the dominant contribution to the waveform, since we are expanding about small dimensionless spin. The nonspinning terms in the waveform are a "background solution" to which we find spin perturbations; we will try to model this background as accurately as possible. In particular, the nonspinning orbital frequency evolution equation is modeled to 3.5PN order with mass-ratio corrections and up to 5.5PN order in the test particle limit. We solve this equation to 8PN order, which artificially extends the series beyond its formal expansion order, so that any differences between our analytic and numerical frequency evolution are exclusively due to spin. We also include nonspinning PN corrections to the time-domain waveform amplitude up to 2.5PN order.

The first-order-in-spin terms in the waveforms are perturbations to the nonspinning waveform background and we model them by expanding in the ratio of the precession to the radiation reaction time scale. We model the spin-orbit precession equations to 2PN order and include first-order-in-spin terms in the frequency evolution equation up to 4PN order with mass ratio corrections. Such precession effects introduce corrections in the waveform phase starting at 3PN order.

The second-order-in-spin terms are even smaller perturbations to the nonspinning background, which we will not systematically include here. For example, the spin precession equations contain spin-spin interactions that we will ignore when analytically solving for the angular momenta. However, we will keep them in the frequency evolution, since they lead to better agreement between our analytical results and numerical solutions. This results in the inclusion of a 2PN second-order-in-spin term in the waveform phase that is evaluated assuming the spins are time independent.

Evolution of the Spin and Orbital Angular Momentum

In this section, we solve for the time evolution of the spin and orbital angular momentum as a series in the dimensionless spin parameters and the ratio of the precession to the radiation reaction time scales. We begin by expanding the precession equations in small spin magnitude. We then solve them by order in this perturbation parameter. At linear order in spin, we solve the precession equations using MSA. We conclude with a comparison between the numerical solution and the analytic approximation to the evolution of the angular momenta.

Precession Equations and the Small Spin Expansion

In the PN approximation, the equations that govern the evolution of the orbital angular momentum $\boldsymbol{L}$ and the spin angular momenta $\boldsymbol{S}_i$ averaged over one orbit are [28, 48]

$$\dot{\boldsymbol{L}} = \omega^2 \sum_{n=0}^{2} \eta^n \omega^{2n/3} \left[C_1^{(\mathrm{n})} \boldsymbol{S}_1 + C_2^{(\mathrm{n})} \boldsymbol{S}_2\right] \times \boldsymbol{L} - \frac{3}{2}\omega^2 [(\boldsymbol{S}_2 \cdot \hat{\boldsymbol{L}})\boldsymbol{S}_1 + (\boldsymbol{S}_1 \cdot \hat{\boldsymbol{L}})\boldsymbol{S}_2] \times \hat{\boldsymbol{L}} - k\boldsymbol{L}\,, \tag{4.1}$$

$$\dot{\boldsymbol{S}}_1 = \omega^2 \sum_{n=0}^{2} C_1^{(\mathrm{n})} \eta^n \omega^{2n/3} \boldsymbol{L} \times \boldsymbol{S}_1 + \omega^2 \left[\frac{1}{2}\boldsymbol{S}_2 - \frac{3}{2}(\boldsymbol{S}_2 \cdot \hat{\boldsymbol{L}})\hat{\boldsymbol{L}}\right] \times \boldsymbol{S}_1\,, \tag{4.2}$$

$$\dot{\boldsymbol{S}}_2 = \omega^2 \sum_{n=0}^{2} C_2^{(\mathrm{n})} \eta^n \omega^{2n/3} \boldsymbol{L} \times \boldsymbol{S}_2 + \omega^2 \left[\frac{1}{2}\boldsymbol{S}_1 - \frac{3}{2}(\boldsymbol{S}_1 \cdot \hat{\boldsymbol{L}})\hat{\boldsymbol{L}}\right] \times \boldsymbol{S}_2\,, \tag{4.3}$$

where we have defined

$$k = \frac{1}{3} a_0 \omega^{8/3} \left[1 + \sum_{i=2}^{11} (a_i + b_i \ln \omega) \omega^{i/3} \right] , \tag{4.4}$$

$$C_A^{(0)} = 2 + \frac{3}{2} \frac{m_B}{m_A} , \tag{4.5}$$

$$C_A^{(1)} = 3 \frac{m_A}{m_B} + \frac{35}{6} + 4 \frac{m_B}{m_A} + \frac{9}{8} \frac{m_B^2}{m_A^2} , \tag{4.6}$$

$$C_A^{(2)} = \frac{27}{4} \frac{m_A^2}{m_B^2} + \frac{31}{2} \frac{m_A}{m_B} + \frac{137}{12} + \frac{19}{4} \frac{m_B}{m_A} + \frac{15}{4} \frac{m_B^2}{m_A^2} + \frac{27}{16} \frac{m_B^3}{m_A^3} , \tag{4.7}$$

where $(A, B) \in \{1, 2\}$, $A \neq B$ and ω is the orbital frequency of the binary, with coefficients (a_i, b_i) given in Appendix A of [62].

The evolution equation for $\boldsymbol{L}$ contains two types of terms that are valid to different PN orders: *conservative* and *dissipative terms*. The dissipative ones are all contained in the last term of Eq. (4.1), which changes the magnitude of $\boldsymbol{L}$, and governs the GW frequency evolution. The conservative terms [the $\boldsymbol{S}_i$-dependent terms in Eq. (4.1) and all terms in Eqs. (4.2)-(4.3)] describe spin-spin and spin-orbit interactions to 2.5PN order. These terms do not change the magnitude of $\boldsymbol{L}$, but only its direction. We work to the highest PN order known (3.5PN) in the spin-orbit interactions [48]. However, since we are carrying out a small spin magnitude analysis we keep here only the leading order spin-spin interaction terms (2PN). We will later ignore this term when deriving an analytic solution for the angular momenta, but include it in the numerical solutions we will compare against, to show that indeed such a term can be neglected.

The evolution of the magnitude of the angular momentum, and thus, the frequency evolution, is controlled by k, which is given in Eq. (4.4) by a sum taken to 5.5PN order. The coefficients (a_i, b_i) listed in Appendix A of [62] are known to 2.5PN order including all spin terms, to 3.5PN to linear-order in spin [48] and to 22PN order

in the test particle limit neglecting spins and BH absorption terms [97, 158, 199, 209]. Since the evolution of nonspinning binaries acts as a background upon which we perturbative expand in $\chi_{1,2}$, we will model the former very accurately, keeping nonspinning terms in (a_i, b_i) to 3.5PN order with η corrections and to 5.5PN order without all η corrections. To linear-order in spin, we keep terms in (a_i, b_i) up to 4PN order, while to quadratic-order we keep terms up to 2PN order. We express the evolution equation (and its solution later on) in terms of (a_i, b_i), so that higher-order PN corrections can be easily incorporated in our analysis by simply modifying these coefficients, when higher PN order terms become available in the future.

For those astrophysically realistic NS binaries that are expected to be detected by advanced ground-based detectors, the dimensionless spin parameter is not expected to exceed $\chi_A \sim 0.1$ [106]. We can, therefore, define the book-keeping parameter ϵ_s as a way to count powers of χ_A. When working to leading order in ϵ_s, we can neglect the spin-spin interactions in Eqs. (4.1)-(4.3), which leads to the precession equations

$$\dot{\boldsymbol{L}} = \omega^2 \sum_{n\geq 0}^{2} \eta^n \omega^{2n/3} \left[C_1^{(\mathrm{n})} \boldsymbol{S}_1 + C_2^{(\mathrm{n})} \boldsymbol{S}_2\right] \times \boldsymbol{L} - k\boldsymbol{L}\,, \tag{4.8}$$

$$\dot{\boldsymbol{S}}_1 = \omega^2 \sum_{n\geq 0}^{2} \eta^n \omega^{2n/3} C_1^{(\mathrm{n})} \boldsymbol{L} \times \boldsymbol{S}_1\,, \tag{4.9}$$

$$\dot{\boldsymbol{S}}_2 = \omega^2 \sum_{n\geq 0}^{2} \eta^n \omega^{2n/3} C_2^{(\mathrm{n})} \boldsymbol{L} \times \boldsymbol{S}_2\,. \tag{4.10}$$

where the coefficients $C_i^{(\mathrm{n})}$ are given in Eqs. (4.5)-(4.7). In Eq. (4.8), we can rewrite $\dot{\boldsymbol{L}} = L\dot{\hat{\boldsymbol{L}}} + \dot{L}\hat{\boldsymbol{L}}$ to separate precession effects from radiation-reaction effects. The magnitude of the spin angular momentum is conserved due to the particular choice of variables [48] (also see the discussion of Chapter 2).

The precession equations will be solved as a function of the independent variable v, a PN expansion parameter defined by

$$v \equiv \omega^{1/3} = \eta \; L^{-1} \,, \tag{4.11}$$

The radiation-reaction equation for the magnitude of the orbital angular moment allows us to write an evolution equation for v

$$\dot{v} = \frac{a_0}{3} v^9 \left\{ 1 + \sum_{i=2}^{11} \left[a_i + 3 b_i \ln(v) \right] v^i \right\} \,, \tag{4.12}$$

where the coefficients (a_i, b_i) are the same as before. Just like velocity or angular frequency, our PN expansion parameter is time dependent, approaching $v \sim 0.3$ by the end of the inspiral.

Having set up the problem, the remainder of this section solves Eqs. (4.8)-(4.10) perturbatively in χ_A. To do so, we perturbatively expand all quantities in ϵ_s:

$$\boldsymbol{A} = \sum_{n=0}^{N} \epsilon_s^n \; \boldsymbol{A}^{(\mathrm{n})} \,, \tag{4.13}$$

where $\boldsymbol{A}$ is any of $\boldsymbol{L}$, $\boldsymbol{S}_1$ or $\boldsymbol{S}_2$, while $\boldsymbol{A}^{(\mathrm{n})}$ is a term proportional to $(\chi_A)^n$. Equation (4.13) is nothing but the mathematical definition of the small-spin expansion, where we will here work to $\mathcal{O}(\epsilon_s)$, i.e. to $N = 1$.

Zeroth-Order-in-Spin Solution

To this order, the NSs' spin angular momenta vanish:

$$\boldsymbol{S}_1^{(0)} = 0 \,, \qquad \boldsymbol{S}_2^{(0)} = 0 \,. \tag{4.14}$$

while the orbital angular momentum evolves according to

$$\dot{\boldsymbol{L}}^{(0)} = -k^{(0)}\boldsymbol{L}^{(0)}\,, \tag{4.15}$$

This equation implies that the angular momentum does not change in direction, but only shrinks in magnitude due to radiation reaction.

We then work in a coordinate system that is adapted to the problem at hand by choosing $\hat{\boldsymbol{z}} = \hat{\boldsymbol{J}}(0)$. Since $\boldsymbol{J}$ is evolving at higher orders in ϵ_s, the $\hat{\boldsymbol{z}}$ axis of the frame will not remain aligned with $\boldsymbol{J}$ at later times. However, to $\mathcal{O}(\epsilon_s^0)$, $\hat{\boldsymbol{J}}^{(0)}$ is not evolving, and thus, $\hat{\boldsymbol{z}} = \hat{\boldsymbol{J}}^{(0)} = \hat{\boldsymbol{L}}^{(0)}$. With this choice of coordinate system, L_x and L_y simply vanish to this order.

The $\hat{\boldsymbol{z}}$ component of the orbital angular momentum satisfies the equation

$$\dot{L}^{(0)}=-\frac{a_0}{3}\frac{\mu^8}{L^{(0)7}}\left\{1+\sum_{i=2}^{11}\left[a_i^{(0)}+3b_i^{(0)}\ln\Big(\frac{\mu}{L^{(0)}}\Big)\right]\frac{\mu^i}{L^{(0)i}}\right\}, \tag{4.16}$$

where recall that $L^{(0)}$ is the magnitude of the orbital angular momentum to order $\mathcal{O}(\epsilon_s^0)$ and the coefficients $a_i^{(0)} = a_i(\boldsymbol{S}_A = 0)$. Since all the coefficients are constant, we can directly integrate the above equation, invert the PN expansion, and obtain $L^{(0)}$ as a function of time. We rewrite Eq. (4.16) in terms of the PN parameter $v^{(0)}$, related to $L_z^{(0)}$ through Eq. (4.11), namely,

$$L_z^{(0)} = \frac{\eta}{v^{(0)}}\,, \tag{4.17}$$

and solve this equation to obtain

$$v^{(0)} = \zeta\left[1+\sum_{i=2}^{11}v_i^{(0)}\zeta^i+\sum_{i=6}^{11}v_i^{\ell,(0)}\zeta^i\ln(\zeta)+\sum_{i=8}^{11}v_i^{\ell^2,(0)}\zeta^i(\ln\zeta)^2+\mathcal{O}(v^{12})\right], \tag{4.18}$$

where we have defined the function of time

$$\zeta \equiv \left[\frac{3}{8a_0(t_c - t)}\right]^{1/8}, \tag{4.19}$$

with t_c the time of coalescence. The PN coefficients $(v_i^{(0)}, v_i^{\ell,(0)}, v_i^{\ell^2,(0)})$ can be obtained from Appendix B of [62] by setting $\boldsymbol{S}_A = 0$. Combining Eq. (4.18) with Eq. (4.17) completes the solution for the time evolution of all angular momenta to $\mathcal{O}(\epsilon_s^0)$.

As explained before, this solution keeps terms beyond 3.5PN order, even though the evolution equation is formally only known to that order. We do so to minimize the difference between the numerical solution to the evolution of the angular momentum and the analytic approximation in Eq. (4.18) in the nonspinning case. In fact, taking this series to 5.5PN order guarantees that the frequency to time mapping is accurate to roughly 10^{-2} Hz during the entire inspiral (from 10Hz up to 400Hz, where finite size effects become important [112,148,180]). Doing so will allow us to isolate any spin precession effects cleanly. Ultimately, however, we will be interested in the frequency-domain waveform, which can be constructed without knowledge of Eq. (4.18).

First-Order-in-Spin Solution

To this order, the orbital and spin angular momenta evolve according to

$$\dot{\boldsymbol{L}}^{(1)} = v^6 \sum_{n=0}^{2} \eta^n v^{2n} \left[C_1^{(\mathrm{n})} \boldsymbol{S}_1^{(1)} + C_2^{(\mathrm{n})} \boldsymbol{S}_2^{(1)}\right] \times \boldsymbol{L}^{(0)} - k^{(0)} \boldsymbol{L}^{(1)} - k^{(1)} \boldsymbol{L}^{(0)}, \tag{4.20}$$

$$\dot{\boldsymbol{S}}_1^{(1)} = v^6 \sum_{n=0}^{2} \eta^n v^{2n} C_1^{(\mathrm{n})} \boldsymbol{L}^{(0)} \times \boldsymbol{S}_1^{(1)}, \tag{4.21}$$

$$\dot{\boldsymbol{S}}_2^{(1)} = v^6 \sum_{n=0}^{2} \eta^n v^{2n} C_2^{(\mathrm{n})} \boldsymbol{L}^{(0)} \times \boldsymbol{S}_2^{(1)}, \tag{4.22}$$

where $\boldsymbol{L}^{(0)}$ is given by Eqs. (4.17) and (4.18). The above equations are easy to decouple: we first use Eqs. (4.21) and (4.22) to solve for $\boldsymbol{S}_1^{(1)}$ and $\boldsymbol{S}_2^{(1)}$, and we then substitute these solutions into Eq. (4.20) to solve for the orbital angular momentum.

Solution for the Spins Without loss of generality, we focus on $\boldsymbol{S}_1^{(1)}$; the solution for $\boldsymbol{S}_2^{(1)}$ can be obtained by an exchange $1 \leftrightarrow 2$. In term of its components, Eq. (4.21) can be written as

$$\dot{S}_{1,x}^{(1)} = -v^6 \sum_{n=0}^{2} \eta^n v^{2n} C_1^{(\mathrm{n})} L S_{1,y}^{(1)} \,, \tag{4.23}$$

$$\dot{S}_{1,y}^{(1)} = v^6 \sum_{n=0}^{2} \eta^n v^{2n} C_1^{(\mathrm{n})} L S_{1,x}^{(1)} \,, \tag{4.24}$$

$$\dot{S}_{1,z}^{(1)} = 0 \,. \tag{4.25}$$

which we rewrite as

$$\frac{d\lambda_1^{(1)}}{d\phi_1} = i\lambda_1^{(1)} \,, \tag{4.26}$$

where we have defined

$$\lambda_1^{(1)} = S_{1,x}^{(1)} + i S_{1,y}^{(1)} \,, \tag{4.27}$$

and the new independent variable

$$\frac{d\phi_1}{dt} = v^5 \mu \sum_{n=0}^{2} \eta^n v^{2n} C_1^{(\mathrm{n})} \,. \tag{4.28}$$

The combined spin evolution equation can now be integrated directly. Doing so, decoupling Eq. (4.27) and using exchange symmetry, we are led to the solution

$$S_{1,x}^{(1)} = S_{1,x}^{(1)}(0) \cos\phi_1 - S_{1,y}^{(1)}(0) \sin\phi_1 \,, \tag{4.29}$$

$$S_{1,y}^{(1)} = S_{1,y}^{(1)}(0) \cos\phi_1 + S_{1,x}^{(1)}(0) \sin\phi_1 \,, \tag{4.30}$$

$$S^{(1)}_{1,z} = S^{(1)}_{1,z}(0)\,, \tag{4.31}$$

$$S^{(1)}_{2,x} = S^{(1)}_{2,x}(0)\cos\phi_2 - S^{(1)}_{2,y}(0)\sin\phi_2\,, \tag{4.32}$$

$$S^{(1)}_{2,y} = S^{(1)}_{2,y}(0)\cos\phi_2 + S^{(1)}_{2,x}(0)\sin\phi_2\,, \tag{4.33}$$

$$S^{(1)}_{2,z} = S^{(1)}_{2,z}(0)\,, \tag{4.34}$$

where ϕ_2 is defined by Eq. (4.28) with $1 \leftrightarrow 2$. This solution amounts to a simple harmonic oscillator with precession frequency $\dot{\phi}_A$.

To complete the calculation, we must solve Eq. (4.28) for the phase angle ϕ_A. Doing so, we find as an expansion in v

$$\frac{\phi_A}{C^{(0)}_A} = \phi_{0,A} - \frac{\eta}{a_0 v^3}\Big[1 + \sum_{i=2}^{8}\phi_{i,A}v^i + \sum_{i=3}^{8}\phi^{\ell}_{i,A}v^i \ln v + \mathcal{O}(v^9)\Big]\,, \tag{4.35}$$

with $A \in \{1,2\}$, where the coefficients $\phi_{i,A}$ are given in Appendix C of [62] and $\phi_{0,A}$ is a constant of integration picked such that $\phi_A(t=0) = 0$ to satisfy Eqs. (4.29-4.34). One could include higher-order PN terms in this expansion. However, we find that truncating it at 4PN order is sufficient to obtain an accurate time-domain waveform phase, i.e. higher-order PN terms induce phase corrections that are smaller than those induced by neglected terms of $\mathcal{O}(\epsilon_s^2)$.

In principle, the coefficients (a_i, b_i) that appear in Eq. (4.35) are given in Appendix A of [62] with all spins set to zero, since $\phi_{i,A}$ should be kept to $\mathcal{O}(\epsilon_s^0)$. However, we will here use the full expressions in Appendix A of [62] for the coefficients (a_i, b_i), with the spin couplings β_i and σ_i evaluated at the initial spin and orbital angular momenta $\boldsymbol{L}(t=0)$ and $\boldsymbol{S}_i(t=0)$. The solution obtained with this substitution differs with the initial $\mathcal{O}(\epsilon_s^0)$ solution by terms of $\mathcal{O}(\epsilon_s^2)$, and is thus equally valid. In practice, we find that using these coefficients leads to better

agreement between the analytical approximation and the numerical solution of the orbital phase.

Solution for the Orbital Angular Momentum The evolution equation [Eq. (4.20)] contains terms that change on two different time scales: a radiation-reaction time scale $T_{\rm rr}$ and a precession time scale $T_{\rm pr}$. The former is associated with the last two terms in Eq. (4.20) and it is defined by

$$T_{\rm rr} \equiv \frac{v}{\dot{v}} \sim \frac{1}{\eta} v^{-8} \,, \tag{4.36}$$

while the latter is associated with the first two terms in Eq. (4.20) and it is defined by

$$T_{\rm pr} \equiv \frac{1}{\omega^2 L} \sim \frac{1}{\eta} v^{-5} \,. \tag{4.37}$$

The ratio of these time scales is $T_{\rm pr}/T_{\rm rr} \sim v^3$, which suggests one should use MSA to solve Eq. (4.20).

We define a new bookkeeping parameter ϵ_p that counts powers of $(T_{\rm pr}/T_{\rm rr})$ and expand all quantities in a bivariate series

$$\boldsymbol{A} = \sum_{n=0}^{N} \sum_{m=0}^{M} \epsilon_s^n \epsilon_p^m \, \boldsymbol{A}^{(n,m)} \,, \tag{4.38}$$

where $\boldsymbol{A}$ is any of $\boldsymbol{L}$, $\boldsymbol{S}_1$ or $\boldsymbol{S}_2$, while $\boldsymbol{A}^{(n,m)}$ is a term proportional to $\chi_{1,2}^n (T_{\rm pr}/T_{\rm rr})^m$. Of course, $\boldsymbol{A}^{(0,m)} \; \forall \; m$ has already been obtained in Eq. (4.18).

In MSA, all quantities must be assumed to depend on all independent time scales, and thus,

$$\boldsymbol{A}^{(n,m)} = \boldsymbol{A}^{(n,m)}(t, \tau) \,, \tag{4.39}$$

where we have defined the long time scale

$$d\tau = \epsilon_p \left(\frac{T_{\rm pr}}{T_{\rm rr}} \right) dt \,. \tag{4.40}$$

The differential operator of Eq. (4.20) is then

$$\frac{d}{dt} = \frac{\partial}{\partial t} + \epsilon_p \left(\frac{T_{\rm pr}}{T_{\rm rr}} \right) \frac{\partial}{\partial \tau} \,. \tag{4.41}$$

The solution for $\boldsymbol{L}^{(1,0)}$ is more easily obtained if we work with the total angular momentum instead. The latter is defined by $\boldsymbol{J}^{(1)} = \boldsymbol{L}^{(1)} + \boldsymbol{S}_1^{(1)} + \boldsymbol{S}_2^{(1)}$ to $\mathcal{O}(\epsilon_s)$, and satisfies the equation

$$\dot{\boldsymbol{J}}^{(1)} = -k^{(0)} \boldsymbol{L}^{(1)} - k^{(1)} \boldsymbol{L}^{(0)} \,. \tag{4.42}$$

This last equation can in turn be expanded in ϵ_p to obtain a bivariate series. Now we can proceed to solve Eq. (4.42) order by order in ϵ_p.

To zeroth order in radiation reaction we have the simple *partial* differential equation

$$\frac{\partial \boldsymbol{J}^{(1,0)}}{\partial t} = 0 \,, \tag{4.43}$$

the solution to which is

$$\boldsymbol{J}^{(1,0)} = [J_x^{(1,0)}(\tau), J_y^{(1,0)}(\tau), J_z^{(1,0)}(\tau)] \,. \tag{4.44}$$

$\boldsymbol{J}^{(1,0)}(\tau)$ is a functions of the long time scale changing over the radiation-reaction time scale, but remaining constant on a precession time scale. Its functional form can only be determined by going to next order in ϵ_p.

To $\mathcal{O}(\epsilon_p)$, the evolution equation for the total angular momentum is

$$\frac{\partial \boldsymbol{J}^{(1,1)}}{\partial t} + \frac{T_{\rm pr}}{T_{\rm rr}} \frac{\boldsymbol{J}^{(1,0)}}{\partial \tau} = -k^{(0,1)} \boldsymbol{L}^{(1,0)} - k^{(1,1)} \boldsymbol{L}^{(0,0)} \,, \tag{4.45}$$

where $\boldsymbol{L}^{(1,0)}$ is to be understood as shorthand for $\boldsymbol{J}^{(1,0)} - \boldsymbol{S}_1^{(1,0)} - \boldsymbol{S}_2^{(1,0)}$. Equation (4.45) is a differential equation for $\boldsymbol{J}^{(1,1)}(t)$, whose solution will grow linearly (a behavior that characterizes a resonance) if sourced by a t-independent term. We eliminate such t-independent terms by requiring that

$$\frac{T_{\rm pr}}{T_{\rm rr}} \frac{dJ_x^{(1,0)}(\tau)}{d\tau} = -k^{(0,1)} J_x^{(1,0)}(\tau) \,, \tag{4.46}$$

$$\frac{T_{\rm pr}}{T_{\rm rr}} \frac{dJ_y^{(1,0)}(\tau)}{d\tau} = -k^{(0,1)} J_y^{(1,0)}(\tau) \,, \tag{4.47}$$

whose solution is

$$\begin{aligned} J_x^{(1,0)} &= J_x^{(1,0)}(0) \exp\left[-\int k^{(0,1)} \left(\frac{T_{\rm rr}}{T_{\rm pr}} \right) d\tau \right] = J_x^{(1,0)}(0) \exp\left[\int \frac{\dot{L}^{(0,1)}}{L^{(0,1)}} dt \right] \,, \\ &= J_x^{(1,0)}(0) \frac{L^{(0,1)}}{L^{(0,1)}(0)} = J_x^{(1,0)}(0) \frac{v^{(0,1)}(0)}{v^{(0,1)}} \,, \end{aligned} \tag{4.48}$$

and $J_y^{(1,0)}$ is obtained by replacing $x \leftrightarrow y$. Collecting all the results and imposing the initial conditions $\hat{\boldsymbol{z}} = \hat{\boldsymbol{J}}(0)$, the solutions for $L_x^{(1,0)}$ and $L_y^{(1,0)}$ are

$$L_x^{(1,0)} = -S_{1,x}^{(1)}(0) \cos\phi_1 + S_{1,y}^{(1)}(0) \sin\phi_1 - S_{2,x}^{(1)}(0) \cos\phi_2 + S_{2,y}^{(1)}(0) \sin\phi_2 \,, \tag{4.49}$$

$$L_y^{(1,0)} = -S_{1,y}^{(1)}(0) \cos\phi_1 - S_{1,x}^{(1)}(0) \sin\phi_1 - S_{2,y}^{(1)}(0) \cos\phi_2 - S_{2,x}^{(1)}(0) \sin\phi_2 \,. \tag{4.50}$$

The z component of the orbital angular momentum to first order in the spins can be obtained in the following way. First, we notice that

$$v \equiv \frac{\eta}{L} = \frac{\eta}{L_z} + \mathcal{O}(\epsilon_s^2)\,. \tag{4.51}$$

Therefore, to this order we have

$$L_z = L_z^{(0)} + \epsilon_s L_z^{(1)} + \mathcal{O}(\epsilon_s^2) = \frac{\eta}{v} + \mathcal{O}(\epsilon_s^2)\,. \tag{4.52}$$

We can further improve on the solution for v as a function of time to $\mathcal{O}(\epsilon_s)$ by revisiting its evolution equation. Equation (4.12) depends on the PN coefficients (a_i, b_i), which in turn depend on the spins only through combinations $\boldsymbol{S}_A \cdot \hat{\boldsymbol{L}}$ and $\boldsymbol{S}_A \cdot \boldsymbol{S}_B$. The spin-spin terms are of $\mathcal{O}(\epsilon_s^2)$ and can thus be neglected, while the spin-orbit terms are constant to $\mathcal{O}(\epsilon_s)$. We can then substitute $\boldsymbol{S}_A \cdot \hat{\boldsymbol{L}} \to S_{A,z}$ in Eq. (4.12), treat all the PN coefficients (a_i, b_i) as constant and integrate to obtain

$$\begin{aligned} v = v^{(0)} + \epsilon_s v^{(1)} + \mathcal{O}(\epsilon_s^2) = \zeta \Bigg[1 + \sum_{i=2}^{11} \left(v_i^{(0)} + \epsilon_s v_i^{(1)}\right) \zeta^i + \sum_{i=6}^{11} \left(v_i^{\ell,(0)} + \epsilon_s v_i^{\ell,(1)}\right) \zeta^i \ln \zeta \\ + \sum_{i=8}^{11} \left(v_i^{\ell^2,(0)} + \epsilon_s v_i^{\ell^2,(1)}\right) \zeta^i (\ln \zeta)^2 + \mathcal{O}(\epsilon_s^2, \zeta^{12}) \Bigg]\,. \end{aligned} \tag{4.53}$$

Recall that $(v_i^{(0)}, v_i^{\ell,(0)}, v_i^{\ell^2,(0)})$ are those in Appendix B of [62] with $\boldsymbol{S}_A = 0$, while the new coefficients $(v_i^{(1)}, v_i^{\ell,(1)}, v_i^{\ell^2,(1)})$ are those in Appendix Bs of [62] that are linear in $\boldsymbol{S}_A$.

However, as in Eq. (4.35), we will here include higher-order terms in ϵ_s to improve the mapping between frequency and time. The coefficients $(v_i^{(0)} + \epsilon_s v_i^{(1)}, v_i^{\ell,(0)} + \epsilon_s v_i^{\ell,(1)}, v_i^{\ell^2,(0)} + \epsilon_s v_i^{\ell^2,(1)})$ will be replaced by those in Appendix B of [62] but with (a_i, b_i) coefficients evaluated at the initial spin and orbital angular momenta $\boldsymbol{L}(t = 0)$

and $\boldsymbol{S}_A(t=0)$. As before, this replacement adds terms of $\mathcal{O}(\epsilon_s^2)$ and higher to the precession phase that improve the accuracy of the analytical solution.

Summary of Results

The final solutions for the orbital angular momentum to first order in spin are

$$L_x = \epsilon_s \left\{-S_{1,x}(0)\cos\phi_1 + S_{1,y}(0)\sin\phi_1 - S_{2,x}(0)\cos\phi_2 + S_{2,y}(0)\sin\phi_2\right\} + \mathcal{O}(\epsilon_s^2, \epsilon_p)\,, \tag{4.54}$$

$$L_y = \epsilon_s \left\{-S_{1,y}(0)\cos\phi_1 - S_{1,x}(0)\sin\phi_1 - S_{2,y}(0)\cos\phi_2 - S_{2,x}(0)\sin\phi_2\right\} + \mathcal{O}(\epsilon_s^2, \epsilon_p)\,, \tag{4.55}$$

$$L_z = \frac{M\mu}{v} + \mathcal{O}(\epsilon_s^2)\,, \tag{4.56}$$

while for the spin angular momenta we find

$$S_{1,x} = \epsilon_s \left[S_{1,x}(0)\cos\phi_1 - S_{1,y}(0)\sin\phi_1\right] + \mathcal{O}(\epsilon_s^2)\,, \tag{4.57}$$

$$S_{1,y} = \epsilon_s \left[S_{1,y}(0)\cos\phi_1 + S_{1,x}(0)\sin\phi_1\right] + \mathcal{O}(\epsilon_s^2)\,, \tag{4.58}$$

$$S_{1,z} = \epsilon_s S_{1,z}(0) + \mathcal{O}(\epsilon_s^2)\,, \tag{4.59}$$

$$S_{2,x} = \epsilon_s \left[S_{2,x}(0)\cos\phi_2 - S_{2,y}(0)\sin\phi_2\right] + \mathcal{O}(\epsilon_s^2)\,, \tag{4.60}$$

$$S_{2,y} = \epsilon_s \left[S_{2,y}(0)\cos\phi_2 + S_{2,x}(0)\sin\phi_2\right] + \mathcal{O}(\epsilon_s^2)\,, \tag{4.61}$$

$$S_{2,z} = \epsilon_s S_{2,z}(0) + \mathcal{O}(\epsilon_s^2)\,, \tag{4.62}$$

The phase angles ϕ_A are given explicitly in Eq. (4.35) in terms of v, which is given explicitly as a function of time in Eq. (4.53). Both of these equations depend on the coefficients (a_i, b_i) which are given in Appendix A of [62], with $[\boldsymbol{L}(t), \boldsymbol{S}_A(t)] \to [\boldsymbol{L}(t=0), \boldsymbol{S}_A(t=0)]$.

Numerical Comparison

We can now show that the bivariate, analytic solution found above is indeed an accurate representation of the full numerical solution. By the latter, we mean the numerical solution to Eqs. (4.1)-(4.3), with k given by Eq. (4.4). Notice that these equations contain spin-spin interactions, i.e. terms of $\mathcal{O}(\epsilon_s^2)$ that are neglected in the analytic solution of the previous sections. We use an adaptive Cash-Karp, fifth-order Runge-Kutta method to solve these equations [55]. We have performed convergence tests to guarantee that the numerical error introduced by the integrator is well controlled and not visible in any of the figures we show.

For the comparisons to follow, we choose a particular system with the following properties:

Test System:

- $m_1 = 1.4M_\odot$ and $m_2 = 1.6M_\odot$, which then implies $\eta = 0.2489$, $\delta m = -0.067$;
- $\chi_1 = 0.1$ and $\chi_2 = 0.1$, which then implies $S_1 \approx 0.196M_\odot^2$, $S_2 \approx 0.256M_\odot^2$;
- $\hat{\boldsymbol{S}}_A(0) = (\cos\phi_{S_A}\sin\theta_{S_A}, \sin\phi_{S_A}\sin\theta_{S_A}, \cos\theta_{S_A})$, where $(\phi_{S_1}, \theta_{S_1}) = (\pi/4, 17\pi/24)$ and $(\phi_{S_2}, \theta_{S_2}) = (\pi/3, -\pi/6)$.
- The source is located at polar angles $(\theta_{\rm d}, \phi_{\rm d}) = (\pi/3, 2\pi/3)$ in the detector frame.

The orbital angular momentum ranges from about $29M_\odot^2$ at 10 Hz to $8.5M_\odot^2$ at 400 Hz, which implies $S_A/L \ll 1$ during the entire inspiral. This system experiences approximately 55 precession cycles from 10Hz to 400 Hz, while accumulating $\sim 14,500$ cycles of GW phase. We choose the integration constants in the analytic solution such that the quantities compared agree at 10 Hz. Although we choose a particular system for the figures to come, the results are representative of all systems we investigated.

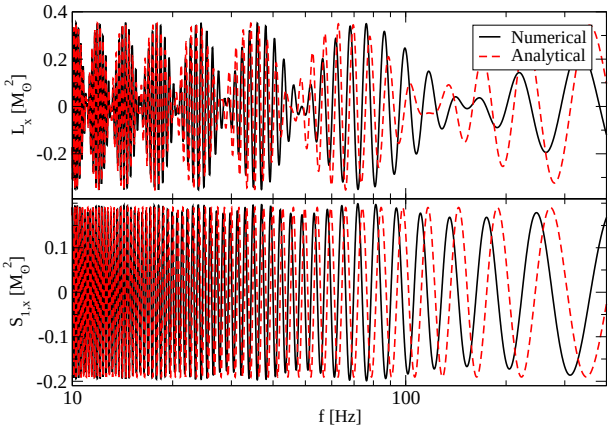

Figure 4.1: Numerical solution (solid black curve) and analytical approximation (red dashed curve) to L_x (top) and $S_{1,x}$ (bottom) as a function of the GW frequency f (twice the orbital frequency).

Figure 4.1 presents the numerical (black solid curve) and analytical (red dashed curve) approximation to the x component of orbital (top) and spin (bottom) angular momentum of NS 1 as a function of the dominant GW frequency (twice the orbital frequency) in units of Hz. The y components present similar behavior. The analytical result tracks the numerical solution closely, coming out of phase by the end of the inspiral. This analytic solution is dramatically better than that which assumes these components simply vanish, as is done when one neglects precession and uses a spin aligned/antialigned approximation. The analytical approximation can, of course, be improved if taken to next order in ϵ_s.

Figure 4.2 shows the numerical (black solid curve) and analytical (red dashed curve) approximation to the z component of the spin angular momentum for NS 1 (top) and NS 2 (bottom) as a function of the GW frequency in Hz. Both solutions start at the same initial value, but the analytical approximation remains constant, while the numerical one oscillates. The amplitude of this oscillation is of $\mathcal{O}(\epsilon_s^2)$, which

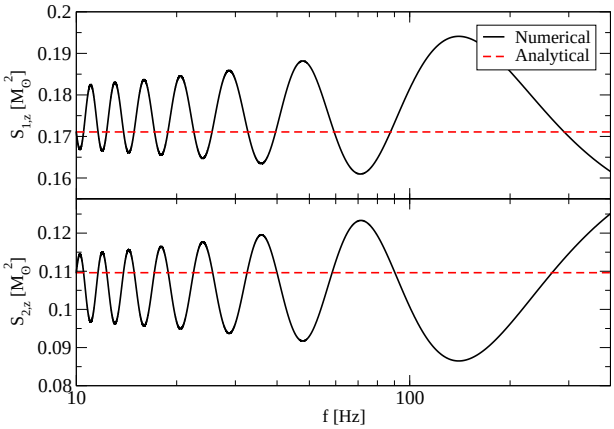

Figure 4.2: Numerical solution (solid black curve) and analytical approximation (red dashed curve) to $S_{1,z}$ (top) and $S_{2,z}$ (bottom) as a function of the GW frequency.

for this system is of $\mathcal{O}(10^{-2})$; the solutions can only be improved by going to next order in spin.

The top panel of Fig. 4.3 presents the numerical (black solid curve) and analytical (red dashed curve) approximation to the z component of the orbital angular momentum as a function of GW frequency. The analytical approximation is so accurate that it cannot be distinguished from the numerical result. For this reason, the bottom panel of Fig. 4.3 shows the absolute value of the fractional difference between the numerical and the analytical result, again as a function of GW frequency. The fractional error never exceeds one part in 10^3.

The construction of time-domain waveforms also requires a mapping between time and frequency. The top panel of Fig. 4.4 shows the numerical (solid black curve) and analytical approximation (red dashed curve) to the evolution of time as a function of GW frequency in Hz; the bottom panel of this figure shows their difference. Again the analytical result tracks the numerical one to a precision better than a few times 10^{-4}.

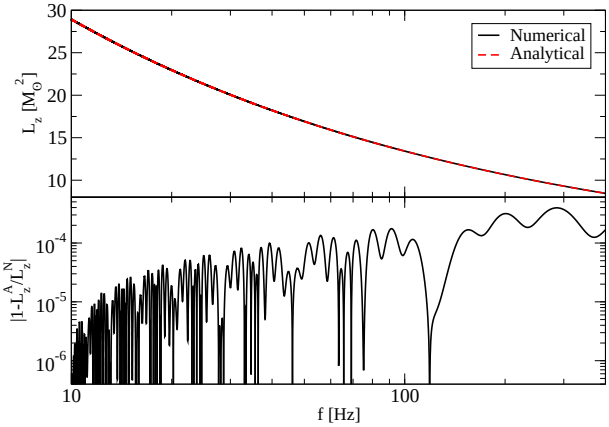

Figure 4.3: Top: Numerical solution (solid black curve) and analytical approximation (red dashed curve) to L_z (top) as a function of the GW frequency. Bottom: Fractional difference of numerical and analytical L_z as a function of GW frequency.

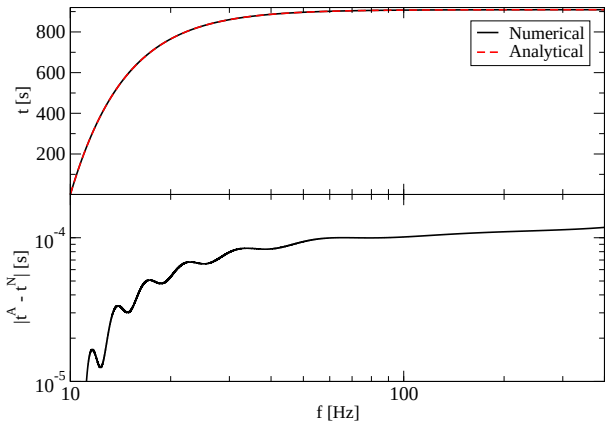

Figure 4.4: Top: Numerical solution (solid black curve) and analytical approximation (red dashed curve) to $t(f)$. Bottom: Absolute value of the difference between the numerical and analytical times as a function of GW frequency.

Time-Domain Waveform

Given the approximate analytical solution to the orbital angular momentum derived in the previous section, we can now construct an analytical approximation to the time-domain GW response function. We begin by defining the basic ingredients that go into the waveform construction. We then derive analytic approximations to the waveform phase, and mode-decompose the response function. We conclude with a comparison between the analytic approximation to the time-domain waveform phase and the phase derived from a numerical solution.

Basics

Let $\hat{\boldsymbol{N}} = [N_x, N_y, N_z]$ be a unit vector that points in the direction of the center of mass of the compact binary relative to the detector. The time-domain response function in the long wavelength approximation can then be written as a sum of harmonics [43]

$$h(t) = \sum_{n\geq 0} (F_+ h_{n,+} + F_\times h_{n,\times}) \,, \tag{4.63}$$

$$h_{n,+} = \mathcal{A}_{n,+}(\iota) \cos n\Phi + \mathcal{B}_{n,+}(\iota) \sin n\Phi \,, \tag{4.64}$$

$$h_{n,\times} = \mathcal{A}_{n,\times}(\iota) \cos n\Phi + \mathcal{B}_{n,\times}(\iota) \sin n\Phi \,, \tag{4.65}$$

where $n \in \mathbb{N}$ is the harmonic number, the inclination angle is

$$\iota \equiv \arccos\left[\hat{\boldsymbol{L}} \cdot \hat{\boldsymbol{N}}\right] \,, \tag{4.66}$$

the antenna pattern functions are given by

$$F_+ = \frac{1}{2}(1 + \cos^2\theta_N) \cos 2\phi_N \cos 2\psi - \cos\theta_N \sin 2\phi_N \sin 2\psi \,, \tag{4.67}$$

$$F_\times = \frac{1}{2}(1+\cos^2\theta_N)\cos 2\phi_N \sin 2\psi + \cos\theta_N \sin 2\phi_N \cos 2\psi\,, \tag{4.68}$$

ψ is the polarization angle

$$\psi = \tan^{-1}\left(\frac{\hat{\boldsymbol{L}}\cdot\hat{\boldsymbol{z}} - (\hat{\boldsymbol{L}}\cdot\hat{\boldsymbol{N}})(\hat{\boldsymbol{z}}\cdot\hat{\boldsymbol{N}})}{\hat{\boldsymbol{N}}\cdot(\hat{\boldsymbol{L}}\times\hat{\boldsymbol{z}})}\right), \tag{4.69}$$

and θ_N, ϕ_N give the location of the GW source on the sky.

The precession of the orbital angular momentum, i.e. of the orbital plane, has two main effects on the waveform [21]:

(i) the inclination and the polarization angles become time dependent and;

(ii) the reference frame used to define the orbital frequency in the orbital plane, $\hat{\boldsymbol{L}}\times\hat{\boldsymbol{N}}$, is no longer constant in time.

These effects will induce corrections to the waveform phase, as well as amplitude modulations. We investigate these effects below.

Waveform Phase

The waveform phase can be decomposed as follows:

$$\Phi = \Phi^{\rm orb} + \delta\phi\,, \tag{4.70}$$

where $\Phi^{\rm orb}$ is the orbital phase and $\delta\phi$ is a precession correction, induced by the changing reference frame. We will refer to the latter as the Thomas precession phase.

The orbital phase can be computed directly from

$$\Phi^{\rm orb} = \int \omega\, dt = \int v^3 \frac{dv}{\dot{v}} = \phi_c - \frac{3}{5a_0 v^5}\left[1 + \sum_{i=2}^{16}\Phi_i^{\rm orb} v^i\right.$$

$$+\sum_{i=5}^{16}\Phi_i^{\rm orb,\ell}v^n\ln v+\sum_{i=12}^{16}\Phi_i^{\rm orb,\ell^2}v^i(\ln v)^2+\mathcal{O}(v^{-17})\Bigg]\,, \tag{4.71}$$

where ϕ_c is a constant of integration (the so-called *phase of coalescence*) and the PN coefficients $(\Phi_i^{\rm orb},\Phi_i^{\rm orb,\ell},\Phi_i^{\rm orb,\ell^2})$ to 8PN order are given in Appendix D of [62]. As before, we extend the series here to 8PN order, so that when spins are zero, the error between this analytical phase and the numerical solution is negligibly small. Doing so will allow us to isolate any dephasings induced by spin.

The Thomas precession phase $\delta\phi$ satisfies the differential equation [21]

$$\delta\dot{\phi}=\frac{1}{L}\frac{\boldsymbol{L}\cdot\hat{\boldsymbol{N}}}{\boldsymbol{L}^2-\left(\boldsymbol{L}\cdot\hat{\boldsymbol{N}}\right)^2}\left(\boldsymbol{L}\times\hat{\boldsymbol{N}}\right)\cdot\dot{\boldsymbol{L}}\,. \tag{4.72}$$

Reference [124] found a uniform asymptotic expansion to the solution of this equation to $\mathcal{O}(\epsilon_s)$, which works both when $\hat{N}_x^2+\hat{N}_y^2=\mathcal{O}(\epsilon_s)$ and when $\hat{N}_x=\mathcal{O}(\epsilon_s)=\hat{N}_y$, namely

$$\delta\phi^{(1)}=-\hat{N}_z\arctan\left[\frac{\hat{N}_xL_z-L_x}{\hat{N}_yL_z-L_y}\right]+\mathcal{O}(\epsilon_s^2)\,, \tag{4.73}$$

where recall that $\boldsymbol{N}$ is constant but $\boldsymbol{L}$ varies on a precession time scale as given by Eqs. (4.54)-(4.56).

We here improve on this solution by including the $\mathcal{O}(\epsilon_s^2)$ secular growth of $\delta\phi$. To do so, we first expand Eq. (4.72) to $\mathcal{O}(\epsilon_s^2)$. The $\mathcal{O}(\epsilon_s^2)$ term depends on the $\boldsymbol{L}$ and $\boldsymbol{S}_A$ solutions found in Eqs. (4.54)-(4.62), and after expanding it in $S_{A,x}/L_z\ll 1\gg S_{A,y}/L_z$ and averaging over a precession cycle, we find

$$\left\langle\delta\dot{\phi}^{(2)}\right\rangle=\frac{1}{4\pi^2}\int_0^{2\pi}\int_0^{2\pi}\delta\dot{\phi}^{(2)}d\phi_1d\phi_2=\frac{1}{2}\frac{\dot{\phi}_1}{L_z^2}S_{1,\perp}^2(0)+1\leftrightarrow 2\,, \tag{4.74}$$

where we have defined $S^2_{1,\perp}(0) = S_{1,x}(0)^2 + S_{1,y}(0)^2$. We can conveniently rewrite this as

$$\left\langle \frac{d}{dv}\delta\phi^{(2)} \right\rangle = \sum_{n=0}^{2} \frac{\eta^n v^{2n} C_1^{(\mathrm{n})}}{2\eta} S^2_{1,\perp}(0) \frac{v^7}{\dot{v}} + 1 \leftrightarrow 2\,, \tag{4.75}$$

Solving this differential equation to 1PN order, we obtain the $\mathcal{O}(\epsilon_s^2)$ secular correction

$$\left\langle \delta\phi^{(2)} \right\rangle = -\frac{5}{64} S^2_{1,\perp}(0) \frac{1}{\eta^2} \frac{1}{v} \left[C_1^{(0)} + \left(a_2 C_1^{(0)} - \eta C_1^{(1)} \right) v^2 \right] + 1 \leftrightarrow 2 + \mathcal{O}(\epsilon_s^3)\,. \tag{4.76}$$

The constant of integration can be absorbed in the constant ϕ_c introduced in Eq. (4.71). We empirically find that it is sufficient to keep terms up to 1PN order in this secular approximation, relative to numerical solutions.

The final expression for the Thomas phase $\delta\phi$ is then the sum of Eq. (4.73) and Eq. (4.76):

$$\delta\phi = \delta\phi^{(1)} + \epsilon_s^2 \left\langle \delta\phi^{(2)} \right\rangle + \mathcal{O}(\epsilon_s^3)\,. \tag{4.77}$$

This expression, of course, is missing the nonsecular evolution of $\delta\phi^{(2)}$, but this cannot be computed without knowledge of the evolution of the angular moment to $\mathcal{O}(\epsilon_s^2)$. We will see later that even without these nonsecular terms, Eq. (4.77) is an excellent approximation to the numerical Thomas phase.

In order to calculate the GW signal that will be measured by a ground-based detector on Earth, we need to work in a frame attached to the arms of the detector. We choose the $\hat{\boldsymbol{z}}_d$ axis to be perpendicular to the detector plane and the $\hat{\boldsymbol{x}}_d$ and $\hat{\boldsymbol{y}}_d$ axes to be aligned with the detector's arms. The subscript d denotes the detector frame, while the subscript s denotes the source frame (see e.g. Fig. 1 in [21]). In this frame, the position of the binary in the sky is given by $\hat{\boldsymbol{N}}_d = [\sin\theta_N \cos\phi_N, \sin\theta_N \sin\phi_N, \cos\theta_N]$. In order to transform vectors from the source frame to the detector frame, we assume

that the binary is oriented in such a way that its total angular momentum at $t = 0$ in the detector frame is given by $\hat{\boldsymbol{J}}_d(0) \equiv [\sin\theta_0 \cos\phi_0, \sin\theta_0 \sin\phi_0, \cos\theta_0]$. Then, the rotation matrix relating the frames is

$$\mathbb{R}_{d\to s} = \begin{bmatrix} \cos\theta_0\cos\phi_0 & \cos\theta_0\sin\phi_0 & -\sin\theta_0 \\ -\sin\phi_0 & \cos\phi_0 & 0 \\ \sin\theta_0\cos\phi_0 & \sin\theta_0\sin\phi_0 & \cos\theta_0 \end{bmatrix} \tag{4.78}$$

We apply this matrix to rotate $\boldsymbol{N}_d$ into $\boldsymbol{N}_s$ and $\hat{\boldsymbol{z}}_d$ into $\hat{\boldsymbol{z}}_s$ when computing the polarization angle and the Thomas precession angle.

Mode Decomposition

The analytical approximations to the Fourier transform that we will employ require that we cast the time-domain response function in the following form

$$h(t) = \mathcal{A}(t) e^{i\Phi_{\rm GW}(t)}\,, \tag{4.79}$$

where $\mathcal{A}(t)$ is a slowly varying amplitude and $\Phi_{\rm GW}(t)$ is a rapidly varying phase. Therefore, we must express all the terms that vary in the orbital or the precessional time scales in terms of exponentials. This includes $\Phi^{\rm orb}$, $\delta\phi$, ι, and ψ because

$$\dot{\Phi}^{\rm orb} \sim \mathcal{O}(v^{-3}), \quad \delta\dot{\phi} \sim \dot{\psi} \sim \dot{\imath} \sim \mathcal{O}(v^{-6})\,, \tag{4.80}$$

and

$$\ddot{\Phi}^{\rm orb} \sim \mathcal{O}(v^{-11}), \quad \delta\ddot{\phi} \sim \ddot{\psi} \sim \ddot{\imath} \sim \mathcal{O}(v^{-11})\,. \tag{4.81}$$

We leave in the amplitude only terms that vary on the radiation reaction time scale.

Expressing the Φ dependence of $h_n(t)$ in Eq. (4.63) and the ψ dependence of antenna pattern functions in Eqs. (4.67) and (4.68) as exponentials, we find [124]

$$h_{n,+} = \frac{1}{2}\left(\mathcal{A}_{n,+} - i\mathcal{B}_{n,+}\right) e^{in\Phi} + \text{c.c.}\,, \tag{4.82}$$

$$h_{n,\times} = \frac{1}{2}\left(\mathcal{A}_{n,\times} - i\mathcal{B}_{n,\times}\right) e^{in\Phi} + \text{c.c.}\,, \tag{4.83}$$

$$F_+ = \frac{1}{2}\left(\mathcal{A}_F + i\mathcal{B}_F\right) e^{2i\psi} + \text{c.c.}\,, \tag{4.84}$$

$$F_\times = \frac{1}{2}\left(\mathcal{B}_F - i\mathcal{A}_F\right) e^{2i\psi} + \text{c.c.}\,, \tag{4.85}$$

where c.c. stands for complex conjugate. The amplitudes

$$\mathcal{A}_F \equiv \frac{1}{2}(1+\cos^2\theta_N)\cos 2\phi_N\,, \tag{4.86}$$

$$\mathcal{B}_F \equiv \cos\theta_N \sin 2\phi_N\,, \tag{4.87}$$

depend only on the slowly varying sky-location angles (θ_N, ϕ_N). The amplitudes $(\mathcal{A}_{n,+}, \mathcal{A}_{n,\times}, \mathcal{B}_{n,+}, \mathcal{B}_{n,\times})$ depend on the rapidly varying ι which we also express in terms of complex exponentials.

Combining all these results and expanding all terms in a Fourier series we obtain [124]

$$h(t) = \frac{\mu v^2}{D_L} \sum_{n\geq 0} \sum_{k\in\mathbb{Z}} \sum_{m=\pm 2} h_{n,k,m}(t)\,, \tag{4.88}$$

where D_L is the luminosity distance and

$$h_{n,k,m}(t) = \mathcal{A}_{n,k,m}(\theta,\phi) e^{i\Phi^{\rm GW}_{nkm}(t)} + \text{c.c.}\,, \tag{4.89}$$

where we have defined

$$\Phi^{\rm GW}_{nkm}(t) \equiv n\Phi^{\rm orb}(t) + n\delta\phi(t) + n\Phi^{\rm log}(t) + k\iota(t) + m\psi(t), \tag{4.90}$$

and

$$\Phi^{\rm log}(t) \equiv \left(6 - 3\eta v^2\right) v^3 \ln v. \tag{4.91}$$

This last term arises when converting certain log-dependent amplitude terms into phase terms [22] and from now on it will be included in $\Phi = \Phi^{\rm orb} + \delta\phi + \Phi^{\rm log}$. The slowly varying amplitudes are given in Appendix E of [124] to 2PN order, while Appendix F of [62] presents the 2.5PN contribution.

Numerical Comparison

We now compare the analytical approximate waveform response of Eq. (4.88) to a numerical one. The latter is computed by numerically solving Eqs. (4.1)-(4.3) for the momenta, with k given by Eq. (4.4), and then inserting these into the response function of Eq. (4.63). The numerical solutions are obtained with the same numerical algorithms discussed previously. Moreover, for the comparisons to follow, we will choose the same test system as before. Additionally, we choose $(\theta_N, \phi_N) = (\pi/3, 2\pi/3)$ for the polar angles of the line-of-sight unit vector in the detector frame, and $(\theta_0, \phi_0) = (2\pi/3, -2\pi/3)$ for the total angular momentum at $t = 0$ in the detector frame. Finally, we align the waveform time offset so that the phases agree at $f_{\rm GW} = 10$ Hz. As already explained, this system is representative of a variety of other systems studied; we postpone a more detailed population study for later.

Figure 4.5 shows the dominant ($n = 2$) harmonic of the total waveform phase as a function of the GW frequency in Hz. The bottom panel shows the phase difference in radians as a function of the GW frequency. The analytical approximation tracks

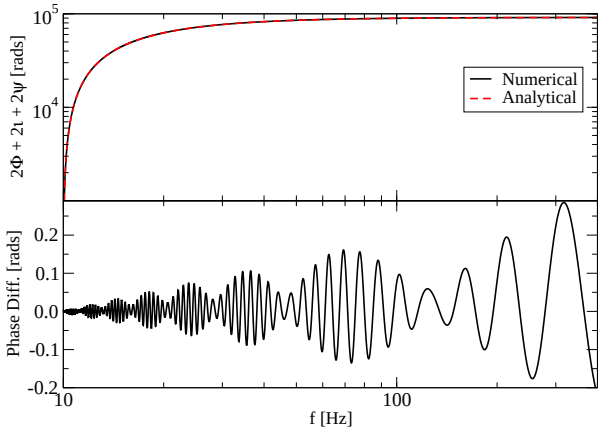

Figure 4.5: Dominant harmonic of the total waveform phase (top) and phase difference (bottom) in radians as a function of GW frequency in Hz, computed numerically (black solid curve) and analytically (red dashed curve).

the numerical solution to better than 0.3 radians over 9×10^4 radians of evolution in the phase. The error is a combination of a secular drift of the mean, superimposed with an oscillation; we shall see below that all of this error is induced by different $\mathcal{O}(\epsilon_s^2)$ effects, which could be recovered if our calculation was carried out to next order.

Figure 4.6 shows different pieces of the dominant ($n = 2$) harmonic of the waveform phase as a function of frequency (top panels), together with the dephasing between analytical and numerical expressions. The top-left plot shows the orbital phase in radians, which is by far the dominant contribution to the total phase. The analytical approximation tracks the numerical result to about 2×10^{-2} radians. The oscillatory features of the phase in Fig. 4.5 are due to the Thomas phase, the inclination angle and the polarization angle, shown in the top-right, bottom-left and bottom-right plots respectively. In all cases the analytical approximation tracks the numerical solution accurately, with dephasing of order a few times 10^{-2}. The Thomas phase, in addition, presents a secular drift, which we can see is accurately matched

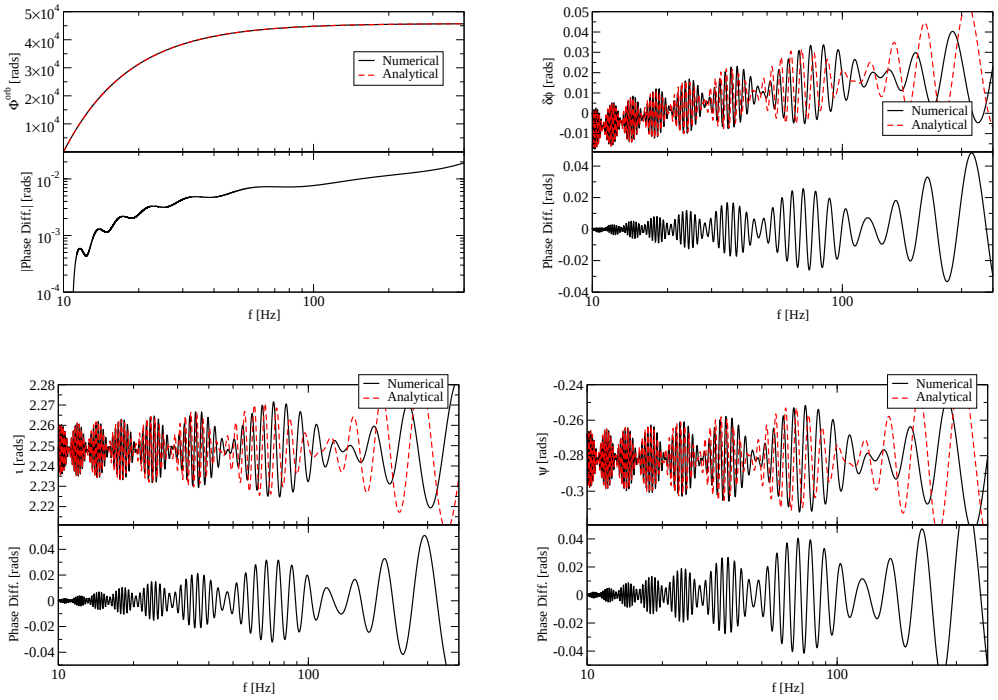

Figure 4.6: Different pieces of the dominant harmonic of the total waveform phase (top panels) in radians and phase difference (bottom panels) in radians and as a function of GW frequency in Hz, computed numerically (black solid curve) and analytically (red dashed curve). These pieces include the orbital phase (top left), the Thomas phase (top right), the inclination angle (bottom left) and the polarization phase (bottom right).

by the analytical approximation due to the correction in Eq. (4.76). We have verified that these errors scale with spin squared, as expected from the fact that the analytic approximation does not consistently account for all $\mathcal{O}(\epsilon_s^2)$ effects.

Frequency-Domain Waveforms

In this section, we construct an analytical approximation to the Fourier transform of the approximate time-domain response function found in the previous section. We begin by defining the basic tools needed and then we apply them to the approximate time-domain response function of Eq. (4.88). We conclude with a comparison of this analytical frequency-domain waveform to the discrete Fourier transform (DFT) of the numerical time-domain waveform used in the previous section.

Basics

Having expressed the time-domain waveform in the desired form, i.e. as a product of a slowly varying amplitude and a rapidly varying phase, we can now Fourier-transform it. The technique that we are going to use is the SPA [37], where the Fourier transform

$$\tilde{h}(f) = \int h(t) e^{2\pi i f t} dt \tag{4.92}$$

is computed by taking into account only the part of the integrand where the integral accumulates the most.

We rewrite the Fourier transform of Eq. (4.92) as a sum of harmonics

$$\tilde{h}(f) = \frac{\mu v^2}{D_L} \sum_{n\geq 0} \sum_{k\in\mathbb{Z}} \sum_{m=\pm 2} \tilde{h}_{n,k,m}(f)\,, \tag{4.93}$$

where

$$\tilde{h}_{n,k,m}(f) = \int \mathcal{A}_{n,k,m} e^{i(2\pi f t + n\Phi + k\iota + m\psi)} dt + \int \mathcal{A}^*_{n,k,m} e^{i(2\pi f t - n\Phi - k\iota - m\psi)} dt \,. \tag{4.94}$$

These integrals are dominated by the regions where the phase is stationary, i.e. where the argument of the exponential is nearly constant. Otherwise, the integrand oscillates rapidly and the integral cancels out by the Riemann-Lebesgue Lemma [37]. Given this and the symmetry properties of Fourier transforms of real signals, the first term of Eq. (4.94) is subdominant for positive frequencies and can be neglected.

The SPA replaces the argument in the exponential of Eq. (4.94) by a Taylor expansion about the stationary point $t_{\rm SP}$ defined by

$$2\pi f = n\dot{\Phi}(t_{\rm SP}) + ki(t_{\rm SP}) + m\dot{\psi}(t_{\rm SP}) \,. \tag{4.95}$$

This approximation works provided the amplitude varies much more slowly than the phase:

$$\left| \frac{\dot{\mathcal{A}}_{n,k,m}}{\mathcal{A}_{n,k,m}} \right| \ll \left| n\dot{\Phi} + ki + m\dot{\psi} \right| \,. \tag{4.96}$$

In the SPA, one must invert Eq. (4.95) to obtain a relation for the orbital frequency, or equivalently the PN expansion parameter v, as a function of the Fourier frequency f. When precession is present, an exact inversion is not possible because Eq. (4.95) is transcendental. We can, however, take Eq. (4.80) into account and approximate the inversion by setting

$$v_n^{\rm SP} = \left(\frac{2\pi f}{n} \right)^{1/3} \equiv v_n \,, \tag{4.97}$$

which is an excellent approximation to the location of the stationary point. We have investigated perturbative corrections to Eq. (4.97) due to precession effects and found that these have a very small effect on the Fourier domain waveform. In fact, this effect is much smaller than other errors already contained in the time-domain waveform.

Waveform Families

With all of this at hand, the full SPA to the Fourier transform of Eq. (4.92) is

$$\tilde{h}_{\rm FSP}(f) = \frac{\eta v^2}{D_L} \sum_{n\geq 0} \sum_{k\in\mathbb{Z}} \sum_{m=\pm 2} \mathcal{A}^{\rm GW}_{n,k,m}(f)\, e^{i\,\Psi_{nkm}(f)}\,. \tag{4.98}$$

The decomposed Fourier amplitude is

$$\mathcal{A}^{\rm GW}_{n,k,m}(f) = \frac{\sqrt{2\pi}\,\mathcal{A}^{*}_{n,k,m}}{\left| n\ddot{\Phi}^{\rm orb}_n + n\delta\ddot{\phi}_n + n\ddot{\Phi}^{\rm log}_n + k\ddot{\iota}_n + m\ddot{\psi}_n \right|^{1/2}}\,, \tag{4.99}$$

where $\mathcal{A}^{*}_{n,k,m}(v_n)$, $\ddot{\Phi}^{\rm orb}_n = \ddot{\Phi}^{\rm orb}(v_n)$, $\delta\ddot{\phi}_n = \delta\ddot{\phi}(v_n)$, $\ddot{\Phi}^{\rm log}_n = \ddot{\Phi}^{\rm log}(v_n)$, $\ddot{\iota}_n = \ddot{\iota}(v_n)$ and $\ddot{\psi}_n = \ddot{\psi}(v_n)$ are to be evaluated at the stationary point of Eq. (4.97). These second time derivatives are presented explicitly in Appendix E of [62].

The decomposed Fourier phase is

$$\Psi_{nkm}(f) = \Psi^{\rm nonprec}_n - n\delta\phi_n - k\iota_n - m\psi_n\,, \tag{4.100}$$

where the nonprecessing Fourier phase is given by

$$\Psi^{\rm nonprec}_n(f) \equiv 2\pi f t_n - n\Phi^{\rm orb}_n - n\Phi^{\rm log}_n + \delta\Psi_n - \frac{\pi}{4}\,. \tag{4.101}$$

The time-frequency function t_n is given in Appendix G of [62], the orbital phase $\Phi^{\rm orb}_n$ in Eq. (4.71), and the log-phase term $\Phi^{\rm log}_n$ in Eq. (4.91) all as a function of v. The

Thomas phase $\delta\phi$ is given in Eqs. (4.73), (4.76) and (4.35), while the inclination and the polarization angles are given by Eqs. (4.66) and (4.69) respectively, all as a function of $\hat{\boldsymbol{L}}$. The angular momentum is given in Eqs. (4.54)-(4.56) in terms of $\phi_{1,2}$, which in turn is given in Eq. (4.35) as a function of v. All of these expressions must be evaluated at the stationary point $v = v_n$ of Eq. (4.97), which is why we included an n subindex.

The nonprecessing Fourier phase can be simplified to

$$\Psi_n^{\rm nonprec}(f) = 2\pi f t_c - n\phi_c - n\Phi_n^{\rm log} + \delta\Psi_n - \frac{\pi}{4} + \frac{3n}{256\eta} v_n^{-5} \times \sum_{i=0}^{16} \left[\Psi_i + \Psi_i^{\ell} \ln v_n + \Psi_i^{\ell^2} (\ln v_n)^2\right] v_n^i + \mathcal{O}(v^{17}) \,, \tag{4.102}$$

where t_c is the time of coalescence, ϕ_c is the phase of coalescence, and the coefficients $(\Psi_i, \Psi_i^{\ell}, \Psi_i^{\ell^2})$ are given in Appendix H of [62] for convenience. As before, recall that formally, the sum in Eq. (4.102) is consistent only up to 3.5PN order, but we here artificially keep terms up to 8PN order, so that any dephasings found when comparing to numerical waveforms are induced purely by spin-precession effects.

The SPA phase correction $\delta\Psi_n$ is the first subleading modification to the SPA condition, i.e. to Eq. (4.96). This term arises by retaining the third time derivative in the Taylor expansion of the argument of the exponential in Eq. (4.94). Reference [84] calculated this correction to leading PN order for the $n = 2$ harmonic in the nonspinning case. We have here extended this to 3PN order beyond leading for arbitrary harmonic number:

$$\delta\Psi_n = \frac{184}{45\,n}\,\eta\,v_n^5 \left[1 + \frac{61}{46} a_2\, v_n^2 + \frac{89}{46} a_3\, v_n^3 + \left(\frac{123}{46} a_4 - \frac{131}{184} a_2^2\right) v_n^4 + \left(-\frac{175}{92} a_3 a_2 + \frac{163}{46} a_5\right) v_n^5 + \left(\frac{147}{46} b_6 + \frac{627}{46} b_6 \ln v_n + \frac{233}{368} a_2^3 - \frac{225}{92} a_4 a_2\right.$$

$$-\frac{227}{184}a_3^2+\frac{209}{46}a_6\Big)\,v_n^6\Big]\ . \tag{4.103}$$

This phase correction improves the agreement between our SPA waveforms and numerical ones for nonspinning systems by as much as an order of magnitude at low frequencies. One could include spin corrections to this equation that, for example, arise from the secular correction to the Thomas phase in Eq. (4.76). We have empirically found, however, that these effects are much smaller than $\mathcal{O}(\epsilon_s^2)$ errors already contained in the time-domain waveform approximant.

The nonspinning ingredients that go into the SPA phase in Eq. (4.100) are artificially of higher PN order than what one is allowed to formally keep. For example, $t_n = t(v_n)$, $\Phi_n^{\rm orb} = \Phi^{\rm orb}(v_n)$ and $\Psi_n^{\rm nonprec} = \Psi^{\rm nonprec}(v_n)$ are kept to 8PN order. As already discussed, whenever possible we keep all such nonspinning ingredients to 8PN order so as to minimize the dephasing for nonspinning inspirals between our SPA waveform and the DFT of the time-domain numerical waveform. Indeed, we find that keeping terms up to this order reduces the dephasing to $\mathcal{O}(10^{-5})$ radians in the nonspinning case. Therefore, any dephasings we show next for spinning systems are exclusively due to spin-precession effects and not due to any disagreement between the SPA and numerical waveforms for the backgroung nonspinning motion.

Restricted PN SPA Now that we have an analytical understanding of the full SPA waveform, we can apply some further approximations to simplify the waveform family. A typical approximation is the *restricted PN* one, where one retains only the leading-order PN terms in the Fourier amplitude, but includes all known PN order terms in the Fourier phase. We calculate this waveform here.

The restricted PN approximation amounts to retaining only the $n = 2$ term in the sum of Eq. (4.98), namely

$$\tilde{h}_{\rm RSP}(f) = \sqrt{\frac{\pi}{\left|\ddot{\Phi}^{\rm orb}\right|}}\frac{\eta v^2}{D_L} e^{i(2\pi f t - 2\Phi^{\rm orb} - 2\Phi^{\rm log} + \delta\Psi^{\rm SP} - \frac{\pi}{4})} \sum_{k=0}^{2}\sum_{m=\pm 2} \mathcal{A}^*_{2,k,m} \left|1 + \frac{\delta\ddot{\phi}}{\ddot{\Phi}^{\rm orb}} + \frac{k}{2}\frac{\ddot{\imath}}{\ddot{\Phi}^{\rm orb}} + \frac{m}{2}\frac{\ddot{\psi}}{\ddot{\Phi}^{\rm orb}}\right|^{-\frac{1}{2}} e^{-i(2\delta\phi + k\iota + m\psi)}\,. \tag{4.104}$$

where $\ddot{\Phi}^{\rm orb}$ is to be substituted with its leading-order PN value $\ddot{\Phi}^{\rm orb} = 96\eta v^{11}/5$, v is to be evaluated at v_2 as given in Eq. (4.97), and we have eliminated the subindex n, since all quantities here are to be evaluated at $n = 2$. To leading PN order the amplitudes $\mathcal{A}^*_{2,k,m}$ are given by Eqs. (E14)-(E16) of [124] with $v = 0$. Notice that, even though we can safely ignore $\ddot{\Phi}^{\rm log} \sim v^{19}$, we *cannot* expand the square root in the second line of Eq. (4.104) because $\ddot{\Phi}^{\rm orb}$ is of the same PN order as $\delta\ddot{\phi}$, $\ddot{\imath}$ and $\ddot{\psi}$.

We can rewrite the above in the more suggestive form

$$\tilde{h}_{\rm RSP}(f) = h_{\rm nonprec}(f)\; h_{\rm prec}(f)\,, \tag{4.105}$$

where we have defined the nonprecessing contribution

$$h_{\rm nonprec} = \sqrt{\frac{5}{96}}\pi^{-2/3}\frac{\eta^{1/2}}{D_L} f^{-7/6} e^{i\Psi_2^{\rm nonprec}(f)}\,, \tag{4.106}$$

and the precession correction

$$h_{\rm prec}(f) = e^{-2i\delta\phi}\sum_{k=0}^{2}\sum_{m=\pm 2}\mathcal{A}^*_{2,k,m} e^{-i(k\iota + m\psi)}\left|1 + \frac{5}{96\eta}v^{-11}\left(\delta\ddot{\phi} + \frac{k}{2}\ddot{\imath} + \frac{m}{2}\ddot{\psi}\right)\right|^{-1/2}, \tag{4.107}$$

where the second time derivatives of the phases are given in Appendix E of [62], we have used that $\ddot{\Phi}^{\rm orb} = 96\eta v^{11}/5$ to leading PN order and ignored $\ddot{\Phi}^{\rm log}$ again, since it is subdominant. The Thomas phase, the inclination angle and the polarization angle can be found in Eqs. (4.73), (4.76), (4.35), (4.66) and (4.69). The amplitudes $\mathcal{A}_{n,k,m}$ can be found in Eqs. (E14)-(E16) of [124] with $v = 0$. The nonprecessing phase is given in Eq. (4.102) with $n = 2$ evaluated at $v_2 = (\pi f)^{1/3}$ with the coefficients $(\Psi_i, \Psi_i^\ell, \Psi_i^{\ell^2})$ given in Appendix H of [62].

Applicability of the SPA

Both the restricted and full SPA waveform families defined above rely on the assumptions of the SPA being valid. In particular, these solutions require that the first nonvanishing time derivative of the argument of the exponential in Eq. (4.94) be the second time derivative. If this is not the case, then the denominator in Eqs. (4.99) or (4.104) would vanish and the SPA amplitude would diverge. When this is the case, a more sophisticated approximation to the generalized Fourier integral is required, such as the method of steepest descent and uniform asymptotics [37, 124].

Whether the second time derivative of the phase vanishes or not depends sensitively on the system considered. The parameters that affect this the most for small spin systems are the angles associated with $\hat{\boldsymbol{N}}$ and those associated with $\hat{\boldsymbol{L}}$. Indeed, as we can see in Appendix E of [62], $\ddot{\Phi}^{\rm orb}$, $\delta\ddot{\phi}$, $\ddot{\imath}$, and $\ddot{\psi}$ are all of the same PN order. However, $\ddot{\Phi}^{\rm orb}$ is of $\mathcal{O}(\epsilon_s^0)$, whereas the others are of $\mathcal{O}(\epsilon_s)$. Furthermore, $\ddot{\Phi}^{\rm orb}$ is always positive whereas the others oscillate around zero. Thus, for the second time derivative of the phase to vanish, a factor of $\mathcal{O}(\epsilon_s^{-1})$ has to multiply $\delta\ddot{\phi}$, $\ddot{\imath}$, or $\ddot{\psi}$ in order for these phases to be comparable to $\ddot{\Phi}^{\rm orb}$. By looking at their expressions, we can see that this is the case when the angle between $\hat{\boldsymbol{N}}$ and $\hat{\boldsymbol{L}}$ is of $\mathcal{O}(\epsilon_s^{-1})$.

To verify this scaling, we carried out a set of Monte Carlo studies by randomizing over all parameters, except for the dimensionless spin parameters that we kept equal for both NSs and fixed for each run. The left panel of Fig. 4.7 shows the probability that the SPA will break down as a function of the spin parameter. For systems with $\chi_A < 0.1$, this probability is smaller than 0.2%. This is because the SPA breaks only when large amounts of precession are present. In particular, the SPA breaks when $\hat{\boldsymbol{N}}$ and $\hat{\boldsymbol{L}}$ are coaligned, so that the system wobbles the most as seen from the detector. This can be seen on the right panel of Fig. 4.7, which shows the probability the SPA will break as a function of the square of the cosine of the inclination angle, where we have fixed both χ_A and ι. For systems with the angle between $\hat{\boldsymbol{N}}$ and $\hat{\boldsymbol{L}}$ larger than 4°, this probability is less than 2%.

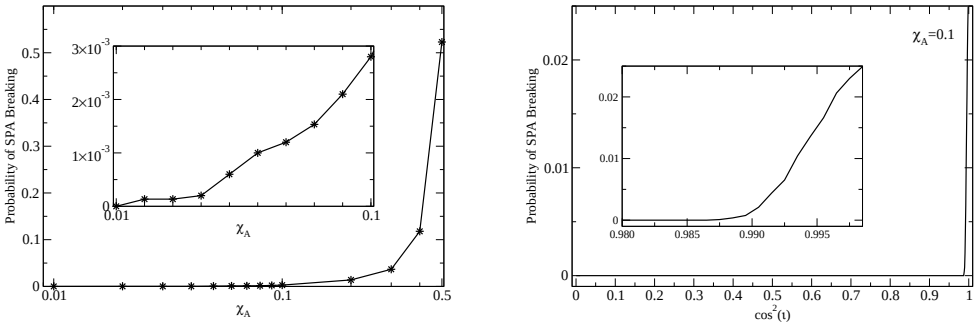

Figure 4.7: Probability that the SPA fails as a function of the dimensionless spin parameter of the system (left) and as a function of the squared of the cosine of the inclination angle (right).

We should note that the systems for which the SPA would break down are precisely those that could lead to a coincident short gamma-ray burst (short GRB) and GW observation. One of the possible progenitors of short GRBs are NS mergers. The electromagnetic observation of such a GRB would require $\hat{\boldsymbol{N}}$ and $\hat{\boldsymbol{L}}$ to be almost exactly aligned or antialigned. In both cases, the SPA waveforms constructed here and elsewhere in the literature may be ineffective at extracting a GW signal. However,

we know in advance which systems will have a failing SPA, and thus, an analysis using our waveforms could switch to the full numerical solution in this (rather small) corner of parameter space. The degree to which this is feasible and effective will be studied in a data analysis framework elsewhere.

Numerical Comparison

Having constructed a full and restricted analytical SPA model for the Fourier transform of the analytical waveform response we now compare it to the purely numerical waveform. The two former ones are given by Eqs. (4.98) and (4.105) respectively. The latter is computed by first applying a window function to the numerical, time-domain waveform, then discretizing this waveform and finally computing the Fourier transform using the FFTW routine. We ensure that the number of points used in the discretized waveform time series is large enough that the Nyquist frequency is at least 5 times larger than the highest frequency of interest in our analysis.

We use a Tukey function, with parameters such that the window varies from zero to unity between $f_{\rm GW,1}$ and $f_{\rm GW,2}$ remains unity until $f_{\rm GW} = 500$ Hz and then falls off to zero between $f_{\rm GW} = 500$ Hz and $f_{\rm GW,3}$, where $f_{\rm GW}$ is the frequency of the dominant harmonic, at twice the orbital frequency. We choose the frequencies $f_{\rm GW,1}$ and $f_{\rm GW,2}$ to be 8.5 Hz and 9.5 Hz respectively when studying waveforms that only include the $n = 2$ harmonic, or 19/7 Hz and 20/7 Hz, respectively, when studying ones that include all harmonics[1]. The window goes to zero at $f_{\rm GW,3}$ which corresponds roughly to twice the orbital frequency of a test particle in the innermost stable circular orbit of a Schwarzschild BH with mass equal to the total mass of the test system. We

[1]When including all harmonics, a smaller starting frequency is required, so that the highest harmonic is included at the beginning of the frequency band, i.e. at 10 Hz.

could have chosen a different maximum frequency, but this would not have changed the results we will show in this section. We investigated a variety of filters, including the Planck-taper function of [154], but found that a Tukey window with the above parameters is optimal for minimizing spectral leakage inside the frequency regime of interest.

For the comparisons that follow, we choose the same test system as before, which is representative of all systems considered, and concentrate on the frequency region $(10, 400)$ Hz. We stop all comparisons at 400 Hz because (i) finite size effects cannot be neglected beyond that frequency and (ii) most of the SNR is contained in this frequency region for binary NS coalescences. The alignment of the time-domain waveforms at 10 Hz does not guarantee alignment in the frequency domain. We have freedom to choose the phase and time of coalescence to perform such an alignment. However, in this section, we will present the waveforms as they are, i.e. without aligning them in frequency space. This implies that the errors shown here are an overestimate of the inaccuracies in our waveforms, i.e. they are conservative.

Dominant Harmonic We begin by focusing on waveforms that contain only the PN dominant harmonic $n = 2$. The numerical waveform will then be the DFT of Eq. (4.88), keeping only the $n = 2$ harmonic. The full SPA waveform is similarly given by Eq. (4.98), keeping only the leading-order $n = 2$ harmonic, while the restricted SPA is given by Eq. (4.105). Notice that the restricted SPA waveform is less accurate than the full SPA waveform because the former keeps only the leading-order PN terms in the amplitude of the SPA Fourier transform.

The left panel of Fig. 4.8 shows the difference in Fourier phase between the DFT and either the SPA (solid black curve), an SPA Fourier phase that assumes aligned spins (dot-dashed green curve), or an SPA that sets all spins to zero (dotted

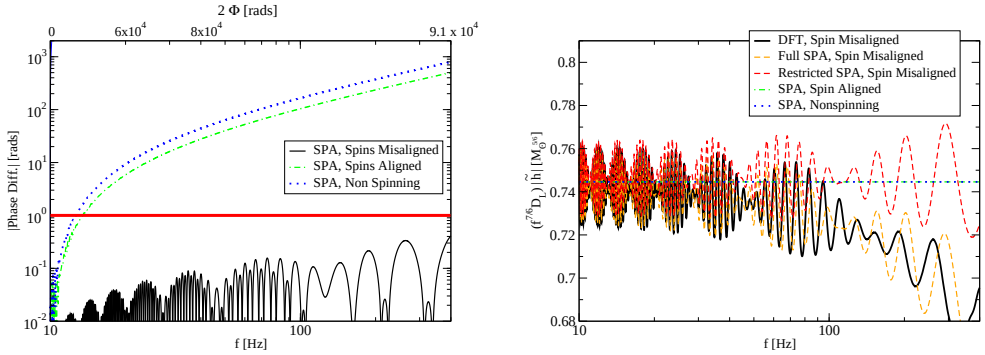

Figure 4.8: Left: Fourier dephasings between the DFT and either the SPA waveform (solid black curve), a spin-aligned SPA waveform (dot-dashed green curve), and a nonspinning SPA waveform (dotted blue curve), together with a reference line (thick red curve) at 1 rad, as a function of frequency on the bottom x axis, and accumulated time-domain GW phase on the top x axis. Right: DFT (solid black curve), full SPA (dashed orange curve), restricted PN SPA (dashed red curve), spin-aligned SPA (dot-dashed green curve), and nonspinning SPA (dotted blue curve) Fourier amplitudes.

blue curve)[2]. The bottom horizontal axis shows the GW frequency, while the top axis shows the accumulated GW phase in the time domain. A horizontal line (thick, solid red curve) at 1 rad is also shown for reference. While the aligned-spin and the nonspinning SPA phases build dephasings of $\mathcal{O}(10^3)$ rads very quickly, the new SPA waveform of this chapter remains in phase to better than 0.4 rads in $\sim 10^5$ rads of GW evolution. This dephasing is dominated by the error in the $t(v)$ function, i.e. the error in the time inversion is roughly 0.2 milliseconds at 400 Hz (see Fig. 4.4), and thus, the error in $2\pi f t(f) \approx 0.5$ rads at 400 Hz.

The right panel of Fig. 4.8 shows the Fourier amplitude multiplied by $(D_L f^{7/6})$ as a function of GW frequency in Hz for the DFT waveform (solid black curve), the full SPA waveform (dashed orange curve), the restricted PN SPA waveform (dashed red cube), the spin-aligned SPA waveform (dot-dashed green curve), and the nonspinning

[2]The latter two waveforms are well-known in the literature [24, 175], and they can be obtained by setting either the spin angular momentum to be aligned with the orbital angular momentum or the spins to zero in Eq. (4.105). In either case there is no precession.

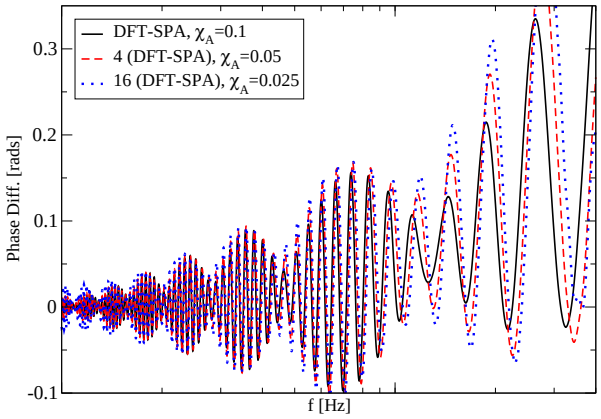

Figure 4.9: Phase difference as a function of GW frequency in Hz. The solid black curve corresponds to the test system with equal spins of $\chi^{(1)} = 0.1$, while the dashed red and dotted blue ones to the same system but with equal spins of $\chi^{(2)} = \chi^{(1)}/2 = 0.05$ and $\chi^{(3)} = \chi^{(1)}/4$, respectively. The latter are multiplied by a factor of $(\chi^{(1)}/\chi^{(2)})^2 = 4$ and $(\chi^{(1)}/\chi^{(3)})^2 = 16$.

SPA waveform (dotted blue curve). The latter two are flat since these waveforms neglect precession altogether. Both the restricted and the full SPA waveforms can capture the precession amplitude oscillations present in the DFT amplitude. The full SPA amplitude, however, does better than the restricted one. Both the restricted and the full SPA amplitudes, however, dephase with respect to the DFT amplitude after roughly 35 cycles of precession oscillations. This dephasing could be eliminated if we extended our results to include all $\mathcal{O}(\epsilon_s^2)$ effects. In practice, this dephasing will induce a systematic error in the determination of spin parameters, but we expect this systematic error to be small.

Figure 4.9 verifies that the error contained in the SPA waveform indeed scales with spin squared. This figure shows the Fourier phase difference as a function of GW frequency between the DFT and either the restricted SPA waveforms with equal spins of $\chi^{(1)} = 0.1$ (solid black curve), $\chi^{(2)} = \chi^{(1)}/2 = 0.05$ (red dashed curve) or $\chi^{(3)} = \chi^{(1)}/4 = 0.025$ (blue dotted curve). The differences computed for the $\chi^{(2)}$

and $\chi^{(3)}$ systems were multiplied by a factor of $(\chi^{(1)}/\chi^{(2)})^2 = 4$ and $(\chi^{(1)}/\chi^{(3)})^2 = 16$ respectively. If the uncontrolled remainder in the SPA is of $\mathcal{O}(\epsilon_s^2)$, we would expect these phase differences to roughly lay on top of each other; Fig. 4.9 verifies this expectation. Notice that the restricted SPA waveform contains errors of $\mathcal{O}(\epsilon_s)$ in the waveform amplitude because of the truncation of $\ddot{\Phi}$ at leading PN order. Therefore, we do not expect the same scaling to be true for the Fourier amplitude of the restricted SPA.

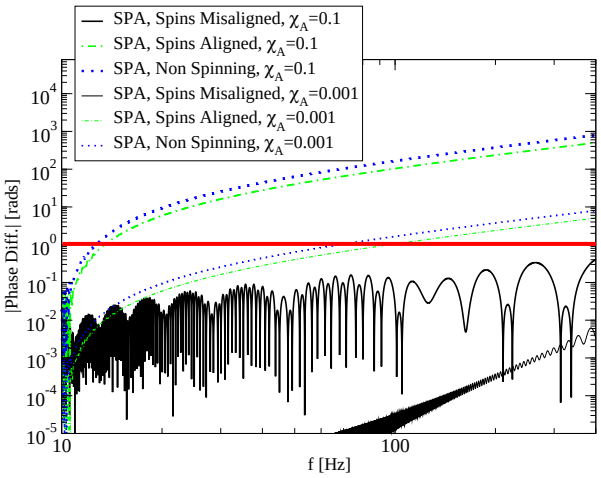

Figure 4.10: Dephasing between the DFT Fourier phase and the SPA (solid black curves), the spins-aligned SPA (dot-dashed green curves) and the nonspinning SPA (dotted blue curves), for systems with dimensionless spins of 10^{-1} (thicker) and 10^{-3} (lighter).

Figure 4.8 suggests that nonspinning and spin-aligned SPA waveforms are inadequate for parameter estimation of precessing and spinning inspirals, but is there a sufficiently small value of $\chi_{1,2}$ for which this would not be the case? Figure 4.10 investigates this question by plotting the dephasing between the DFT and either the SPA (solid black curves), the spin-aligned SPA (dot-dashed green curves) or the nonspinning SPA (dotted blue curves) waveforms, for systems with equal spins of 0.1 (thicker curves) and spins of 10^{-3} (thinner curves). Only when the spins become

smaller than $\mathcal{O}(10^{-3})$ does the dephasing of the spin-aligned or nonspinning SPA waveforms become comparable to 1 rad. This suggests that parameter estimation systematics would be introduced if we used spin-aligned waveforms to analyze precessing signals when the dimensionless spin magnitude of the latter exceeds 10^{-3}.

Full Waveform We now focus on waveforms that include all harmonics and PN amplitude corrections. The numerical waveform will then be the DFT of Eq. (4.88), keeping all known modes, i.e. up to $n = 7$. This then implies that numerical solutions to the time evolution of the angular momenta must be started at sufficiently low orbital frequency, such that the highest harmonic ($n = 7$) contributes at 10 Hz. The full SPA waveform is similarly given by Eq. (4.98), keeping all known harmonics, while the restricted SPA is again given by Eq. (4.105).

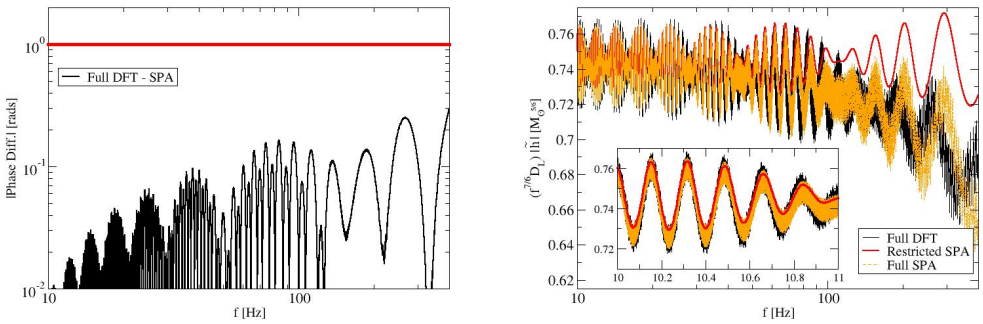

Figure 4.11: Left: difference between the full DFT and SPA Fourier phases (solid black curve) as a function of GW frequency in Hz, together with a reference line (thick red curve) at 1 rad. Right: full DFT (solid black curve), full SPA (dashed orange curve), and restricted PN SPA (solid red curve) Fourier amplitudes.

The left panel of Fig. 4.11 shows the difference between the full DFT and SPA (solid black curve) Fourier phases as a function of GW frequency in Hz. For reference, we include a horizontal (thick red) curve at 1 rad. The dephasing is still roughly 10^{-1} rads, as we found earlier when looking at a single harmonic.

The right panel of Fig. 4.11 shows the full DFT (black solid curve), the full SPA (dashed orange curve) and the restricted SPA (solid red curve) Fourier amplitudes, normalized to $f^{7/6}D_L$, as a function of GW frequency in Hz. The inset shows a zoom to the region close to 10 Hz. This figure presents a number of interesting features. First, the full DFT and SPA amplitude curves are thick. This thickness is not due to numerical noise, but rather due to the beating of higher harmonics that induce high frequency oscillations. This feature is not present in the restricted SPA amplitude, as this neglects higher amplitude harmonics all together. Second, the average of the full DFT Fourier amplitude tends to decrease with frequency. This feature is captured by the full SPA amplitude, because it includes high PN order effects that induce this trend. The restricted SPA amplitude, however, cannot recover this trend, since it retains the lowest PN order terms only. Third, in spite of not perfectly matching the full DFT waveform, the restricted SPA amplitude does a superb job at catching the initial oscillations of the full DFT up to roughly 100 Hz. This is important because most of the power accumulates between 10 and 100 Hz for NS binary inspirals.

Data Analysis Comparisons

In this section we perform a more detailed data analysis comparison for a variety of different systems. This comparison will be based on the *faithfulness* measure, namely

$$F \equiv \max_{t_c,\phi_c} \frac{(h_1 \mid h_2)}{\sqrt{(h_1 \mid h_1)\,(h_2 \mid h_2)}}\,, \tag{4.108}$$

where $h_{1,2}$ are different waveforms with the *same* physical parameters. The inner-product is defined in the usual way:

$$(h_1 \mid h_2) \equiv 4\Re \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1 \tilde{h}_2^*}{S_n}\, df\,, \tag{4.109}$$

where $\Re[\cdot]$ is the real part operator, $(f_{\min}, f_{\max})$ are the boundaries of the detector's sensitivity band, S_n is the detector's spectral noise density, and $\tilde{h}$ denotes the Fourier transform of h. Given this definition, the *fitting factor* FF is nothing but the faithfulness maximized over all template parameters, which then implies that $FF \geq F$.

We concentrate here on NS binary inspirals, and thus, on sources suitable for detection with ground-based instruments, such as LIGO. When calculating overlaps through Eq. (5.79), we choose $f_{\min} = 10$ Hz and $f_{\max} = 10$ kHz, with observation times of about 3×10^5 seconds since the lowest harmonic evolves from $f_{\rm orb} = 10/7$ Hz to coalescence. We here employ an aLIGO noise curve given by [69]

$$S_n(f) = 10^{-49}\left\{ \bar{f}^{-4.14} - \frac{5}{\bar{f}} + 111\left[1 - \bar{f}^2\left(1 - \frac{1}{2}\bar{f}^2\right)\right]\left(1 - \frac{1}{2}\bar{f}^2\right)^{-1}\right\}, \tag{4.110}$$

where we have defined the dimensionless frequency $\bar{f} = f/(215\text{ Hz})$. We use a Tukey window in all our waveforms. The Tukey window is sufficiently slowly varying that we can include it in the amplitudes of our SPA waveforms.

The image of the faithfulness measure is in the interval $[-1, 1]$; it quantifies how well waveforms agree with each other, with unity representing perfect agreement. All integrations are done numerically, with errors of $\mathcal{O}(10^{-5})$; thus, we consider that a match of $F = 0.9999$ is consistent with unity. Conventionally, a fitting factor about 97% is generally considered to be sufficient for detection. Therefore, a faithfulness of

98% certainly implies a fitting factor of at least 98%, which is also good enough for detection.

The faithfulness measure will be evaluated using a full DFT waveform as the signal and either the full SPA or the restricted SPA as the template. That is, the full DFT waveform will be our reference waveform, to which the other two template families will be compared.

Before proceeding,we add one last word of caution. Faithfulness comparisons are *conservative* because the match is not maximized over physical parameters, such as the total mass, mass ratio, spin magnitudes, or angles. Higher matches would indeed be obtained if we allowed the templates to have different physical parameters from the signal, as one does in parameter estimation. Such higher matches, of course, will come at the cost of a systematic bias in the recovered parameters.

Instead of working with a particular system, we perform a Monte-Carlo simulation with 1000 points in parameter space, with all system parameters randomized, except for the dimensionless spin magnitudes, which will be set to be equal to each other and constant along each run. We consider systems with individual masses in the range $(1.2, 2)\ M_{\odot}$, appropriate for NSs, with a flat distribution in log space. The distribution of unit vectors is chosen to be uniform on the sphere.

Figure 4.12 shows the median match and 68% (1σ) interval regions between the full DFT-full SPA waveform (red solid curve), the full DFT-restricted SPA waveform (green dot-dashed curve), the full DFT-full spin-aligned SPA waveform (blue dashed curve), and the full DFT-full nonspinning SPA waveform (magenta dotted curve) as a function of spin magnitude χ_A. The faithfulness of the new precessing waveform families is above 99% for all systems with $\chi < 0.1$, which is dramatically better than the faithfulness when using spin-aligned or nonspinning templates. In fact, this figure suggests that the latter two could introduce serious systematic errors in parameter

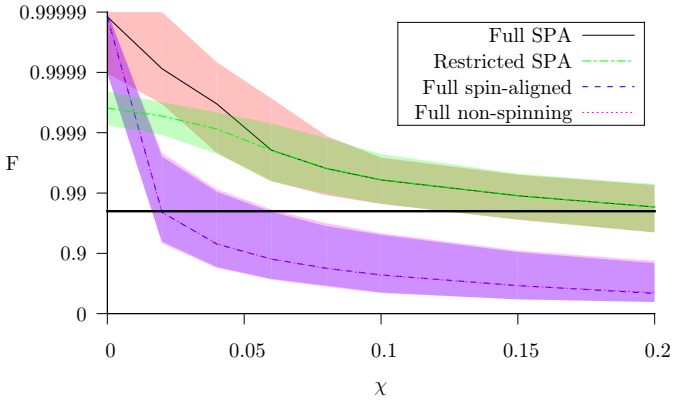

Figure 4.12: Mean faithfulness of the full SPA family (red, solid line), the restricted PN SPA family (green dot-dashed line), the full spin-aligned SPA family (blue dashed line) and the full nonspinning SPA family (red dotted line) against numerical waveforms for 1000 different systems with randomized parameters. The shaded areas show the faithfulness regions for 68% of the systems considered. For reference, the black solid line corresponds to a faithfulness of 97%.

estimation. The performance of the full and restricted SPA families is very similar because their difference is dominantly due to odd amplitude harmonics, which are proportional to the dimensionless mass difference, a small number for binary NSs. This results in higher harmonic corrections being unimportant when the faithfulness reaches $F \gtrsim 0.999$.

Apart from the spin magnitude, he faithfulness might also depend on other parameters. The two that are perhaps most important are the component masses; in this case, m_1 and m_2 have a rather limited range because we are considering NS binaries. Figure 4.13 presents contour plots of the faithfulness in the (m_1, m_2) plane for systems with equal dimensionless spin magnitudes of 0.04 (left panel) and 0.1 (right panel). We considered 1000 systems with all parametrized randomized, except for the spin magnitudes. The value of the faithfulness at any given point is the median value inside a circle centered at that point with radius $0.1M_\odot$. The faithfulness is largest along the equal-mass symmetry line where precession is suppressed.

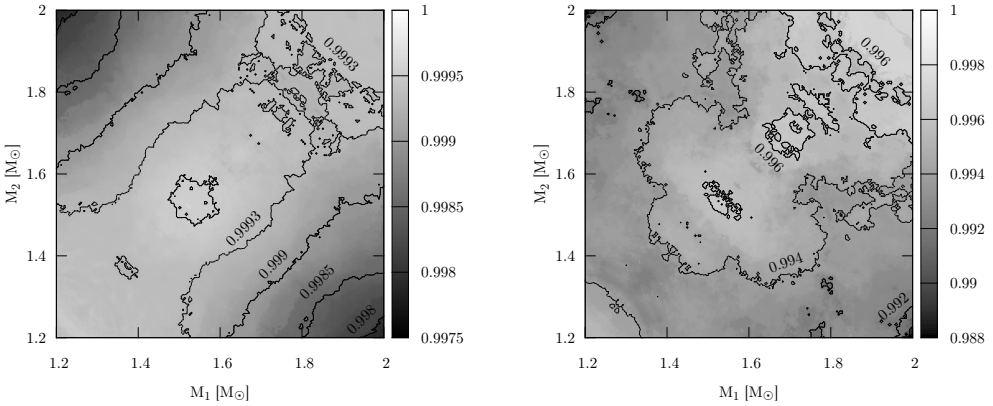

Figure 4.13: Faithfulness contour plot in the (m_1, m_2) plane for systems with equal spin magnitudes of 0.04 (left panel) and 0.1 (right panel). Each point corresponds to the median faithfulness inside a circle centered at the given point with a radius of $0.1M_{\odot}$.

Beyond the Early Inspiral

The starting frequency of our numerical comparisons corresponds to the beginning of the LIGO sensitivity band, but the ending frequency (400 Hz) does not correspond to the end of this band or to the merger frequency. Rather, it corresponds to the frequency at which finite-size effects can no longer be neglected [112, 148, 180]. In this section we summarize how such effects could be taken into account, extending the waveforms obtained here beyond 400Hz.

The finite-size effects give rise to two types of deformations: (i) multipole and, (ii) tidal. The former are described by the NS's multipole moments, which measure how much the object is deformed away from sphericity. The latter are caused by the companion's external field and they are quantified by the tidal Love number [42, 77, 110] that characterizes the deformability of the NS.

The leading-order multipole effect is the quadrupole one which is an $\mathcal{O}(\epsilon_s^2)$ effect entering the GW as a 2PN phase correction given by [157, 172]

$$\delta\Psi_n^Q = \frac{75n}{256v_n}\left(\frac{Q_1}{m_1^2 m_2} + \frac{m_1}{m_2}\chi_1^2\right)\left[3\left(\hat{\boldsymbol{S}}_1\cdot\hat{\boldsymbol{L}}\right)^2 - 1\right] + 1 \leftrightarrow 2\,, \tag{4.111}$$

where recall that $v_n = (2\pi f/n)^{1/3}$ and n is the harmonic number. In the above equation, Q_A is the quadrupole moment of each binary component[3]. In principle there are higher-order multipole deformations that affect the waveform phase, but these are proportional to higher powers of spin, and thus, negligible.

On the other hand, the tidal deformations caused by the gravitational field of the companion, will result in a 5PN phase correction [95]

$$\delta\Psi_n^\lambda = -\frac{9n}{32\eta}\left(1 + 12\frac{m_2}{m_1}\right)\lambda_1 v_n^5 + 1 \leftrightarrow 2\,, \tag{4.112}$$

where λ_A is the tidal Love number of the Ath binary component. Corrections to the above equation up to 2.5PN can be found in [78, 222].

Even though we can use Eqs. (4.111) and (4.112) to minimize the error induced by the point-particle approximation, we still cannot extend the region of validity of inspiral PN waveforms beyond a certain frequency. For sure one cannot use simple inspiral waveform beyond the frequency at which the two NSs touch. This contact frequency can be approximated by that at which the NS's separation is equal to the sum of their radii, neglecting tidal deformations. In Table 4.1 we give an estimate of that frequency for our test system for various equations of state (EoS). For all EoS considered here, the regime we study can be considered as the early inspiral,

[3]As in Ref. [239], Q_A is normalized such that Eq. (4.111) gives the contribution to the phase from deviations of the quadrupole moment from the BH value.

EoS	R_1(km)	R_2(km)	f(Hz)
APR [12]	12.2	12.2	1425
SLy [87]	11.4	11.6	1553
LS220 [135]	12.7	13.4	1292
Shen [196]	14.5	14.9	1089

Table 4.1: Frequency at which the separation of the two NSs is equal to the sum of their radii for various EoS.

where finite-size effects have not yet started to have an observable signature in the waveforms.

Beyond this frequency the system enters the highly nonlinear and dynamical merger phase, the outcome of which is still rather uncertain [15, 90, 102, 115, 122, 142, 200]. Depending on the EoS and the mass ratio, the merger remnant might be a hypermassive NS, which then collapses to a BH, or it could simply collapse directly to a BH. In either case, nonlinear dynamics play an important role and would have to be mimicked in some phenomenological way if one wishes to construct an effective analytical template.

Conclusions

In this chapter we compute computationally inexpensive, analytical waveforms for spinning binaries with small spins. We do so by first obtaining an analytic perturbative solution to the spin precession equations and, then, by analytically Fourier transforming the resulting time domain waveform. The waveform family presented here is a definite step towards modeling spinning binaries and constructing waveforms that can be used for parameter estimation. Figure 4.12 has explored how efficient this waveform family is by calculating the faithfulness between full numerical

waveforms and analytical ones, but a full data analysis parameter estimation study is left for Chapter 6.

The waveform family computed in this chapter fails to capture the strong precession effects induced by rapidly spinning compact objects with arbitrary spin orientations, like BHs. Astrophysical BHs can easily have dimensionless spin magnitudes larger than 0.1, thus violating our small spins approximation. One could extend our calculation to next order in $\mathcal{O}(\epsilon_s)$ and study whether the extended solution is now accurate enough to model moderately spinning BHs. Instead of improving the small-spin model presented here, in Chapter 5 we present a GW model that is valid for arbitrary spins magnitudes and orientations.

GW TEMPLATES FOR GENERIC BINARIES

In this chapter we utilize two recent breakthroughs to construct waveforms for spin-precessing systems with arbitrary spin magnitudes and orientations with MSA.

The first breakthrough in the modeling of generic spin-precessing binaries was by Kesden, et al. [99, 100, 121]. Neglecting radiation reaction, the authors found an *exact* solution to the precession equations that govern the evolution of the orbital and the spin angular momenta of the binary. By identifying certain constants of the precessional motion, they were able to express all angular momenta as functions of the total spin magnitude, which satisfies an ordinary differential equation. We here solve this differential equation analytically and obtain an exact solution to the precession equations in the absence of radiation reaction. We then use MSA to introduce radiation reaction perturbatively as an expansion in the ratio of the precession to the radiation reaction timescale. With this at hand, we obtain time-domain waveforms in terms of the parameters of the system only.

The second breakthrough in the modeling of generic spin-precessing binaries was by Klein, et al. [125], and tackles the failure of the SPA. The authors introduced the *shifted uniform asymptotics* (SUA) method where the waveform is decomposed into Bessel functions, the Fourier integral is evaluated term by term in the SPA, and then resumed using the exponential shift theorem. The result is a closed-form analytic expression for the gravitational wave in the frequency domain as a series of time-domain waveforms evaluated at shifted stationary times. Unlike previous approaches, both the time- and frequency-domain waveforms that we obtain are valid for arbitrary mass ratios, arbitrary spin magnitudes and arbitrary spin orientations.

Closed-form expressions for the waveforms have several advantages. From a theoretical standpoint, analytic solutions shed light on the physical processes at play,

the structure of the resultant signal, and the transition through different resonant states. From a practical standpoint, analytic solutions are in general faster to evaluate, avoiding costly numerical integrations and discrete Fourier transforms. Estimating the computational gain from closed-form, analytic expressions relative to numerical ones is not straightforward since it depends heavily on the implementation. However, the closed-form, analytic frequency-domain waveforms computed here has the potential to be much faster than any numerical waveform, while still encompassing all precessional effects.

Spin and Angular Momentum Evolution

A quasicircular binary system consisting of generic spinning compact objects is subject to spin-orbit and spin-spin interactions that force all angular momenta to precess. Averaging over one orbit, the precession equations governing the conservative evolution of the orbital and spin angular momenta are [28, 52, 211]

$$\begin{aligned}\dot{\hat{\boldsymbol{L}}} = &\left\{\left(2+\frac{3}{2}q\right) - \frac{3}{2}\frac{v}{\eta}\left[(\boldsymbol{S}_2+q\boldsymbol{S}_1)\cdot\hat{\boldsymbol{L}}\right]\right\} v^6\left(\boldsymbol{S}_1\times\hat{\boldsymbol{L}}\right)\\ &+\left\{\left(2+\frac{3}{2q}\right) - \frac{3}{2}\frac{v}{\eta}\left[\left(\boldsymbol{S}_1+\frac{1}{q}\boldsymbol{S}_2\right)\cdot\hat{\boldsymbol{L}}\right]\right\} v^6\left(\boldsymbol{S}_2\times\hat{\boldsymbol{L}}\right)+\mathcal{O}(v^7),\end{aligned} \tag{5.1}$$

$$\dot{\boldsymbol{S}}_1 = \left\{\eta\left(2+\frac{3}{2}q\right) - \frac{3v}{2}\left[(q\boldsymbol{S}_1+\boldsymbol{S}_2)\cdot\hat{\boldsymbol{L}}\right]\right\} v^5\left(\hat{\boldsymbol{L}}\times\boldsymbol{S}_1\right)+\frac{v^6}{2}\boldsymbol{S}_2\times\boldsymbol{S}_1+\mathcal{O}(v^7), \tag{5.2}$$

$$\dot{\boldsymbol{S}}_2 = \left\{\eta\left(2+\frac{3}{2q}\right) - \frac{3v}{2}\left[\left(\frac{1}{q}\boldsymbol{S}_2+\boldsymbol{S}_1\right)\cdot\hat{\boldsymbol{L}}\right]\right\} v^5\left(\hat{\boldsymbol{L}}\times\boldsymbol{S}_2\right)+\frac{v^6}{2}\boldsymbol{S}_1\times\boldsymbol{S}_2+\mathcal{O}(v^7). \tag{5.3}$$

Radiation reaction drives the evolution of the magnitude of the orbital angular momentum, leaving the magnitude of the spin angular momenta unaltered to our current knowledge of the PN expansion and ignoring all energy and angular

momentum flux through BH horizons [63, 173]. The magnitude L is related to the evolution of the orbital frequency ω, and the PN expansion parameter v, leading to

$$\dot{v} = \frac{v^9}{3} \frac{1}{\sum_{n=0}^{7} \left[g_n + 3 g_n^{\ell} \ln(v) \right] v^n}. \tag{5.4}$$

The coefficients $\{g_n, g_n^{\ell}\}$ are functions of the symmetric mass ratio and inner products of the angular momenta, given in Appendix A of [61].

Equations (5.1)-(5.3) describe the *conservative dynamics*, while Eq. (5.4) describes the *dissipative dynamics*. The former models the spin-spin and spin-orbit interactions, that change only the direction of $\boldsymbol{L}, \boldsymbol{S_1}$ and $\boldsymbol{S_2}$. We use only the leading PN order expressions in each interaction[1]. We do not use higher PN order corrections because the spin-spin and spin-cubed terms have not been fully calculated for generic precessing orbits yet [46, 178]. In principle, we could have included the spin-orbit corrections. However, as explained later, our solution makes use of a certain quantity [177] that is conserved by the leading-order in spin-orbit and spin-spin interactions precession equations. If we use partial precession equations (including spin-obit but not spin-spin corrections) it is not clear if we can modify this quantity to remain conserved. Once the spin-spin and spin-cubed terms have been fully calculated we can revisit this problem.

The dissipative dynamics govern the GW frequency evolution by changing the magnitude of the Newtonian orbital angular momentum $L = \eta/v$. This equation is known to 2.5PN order in all spin interactions [178], 4PN in linear-in-spin terms [44, 47, 48, 92, 150, 178] and 22PN order in the point particle limit, neglecting spins and BH absorption effects [97, 158, 199, 209]. In our analysis we keep terms in Eq. (5.4) to 3.5PN order since this is the highest complete PN order, ignoring spin-spin terms.

[1]Spin-orbit corrections can be found in [48, 94], spin-spin in [46], and spin-cubed in [149].

In this case, we can easily include partial PN terms in radiation reaction to make the evolution more accurate. When the 3PN spin-spin term has been fully calculated for precessing orbits [46] we can include it in our model.

Conservative and dissipative equations evolve on distinct timescales. The former evolve on the *precession timescale*

$$T_{\rm pr} \equiv \frac{|\boldsymbol{S}_1|}{|\dot{\boldsymbol{S}}_1|} \sim v^{-5}, \tag{5.5}$$

while the later evolve on the *radiation reaction timescale*

$$T_{\rm rr} \equiv \frac{v}{\dot{v}} \sim v^{-8}. \tag{5.6}$$

The ratio $T_{\rm pr}/T_{\rm rr} \sim v^3$ is small in the inspiral and thus a natural expansion parameter.

Recently, Kesden et. al. [121] found an exact solution to the precession equations [Eqs. (5.1)-(5.3)] ignoring radiation reaction [Eq. (5.4)]. This solution can be used to "precession-average" the full precession equations with radiation reaction (analogously to orbit-averaging). The final precession-averaged equations depend only on quantities that vary on the radiation reaction timescale, and can be numerically integrated with a larger step size [100].

Here we take a different approach. Rather that precession-averaging Eqs. (5.1)-(5.3) and numerically accounting for Eq. (5.4), we make explicit use of the fact that $T_{\rm pr}/T_{\rm rr} \sim v^3$ to solve the precession equations analytically. We use a perturbation theory technique known as multiple scale analysis and treat radiation reaction as a slowly-evolving perturbation on top of precession. This approach allows us to find a solution to the full set of Eqs. (5.1)-(5.4) as an expansion in $T_{\rm pr}/T_{\rm rr}$.

Analytic Solution to the Precession Equations without Radiation Reaction

Ignoring radiation reaction, the precession equations can be solved analytically by making use of certain conserved quantities of the system. Below we review and complete the solution first presented in [121].

A precessing binary has a total of 9 degrees of freedom arising from the 3 components of 3 Newtonian vectors $(\boldsymbol{L}, \boldsymbol{S_1}, \boldsymbol{S_2})$. The precession equations lead to 7 conserved quantities, reducing the degrees of freedom to 2. Of the remaining degrees of freedom, one is associated with the choice of a coordinate system, while the other corresponds to a dynamical quantity that changes with time. This dynamical quantity is chosen to be the magnitude of the total spin angular momentum $S = |\mathbf{S_1} + \mathbf{S_2}|$.

The conserved quantities are $\boldsymbol{\lambda} \equiv (S_1, S_2, L, J, \hat{\boldsymbol{J}}, \xi)$: the magnitudes of the spin angular momenta, the magnitude of the total angular momentum, the magnitude and direction of the total angular momentum, and the mass weighted effective spin [177]

$$\xi \equiv (1+q)\boldsymbol{S}_1 \cdot \hat{\boldsymbol{L}} + (1+q^{-1})\boldsymbol{S}_2 \cdot \hat{\boldsymbol{L}}. \tag{5.7}$$

In the effective-one-body formalism, ξ corresponds to the projection of the spin angular momentum of the body at the center of mass onto the orbital angular momentum. Once the system is allowed to evolve under radiation reaction, S_1, S_2, and ξ are still conserved, while L, J and $\hat{\boldsymbol{J}}$ evolve on the radiation reaction timescale.

In the remainder of this section we use these 7 conserved quantities to geometrically solve for the 9 components of the angular momenta as a function of S in a specific coordinate system. We then complete the solution for the angular momenta as a function of time by solving a differential equation to determine $S(t)$.

Precession in a non Inertial Frame

The identification of $\hat{\boldsymbol{J}}$ as a conserved quantity suggests a coordinate frame where $\hat{\boldsymbol{z}} = \hat{\boldsymbol{J}}$ (see Fig. 5.1). We further pick the $x - y$ plane to be precessing around $\hat{\boldsymbol{z}}$ (a non inertial frame), following the precession of the orbital angular momentum which is chosen to be in the $x - z$ plane, at an angle

$$\cos\theta_L = \hat{\boldsymbol{J}} \cdot \hat{\boldsymbol{L}} = \frac{J^2 + L^2 - S^2}{2JL}, \tag{5.8}$$

from the $\hat{\mathbf{z}}$ axis. This allows us to express $\boldsymbol{L}$ as

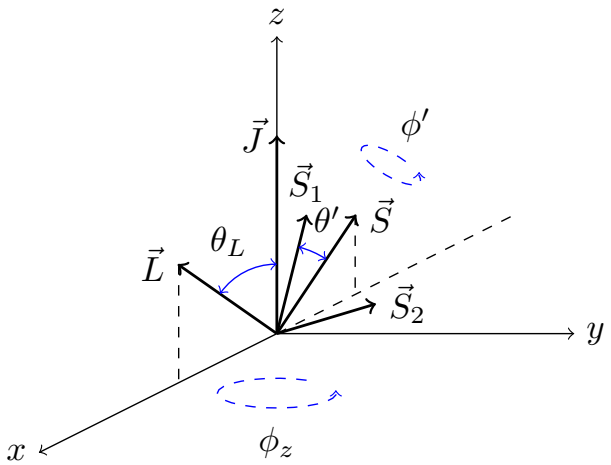

Figure 5.1: Initial configuration of the angular momenta in a non inertial frame precessing around $\hat{\boldsymbol{z}}$.

$$\boldsymbol{L}(S;\boldsymbol{\lambda}) = L\,[\sin\theta_L, 0, \cos\theta_L]. \tag{5.9}$$

The total spin angular momentum then is

$$\boldsymbol{S}(S;\boldsymbol{\lambda}) = \boldsymbol{J} - \boldsymbol{L} = [-L\sin\theta_L, 0, J - L\cos\theta_L]. \tag{5.10}$$

In another frame with $\hat{\boldsymbol{z}}' = \hat{\boldsymbol{S}}, \hat{\boldsymbol{y}}' = \hat{\boldsymbol{y}}$, and $\hat{\boldsymbol{x}}' = \hat{\boldsymbol{y}}' \times \hat{\boldsymbol{z}}'$, we define angles (θ', ϕ') (see Fig. 5.1) such that

$$\boldsymbol{S}_{\mathbf{1}}' = S_1[\sin\theta'\cos\phi', \sin\theta'\sin\phi', \cos\theta']. \tag{5.11}$$

Using the definition of ξ given in Eq. (5.7) we get

$$\cos\theta' = \hat{\boldsymbol{S}}_{\mathbf{1}} \cdot \hat{\boldsymbol{S}} = \frac{S^2 + S_1^2 - S_2^2}{2SS_1}, \tag{5.12}$$

$$\cos\phi' = \frac{(J^2 - L^2 - S^2)[S^2(1+q)^2 - (S_1^2 - S_2^2)(1-q^2)] - 4qS^2L\xi}{(1-q^2)A_1A_2A_3A_4}, \tag{5.13}$$

where

$$A_1 = \sqrt{J^2 - (L-S)^2}, \tag{5.14}$$

$$A_2 = \sqrt{(L+S)^2 - J^2}, \tag{5.15}$$

$$A_3 = \sqrt{S^2 - (S_1 - S_2)^2}, \tag{5.16}$$

$$A_4 = \sqrt{(S_1 + S_2)^2 - S^2}. \tag{5.17}$$

In the original unprimed system

$$\boldsymbol{S}_{\mathbf{1}}(S; \boldsymbol{\lambda}) = \mathbb{R}(\hat{\boldsymbol{y}}, \theta_S)\boldsymbol{S}_{\mathbf{1}}', \tag{5.18}$$

where $\mathbb{R}(\hat{\boldsymbol{y}}, \theta_S)$ is a rotation around $\hat{\boldsymbol{y}}$ by an angle θ_S and

$$\cos\theta_S = \hat{\boldsymbol{S}} \cdot \hat{\boldsymbol{J}} = \frac{J^2 + S^2 - L^2}{2JS}. \tag{5.19}$$

Once we have $\boldsymbol{S_1}$ in the original unprimed system, then

$$\boldsymbol{S_2}(S;\boldsymbol{\lambda}) = \boldsymbol{J} - \boldsymbol{L} - \boldsymbol{S_1}. \tag{5.20}$$

Equations (5.9), (5.18), and (5.20) determine the angular momenta in a non inertial frame as a function of S up to the sign of $\sin\phi'$, which we will tackle shortly.

At this point, the various orbital angular momenta have been written in a non inertial frame in terms of S using purely geometrical arguments. The evolution equation of S can be derived from Eqs. (5.1)-(5.3):

$$\left(\frac{dS^2}{dt}\right)^2 = -A^2\left(S^6 + BS^4 + CS^2 + D\right). \tag{5.21}$$

where the coefficients A, B, C, D depend only on quantities that change on the radiation reaction timescale. Their explicit form is given in Appendix B of [61]. The roots of the polynomial on the right-hand side of Eq. (5.21) have a simple interpretation. When S^2 is equal to one of the roots, its derivative is zero. Therefore, two of the roots are the maximum S_+^2 and the minimum S_-^2 of S^2. The third root S_3^2 does not correspond to any physically interesting scenario; in fact, it is negative for most systems[2].

Making explicit use of the roots of the polynomial, we can rewrite Eq. (5.21) as

$$\left(\frac{dS^2}{dt}\right)^2 = -A^2(S^2 - S_+^2)(S^2 - S_-^2)(S^2 - S_3^2). \tag{5.22}$$

[2]In the most generic case, a third order polynomial with real coefficients can have complex roots. However, we argue that this is an unphysical scenario. Unless two of the roots are real, S^2 will increase or decrease with no bound. If two roots of a third order polynomial with real coefficients are real, then the third root must be real too.

The solution to this equation is

$$S^2 = S_+^2 + (S_-^2 - S_+^2)\,\mathrm{sn}^2(\psi, \mathrm{m}) \tag{5.23}$$

where sn is a Jacobi Elliptic function (see Section 16 of [4] for a detailed introduction to the Jacobi Elliptic functions, and [51] for a physics-oriented approach), ψ is its phase, and $m \in [0, 1]$. When $m = 0$, sn reduces to a sine, while for $m = 1$ it gives a hyperbolic tangent. The period of S^2 is $2K(m)$, where $K(m)$ is the complete elliptic integral of the first kind. The phase and the parameter m are given by

$$\frac{d\psi}{dt} = \frac{A}{2}\sqrt{S_+^2 - S_3^2} \tag{5.24}$$

and

$$m = \frac{S_+^2 - S_-^2}{S_+^2 - S_3^2}\,. \tag{5.25}$$

Clearly, this solution requires that $S_+^2 \neq S_3^2$, which is almost always the case because S_+^2 and S_3^2 are defined to be the largest and smallest roots respectively. The only possible case when $S_3^2 = S_+^2$ is when $S_+^2 = S_-^2$, but then S^2 is constant in the first place and there is no precession. The phase ψ can be obtained by noticing that $\dot{\psi}$ is constant if we ignore radiation reaction, so that

$$\psi = \frac{A}{2}\sqrt{S_+^2 - S_3^2}\,t. \tag{5.26}$$

The final ingredient we need in order to have a complete expression for all angular momenta as function of time in a non-inertial frame precessing around $\hat{\boldsymbol{z}}$ is the sign

of $\sin\phi'$. Equation (5.11) implies that

$$\mathrm{sign}(\sin\phi') = \mathrm{sign}(\boldsymbol{S_1}\cdot\boldsymbol{y'}), \tag{5.27}$$

which after some algebra can be shown to be equivalent to

$$\mathrm{sign}(\sin\phi') = \mathrm{sign}\left[(\hat{\boldsymbol{L}}\times\boldsymbol{S_1})\cdot\boldsymbol{S_2}\right] = \mathrm{sign}\left(-\frac{dS^2}{dt}\right) = \mathrm{sign}\left[\mathrm{sn}(\psi,\mathrm{m})\mathrm{cn}(\psi,\mathrm{m})\right], \tag{5.28}$$

where $\mathrm{cn}(\psi,\mathrm{m})$ is another Jacobi Elliptic function and in the last equality we have used Eq. (5.23).

Precession in an Inertial Frame

All angular momenta so far have been expressed in a non inertial frame that precesses around $\hat{\boldsymbol{J}}$. An Euler rotation of $\boldsymbol{L}, \boldsymbol{S_1}$, and $\boldsymbol{S_2}$ around $\hat{\boldsymbol{z}}$ by some angle ϕ_z and substitution into the precession equations yields [121]

$$\frac{d\phi_z}{dt} \equiv \Omega_z = \frac{J}{2M}v^6\left\{1+\frac{3}{2\eta}(1-\xi v) - \frac{3(1+q)}{2qA_1^2A_2^2}(1-\xi v)\left[4(1-q)L^2(S_1^2-S_2^2)\right.\right.$$
$$\left.\left.-(1+q)(J^2-L^2-S^2)(J^2-L^2-S^2-4\eta L\xi)\right]\right\}. \tag{5.29}$$

The precession angle ϕ_z changes on the precession timescale through S and on the radiation reaction timescale through J and L. We recast it in the form

$$\frac{\dot{\phi}_z}{J} = a + \frac{c_0 + c_2\,\mathrm{sn}^2(\psi,m) + c_4\,\mathrm{sn}^4(\psi,m)}{d_0 + d_2\,\mathrm{sn}^2(\psi,m) + d_4\,\mathrm{sn}^4(\psi,m)}, \tag{5.30}$$

where a, the d_i's and the c_i's are quantities that evolve on the radiation reaction timescale only. Their explicit form is given in Appendix B of [61]. Now $\dot{\phi}_z$ can be

integrated exactly in the absence of radiation reaction to give

$$\frac{\phi_z}{J} = A_\phi \frac{\psi}{\dot{\psi}} + iB_\phi \frac{F[i\sinh^{-1}(\mathrm{sc}(\psi,m)),1-m]}{\dot{\psi}} + iC_\phi \frac{\Pi[n_c, i\sinh^{-1}(\mathrm{sc}(\psi,m)),1-m]}{\dot{\psi}} + iD_\phi \frac{\Pi[n_d, i\sinh^{-1}(\mathrm{sc}(\psi,m)),1-m]}{\dot{\psi}}, \tag{5.31}$$

where $\dot{\psi}$ is given by Eq. (5.24), F is the elliptic integral of the first kind, Π is the elliptic integral of the third kind, and sc is a Jacobi elliptic function. The quantities $A_\phi, B_\phi, C_\phi, D_\phi, n_c, n_d$ are functions of $\{a, c_i, d_i\}$, and they are constant in the absence of radiation reaction. They are given in Appendix B of [61].

This concludes the solution to the precession equations in the absence of radiation reaction in a frame where $\hat{\boldsymbol{J}} = \hat{\boldsymbol{z}}$. In summary, at some initial time:

- The orbital angular momentum $\boldsymbol{L}$ is given by Eq. (5.9), which depends on the angle θ_L given in Eq. (5.8). The latter depends on S, which varies on the precession timescale as described in Eq. (5.23);
- The spin angular momentum of the heavier body $\boldsymbol{S_1}$ is given in Eq. (5.18), which depends on the angle θ_S given in Eq. (5.19) as well as on $\boldsymbol{S_1'}$ given in Eq. (5.11) in terms of the angles (θ', ϕ') of Eqs. (5.12) and (5.13). All of these depend on S, which again is described by Eq. (5.23);
- The spin angular momentum of the lighter body $\boldsymbol{S_2}$ is given by Eq. (5.20), which depends on $\boldsymbol{L}$ and $\boldsymbol{S_1}$ described above.

The full precessional motion of these angular momenta in an inertial frame is obtained by rotating them around $\hat{\boldsymbol{z}}$ by ϕ_z, given in Eq. (5.31).

Addition of Radiation Reaction

The exact solution to the precession equations obtained in the previous section is valid only in the absence of radiation reaction. The problem of including radiation reaction admits a perturbative solution owing to its two distinct timescales: radiation reaction unfolds on a much longer timescale than precession. This natural separation of timescales allows us to treat radiation reaction as a slow perturbation of the more rapid precession, a technique formally known as multiple scale analysis [37].

In MSA, every quantity is expanded in the ratio of the two distinct timescales. In our case, we expand in the ratio of the precessional timescale $T_{\rm pr}$ to the radiation reaction timescale $T_{\rm rr}$; radiation reaction is a 1.5PN effect on top of precession. This is not the first application of MSA to the precession problem. In fact, the precession equations we started with are orbit-averaged, which would be the first term in an MSA expansion about the ratio of the fast orbital timescale to the precession timescale.

Choice of an Inertial Frame

The precession solution of the previous section was built around the assumption that $\hat{\boldsymbol{J}}$ is conserved and aligned with $\hat{\boldsymbol{z}}$. Our first task when adding radiation reaction is to check whether this remains true. If it does, then the functional form of Eqs. (5.9), (5.18), and (5.20) holds, since they were derived solely on geometrical arguments.

Radiation reaction does not strictly conserve the direction of the total angular momentum. However, it has been argued [21] that in the context of simple precession ($\boldsymbol{S_2} = 0$) the variation of $\hat{\boldsymbol{J}}$ in a precession cycle averages out. Here we show that this is approximately true for generic precession as well.

Equations (5.1)-(5.3) imply

$$\dot{\boldsymbol{J}} = \dot{L}\hat{\boldsymbol{L}}, \tag{5.32}$$

and after some algebra we can rewrite this as

$$\dot{\hat{\boldsymbol{J}}} = \frac{\dot{L}}{JL}\boldsymbol{L} - \frac{\dot{J}}{J^2}\boldsymbol{J}. \tag{5.33}$$

Averaging over ϕ_z we find

$$\left\langle \dot{J}_x \right\rangle_{\phi_z} = \left\langle \frac{\dot{L}}{J} \sin\theta_L \cos\phi_z \right\rangle_{\phi_z}, \tag{5.34}$$

$$\left\langle \dot{J}_y \right\rangle_{\phi_z} = \left\langle \frac{\dot{L}}{J} \sin\theta_L \sin\phi_z \right\rangle_{\phi_z}, \tag{5.35}$$

$$\left\langle \dot{J}_z \right\rangle_{\phi_z} = 0. \tag{5.36}$$

This averaging induces an error in $\dot{\hat{\boldsymbol{J}}}$ that is $\mathcal{O}(T_{\rm pr}/T_{\rm rr})$, or $\dot{\hat{\boldsymbol{J}}} - \langle\dot{\hat{\boldsymbol{J}}}\rangle_{\phi_z} \sim v^3$. At this order, we can treat $\dot{L}$ as a constant, since the spin couplings in Eq. (5.4) first enter at $\mathcal{O}(v^3)$. They are therefore of the same order as the averaging error, and can be neglected.

Working to this order we have

$$\begin{aligned}\left\langle \dot{J}_{x,y} \right\rangle_{\phi_z} &\sim \langle \sin\theta_L \cos\phi_z \rangle_{\phi_z} \\ &\sim \left\langle \sqrt{1 - \left(\frac{J^2 + L^2 - S^2}{2JL}\right)^2} \cos\phi_z \right\rangle_{\phi_z}.\end{aligned} \tag{5.37}$$

Since $L \sim \mathcal{O}(v^{-1})$, $J \sim \mathcal{O}(v^{-1})$, and $S \sim \mathcal{O}(v^0)$, a PN expansion yields schematically

$$\left\langle \dot{J}_{x,y} \right\rangle_{\phi_z} \sim \left\langle \cos\phi_z \right\rangle_{\phi_z} + v^2 \left\langle S^2 \cos\phi_z \right\rangle_{\phi_z} + \mathcal{O}(v^4). \tag{5.38}$$

The first term vanishes, while the second is of higher PN order and we neglect it. This situation is different from simple precession. In the latter the averaging out of $\dot{\hat{\boldsymbol{J}}}$ is exact, while here it requires a PN expansion. We therefore expect this result to become less and less accurate as the binary approaches merger.

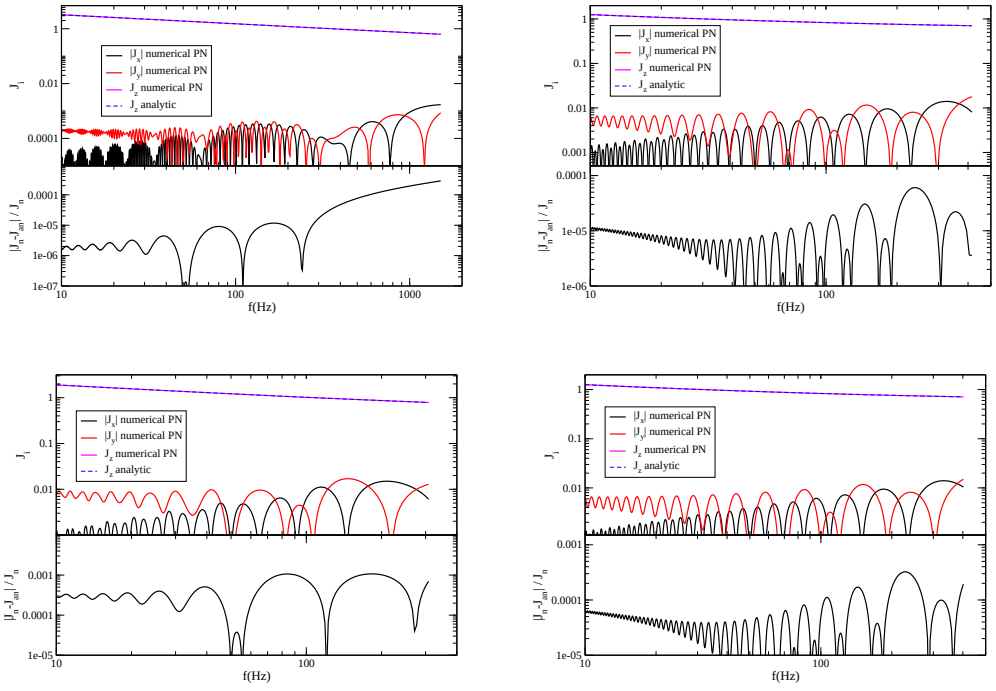

Figure 5.2: (Top Panel) Comparison between the numerical PN and the analytic components of the total angular momentum as a function of the GW frequency for the NSNS (Top Left), the BHNS (Top Right), the BHBH (Bottom Left), and the HSNSBH (Bottom Right) system of Table 5.1. (Bottom Panel) Fractional error between the magnitude of the total angular momentum obtained numerically and analytically.

	NSNS	BHNS	BHBH	HSNSBH
m_1	$1.6M_\odot$	$10M_\odot$	$10M_\odot$	$10M_\odot$
m_2	$1.4M_\odot$	$1.4M_\odot$	$5M_\odot$	$1.4M_\odot$
$\cos\theta_L$	1	1	1	1
ϕ_L	0	0	0	0
$\cos\theta_1$	0.5	0.5	0.5	0.5
ϕ_1	1.2	1.2	1.2	1.2
χ_1	0.08	0.7	0.7	0.7
$\cos\theta_2$	0.7	0.7	0.7	0.7
ϕ_2	2.5	2.5	2.5	2.5
χ_2	0.1	0.1	0.6	0.6

Table 5.1: Parameters of the systems we use for comparisons of our analytic solution to the numerical solution to the PN precession equations. All parameters are defined at 50Hz and in a frame were the orbital angular momentum is aligned with the z axis. The angles θ_L and ϕ_L are the polar angles of $\boldsymbol{L}$, while θ_A and ϕ_A are the polar angles of $\boldsymbol{S}_A$ in a frame were the z axis is aligned with the orbital angular momentum.

The above calculation implies that $\langle \dot{\hat{\boldsymbol{J}}} \rangle_{\phi_z} = 0$; radiation reaction changes the magnitude of $\boldsymbol{J}$ while leaving its direction approximately constant. The components J_x and J_y are expected to oscillate with an amplitude much smaller than J_z without exhibiting any secular growth. Figure 5.2 tests the validity of this statement. We select 4 systems with typical parameters as expected for NSNS, BHNS, BHNS, and HSNSBH binaries (see Table 5.1) and plot the components of $\boldsymbol{J}$ obtained by numerically solving Eqs. (5.1)-(5.4) as a function of the GW frequency f. In all cases J_x and J_y are at least 2 orders of magnitude smaller that J_z and oscillate around 0, with no signs of secular growth.

Based on this result we can build a solution to the precession equations including radiation reaction in the inertial frame introduced in the previous section. That is, we neglect any variation in the direction of $\hat{\boldsymbol{J}}$ and align it with $\hat{\boldsymbol{z}}$. This choice of frame automatically means that the functional form of Eqs. (5.9), (5.18), and (5.20) for the

orbital and spin angular momenta respectively is still valid, since they were derived on purely geometric arguments. On the contrary, any quantity that was derived based on Eqs. (5.1)-(5.3) needs to be revisited and recalculated by taking Eq. (5.4) into account. This involves the remaining 5 conserved quantities of precession (S_1, S_2, L, J, ξ), Eq. (5.22) for the magnitude of the total spin angular momentum, and Eq. (5.30) for the precession angle.

Constants of the Precessional Motion

In principle, the constants of the precessional motion need not remain constant when radiation reaction is invoked. The magnitudes of the two spin angular momenta S_1 and S_2, and the mass weighted effective spin ξ remain constant under radiation reaction to the PN order we work here and ignoring horizon absorption. The magnitude of the orbital angular momentum L is updated by definition through $L = \eta/v$. The magnitude of the total angular momentum J depends on L and also changes under radiation reaction. The evolution equation for J averaged over one period of $S(t)$ is [121]

$$\left\langle \frac{dJ}{dL} \right\rangle_{\rm pr} = \frac{J^2 + L^2 - \langle S^2 \rangle_{\rm pr}}{2JL}. \tag{5.39}$$

This can be integrated exactly to yield

$$J^2 = L^2 + \frac{2c_1}{\eta} L - L \int \frac{\langle S^2 \rangle_{\rm pr}}{L^2} dL, \tag{5.40}$$

where here and in what follows J is approximated by its precession average, and c_1 is an integration constant. As we will show below, $\langle S^2 \rangle_{\rm pr}$ is constant when ignoring

high-order PN effects, and the integral of Eq. (5.40) can be calculated to give

$$J^2 = L^2 + \frac{2c_1}{v} + \left\langle S^2 \right\rangle_{\rm pr} + \mathcal{O}(v). \tag{5.41}$$

The quantity $\left\langle S^2 \right\rangle_{\rm pr}$ can be computed from Eq. (5.23):

$$S^2_{\rm av} \equiv \left\langle S^2 \right\rangle_{\rm pr} = \frac{1}{m}\left[(m-1)S^2_+ + S^2_- + \frac{E(m)}{K(m)}\left(S^2_+ - S^2_-\right)\right], \tag{5.42}$$

where $K(m)$ and $E(m)$ are the complete elliptic integrals of the first and second kind respectively. PN expanding S^2_+ and S^2_- around their initial value we find

$$S^2_\pm = S^2_{\pm,0} + \mathcal{O}\left(v\right), \qquad S^2_3 = \mathcal{O}\left(v^{-2}\right), \tag{5.43}$$

which together with Eq. (5.25) yields

$$m = \mathcal{O}\left(v^2\right), \tag{5.44}$$

and

$$S^2_{\rm av} = \frac{1}{2}\left(S^2_{+,0} + S^2_{-,0}\right) + \mathcal{O}\left(v^2\right). \tag{5.45}$$

In the above expressions $S^2_{\pm,0}$ are the roots computed from the initial conditions.

Combining the result for J obtained here and keeping $\hat{\boldsymbol{J}}$ aligned with $\hat{\boldsymbol{z}}$, our analytic approximation for the total angular momentum is

$$\boldsymbol{J} = [0, 0, J]. \tag{5.46}$$

To verify that this approximate $\boldsymbol{J}$ stays close to the numerical PN solution we plot it in Fig. 5.2 as a function of the GW frequency for our three study systems. The analytic J_x and J_y are identically zero, so we omit them. The bottom panel shows the fractional error in the magnitude of the total angular momentum when approximated by Eq. (5.41). The maximum discrepancy in the magnitude J is of $\mathcal{O}\left(10^{-2}\right)$ in those particular examples indicating both that Eq. (5.41) is accurate and that setting J_x and J_y equal to zero is justified.

Magnitude of the Total Spin Angular Momentum

Once radiation reaction is included, Eq. (5.22) for the magnitude of the total spin angular momentum needs to be solved with MSA. We first explicitly separate the timescales by writing $S^2(t) = S^2(t_{\rm pr}, t_{\rm rr})$, where $t_{\rm pr}$ denotes variation on the precession timescale, while $t_{\rm rr} = \epsilon t_{\rm pr}$ denotes variations on the radiation reaction timescale, with ϵ a bookkeeping parameter.

Expanding S^2 as

$$S^2(t_{\rm pr}, t_{\rm rr}) = \sum_{n\geq 0} \epsilon^n S_n^2(t_{\rm pr}, t_{\rm rr}). \tag{5.47}$$

and substituting this expression into Eq. (5.22), at leading order in ϵ, we recover Eq. (5.22) for $S_0^2(t_{\rm pr}, t_{\rm rr})$ with the time derivative taken on the precession timescale $t_{\rm pr}$:

$$\left(\frac{\partial S_0^2}{\partial t_{\rm pr}}\right)^2 = -A^2(t_{\rm rr})[S_0^2(t_{\rm pr}, t_{\rm rr}) - S_+^2(t_{\rm rr})][S_0^2(t_{\rm pr}, t_{\rm rr}) - S_-^2(t_{\rm rr})][S_0^2(t_{\rm pr}, t_{\rm rr}) - S_3^2(t_{\rm rr})].$$

The solution to this differential equation is similar to Eq. (5.23), except that quantities that were previously constant are now promoted to functions of $t_{\rm rr}$:

$$S_0^2 = S_+^2(t_{\rm rr}) + \left[S_-^2(t_{\rm rr}) - S_+^2(t_{\rm rr})\right] \mathrm{sn}[\psi(t_{\rm pr}, t_{\rm rr}), m(t_{\rm rr})], \tag{5.48}$$

where $S_+^2(t_{\rm rr})$, $S_-^2(t_{\rm rr})$, and $m(t_{\rm rr})$ now depend on time through $L(t_{\rm rr})$ and $J(t_{\rm rr})$.

The angle $\psi(t_{\rm pr}, t_{\rm rr})$ satisfies

$$\frac{d\psi}{dt} = \frac{A(t_{\rm rr})}{2}\sqrt{S_+^2(t_{\rm rr}) - S_3^2(t_{\rm rr})}. \tag{5.49}$$

where we keep terms of $\mathcal{O}(\epsilon)$ by taking the derivative with respect to t rather than $t_{\rm pr}$. We can integrate this equation using a PN integration, i.e. expanding it in powers of v and integrating term by term. The result is

$$\psi = \psi_0 - \frac{3g_0}{4}\delta m\, v^{-3}\left(1 + \psi_1 v + \psi_2 v^2\right), \tag{5.50}$$

where ψ_0 is an integration constant, and the constants ψ_1, ψ_2 are given in Appendix C of [61]. We find that expanding Eq. (5.50) to relative 1PN order suffices.

We test this solution for S in Fig. 5.3 by plotting the numerical PN, analytic, and hybrid magnitude of the total spin angular momentum S as a function of the GW frequency for the 4 systems we study. The hybrid S is obtained through Eq. (5.48) but with a numerical solution to Eq. (5.49). For all systems, the amplitude of S shows excellent agreement with the numerical PN results, which is controlled by the roots S_+^2 and S_-^2. For the NSNS and BHBH systems, the analytic phase ψ also shows very good agreement with the numerical PN result, although the dephasing for the BHNS and HSBHNS systems is about 2 cycles. However, both systems are dominated by the spin of the BH, making the motion close to that of simple precession; the variation

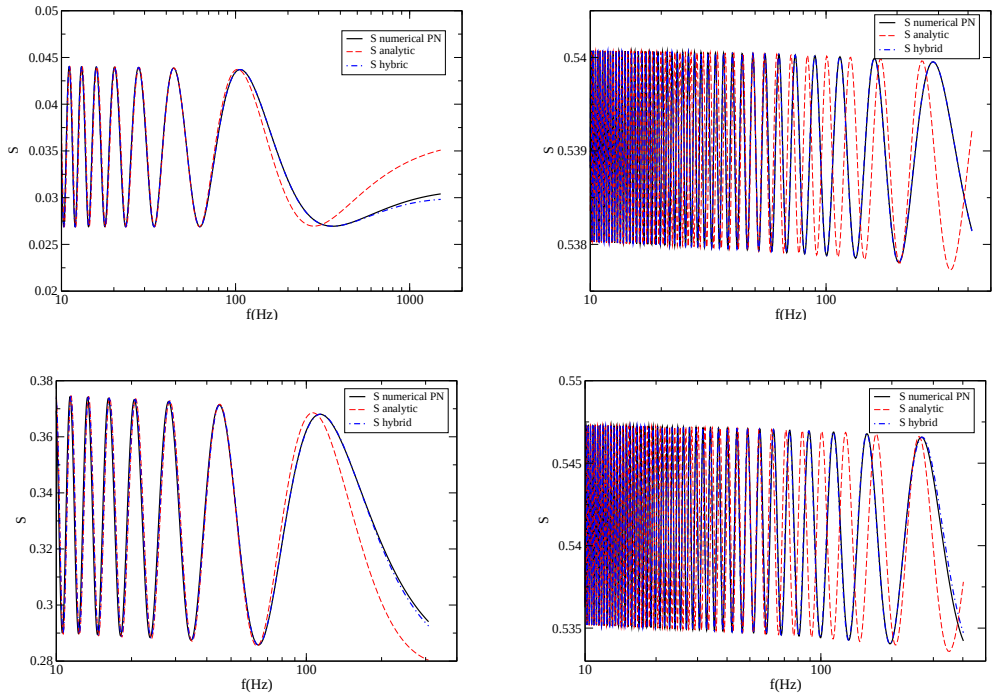

Figure 5.3: Comparison between the numerical PN, the analytic, and the hybrid magnitude of the total spin angular momentum as a function of the GW frequency for the NSNS (Top Left), the BHNS (Top Right), the BHBH (Bottom Left), and the HSNSBH (Bottom Right) system of Table 5.1.

in S is very small as demonstrated by the scale of the y axis of the right panels of Fig. 5.3 and this dephasing should not affect the emitted waveform considerably.

On the other hand, the phase of the hybrid S is always in excellent agreement with the numerical solution, indicating that if we do indeed need an improved solution in the future[3] we can obtain it by carrying out the expansion of Eq. (5.50) to higher order.

Precession Angle

The final quantity that needs to be recalculated to account for radiation reaction is the precession angle. Its derivative, given in Eq. (5.30), depends both on the precession and the radiation reaction timescale, so it requires a MSA treatment.

We write

$$\frac{d\phi_z}{dt} = \Omega_z[S(t), L(t), J(t)] = \Omega_z[S(t_{\rm pr}, t_{\rm rr}), L(t_{\rm rr}), J(t_{\rm rr})], \tag{5.51}$$

and expand the precession angle as

$$\phi_z(t_{\rm pr}, t_{\rm rr}) = \epsilon^{-1}\phi_{z,-1}(t_{\rm pr}, t_{\rm rr}) + \phi_{z,0}(t_{\rm pr}, t_{\rm rr}) + \mathcal{O}(\epsilon)\,. \tag{5.52}$$

The reason ϕ_z includes a term of $\mathcal{O}(\epsilon^{-1})$ is because the binary precesses even in the absence of radiation reaction.

Solving Eq. (5.51) order by order in ϵ, we find to $\mathcal{O}(\epsilon^{-1})$

$$\frac{1}{\epsilon}\frac{\partial\phi_{z,-1}}{\partial t_{\rm pr}} = 0\,, \tag{5.53}$$

[3]For example, if and when LIGO's sensitivity increases, so will its requirement for more accurate waveforms.

which means $\phi_{z,-1} = \phi_{z,-1}(t_{\rm rr})$. To next order, we find

$$\frac{\partial \phi_{z,-1}}{\partial t_{\rm rr}} + \frac{\partial \phi_{z,0}}{\partial t_{\rm pr}} = \Omega_z(t_{\rm pr}, t_{\rm rr})\,, \tag{5.54}$$

and averaging over $t_{\rm pr}$ we find

$$\frac{d\phi_{z,-1}}{dt_{\rm rr}} = \langle \Omega_z \rangle_{\rm pr}(t_{\rm rr}), \tag{5.55}$$

where we set $\langle \partial\phi_{z,0}/\partial t_{\rm pr}\rangle_{\rm pr} = 0$ to cancel secular terms. Equation (5.55) can be solved with a PN integration. Going back to Eq. (5.54) we get

$$\frac{\partial \phi_{z,0}}{\partial t_{\rm pr}} = \Omega_z(t_{\rm pr}, t_{\rm rr}) - \langle \Omega_z \rangle_{\rm pr}(t_{\rm rr}). \tag{5.56}$$

Integrating the first term on the right hand side of Eq. (5.56) we recover Eq. (5.31) for ϕ_z in the absence of radiation reaction. Integrating the second term is straightforward.

The full solution for ϕ_z is then

$$\phi_z = \phi_{z,-1} + \phi_{z,0} + \mathcal{O}(\epsilon), \tag{5.57}$$

where

$$\phi_{z,-1} = \int \langle \Omega_z \rangle_{\rm pr}(t_{\rm rr})\, dt_{\rm rr}, \tag{5.58}$$

$$\phi_{z,0} = \int \Omega_z(t_{\rm pr}, t_{\rm rr})\, dt_{\rm pr} - \int \langle \Omega_z \rangle_{\rm pr}(t_{\rm rr})\, dt_{\rm pr}. \tag{5.59}$$

The meaning of each term in the MSA expansion is clear. The first term $\phi_{z,-1}$ is averaged over the fast (relative to radiation reaction) precession timescale, and then

integrated over radiation reaction. The next term $\phi_{z,0}$ is a first order correction to this precession averaging.

Leading Order MSA The leading order MSA term is defined in Eq. (5.55) which to first order in ϵ is equivalent to

$$\left\langle \frac{d\phi_z}{dt} \right\rangle_{\rm pr} = \langle \Omega_z \rangle_{\rm pr} \,. \tag{5.60}$$

The average of Ω_z can be obtained by taking the difference between Eq. (5.31) evaluated at $\psi = 0$ and at $\psi = 2K(m)$, where recall that $K(m)$ is the complete elliptic integral of the first kind. However, for reasons explained in Appendix A, we prefer to use Eq. (5.30) and find an alternative way of calculating $\left\langle \dot{\phi}_z \right\rangle_{\rm pr}$. We write

$$\begin{gathered}
\frac{\dot{\phi}_z}{J} - a \equiv \dot{\phi}_z^{\rm red.} = \frac{c_0 + c_2\,{\rm sn}^2(\psi,m) + c_4\,{\rm sn}^4(\psi,m)}{d_0 + d_2\,{\rm sn}^2(\psi,m) + d_4\,{\rm sn}^4(\psi,m)} \Rightarrow \\
[d_0 + d_2\,{\rm sn}^2(\psi,m) + d_4\,{\rm sn}^4(\psi,m)]\dot{\phi}_z^{\rm red.} = c_0 + c_2\,{\rm sn}^2(\psi,m) + c_4\,{\rm sn}^4(\psi,m) \Rightarrow \\
d_0 \left\langle \dot{\phi}_z^{\rm red.} \right\rangle_{\rm pr} + d_2 \left\langle {\rm sn}^2(\psi,m)\dot{\phi}_z^{\rm red.} \right\rangle_{\rm pr} + d_4 \left\langle {\rm sn}^4(\psi,m)\dot{\phi}_z^{\rm red.} \right\rangle_{\rm pr} \\
= c_0 + c_2 \left\langle {\rm sn}^2(\psi,m) \right\rangle_{\rm pr} + c_4 \left\langle {\rm sn}^2(\psi,m) \right\rangle_{\rm pr}
\end{gathered}$$

where on the third line we average over precession.

Unfortunately, no closed form expressions exist for $\left\langle {\rm sn}^2(\psi,m) \right\rangle_{\rm pr}$ and $\left\langle {\rm sn}^4(\psi,m) \right\rangle_{\rm pr}$ for arbitrary m. We can, however, calculate these averages as an expansion in $m \ll 1$ since, as already discussed, $m \sim \mathcal{O}(v^2)$. We could in principle retain high order in m terms in this expansion, but in practice we find that working to leading order in m suffices. Expanding the above expression to leading order in $m \ll 1$, we find

$$d_0 \left\langle \dot{\phi}_z^{\rm red.} \right\rangle_{\rm pr} + d_2 \left\langle \dot{\phi}_z^{\rm red.} \sin^2\psi \right\rangle_{\rm pr} + d_4 \left\langle \dot{\phi}_z^{\rm red.} \sin^4\psi \right\rangle_{\rm pr} = c_0 + \frac{1}{2}c_2 + \frac{3}{8}c_4 \Rightarrow$$

$$d_0 \left\langle \dot{\phi}_z^{\text{red.}} \right\rangle_{\text{pr}} + d_2 D_2 \left\langle \dot{\phi}_z^{\text{red.}} \right\rangle_{\text{pr}} + d_4 D_4 \left\langle \dot{\phi}_z^{\text{red.}} \right\rangle_{\text{pr}} = c_0 + \frac{1}{2} c_2 + \frac{3}{8} c_4 \Rightarrow$$
$$\left\langle \dot{\phi}_z \right\rangle_{\text{pr}} = J\Big(a + \frac{c_0 + \frac{1}{2} c_2 + \frac{3}{8} c_4}{d_0 + d_2 D_2 + d_4 D_4}\Big), \tag{5.61}$$

where we have defined

$$D_2 \equiv \frac{\left\langle \dot{\phi}_z^{\text{red.}} \sin^2 \psi \right\rangle_{\text{pr}}}{\left\langle \dot{\phi}_z^{\text{red.}} \right\rangle_{\text{pr}}} = \frac{\left\langle \frac{c_0 + c_2 \sin^2 \psi + c_4 \sin^4 \psi}{d_0 + d_2 \sin^2 \psi + d_4 \sin^4 \psi} \sin^2 \psi \right\rangle_{\text{pr}}}{\left\langle \frac{c_0 + c_2 \sin^2 \psi + c_4 \sin^4 \psi}{d_0 + d_2 \sin^2 \psi + d_4 \sin^4 \psi} \right\rangle_{\text{pr}}}, \tag{5.62}$$

$$D_4 \equiv \frac{\left\langle \dot{\phi}_z^{\text{red.}} \sin^4 \psi \right\rangle_{\text{pr}}}{\left\langle \dot{\phi}_z^{\text{red.}} \right\rangle_{\text{pr}}} = \frac{\left\langle \frac{c_0 + c_2 \sin^2 \psi + c_4 \sin^4 \psi}{d_0 + d_2 \sin^2 \psi + d_4 \sin^4 \psi} \sin^4 \psi \right\rangle_{\text{pr}}}{\left\langle \frac{c_0 + c_2 \sin^2 \psi + c_4 \sin^4 \psi}{d_0 + d_2 \sin^2 \psi + d_4 \sin^4 \psi} \right\rangle_{\text{pr}}}. \tag{5.63}$$

The quantities D_2 and D_4 are functions of v and can be calculated exactly. For reasons explained in Appendix A we do not wish to use these full expressions, but rather we keep the quantities D_2 and D_4 constant and set them equal to their leading PN order expressions.

We can now integrate the right hand side of Eq. (5.61) by first PN expanding it. However, we find it more convenient to factor J out of $\langle \Omega_z \rangle_{\text{pr}}$ and PN expand the remaining terms. We do so to avoid artificial divergences in the small mass ratio limit arising from expanding around essentially η/v; see Appendix A. We, then, have to perform an integral of the form

$$\phi_{z,-1} = \int \frac{J}{\xi^3} \sum_{n=0}^{5} \langle \Omega_z \rangle^{(n)} \, v^n d\xi, \tag{5.64}$$

where the coefficients $\langle \Omega_z \rangle^{(n)}$ are given in Appendix D of [61]. This integral can be directly calculated to give

$$\phi_{z,-1} = \sum_{n=0}^{5} \langle \Omega_z \rangle^{(n)} \phi_z^{(n)} + \phi_{z,-1}^0, \tag{5.65}$$

where $\phi_z^{(n)}$ are functions given in Appendix D of [61] and $\phi^0_{z,-1}$ is an integration constant.

Correction to MSA The first-order correction to MSA is given in Eq. (5.59). The solution to the first integral is Eq. (5.31) where we set $m = 0$. The second integral is trivial since $\langle \Omega_z \rangle_{\rm pr}$ does not depend on the precession timescale $t_{\rm pr}$, and the result is $\langle \Omega_z \rangle_{\rm pr} t_{\rm pr}$. In that expression, we choose for convenience to substitute $t_{\rm pr} = \psi/\dot{\psi}$.

Collecting all the elements together, the correction to the precession phase is given by

$$\phi_{z,0} = \frac{C_\phi}{\dot{\psi}} \frac{\sqrt{n_c}}{n_c - 1} \arctan\left[\frac{(1-\sqrt{n_c})\tan\psi}{1+\sqrt{n_c}\tan^2\psi}\right] + \frac{D_\phi}{\dot{\psi}} \frac{\sqrt{n_d}}{n_d - 1} \arctan\left[\frac{(1-\sqrt{n_d})\tan\psi}{1+\sqrt{n_d}\tan^2\psi}\right], \tag{5.66}$$

where $\dot{\psi}$ is given in Eq. (5.24), ψ is given in Eq. (5.50) and C_ϕ, D_ϕ, n_c and n_d are functions of v given in Appendix B of [61].

Comparisons In Fig. 5.4 we plot the numerical PN and analytic solutions for ϕ_z with and without the MSA corrections. The small oscillations of the numerical PN phase are reproduced by the analytic phase with MSA corrections. These oscillations are more pronounced for the NSNS and BHBH systems where both spins contribute significantly to the dynamics. The bottom panel shows the error in the precession phase with and without MSA corrections.

Building the Waveform

Using the solution for the angular momenta described above, we can calculate an analytic time-domain waveform for generic precessing binaries. The gravitational wave signal emitted by a precessing binary system as observed in an interferometric

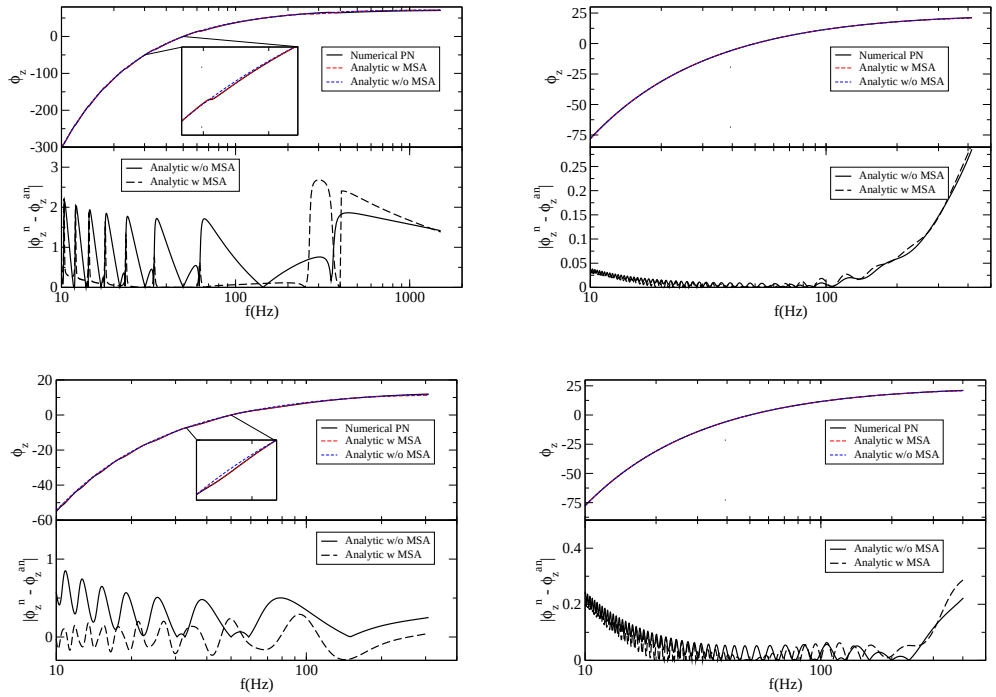

Figure 5.4: (Top Panel) Comparison between the numerical PN and the analytic precession phase as a function of the GW frequency for the NSNS (Top Left), the BHNS (Top Right), the BHBH (Bottom Left), and the HSNSBH (Bottom Right) system of Table 5.1. (Bottom Panel) Error in ϕ_z with and without the MSA corrections.

detector is [21, 24, 43, 49, 144]:

$$h(t) = F_+ h_+ + F_\times h_\times, \tag{5.67}$$

where

$$F_+ = \frac{1}{2}\left(1 + \cos^2\theta_N\right)\cos 2\phi_N \cos 2\psi - \cos\theta_N \sin 2\phi_N \sin 2\psi, \tag{5.68}$$

$$F_\times = \frac{1}{2}\left(1 + \cos^2\theta_N\right)\cos 2\phi_N \sin 2\psi + \cos\theta_N \sin 2\phi_N \cos 2\psi, \tag{5.69}$$

are the antenna pattern functions, $h_{+,\times}$ are the GW polarization states, (θ_N, ϕ_N) are the polar angles of $\hat{\boldsymbol{N}}$ in a frame tied to the arms of the detector with $\hat{\boldsymbol{z}}'$ the normal to the detector plane, and ψ is given by

$$\psi = \arctan\left[\frac{\left(P_N \hat{\boldsymbol{J}}\right)\cdot \hat{\boldsymbol{z}}'}{\left(\hat{\boldsymbol{N}} \times \hat{\boldsymbol{J}}\right)\cdot \hat{\boldsymbol{z}}'}\right], \tag{5.70}$$

where P_N acts as a projection along $\hat{\boldsymbol{N}}$.

The polarization states can be decomposed into a spin-weighted spherical harmonic basis [24, 49, 144]

$$h_+ - i h_\times = \sum_{l\geq 2}\sum_{m=-l}^{l} H^{lm}(\theta_s, \phi_s) e^{-im\Phi}, \tag{5.71}$$

where

$$\Phi = \phi_{\rm orb} - 3v^3(2 - \eta v^2)\ln v\,, \tag{5.72}$$

and (θ_s, ϕ_s) are the spherical angles of $\hat{\boldsymbol{N}}$ in a frame where $\hat{\boldsymbol{J}}$ is along the z-axis, $\phi_{\rm orb}$ is the orbital phase, and

$$H^{lm} = h^{lm} \sum_{m'=-l}^{l} D^{l}_{m',m}(\phi_z, \theta_L, \zeta)\, {}_{-2}Y_{lm'}(\theta_s, \phi_s), \tag{5.73}$$

where the amplitudes $h^{lm}(\iota)$ can be found in [43], $D^{l}_{m,m'}$ are the Wigner D-matrices, ${}_sY_{lm}$ are the spin-weighted spherical harmonics, the angles θ_L and ϕ_z are the spherical angles of $\hat{\boldsymbol{L}}$ in the same frame as θ_s and ϕ_s are defined, and ζ satisfies $\dot{\zeta} = \dot{\phi}_z \cos\theta_L$. In order to solve for ζ we can employ the same techniques as for ϕ_z, namely MSA. An explicit expression for ζ is given in Appendix F of [61].

The above prescribe a waveform $h(t)$ in the time domain. To compute its Fourier transform, we use the shifted uniform asymptotics method devised in [125] and write

$$\tilde{h}(f) = \sqrt{2\pi} \sum_{m\geq 1} T_m e^{2\pi i f t_m - m\Phi - \pi/4} \sum_{l\geq 2} \sum_{k=-k_{\max}}^{k_{\max}} \frac{a_{k,k_{\max}}}{2-\delta_{k,0}} \mathcal{H}_{lm}(t_m + kT_m), \tag{5.74}$$

where t_m and T_m are defined by

$$2\pi f = m\dot{\Phi}(t_m), \tag{5.75}$$

$$T_m = \frac{1}{\sqrt{m\ddot{\Phi}(t_m)}}, \tag{5.76}$$

$$\begin{aligned}\mathcal{H}_{lm} = {} & \frac{1}{2}(F_+ + iF_\times) \sum_{m'=-l}^{l} h^{lm} D^{l}_{m',m}(\phi_z, \theta_L, \zeta)\, {}_{-2}Y_{lm'}(\theta_s, \phi_s) \\ & + \frac{1}{2}(F_+ - iF_\times) \sum_{m'=-l}^{l} h^{l,-m} D^{l}_{m',-m}(\phi_z, \theta_L, \zeta)\, {}_{-2}Y_{lm'}(\theta_s, \phi_s),\end{aligned} \tag{5.77}$$

and the constants $a_{k,k_{\rm max}}$ satisfy the linear system

$$\frac{(-i)^p}{2^p p!} = \sum_{k=0}^{k_{\rm max}} a_{k,k_{\rm max}} \frac{k^{2p}}{(2p)!}, \tag{5.78}$$

for $p \in \{0, \ldots, k_{\rm max}\}$. In this expression, Eq. (5.75) expresses the stationary time t_m as a function of the frequency f. For a static detector, $\mathcal{H}_{lm}$ depends on time (and thus on frequency) only through ϕ_z, θ_L, and ζ.

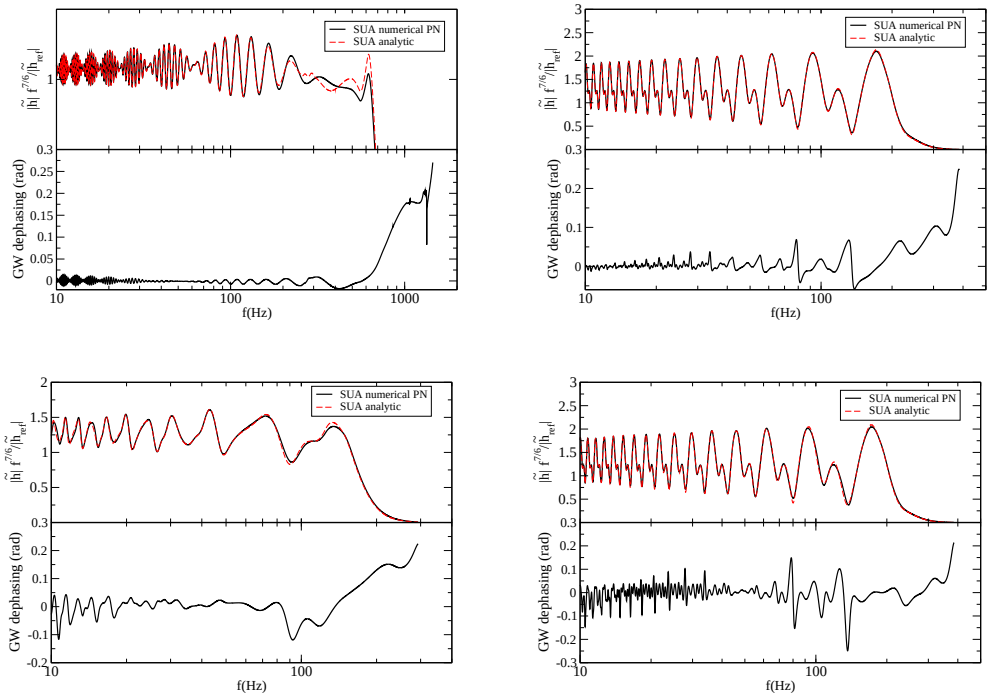

Figure 5.5: (Top Panel) Amplitude of the GW including only the dominant ($\ell = 2, m = 2$) harmonic as a function of the GW frequency for the numerical PN and analytic SUA waveforms for the NSNS (Top Left), the BHNS (Top Right), the BHBH (Bottom Left), and the HSNSBH (Bottom Right) system of Table 5.1. The reference amplitude $|\tilde{h}_{\rm ref}|$ is the numerical PN SUA amplitude at 50Hz. (Bottom Panel) GW dephasing between the numerical PN and analytic SUA waveforms.

Figure 5.5 compares the frequency domain GWs for the 4 systems of Table 5.1 using only the leading ($\ell = 2, m = 2$) harmonics of Eq. (5.74). The two waveforms

are computed with the numerical solution to the PN precession equations and with the analytic solution described previously. Both waveforms are Fourier-transformed with SUA, allowing us to assess the effect of our new analytic solution to the GW amplitude and phase. The agreement between the wave amplitudes is excellent over a wide range of frequencies, while the dephasing between the two waveforms never exceeds 0.3rad, even for our BHBH system. This figure serves as a first indication of the accuracy of our model to accurately capture generic precessing features in GWs.

Waveform Comparison

In order to have a more complete picture of our waveform's ability to model generic systems, we carry out a Monte Carlo study randomizing over the 15 parameters describing a quasicircular compact binary waveform. For the randomization, we draw the components' masses form a flat distribution in log space between $[1, 2.5]M_{\odot}$ for NSs and $[2.5, 20]M_{\odot}$ for BHs, while the components' spin magnitudes are uniformly distributed in $[0, 0.1]$ for NSs and $[0, 1]$ for BHs. We selected seemingly low black hole masses in order to select systems for which the inspiral part is the most important. Indeed, those are the ones for which the accurate modelling of the precession effects are the most challenging, due to the increased number of precession cycles that low masses entail. All directions (spin, sky location, orbital angular momentum) are drawn uniformly on a unit sphere. The phase of coalescence is assumed to be uniform in $[0, 2\pi]$, while the time of coalescence and the distance are fixed at 10^5 seconds and 100Mpc respectively.

The large number of systems simulated can only be analyzed through some appropriate and efficient statistic; we use the *faithfulness* (or *match*) defined as

$$F \equiv \max_{t_c, \phi_c} \frac{(h_1 \mid h_2)}{\sqrt{(h_1 \mid h_1)\,(h_2 \mid h_2)}}. \tag{5.79}$$

The faithfulness is calculated between two waveforms h_1 and h_2 with the same physical parameters, but maximized over any un-physical parameters: the time t_c and phase ϕ_c of coalescence. As such, it is a good estimator of a model's suitability for parameter estimation. The faithfulness always falls between -1 and 1, with the latter indicating perfect agreement between the waveforms.

Unlike *fitting factors*[4], selecting a value for the faithfulness that is 'good enough' is not straightforward. The nominal fitting factor threshold of 0.965 corresponds to a 10% drop in detection rates. On the other hand, a faithfulness threshold should be translatable to a requirement about parameter estimation accuracy: the systematic mismodeling error should be smaller than the statistical measurement error. The latter depends on the signal-to-noise ratio (SNR) of the signal, while the former does not, meaning that any faithfulness threshold should take the strength of the signal into account. In Appendix F we calculate the faithfulness threshold as a function of the SNR and find that for an SNR of 10(25)[50], a faithfulness of 0.96(0.9936)[0.9984] suffices for accurate parameter estimation. Led by the SNR of the first detected GW, we set our faithfulness threshold to 0.994.

In our study h_1 is a waveform calculated by numerically solving the precession equations, while h_2 uses our new analytic solution. Both waveforms are Fourier-transformed with the SUA method, justified by [125] where it was shown that SUA induces a negligible loss of faithfulness compared to a discrete Fourier transform. The use of SUA in both waveforms allows us to isolate the effect of our new solution: any mismatch is solely caused by the solution to the precession equations described in this paper.

[4]A fitting factor is the faithfulness maximized over all model parameters, quantifying a model's suitability for detection.

The inner product in Eq. (5.79) is defined in the usual way

$$(h_1 \mid h_2) \equiv 4\Re \int_{f_{\rm min}}^{f_{\rm max}} \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)}\, df\,, \tag{5.80}$$

where $f_{\rm min} = 10$Hz is aLIGO's lower frequency cutoff, $f_{\rm max}$ is the frequency that corresponds to an orbital separation of $6M$, and $S_n(f)$ is aLIGO's design zero-detuning, high power noise spectral density [201].

Figure 5.6 shows the distributions of $1 - F$ for 4 sets, each containing $10,000$ systems. The first 3 sets contain systems with masses and spins corresponding to NSNS, BHNS, and BHBH systems respectively. The fourth set contains an additional, less astrophysically motivated but useful to test our model in the most challenging setting, type of system where both masses were drawn from a log-flat distribution ranging from $1M_\odot$ to $20M_\odot$, and both spin magnitudes uniformly distributed in [0, 1]. We study 2 different types of waveforms: *full* waveforms (FWF) contain all the known harmonics in Eq. (5.74), while *restricted* waveforms (RWF) contain only the dominant $(\ell = 2, m = 2)$ harmonic.

The agreement between our analytical waveform and the numerical PN one is excellent for a wide range of parameters. In the NSNS case we find that only 0.06%(0.3%) of the systems have a faithfulness below the 0.965(0.994) for both waveforms, while for BHNS systems, this number is 0.33%(1.6%) for both waveforms. The percentage of systems below the nominal faithfulness threshold is increased to 1.85%(10.4%) (RWF) and 1.85%(10.7%) (FWF) in the case of BHBHs and 1.14%(9.2%) (RWF) and 1.26%(9.9%) (FWF) for the $4^{\rm th}$ generic set. This increase is not unexpected, since, precessional feature are more pronounced, and hence more difficult to model, when the spins are large and the masses different.

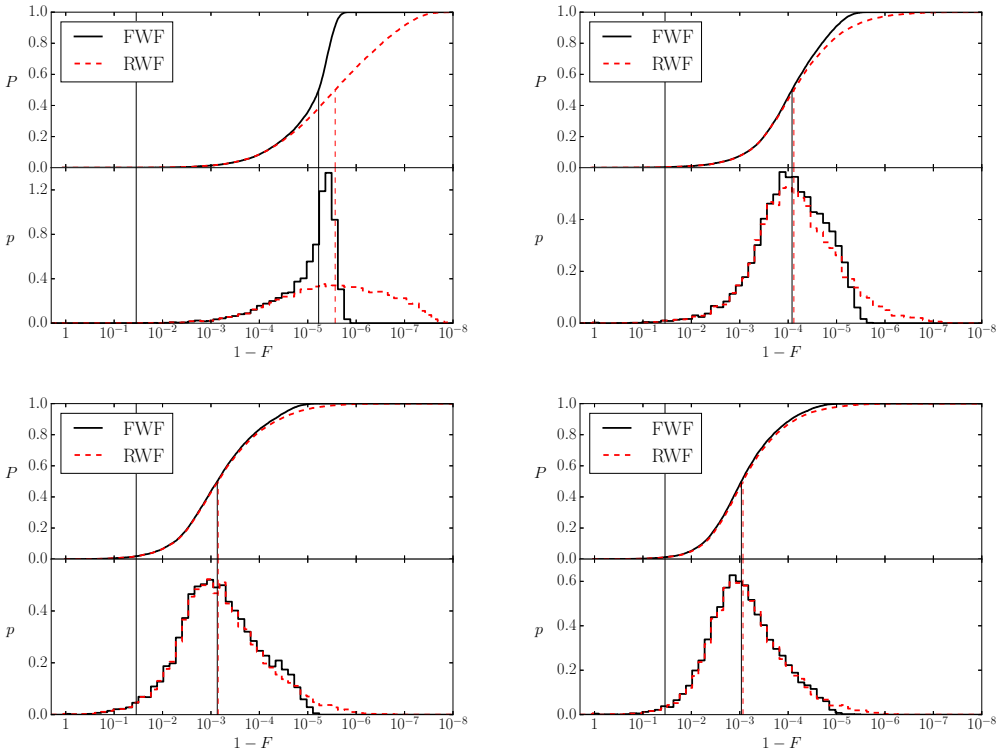

Figure 5.6: Distribution of $1 - F$ for NSNS (Top Left), BHNS (Top Right), BHBH (Bottom Left) and the 4^{th} generic set containing all masses and spins (Bottom Right) for waveforms with full harmonic content (solid black line) and waveforms restricted to the leading ($\ell = 2, m = 2$) mode (dashed red line). Top panels show cumulative distribution functions, and bottom panels give the correspondig probability distribution function.

The majority systems with faithfulnesses below 0.96 are explained by the effect described in Appendix A. For them the orbital angular momentum becomes approximately (anti)aligned with the total spin angular momentum at some point in the evolution of the systems. In this case the PN expansion of Eq. (5.61) becomes ill-defined. Our specific choice for the values of D_2 and D_4 in Eqs. (5.62) and (5.63) to some extent ameliorates this problem, yet it does not fully solve it. We have explored many choices for D_2 and D_4, some leading up to 8% of systems with faithfulnesses below 0.96 in the BHBH case. The particular values for D_2 and D_4 we employ in our model [Eqs. (A.3) and (A.4)] yield the best results among all the values we tested. The remaining low F systems for which $\boldsymbol{S}$ and $\boldsymbol{L}$ do not become (anti)aligned all are approximately edge-on ($\boldsymbol{L} \cdot \boldsymbol{N} \sim 0$), a configuration which maximizes precession and causes small inaccuracies in the model to become noticeable.

Besides systems with very low faithfulnesses, we also encounter a number of systems with overlaps between 0.96 and 0.994, which can be modeled accurately only for low SNR signals. The unfaithfulness of these systems can be attributed to the various approximations we have used in our model construction: PN inaccuracies in ψ (Eq. (5.50)), PN or MSA inaccuracies in ϕ_z, the total angular momentum e.t.c. The analytic nature of our waveform model allows us to track all these approximations and study how they influence model accuracy, while having a clear-cut way to improve them if a more faithful waveform model is required in the future. This is, perhaps, the most attractive feature of having analytic control over complicated processes like spin-precession.

Conclusions

We have constructed the first closed-form fully analytic GW template in the frequency domain that can accurately model quasicircular systems of generic masses,

spin magnitudes, and spin orientations in the inspiral phase. We expand the exact solution to the precession equations in the absence of radiation reaction derived by Kesden et al. [121] to include radiation reaction using elements from *multiple scale analysis*. This allows us to derive the first closed-form time-domain GW model valid for generic inspirals. We then use the method of *shifted uniform asymptotics* to transform this waveform from the time domain to the frequency domain.

The resultant waveform is ideal for extracting parameters from generically precessing quasicircular inspirals as demonstrated by a Monte Carlo study of $40,000$ system; only $0.8\%(5.4\%)$ of them had a faithfulness with a numerical PN waveform that solves the precession equations numerically below $0.965(0.006)$. The remaining inaccuracies of our model can be mapped back to specific assumptions we made while solving the spin-precession equations including radiation reaction. Analytical understanding of all these assumptions and the elements that enter our waveform construction enable us to improve the accuracy of our model if deemed necessary when more sensitive GW detector networks become available.

DETECTION AND PARAMETER ESTIMATION OF GWS FROM SPIN-PRECESSING BINARIES

In this chapter, we study the performance of various analytical frequency-domain templates for detection and parameter estimation of GWs from spin-precessing, quasicircular, compact binary inspirals. We begin by assessing the extent to which nonspinning, spin-aligned, and the new (analytical, frequency-domain, small-spin) double-precessing frequency-domain templates constructed in Chapter 4 can be used to detect signals from such systems. For effective, dimensionless spin values above 0.2, the use of nonspinning or spin-aligned templates for detection purposes will result in a loss of up to 30% of all events, while in the case of the double-precessing model, this never exceeds 6%. Moreover, even for signals from systems with small spins, nonspinning and spin-aligned templates introduce large biases in the extracted masses and spins. The use of a model that encodes spin-induced precession effects, such as the double-precessing model constructed in this dissertation, improves the mass and spin extraction by up to an order of magnitude. The additional information encoded in the spin-orbit interaction is invaluable if one wishes to extract the maximum amount of information from GW signals.

Introduction

The recent detection of GWs highlights the importance of efficient data analysis, a nontrivial task for signals that are deeply buried in detector noise. The most efficient way to extract and analyze such signals is through template filters: analytical or numerical models for the response of the detectors to impinging GWs. The templates are functions of a parameter vector $\vec{\theta}$ that characterizes the GW emitting system, and its position relative to the Earth. Parameter estimation consists of finding the

components of $\vec{\theta}$ that best fit the signal, as well as their error. Not surprisingly, the efficiency of such an analysis is highly dependent on the accuracy of the template model itself [74, 186].

Having established the accuracy of the analytical, small-spin, double-precessing model (in this chapter, we will refer to it as just double-precessing) relative to a purely numerical PN model in Chapter 4, we now study how good the former is at detecting and estimating the parameters of signals in noise. The double-precessing model is expected to be able to recover more information from precessing signals, because it can capture the amplitude and phase modulations induced by precession, and thus, break degeneracies that are present in the absence of precession. We find that this is indeed the case: the precessing model breaks degeneracies between the mass ratio and the spin magnitudes [73], allowing for a much better estimation of both quantities, by up to an order of magnitude. The improvement in parameter estimation is such that the precessing model can distinguish between NSs and BHs in the mass gap, *even for nonspinning signals* [60]. This result is in contrast to the conclusions one would arrive at if using spin-aligned templates that lack precession effects [106].

The idea that spin precession can significantly improve parameter extraction is by no means new. Vecchio [218] was the first to show that spin-precession effects improve parameter extraction in the context of LISA sources. The restricted 1.5PN simple-precession model he considered [21] was later extended to 2PN order through numerical PN waveforms by Lang and Hughes [131], who reached similar conclusions. Klein et al. [126] included higher harmonics and showed that parameter extraction was further improved. Concerns that binaries in gas rich environments tend to have partially aligned spins, prompted Lang et al. [132] to study partially aligned models; they found that restricting precession degrades parameter extraction significantly, but the inclusion of higher harmonics improves extrinsic parameter extraction again. A

similar result was recently found by O'Shaughnessy et al. [161]. In another recent paper, Vitale et al. [223] performed an extensive search of the parameter space and found that parameter extraction is improved when precessional effects are maximized, i.e. when the binary is observed edge on. The results of this and the following chapter verify the above results and further demonstrate that the more accurate double-precessing model improves detection rates and parameter estimation for NS binaries so much so that it enables distinguishing between NSs and BHs and measure NS spins. The above comparison excludes the numerical PN templates, since they are slower by about a factor of 10^2 or more than the double-precessing model [125], a fact that makes them prohibitive for parameter estimation studies.

Regarding detection, we study the efficiency of nonspinning, spin-aligned and double-precessing templates at detecting a numerical PN model of GWs emitted by generically precessing, spinning binaries with arbitrary spin magnitudes. We address this by calculating the so-called *fitting factor* [see Eq. (6.4)]: the normalized, noise-weighted inner product [see Eq. (6.2)] between a model and the signal, maximized over all template parameters. Such a measure is ideal to estimate how good a model is at recovering as much of the signal as possible at the expense of distorting the recovered parameters. This measure is above 98%, corresponding to a 6% drop in detection rate, when using the spin-aligned and the small-spin, double-precessing templates for *all* NS binaries with astrophysically realistic spins [147]. In the spin-aligned model, however, this large fitting factor comes at the expense of large biases in the extracted masses and spins. Binary BHs can have much larger spin magnitudes than NSs, and thus, the nonspinning and the spin-aligned models reach fitting factors above 98% only for $\chi_A < 0.4$. The double-precessing model reaches fitting factors above this threshold for all $\chi_A < 1$, at the expense of large parameter biases.

We then consider the efficiency of these templates in parameter estimation, focusing on spin detectability and the accuracy in parameter extraction. In particular, we study what SNR and what injected spin parameter we can claim that a NS binary signal was produced by spinning NSs. If we can claim the signal corresponds to such a spinning binary, we can then address how well the spin magnitudes can be measured, again as a function of SNR and injected spin parameter. We tackle these questions in a Bayesian framework [69, 70, 140, 185, 186], where we inject a small-spin, double-precessing signal and search for it through Markov-Chain Monte Carlo (MCMC) methods with either a spin-aligned template or a small-spin, double-precessing template. Such MCMC methods allow us to not only find the best fit parameters $\vec{\theta}_{\rm best}$, but also to construct their posterior probability distribution, as well as to determine which template model is best supported by the data.

One may be concerned that the small-spin, double-precessing template should not be used to estimate the statistical accuracy with which parameters can be inferred or alternative models distinguished, given that for higher spin values there will be systematic bias when recovering the true GWs we expect from nature. However, as we show in Appendix B, systematic and statistical errors are independent for small model deviations. Thus, the analytic double-precessing model can be used for reliable Bayesian inference and model selection.

The first parameter estimation question we tackle is that of spin detectability, which is a model selection problem [1, 80, 104, 136]: given a signal, we wishes to determine which of two competing models ("the signal was produced by a spinning binary" versus "the signal was produced by a nonspinning binary") is best supported by the data. We address this problem by calculating the *Bayes factor* (BF), which provides an estimate of how well a model fits the data compared to another model. Since we are dealing with *nested models* (models which reduce to each other when

a subset of their parameters $\vec{\theta}$ acquire certain values), the BF can be calculated through the *Savage-Dickey density ratio* [221]: the ratio of the prior to the posterior evaluated at vanishing spins. We find that the data prefer the small-spin, double-precessing model over the nonspinning model at dimensionless spin magnitudes larger than roughly 0.02 for SNR 10 with aLIGO [60] and 0.01 for SNR 30 with LIGO3 [6]. On the other hand, use of the spin-aligned model increases the spin detection threshold to roughly 0.05 and 0.02 respectively.

The second parameter estimation question we address is that of accuracy in parameter extraction. Given a small-spin double-precessing signal, we determine the best-fit parameters and their 1-σ confidence region (the smallest area in parameter space that contains 68% of the posterior weight) for either a spin-aligned or a double-precessing model. We find that the small-spin, double-precessing templates can measure masses and spins roughly 1 order of magnitude better than spin-aligned templates. This is because even a small amount of precession is sufficient to greatly deteriorate the likelihood of a double-precessing template, while a spin-aligned template cannot access this extra structure. This structure breaks degeneracies between the mass ratio and the spin magnitudes, allowing for a better measurement of both quantities. We show that these result are insensitive to the specific choice of spin priors: uniform over spin magnitudes and uniform over spin orientations on a 2-sphere. The improvement in parameter estimation is so dramatic that one should be able to distinguish between NS binaries and BH binaries purely from the detection of GWs during the inspiral phase. We expand on this last point in Chapter 7.

Waveform Models

We consider BHBH binaries and NSNS binaries with masses m_1 and m_2 (where $m_1 \geq m_2$) and spin angular momentum magnitudes S_1 and S_2 respectively in

adiabatically evolving, quasicircular orbits in the inspiral phase. GWs emitted from such a system induce a signal on ground-based detectors described by the parameter vector

$$\vec{\theta} =(\mathcal{M}, m, \cos\theta_N, \phi_N, D_L, \cos\theta_L, \phi_L, t_c, \phi_c, \cos\theta_1, \phi_1, \chi_1, \cos\theta_2, \phi_2, \chi_2), \tag{6.1}$$

where $\mathcal{M} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$ is the chirp mass, $m = m_1 + m_2$ is the total mass, $\cos\theta_N$ and ϕ_N are sky location angles, D_L is the luminosity distance, $\cos\theta_L$ and ϕ_L are angles that describe the direction of the initial orbital angular momentum vector, $\cos\theta_A$ and ϕ_A are angles that describe the direction of the initial spin angular momentum vectors, with $\chi_A = S_A/m_A^2$ the dimensionless spin magnitude, for the Ath binary component. All angles are measured in a geocentric frame [16].

We consider five different waveform models: one that is purely numerical; two versions of the analytical, small-spin, double-precessing model; one version of the analytical spin-aligned model; and one version of the analytical nonspinning model. When considering detection, we use the numerical PN model as the signal and the other four models as templates. When considering parameter estimation, we use one of the double-precessing models as the signal, and a subset of the other analytical models as the template. We describe each of these models below.

Numerical PN model. This model is constructed by first solving the most accurate PN spin-precession equations numerically (see e.g. [62, 124]), and then Fourier-transforming the numerical PN time-domain response function through a discrete Fourier transform. We stress that this model is constructed by solving PN ordinary differential equations numerically, similarly to the SpinTaylorT4 [1, 54], or effective-one-body models [53,75]. Therefore, we regard it as a numerical PN model, in contrast to the closed-form analytical models we describe below, and the full numerical

relativity based models of [10, 85, 165, 188]. Sometimes in the literature this model is referred to as 'semianalytical'.

Double-precessing models [62]. The precession equations are solved by separating the three intrinsic time scales of the problem: the *orbital* time scale, which is much shorter than the *precession* time scale, which is much shorter than the *radiation-reaction* time scale. The resultant orbital precession equations are then expanded in $\chi_A \ll 1$ and in the ratio of the different time scales. Such a *multiple-scale analysis* treatment results in an analytical solution for the temporal evolution of the orbital and the spin angular momenta, valid to first order in χ_A and in the ratio of the precession to the radiation-reaction time scale. This solution can then be used to construct a time-domain response function that is Fourier-transformed through the *stationary-phase approximation* [84, 245]. The model was constructed in Chapter 4 of this thesis and comes in two versions:

(i) *Full, double-precessing:* both the Fourier amplitude and phase are kept to high PN order [see Eqs. (105), (106), (107) of [62].]

(ii) *Restricted, double-precessing:* the Fourier amplitude is kept only to leading PN order, while all known PN corrections are kept in the Fourier phase [see Eqs. (98), (99), (100) of [62].]

Restricted spin-aligned model. This waveform is constructed by assuming the spin angular momenta are *exactly* aligned with the orbital angular momentum. Such an alignment prevents the system from precessing, thus rendering the spin-precession equations simple to solve [43]. One then solves the evolution equation for the orbital frequency and phase through a PN expansion, which allows the construction of a time-domain response function. The latter is Fourier-transformed through the SPA. We here consider a restricted model, where only the leading PN order term is kept in the

Fourier amplitude, while the Fourier phase is kept to 3.5PN order. The performance of these templates has been studied in numerous papers [1, 7, 41, 106]; most of them conclude that, even though spin-aligned templates might be good enough for detection of NSNS binaries, they lead to large biases when used in parameter estimation.

Restricted nonspinning model. This waveform is derived assuming that the binary components have no spin angular momenta. The temporal evolution of the orbital frequency is obtained analytically through a PN expansion, which is then used to construct a time-domain response. The latter is Fourier transformed through the SPA. We here focus on a restricted version of this waveform, where we keep the Fourier amplitude to leading PN order, but the Fourier phase is kept to 3.5PN order. Such waveforms have been studied extensively in the literature [11, 23, 143, 156, 162, 184, 194, 219], mainly as detection templates, despite their inherent inability to measure spins.

When studying parameter estimation, we will be in part interested in the errors associated with the extraction of parameters. There are two main types of errors in parameter recovery: *systematic errors* and *statistical errors*. Systematic errors are associated with a shift in the peak of the posterior distribution of the recovered parameter away from the injected value; such errors can be produced by inaccuracies in the template model. Statistical errors are associated with the inherent width of the posterior distribution; such an error is produced by the signals possessing a finite SNR. As we show explicitly in Appendix B, systematic and statistical errors are independent to first order in the inaccuracies of the model, and thus, we will study them separately.

When studying systematic errors, we will randomize over all model parameters, while when focusing on statistical errors, we will select a few characteristic systems. By doing so, we isolate the effects of SNR, injected spin parameter and detector

	m_1	m_2	χ_1	χ_2	$\cos\theta_N$	ϕ_N	ι	κ
1	1.43	1.23	*	*	-0.11	3.71	63°	30°
2	1.43	1.23	0.04	0.04	-0.11	3.71	63°	*
3	*	*	0.04	0.04	-0.11	3.71	63°	30°
4	*	*	0	0	-0.11	3.71	63°	30°

Table 6.1: Summary of the systems used in the parameter estimation analysis. The masses are in units of solar masses, ι is the angle between the orbital angular momentum and the line of sight at GW frequency 70Hz, and κ is the opening angle between the orbital angular momentum and the total spin angular momentum again at 70Hz. The asterisk denotes the parameters that are varied.

(aLIGO or LIGO 3) on parameter recovery. The three systems we work with are characterized by the parameters in Table 6.1.

Detection

In GW astronomy, there are two measures that estimate the extent to which two models are similar to each other: the *faithfulness* and the *fitting factor* [76][1] . Both measures depend on the noise-weighted inner product between two models for the response function, h_1 and h_2:

$$(h_1 \mid h_2) \equiv 4\Re \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)}\, df\,, \tag{6.2}$$

where the overhead tilde stands for the Fourier transform, $\Re[\cdot]$ is the real part operator, $(f_{\min}, f_{\max})$ are the limits of integration, and $S_n(f)$ is the detector's spectral noise density; we here use the high-power, zero-detuned $S_n(f)$ of aLIGO [201].

The faithfulness is a measure of how good a template is at recovering a signal with the *same* parameters, and thus, how efficient the model is at parameter recovery.

[1]The faithfulness was defined and discussed in Chapter 4. We repeat part of this discussion here for completeness and to emphasize its relation to the fitting factor.

It is defined as

$$F \equiv \max_{t_c,\phi_c} \frac{(h_1 \mid h_2)}{\sqrt{(h_1 \mid h_1)\,(h_2 \mid h_2)}}, \tag{6.3}$$

where the inner product is maximized only over the time of coalescence t_c and the phase of coalescence ϕ_c.

The fitting factor is a measure of how good a template is at recovering a signal regardless of biasing parameter recovery. It is defined through

$$FF \equiv \max_{\vec{\theta}} \frac{(h_1 \mid h_2)}{\sqrt{(h_1 \mid h_1)\,(h_2 \mid h_2)}}\,, \tag{6.4}$$

where the inner product is maximized over all parameters. In general, the highest FF is achieved between models h_1 and h_2 that have *different* parameters.

In order to obtain reliable estimates for these two measures that are independent of the specific system considered, we create a random distribution of systems through Monte Carlo (MC) methods. The mass distribution is chosen to be flat in log space, with boundaries chosen depending on the class of system considered: for NS binaries, we choose the range $[1, 2.5]M_\odot$, while for BH binaries, we choose the range $[5, 10]M_\odot$. All vector directions are chosen uniformly on the sphere. We present our results as a function of the symmetric dimensionless spin parameter $\chi_s \equiv (\chi_1 + \chi_2)/2$ since it gives a measure of how applicable the small-spin approximation of Chapter 4 is for the particular system studied. Below we study BH binaries separately from NS binaries, since the latter are expected to have comparable masses and small spin magnitudes, while the former are not.

NSNS Binaries

NS binaries that enter the sensitivity band of ground-based detectors are not expected to have large spin magnitudes. This is because although NSs can be spun up by accretion, they spin down due to magnetic breaking. By the time they have spiraled into each other sufficiently to be emitting GWs detectable by ground-based detectors, their spin magnitudes are not expected to exceed $\chi_A = 0.2$ [147]. This fact makes NSNS binaries an ideal candidate for the small-spin, double-precessing model of Chapter 4.

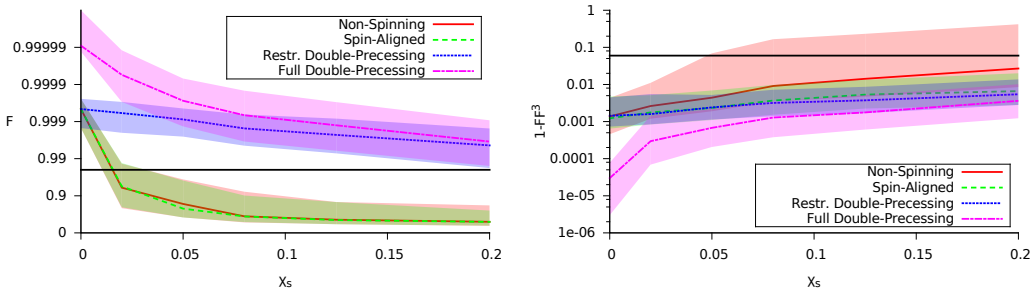

Figure 6.1: Median faithfulness (left panel) and median drop in detection rates (right panel) between a numerical PN waveform and a full double-precessing waveform (magenta dot-dashed line), a restricted double-precessing waveform (blue dotted line), a restricted spin-aligned waveform (green dashed line), and a restricted nonspinning waveform (red sold line) for NSNS binaries as a function of the symmetric spin. The shaded areas give the 1-σ confidence regions and the black solid line represents the 98% threshold. In the case of detection rates this threshold corresponds to the loss of 6% of all events.

Figure 6.1 shows the faithfulness and (one minus the cube of) the fitting factor for NSNS binaries between the numerical PN model and all the analytic models, as a function of the symmetric dimensionless spin parameter. Since the recovered SNR of a source scales as the fitting factor, $1 - \mathrm{FF}^3$ gives an estimate of the reduction of the volume accessible to the detectors due to model inaccuracies. In other words, when the fitting factor drops, the source needs to be closer to earth to give the same

SNR value and be detectable. For this reason we interpret $1 - \mathrm{FF}^3$ as the drop in overall expected detection rates of aLIGO/AdV, which are highly uncertain to begin with [3].

Each point in χ_s is computed by averaging over 2000 random systems (600 for the full double-precessing model due to computational restrictions) with masses in $[1, 2.5]M_\odot$. The lower limit of integration is $f_{\rm min} = 10$Hz, the frequency at which GWs enter the aLIGO band. The upper limit of integration is $f_{\rm max} = 400$Hz, while the system is still in the inspiral phase, in order to avoid finite size effects that enter above this frequency [112, 148, 180].

Three primary conclusions can be drawn from these plots. First, the faithfulness stays above the nominal 98% threshold when using the double-precessing models for all spins considered, while it drops below this threshold for the nonspinning and spin-aligned system above $\chi_s = 0.02$. This indicates that only the double-precessing models can be considered as reliable parameter estimation templates. Second, the fitting factor is above the 98% threshold, corresponding to a loss of event rate smaller than 6%, for all models. As expected from previous results, the nonspinning and the spin-aligned models can serve as detection templates for slowly spinning systems, like NS binaries. Third, we find similar fitting factors when using the restricted and the full double-precessing models. This implies that the restricted model is sufficient for parameter estimation studies.

Comparing the two panels of Fig. 6.1 we see how the spin-aligned and nonspinning templates are able to distort their parameters to achieve a better overlap with the numerical PN model. In Fig. 6.2 we plot the bias that such shifting induces on the chirp mass, the total mass, and the absolute value of the effective spin parameter [the symmetric spin combination projected onto the orbital angular moments; see Eq. (6.12)]. Clearly, if the nonspinning or the spin-aligned waveforms are used for

parameter estimation, the resulting parameter bias will be significant (about an order of magnitude larger than from the double-precessing models), and the systematic error will most likely dominate the total error.

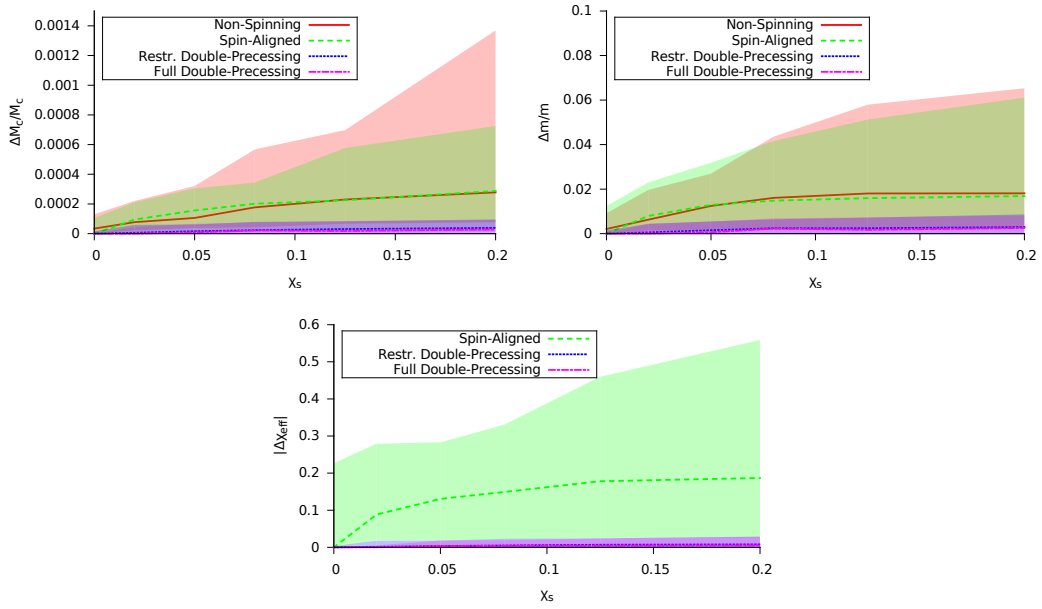

Figure 6.2: Median parameter bias for the chirp mass (top left), the total mass (top right), and the absolute value of the dimensionless effective spin parameter (bottom) for full double-precessing waveforms (magenta dot-dashed line), restricted double-precessing waveforms (blue dotted line), restricted spin-aligned waveforms (green dashed line), and restricted nonspinning waveforms (red sold line) for NSNS binaries as a function of the symmetric spin parameter.

BHBH Binaries

Unlike NSs, there is no astrophysical reason to limit the spin magnitude of BHs (other than cosmic censorship, $\chi \leq 1$). One may thus expect the small-spin, double-precessing model of Chapter 4 to perform badly when attempting to detect highly spinning signals. However, we find this not to be the case, due to the ability of the double-precessing model to shift its 15 parameters in order to recover as much of the signal as possible.

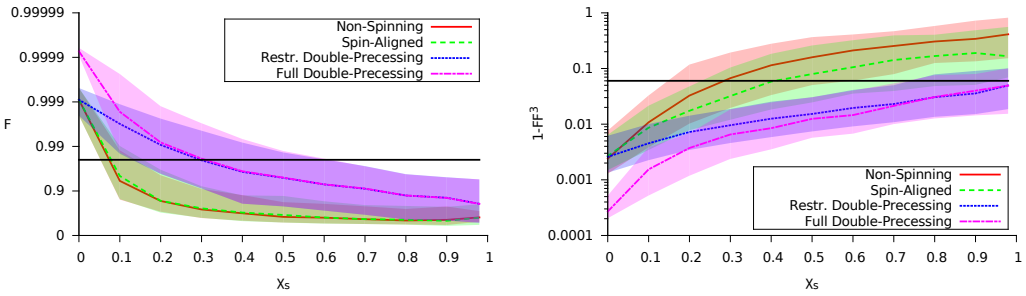

Figure 6.3: Median faithfulness (left panel) and median drop in detection rates (right panel) between a numerical PN waveform and a full double-precessing waveform (magenta dot-dashed line), a restricted double-precessing waveform (blue dotted line), a restricted spin-aligned waveform (green dashed line), and a restricted nonspinning waveform (red sold line) for BHBH binaries as a function of the symmetric spin. The shaded areas give the 1-σ confidence regions and the black solid line represents a value of 98% (corresponding to a 6% drop in detection rates).

Figure 6.3 shows the faithfulness and (one minus the cube of) the fitting factor for BHBH binaries between the numerical PN model and all the analytic models as a function of χ_s. Each point in χ_s is computed by averaging over 6000 random systems (1100 for the full double-precessing model due to computational restrictions) with masses in $[5, 10]M_\odot$. The lower limit of integration is again set at $f_{\rm min} = 10$Hz. However, since GWs emitted by BHs do not have any finite size effects, we extend the integration to the frequency corresponding to GWs emitted by a test particle at the innermost stable circular orbit (ISCO) of a Schwarzschild BH, i.e. $f_{\rm max} = 6^{-3/2}/(\pi m)$.

For spins larger than about $\chi_s = 0.3$ the faithfulness drops below 98%; a first order expansion in the spins seems inadequate in capturing the strong precessional effects present in binaries with large spins. However, if the waveforms are allowed to adjust their parameters to fit the signal, they perform significantly better in detecting sources with large spin magnitudes. The right panel of Fig. 6.3 shows the drop in detection rates for all analytical models. Now the double-precessing waveforms obtain overlaps greater than 98% for all dimensionless spin magnitudes. The nonspinning

and spin-aligned templates perform adequately for spins only up to 0.3 and 0.4 respectively. Clearly, the two double-precessing models are the only reliable detection templates for highly spin-precessing BHBH binaries of all models considered here.

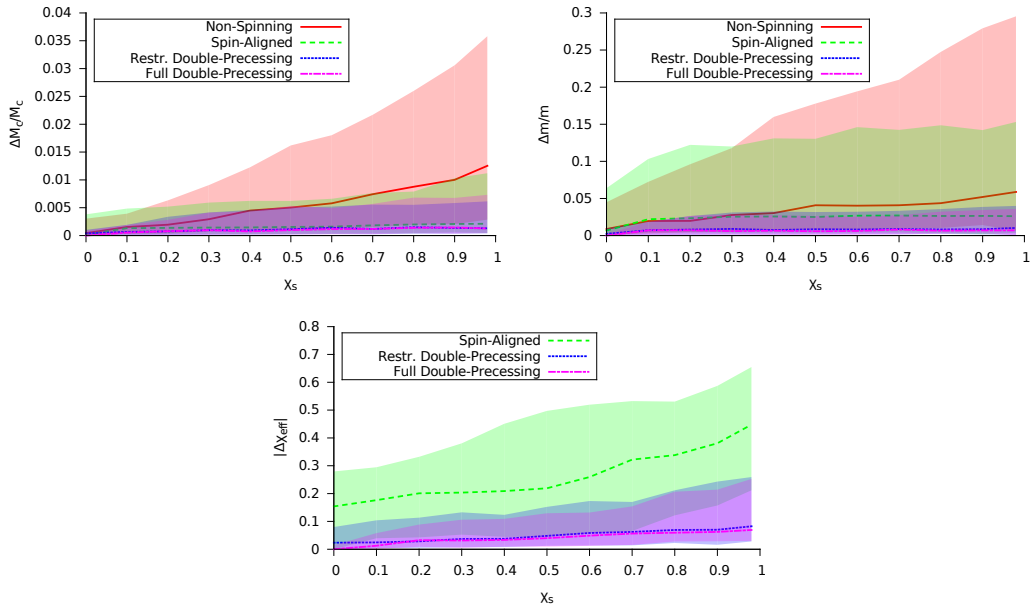

Figure 6.4: Median parameter bias for the chirp mass (top left), the total mass (top right), and the effective spin (bottom) for full double-precessing waveforms (magenta dot-dashed line), restricted double-precessing waveforms (blue dotted line), restricted spin-aligned waveforms (green dashed line), and restricted nonspinning waveforms (red sold line) for BHBH binaries. The shaded areas give the 1-σ confidence regions.

The large difference between faithfulness and fitting factor shows that even the double-precessing waveforms have to adjust their parameters significantly to achieve high overlaps with the numerical PN waveforms. Figure 6.4 shows the bias in the chirp mass, the total mass, and the absolute value of the effective spin [Eq. (6.12)]. The bias induced by using double-precessing templates is about an order of magnitude smaller than that incurred when using nonspinning or spin-aligned templates. Nonetheless, even the double-precessing templates induce a significant bias, making them unsuitable for parameter estimation of BHBH binaries [1, 216].

Likelihood as a Function of Mass

Maximizing the fitting factor reduces to maximizing the *likelihood* [the importance of which will become more evident when we carry out parameter estimation; see also Eq. (6.8)]. The efficiency of any maximization algorithm is highly dependent on our understanding of the behavior of the likelihood surface, allowing us to propose better jumps that find the peak and explore it more efficiently. It is, therefore, important to study how precession affects the likelihood surface.

Figure 6.5 shows the log of the likelihood maximized over the time of coalescence, the phase of coalescence, and the luminosity distance as a function of the chirp mass (left panel) and as a function of the total mass (right panel) for a BHBH system with $\mathcal{M} = 7.23M_\odot$, $M = 16.8M_\odot$, and $\chi_1 = \chi_2 = 0.5$. The left panel shows a strong preference for the injected value; it is unlikely that a different chirp mass will be recovered. Indeed, when we studied parameter biases in Figs. 6.2 and 6.4, we found that the chirp mass is biased by about 0.01% for NSNS binaries and 0.3% for BHBH binaries, depending on the spin of the injection.

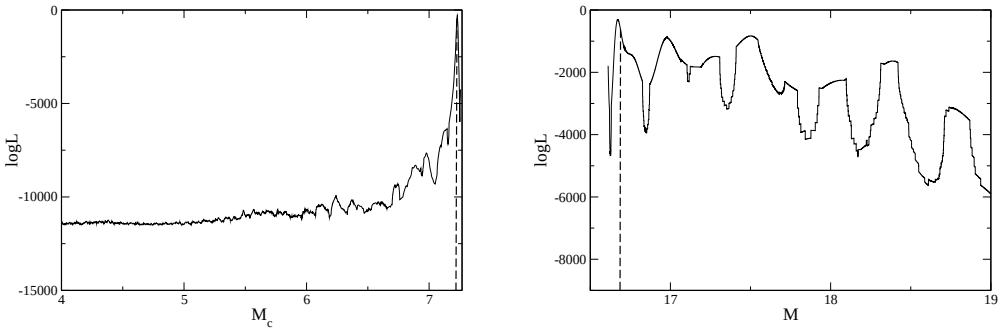

Figure 6.5: Log likelihood as a function of the chirp mass (left) and of the total mass (right) for a BHBH system. The vertical dashed lines correspond to the injected values $\mathcal{M} = 7.23M_\odot$ and $M = 16.7M_\odot$.

On the other hand, the right panel of Fig. 6.5 shows a completely different dependence of the likelihood on the total mass. We see that the log of the likelihood presents a series of peaks with comparable heights. Therefore, as the other parameters in the model are varied from their injected value, one of the secondary peaks might become the primary one, resulting in a higher fitting factor. In fact, this is the case for the system presented here. The injected value for the total mass corresponds to the first peak around $M = 16.8M_{\odot}$. However, the recovered, or best fit, value for the total mass corresponds to the second peak at around $M = 16.9M_{\odot}$. By appropriately adjusting its parameters, the template model managed to find a better fit to the signal than by using the signal's parameters, resulting in a fractional systematic error of about 0.5% for the recovered total mass. This is also verified by the top, right panel of Fig. 6.4.

The series of peaks in log likelihood suggests that to map the likelihood surface sufficiently one should propose jumps between peaks, so that the Markov chains do not get stuck in a local maximum. To do so, we used the log likelihood as a function of the total mass maximized over t_c, ϕ_c and D_L as an additional jump proposal, where all other parameters were held fixed, and the new total mass point was drawn from this distribution through rejection sampling. These jumps ensure that all peaks are explored adequately and the one with the maximum likelihood is selected.

An interesting consequence of the behavior of the likelihood surface is related to the *theoretical bias* [74]. The latter is defined as the mismodeling error in parameter recovery induced by inaccuracies in the template model, e.g. due to truncation of the PN series. One semianalytic estimate of this error can be obtained by modeling the likelihood surface as a single peak of finite width [74]:

$$\Delta_{th}\theta^i = \left(\Gamma^{-1}(\theta_{tr})\right)^{ij} \left(\partial_j h_{\rm SPA}(\theta_{tr}) \mid h_{\rm DFT}(\theta_{tr}) - h_{\rm SPA}(\theta_{tr})\right), \tag{6.5}$$

where $h_{\rm DFT}$ is the true signal (in our case, the numerical PN model), $h_{\rm SPA}$ is the "incorrect" template that is used (in our case, any of the analytical models), and $(\Gamma^{-1})^{ij}$ is the inverse of the Fisher information matrix

$$\Gamma_{ij} = (\partial_i h \,|\, \partial_j h)\,, \tag{6.6}$$

where ∂_i denotes differentiation with respect to the ith parameter. All quantities are evaluated at the injected parameters θ_{tr}.

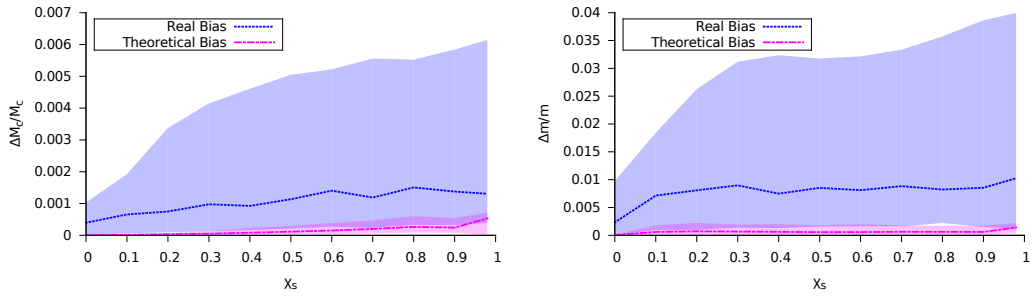

Figure 6.6: Real bias (blue dotted line) and theoretical bias (magenta dot-dashed line) for the chirp mass (left panel) and the total mass (right panel) for a BHBH binary as a function of the injected spin. The shaded regions give the 1-σ confidence intervals.

This estimate of the mismodeling error due to theoretical bias is indeed approximately correct when the likelihood surface is single peaked, but it can grossly underestimate the biases when the surface is multipeaked. Figure 6.6 shows the real error, as estimated from the posterior distribution, and the mismodeling error, as estimated with Eq. (6.5), for the chirp mass (left panel) and the total mass (right panel) as a function of the injected χ_s, given a BHBH binary signal. For the total mass, Eq. (6.5) (roughly the width of the first peak) underestimates the true bias (roughly the distance between peaks) by an order of magnitude. As a further

verification of this, we restricted the total mass range around the primary peak of the log likelihood and found that the real bias agreed with the theoretical bias.

Parameter Estimation

Parameter estimation in the Bayesian framework gives the posterior belief that a certain model with a parameter vector $\vec{\theta}$ describes the data d. The posterior distribution is calculated through Bayes' theorem

$$p(\vec{\theta}|d) = \frac{p(d|\vec{\theta})p(\vec{\theta})}{p(d)}, \tag{6.7}$$

where $p(\vec{\theta}|d)$ is the posterior belief, $p(\vec{\theta})$ is the prior belief on the parameters, $p(d)$ is the evidence (here, an irrelevant normalization factor), and $p(d|\vec{\theta})$ is the *likelihood* that data d were produced by a model with parameters $\vec{\theta}$. For the prior, we choose uniform distributions in the allowed region of parameter space. In GW studies, the likelihood is the noise model, which we here assume to be Gaussian and stationary

$$p(d|\vec{\theta}) \sim \exp\left[-\frac{1}{2}\left(s-h \mid s-h\right)\right], \tag{6.8}$$

with s the detectors' output and h the template model. The $15-$dimensional posterior distribution is sampled though an MCMC algorithm. The posterior distribution for each parameter is then obtained by marginalizing over all other parameters.

Unlike in the previous section, where we were interested in maximizing the likelihood to recover the maximum overlap between the signal and the model, we are now interested in the likelihood surface itself. For that reason we use the restricted double-precessing waveform as the signal and recover it with the spin-aligned model and the double-precessing model as templates. A very wide likelihood surface results

in poor parameter extraction, while a peaked likelihood results in small errors for the recovered parameters.

Henceforth, we assume GW detections with the following three-detector network configurations: (i) two aLIGO detectors and one AdV detector with network SNRs of 10 and 20, and (ii) three detectors with the LIGO3 noise model [6] and SNRs of 30 and 60. We concentrate on observations of the characteristic systems described in Table 6.1.

Furthermore, since we are dealing with NS binaries, we stop our analysis at a GW frequency of 400Hz in order to avoid finite size effects [112, 148, 180]. Extending our analysis beyond this frequency would only serve to strengthen our results for the following reasons: (i) for a given GW source at a fixed distance, the inclusion of the late inspiral, plunge and merger, increases the SNR, which leads naturally to an improvement in parameter estimation; (ii) the finite size effects that NSs experience can provide useful information in mass extraction and distinguishing between NSs and BHs; and (iii) electromagnetic counterparts from the merger phase can aid in differentiating between NSs and BHs. Thus, from this standpoint, our parameter estimation results could be thought of as conservative.

Distinguishing between
Nonspinning and Spinning Binaries

Given a GW detection, a particularly important follow-up question is whether the signal was produced by a spinning binary or not. In this section, we address this issue by examining whether the restricted, double-precessing model can be used to distinguish between spinning and non-spinning signals. We do so by calculating the *Bayes factor*, in the case of uninformative flat priors the betting odds, in favor of the spinning model. If the BF is less than one, then the non-spinning model is

preferred and we cannot conclude from the data that the binary components have nonzero spin magnitudes.

When considering nested models, i.e. models that reduce to each other when a subset of the parameters in one of them acquire certain values, the BF reduces to the *Savage-Dickey density ratio*, which is given by

$$\mathrm{BF} = \frac{p(\chi_1 = 0, \chi_2 = 0)}{p(\chi_1 = 0, \chi_2 = 0|d)}\,, \tag{6.9}$$

the ratio of the prior belief that the spins were zero to the posterior belief that the spins are zero. In Appendix C, we derive this result for models that differ by multiple parameters, some of which do not contribute to the likelihood unless others are nonzero, e.g. the spin angle parameters do not matter if the spin magnitude is zero.

Although the prior can be easily evaluated at $(\chi_1, \chi_2) = (0, 0)$ since it is uniform, the posterior is much more difficult to calculate. As already mentioned, the process of determining the $2D$ posterior $p(\chi_1, \chi_2|d)$ involves marginalizing over all other parameters. This is done by dividing the (χ_1, χ_2) space into bins of size $d\chi_1 = d\chi_2$ and counting how many times the Markov chains visit each corresponding bin. The value of the posterior at $(\chi_1 = 0, \chi_2 = 0)$ is proportional to the number of samples in the first bin $(0, d\chi_1, 0, d\chi_2)$.

Clearly, the result depends sensitively on the number of bins used, or equivalently, on the size of each bin. There are two main sources of error in this calculation. If the size of the bins is too small, there will not be enough samples in each of them to give a statistically reliable result, i.e. there are large root n errors, where n is the number of samples in the bin. A very large bin size, on the other hand, will result in an inaccurate estimate of the value of the posterior at $(\chi_1, \chi_2) = (0, 0)$.

In order to reduce root n error, we need more samples in the first bin, which we achieve through a two-stage analysis. In the first stage, the *pilot run*, we obtain N_1 samples from the full posterior distribution. Given that, one can estimate the rectangle $(0, \chi_1^{\rm max}, 0, \chi_2^{\rm max})$ that contains n_1 samples (n_1 chosen to be $\sim$ 10% of N_1). In the second stage, the *focused run*, we carry out an analysis with a flat prior in $(0, \chi_1^{\rm max}, 0, \chi_2^{\rm max})$ and a zero prior elsewhere, effectively forcing the chains to visit points close to zero spin magnitude, and thus reducing the statistical root n fluctuations. The focused run results in a total of N_2 points, n_2 of which are in $(0, d\chi_1, 0, d\chi_2)$, where recall that $d\chi_1 = d\chi_2$ is the size of the bins. Figure 6.7 gives an illustration of this procedure. The value of the normalized $2D$ posterior at vanishing

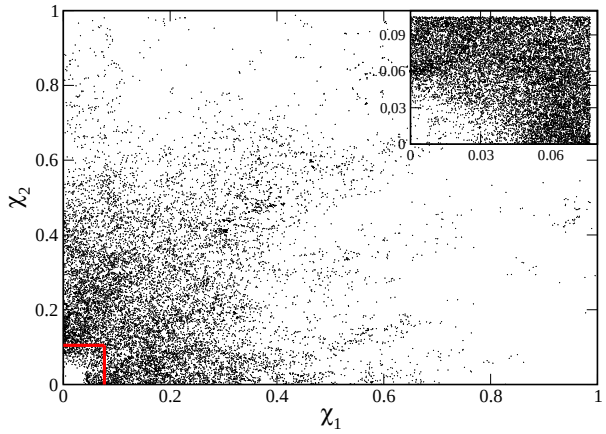

Figure 6.7: Example 2D scatter plots for $\chi_1 - \chi_2$ for the pilot run (main plot) and the focused run (inset) to illustrate the two-stage analysis. The red box in the pilot run indicates the size of the region $(0, \chi_1^{\rm max}, 0, \chi_2^{\rm max})$ of the focused run. This region contains $\sim$ 10% of the total points of the pilot run.

spins is then

$$p(\chi_1 = 0, \chi_2 = 0|d) = \frac{n_2}{N_2}\frac{n_1}{N_1}\frac{1}{d\chi_1 d\chi_2}\,, \tag{6.10}$$

while the fractional error from this procedure can be estimated through

$$\sqrt{\frac{1}{n_1}+\frac{1}{n_2}}\,. \tag{6.11}$$

Having ensured that there are enough samples in the first bins, we still need to choose a bin size that provides an accurate estimate of the value of the posterior at zero spin magnitude. We do so by plotting the posterior in the first bin $(0, d\chi_1, 0, d\chi_2)$ as a function of the bin size $d\chi_1 = d\chi_2$. From this plot, we choose the points $\{b_i\}$ that satisfy the two following requirements: (i) the bin size is not comparable to the injected spin value and (ii) there are at least 30-50 samples in the first bin. Each point has an error bar $\{b_i^{\rm min}, b_i^{\rm max}\}$ calculated through Eq. (6.11). The BF is, then, given by the average of these points with an error bar $\{\min(b_i^{\rm min}), \max(b_i^{\rm max})\}$. The convergence of this procedure, and the accuracy of the error estimates, were checked by performing multiple runs with different random number seeds for a few examples. These multiple runs produced consistent results, with a spread in values that agreed with the error estimates.

Before we can discuss distinguishability between spinning and non-spinning systems, we must understand how spin enters the waveform templates. For systems with similar component masses, spin first enters through the effective spin parameter

$$\chi_{\rm eff} = \frac{\vec{\chi}_1 \cdot \hat{L} + \vec{\chi}_2 \cdot \hat{L}}{2}. \tag{6.12}$$

Not surprisingly, this is the parameter that can be extracted most accurately, just like the chirp mass is measured more accurately than the symmetric mass ratio. In this case, however, a measurement of $\chi_{\rm eff}$ only provides information about the component of the spin angular momentum along the orbital one. Measuring the perpendicular

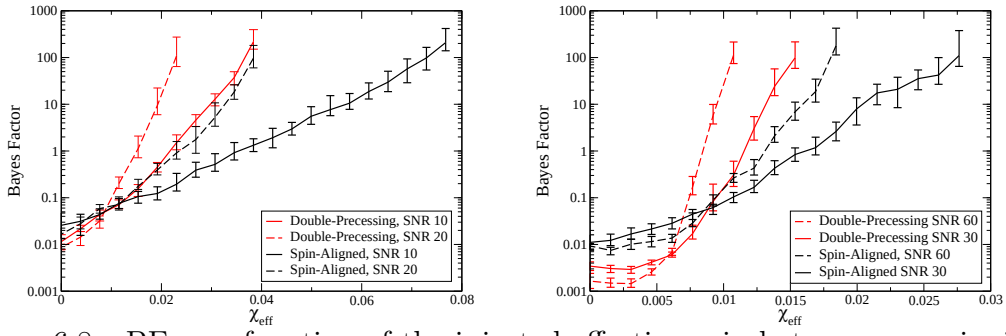

Figure 6.8: BF as a function of the injected effective spin between nonspinning and spinning models for System 1 of Table 6.1 and for spin-aligned (black) and double-precessing (red) templates, assuming an injection with SNR 10 (solid) and 20 (dotted) in an aLIGO-AdV network (left) and SNR 30 (solid) and 60 (dashed) as seen by LIGO3 (right).

components of the spin angular momentum would require measuring the cone of precession, which is difficult with the SNRs expected with aLIGO.

We tackle the distinguishability of spinning and non-spinning systems as a model selection problem [69, 80, 104, 185, 212]. We inject system 1 of Table 6.1, where we vary $\chi_1 = \chi_2$ and recover it with the double-precessing and the spin-aligned model. Figure 6.8 shows the BF between non-spinning and spinning models as a function of the $\chi_{\rm eff}$ for the aLIGO/AdV system (left) and for LIGO3 (right). BF > 1 indicates that the data favors the spinning model. For the same SNR the double-precessing template correctly identifies the signal as produced by a spinning source at a lower value of $\chi_{\rm eff}$ than the spin-aligned model, while both models correctly identify a non-spinning signal ($\chi_{\rm eff} = 0$ case).

The effective spin parameter is the appropriate variable to use when studying spin detectability. This is because it is $\chi_{\rm eff}$ which enters to leading PN order in the evolution of the GW phase. We can demonstrate the validity of this argument by calculating the BF in favor of the spinning model as $\chi_{\rm eff}$ is increased in two different

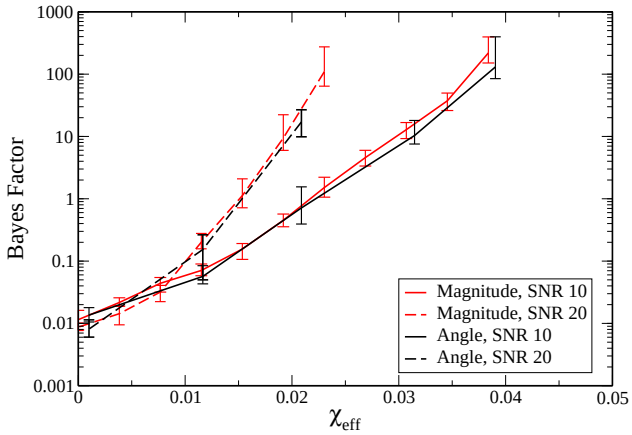

Figure 6.9: BF as a function of the injected effective spin between nonspinning and spinning models with double-precessing templates, where the effective spin has been updated through a change in the spin magnitudes (red lines, System 1 of Table 6.1) and a change in the spin angles (blue lines, System 2 of Table 6.1). The injected signal has a SNR of 10 (solid lines) and 20 (dashed lines) and is measured by aLIGO.

ways: (i) by increasing the value of the injected dimensionless spin parameters $\chi_1 = \chi_2$ (System 1 in Table 6.1), and (ii) by decreasing the angle between the spin and the orbital angular momenta (System 2 in Table 6.1). Figure 6.9 shows the BF in favor of the spinning model for the double-precessing model calculated in both ways. Both approaches give similar results, demonstrating that the model depends indeed on $\chi_{\rm eff}$ and not on the individual spin magnitudes and orientations.

Accuracy of Recovered Parameters

The double-precessing model can break degeneracies between the spins and the masses, improving the accuracy of mass extraction significantly, as compared to the spin-aligned model. This phenomenon has deep implications for astrophysical inference with ground-based detectors and merits further discussion and explanation. We do so in Chapter 7. Here we focus on another aspect of the double-precessing model: its ability to measure spins.

Once the components' masses have been accurately recovered (see the following chapter), what is left is recovering the spin magnitudes themselves. We quantify our belief on the recovered, effective spin parameter by estimating the minimum interval in the $\chi_{\rm eff}$ space that contains 68% of the posterior distribution. This interval corresponds to the 1-σ confidence region for a Gaussian distribution. Figure 6.10 shows the value of $\chi_{\rm eff}$ at the peak of the posterior for System 1 of Table 6.1 as a function of the injected effective spin value for signals detected with aLIGO (top panels) and LIGO3 (bottom panels) with the double-precessing model (left panels) and spin-aligned model (right panels) and SNR of 10 and 20 for aLIGO, and 30 and 60 for LIGO3. We indicate the 68% confidence area with shaded regions.

As expected, the distributions are peaked closer to the injected value and the error bars decrease as the SNR increases. The error bars always include the injected value. For an aLIGO detection with a SNR of 10, one could determine $\chi_{\rm eff}$ with a confidence of about ± 0.02. The use of spin-aligned templates deteriorates this by an order of magnitude; the error in $\chi_{\rm eff}$ is now about ± 0.2. We find similar results for signals detected by LIGO3. This time, the double-precessing model can place error bars of about ± 0.01 at SNR 30, while the spin-aligned model achieves an accuracy of about ± 0.1 only at the same SNR. For the same model the accuracy is essentially independent on the injected spin value, and depends mainly on the SNR value.

Conclusions

In this chapter we carried out a full Bayesian study of the performance of various analytical templates for detection and parameter estimation of double-precessing, compact binary inspirals. We considered the usual nonspinning and spin-aligned models, as well as the new analytical (small-spin), double-precessing model we constructed in Chapter 4.

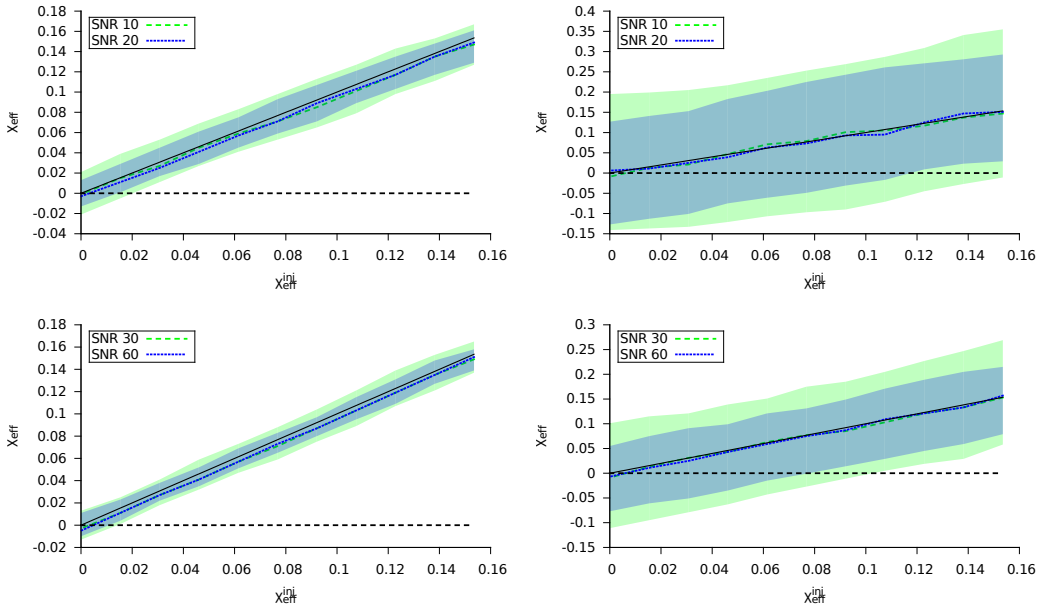

Figure 6.10: Maximum posterior value for effective spin for System 1 of Table 6.1 as a function of its injected value. The top panels correspond to a signal measured by aLIGO, while the bottom ones are for LIGO3. The left panels show signals recovered with the double-precessing templates, while the right ones with the spin-aligned ones. The shaded regions indicate the minimum interval that contains 68% of the posterior distribution. The black dashed line indicates $\chi_{\rm eff} = 0$, while the black solid line gives the $\chi_{\rm eff} = \chi_{\rm eff}^{\rm inj}$ curve.

We found that even though the non-spinning and spin-aligned models can be used for the detection of NSNS binaries (symmetric spin parameters up to 0.2), they are inadequate for BHBH binaries (arbitrary symmetric spin parameters). Furthermore, they induce mismodeling biases that make any use of them for parameter estimation purposes prohibitive. On the other hand, the spin-precessing model can achieve fitting factors above the 98% threshold for all symmetric spin values and lead to a reduction of the systematic bias in mass and spin of at least an order of magnitude [126, 131, 218, 223]. This enables us to use it for parameter estimation of NSNS binary systems.

The parameter extraction analysis is carried through a search of the likelihood surface with a MCMC technique, which is better suited for sampling complicated and weak signals than Fisher information matrix estimates [26, 27, 159, 183, 214, 215, 217, 224]. We find that the double-precessing template can not only lead to the detection of spins, but also to the measurement of the effective spin parameter to high accuracy compared to an alternative spin-aligned model.

The results presented here demonstrate the importance of precessional effects in the analysis of precessing systems, even when the spin magnitudes and angles are small. Failure to accurately include them will lead to a significant loss in the volume accessible to GW detectors. Given the already low detection rates expected, such a reduction might lead to erroneous astrophysical conclusions. Apart from the detectability issue, the use of precessing templates in data analysis can lead to answers to many astrophysically important questions, like mass and spin distributions of astrophysical objects.

BREAKING THE BH–NS DEGENERACY

The GW signals from compact binary merger events can be used to study the mass and spin distribution of stellar remnants, and provide information about BH horizons and the material properties of NSs. However, it has been suggested that degeneracies in the way that the star's mass and spin are imprinted in the waveforms may make it impossible to distinguish between BHs and NSs. In this chapter we show that the precession of the orbital plane due to spin-orbit coupling breaks the mass-spin degeneracy, and allows us to distinguish between standard NSs and alternative possibilities, such as BHs or exotic NSs with large masses and spins.

Introduction

The final stages of the inspiral and merger proceed differently for BHs and NSs, and in principle, this should allow us to identify the make-up of the system from the GW signal alone. However, the number of GW cycles in the signal and the aLIGO/AdV sensitivity fall off rapidly with increasing frequency, meaning that there is very little information past $\sim$ 500 Hz (less than 2% of the SNR). Probes of BH physics and the equation of state of NSs will likely require multiple detections [79,80]. An electromagnetic counterpart to the GW signal, such as a short-hard gamma-ray burst or an associated kilonova/macronova emission [155], would indicate that at least one of the bodies was a NS, but beaming effects or the luminosity of the signal may make detecting a counterpart difficult for the majority of mergers [2,3]. Absent a counterpart, we must rely on the early inspiral to extract information about the make-up of the binary, which poses a challenge since finite size effects are completely negligible during inspiral [180]. All we have to go on to decide the composition of the binary are the values of the masses and spins inferred from the inspiral signal.

General arguments based on stability and causality limit the mass and spin of NSs to the range $M \in [0.1, 3.2]M_\odot$ for the mass and $\chi \in [0, 0.7]$ for the dimensionless spin magnitude, $\chi \equiv |\vec{S}|/M^2$, where $\vec{S}$ is the spin angular momentum [134, 182, 236]. Realistic equations of state yield a tighter mass range $M \in [1.0, 2.5]M_\odot$. The observed range of masses and spins is somewhat tighter [134, 164]: $M \in [1.0, 2.0]M_\odot$, $\chi \in [0, 0.3]$. The old NSs that merge are expected to have spun down by magnetic breaking to the point where the maximum spin is much lower, $\chi \lesssim 0.05$, than in the general NS population [147]. Furthermore, the standard isolated NSNS binary formation scenario ensures that after every common envelope phase (that tends to align the spins) follows a supernovae kick that misaligns the spins (unless the kick is in the orbital plane, though there is evidence that this is not the case [119]). Thus, we adopt the definition that *normal* NSs seen by aLIGO/AdV have $M \in [1, 2.5]M_\odot$ and $\chi \leq 0.05$, and term NSs with larger masses or spins *exotic*. Einstein's theory of gravity allows BHs to have spin in the range $\chi \in [0, 1]$ with any mass. X-ray and GW observations have identified stellar remnant BHs with $M \in [3.6, 36]M_\odot$ and $\chi \in [0, 1]$. There is currently some debate as to the existence of a mass gap between NSs and BHs [36, 91, 163], but for the purpose of determining whether a normal NS could be misidentified as a BH or an exotic NS, the existence of a gap is moot.

The PN approximation can be used to construct a model of the GWs emitted during inspiral. The combination of such a GW model with a model for the instrument response yields templates for the signals as seen by the detector. Subtracting the model from the data produces a residual, and demanding that the residual is consistent with a model for the instrument noise defines a likelihood function. From this function and our prior knowledge we can derive a posterior distribution for the model parameters that are consistent with the observed data. It often happens that

there are strong correlations between these parameters, limiting our ability to measure each parameter individually.

Recent work [106] has suggested that the correlation between mass and spin [72, 73] may make it impossible to distinguish between a NSNS binary and a NSBH or a BHBH binary. This result hinges on a simplified waveform model that assumes that the spin and orbital angular momenta are perfectly aligned, and thus, spin-orbit induced precession [28, 48] is absent. However, we have no reason to expect the spin and orbital angular momenta to be aligned in stellar remnant binaries. Indeed, the NS binaries observed at much longer orbital periods are far from aligned and are precessing [50, 113, 228]. It has been hypothesized [25, 106] that spin precession would not significantly alter the conclusions drawn using spin-aligned waveforms. We have tested this hypothesis and found, as first suggested by [72], that spin precession adds additional richness to the signals that almost completely breaks the mass-spin degeneracy, producing an order-of-magnitude improvement in the extraction of the individual masses and spins, which allows us to distinguish between NSs and BHs. We show that normal NS binaries will not be mistaken for BHs or exotic NSs, but we cannot rule out the possibility that some exotic NSs or low mass/low spin BHs may be misidentified as normal NSs.

Methodology

We employ Bayesian inference [1, 70, 140, 212] to quantify the astrophysical information can be extracted from a GW detection. In particular, when comparing models, we compute the *Bayes Factor*, which is the ratio of the evidence for one model to that for another. We compute BFs through MCMC techniques, as described by [1, 70, 140]. We consider only the inspiral phase, from 10 Hz up to 400 Hz, at which point NS tidal deformations can no longer be neglected; extending the analysis

beyond 400Hz would only strengthen the results obtained here. With these tools, and assuming a GW detection, we address the following questions:

1. Can we distinguish between NSNS binaries and low-mass, small-spin NSBH binaries only by the inspiral portion of the waveform?

2. Is the mass uncertainty large enough to lead to a false detection of astrophysically "exotic" NSs?

To answer these we need a waveform template that accurately models the GWs emitted during the quasicircular inspiral of spin-precessing, compact binaries. Previous studies were limited to spin-aligned or antialigned systems [25, 106], as until recently, these were the only systems for which fast, closed-form frequency domain waveforms were available (numerical time-domain templates are available, but their high computational cost limits their utility, see the discussion in [59]). However, we derived an analytical model for precessing systems in Chapter 4 of this dissertation using elements of MSA [62, 111, 124] (see also [144]) and the SPA method [84, 245], leading to small-spin *double-precessing* templates.

Each GW model incorporates spin effects in a different way and has, thus, a different spin prior. For the double-precessing model, we use uniform priors on the spin magnitudes and uniform priors on the unit sphere for the spin angles. For the spin-aligned model we again use uniform priors on the spin magnitudes, but delta functions about (anti)alignment with the orbital angular momentum for the spin angles. Clearly, the prior used in the double-precessing case is the most generic one since it assumes the least amount of prior information about the signal. Furthermore, a prior favoring spin alignment is not supported by astrophysical data. All the models use uniform priors on the masses.

Distinguishing between NSs and BHs

We simulated four *non-spinning* signals based on System 4 of Table 6.1 and recovered them with non-spinning, spin-aligned [9, 24, 132, 175], and the double-precessing models [62]. Figure 7.1 shows a 2D scatter plot of points in the (m_1, m_2) plane (with $m_1 \geq m_2$) that belong in the 90% probability quantile of the posterior distributions. The points are clustered along lines of constant chirp mass, with the extent of the lines determined by how well the dimensionless, symmetric mass ratio η is determined. For waveforms with spin, the degeneracy between spin and mass ratio enlarges the 90% confidence region.

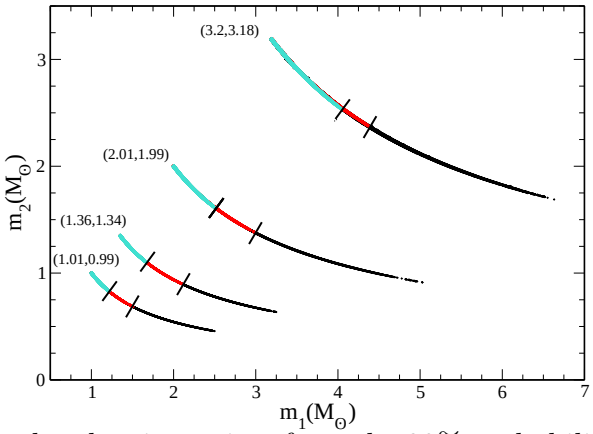

Figure 7.1: Scatter plot showing points from the 90% probability quantile in (m_1, m_2) for non-spinning signals with different masses of SNR 10 extracted with non-spinning (turquoise), spin-aligned (black), and double-precessing (red) templates. The injected masses are given on the top-left of each posterior. The posteriors overlap from the equal mass boundary to the short lines that cut across the scatter plots indicating the separation between the different posteriors in the direction orthogonal to the chirp mass. The use of double-precessing templates leads to more accurate mass extraction that spin-aligned templates.

How well the mass ratio can be measured depends on the particular model used. Non-spinning templates lead to the smallest spread in the recovered masses,

but at the cost of large systematic biases when one considers astrophysical realistic spin-precessing signals. Spin-aligned templates measure the mass ratio with a larger spread, due to degeneracies between masses and the spins. The inclusion of spin-precession partially breaks this degeneracy, translating into an improvement in the accuracy of the mass extraction that resembles what one would obtain with non-spinning templates. Similar results are shown in the left panel of Fig. 7.2 for the spin-prcessing System 3 in Table 6.1 at SNR 10. The right panel of Fig. 7.2 gives the 90% probability quantile in the (χ_m, m_1), where $\chi_m \equiv (m_1\hat{\chi}_1\cdot\hat{L}+m_2\hat{\chi}_2\cdot\hat{L})/(m_1+m_2)$, for one of the systems of the left panel. The results obtained for spin-precessing injections are stronger than for nonspinning ones.

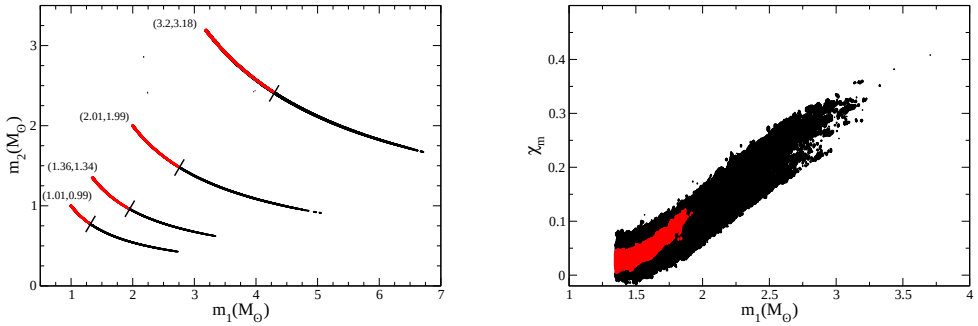

Figure 7.2: (Left panel) Same as Fig. 7.1 for System 3 of Table 6.1, but without the non-spinning template. (Right panel) Scatter plot showing the 90% probability quantile in the (χ_m, m_1) plane for the $(1.36, 1.34)M_\odot$ system from the left panel. The mass-spin correlation is far more pronounced for the spin-aligned model (black) than for the double-precessing (red) waveform model.

Figures 7.1 and 7.2 demonstrate that the use of the double-precessing model results in a significantly improved parameter extraction accuracy *even for nonspinning injections*. This might seem counterintuitive, because one may expect that increasing the dimensionality of the parameter space *without* increasing the complexity of the data (as is the case with the nonspinning injection) deteriorates measurement accuracy. Indeed, we find this to be the case when using the spin-aligned model to

extract a nonspinning signal; the two extra spin parameters, the spin magnitudes, introduce degeneracies with the individual masses that degrade the accuracy of mass extraction.

Following that reasoning, one may expect that the double-precessing model would perform even worse, since it has four more parameters than the spin-aligned model: the spin angles. This is *not* the case for the following reason. The spin angles offer the model more ways to leave the region of parameter space where the mass-spin degeneracies are more pronounced. When this occurs, the likelihood calculated between the signal and the double-precessing model deteriorates severely, leading to the rejection of the proposed jumps that have large masses.

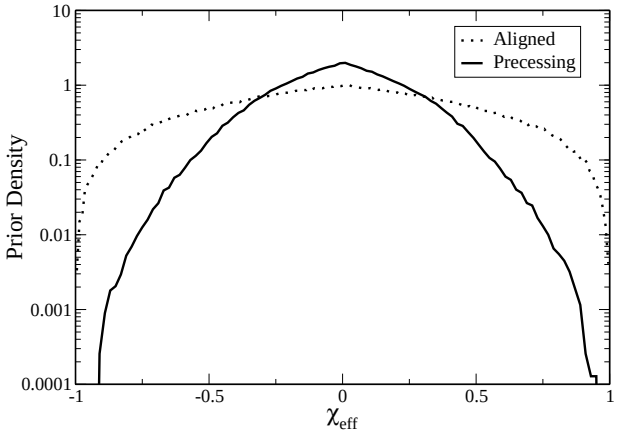

Figure 7.3: Prior distribution for the effective spin for the priors we have used in the two different models: uniform spin magnitudes and uniform priors on the unit sphere for all direction angles for the precessing model (solid line), and uniform spin magnitudes for the aligned model (dotted line). In the range of interest $[-0.5, 0.5]$ the priors differ by less than a factor of $\sim 3-4$, demonstrating that it is not a difference in priors that results in the increased measurement accuracy of the double-precessing model.

The tendency of the double-precessing model to leave the mass-spin degenerate region of parameter space is *not* a result of the choice of prior. In the spin-aligned case, we chose uniform priors on the spin magnitudes in the range $[0, 1]$. In the double-precessing model, we chose uniform priors on the spin magnitudes in the range $[0, 1]$ and uniform priors on the unit sphere for all direction angles. To determine the influence of these choices on our results we imposed the precessing $\chi_{\rm eff}$ prior on the spin-aligned model. We, indeed, found that the results of Fig. 7.2 are not noticeably modified, demonstrating that it is not the choice of prior that enhances the performance of the double-precessing model. This result should not be surprising if one compares the two priors. Figure 7.3 shows the prior distribution for $\chi_{\rm eff}$ for the two models considered here. The two priors on $\chi_{\rm eff}$ show the same qualitative behavior in the region $[-0.5, 0.5]$, the range of interest here (see also the right panel of Fig. 7.2).

If this behavior cannot be attributed to the prior, it must be attributed to the likelihood: the extra freedom in the spin orientation of the precessing model makes it less likely for systems with large masses or spins to match the signal. Figure 7.4 illustrates this through the dephasing between one of the systems of Fig. 7.1 and a system whose mass and spin magnitude are in the 90% probability quantile of the spin-aligned model but *not in that of the precessing one*, for different angles between the total spin and the orbital angular momentum. The dephasing induced by the spin-aligned model is below 1 radian for a wide range of frequencies $[20, 400]$Hz despite the high value of χ_2, a manifestation of the mass-spin degeneracy. On the other hand, the double-precessing model results in a large dephasing, and hence a low likelihood bad fit, which breaks the mass-spin degeneracy.

The fact that even a very small transverse spin can have such a big effect on data analysis can be understood as follows. Spin-alignment introduces a very strong

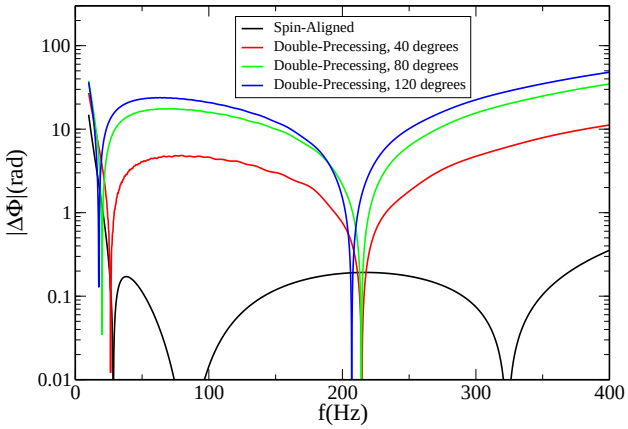

Figure 7.4: Phase difference between the nonspinning system $(m_1, m_2) = (1.36, 1.34)M_\odot$ of Fig. 7.1 and a system that belongs in the 90% probability quantile of the spin-aligned model with $(m_1, m_2) = (2.51, 0.79)M_\odot$ and $(\chi_1, \chi_2) = (0.04, 0.82)$ (black line). Keeping the masses and the spin magnitudes of the second system fixed, we misalign the spins and plot the phase difference between the initial nonspinning system and the new precessing system for 40° (red line), 80° (green line), and 120° (blue line) between the total spin and the orbital angular momentum at 10Hz.

correlation between the masses and the spins. As a result, the parameter covariance matrix is nearly singular. The near singularity of the covariance matrix means that very small changes in the waveforms can have significant effects on parameter estimation. Even the small amount of precession expected for NS binaries is sufficient to alter the mass-spin correlation and lead to very different parameter estimation results.

These results indicate that the increased accuracy we achieve with the double-precessing model is *not* a consequence of our choice of parameter priors, but rather it is due to the likelihood itself and its dependence on the precession features of the signal. The latter offer more ways for the double-precessing model to produce large mismatches with the injected signal, when it has to select the additional four spin angle parameters. As a consequence, the large-spin/large-mass points tend to

give lower likelihoods. The double-precessing model prefers to stay in the region of parameter space that fits the injected parameter, rather than wander off into these regions of lower likelihood. Effectively, in non-precessing models, a change in mass ratio results in a change of the rate of monotonic increase in the phase and the amplitude of the GW. The same can be achieved through a change in spin magnitude, due to the mass-spin degeneracy. On the other hand, precessional effects introduce phase and amplitude modulations that cannot be reproduced by a change in mass ratio.

Distinguishing between Normal and Exotic NSs

Imagine we have detected a GW produced by a NSNS binary. The double-precessing model has enabled us to correctly identify it as consisting of NSs. But, are the remaining parameter uncertainties enough to lead to an erroneous inference that we have detected a NS with parameters outside those expected from astrophysical models? We define a *normal* NS binary as one with $m_{1,2} \in [1, 2.5]M_{\odot}$ and $\chi_{\rm eff} \in [-0.05, 0.05]$ and an *exotic* NS as one that is not normal. We could have chosen different values for the boundaries in $m_{1,2}$ and $\chi_{\rm eff}$, but these are consistent with current astrophysical considerations, and the results would not qualitatively change if we picked other values. Notice that a $\chi_{\rm eff}$ in that range does not guarantee $\chi_{1,2} \leq 0.05$, due to the effect of the projection along the orbital angular momentum. Nonetheless, a detection of a system with $\chi_{\rm eff} \geq 0.05$ would unambiguously imply that the system possesses at least one $\chi \geq 0.05$.

Figure 7.5 shows the BF in favor of the exotic NS model for a precessing signal as a function of $\chi_{\rm eff}$ for different SNRs using the double-precessing model. Regardless of the SNR of the signal, a normal NS is always recovered as such. There exists, however, a window in parameter space (signals with SNRs of 10 and $\chi_{\rm eff} \in [0.05, 0.07]$) that

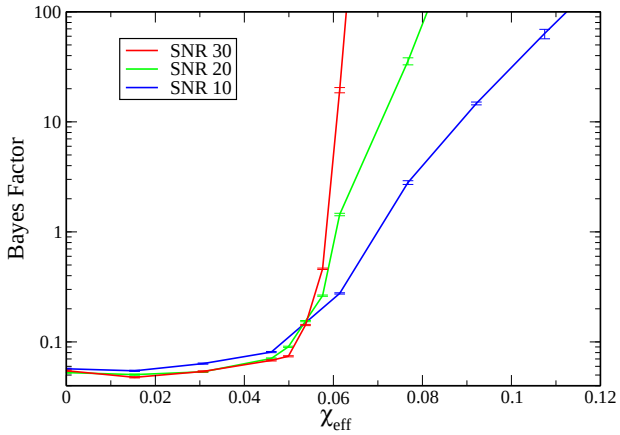

Figure 7.5: BF in favor of the exotic NS model as a function of the effective spin for different SNR values. The simulated signal is produced by System 1 of Table 6.1.

could lead to the characterization of the system as normal, when in reality it was exotic.

Conclusions

We showed that the inclusion of spin-precession in waveform templates breaks the degeneracy between the system's individual masses and spins, and allows us to distinguish between NSNS binaries and low-mass, small-spin NSBH or BHBH binaries. Moreover, even for signals with modest SNR, we can distinguish between "normal" and "exotic" NSs. These results open the door to population studies with the first GW detections, as well as coincident studies between the electromagnetic detection of short gamma-ray bursts and GWs. Indeed, if such a coincident observation is made, being able to identify the source from purely GW observations as a NS binary, a mixed binary or a BH binary would prove invaluable.

The results presented here are subject to several assumptions. First, the noise is assumed to be stationary and Gaussian, while in reality this may not be the case.

Proper noise modeling along the lines of [71, 141] will help to restore performance to levels close to the ideal case. Second, calibration errors and non-stationary drifts in the noise spectrum should be marginalized over, but these mostly impact the amplitude parameters, and only have a small impact on the spin measurement. Third, the waveform model inaccuracies do not affect our estimates of the statistical errors at leading order [59], so our conclusions will apply to more accurate waveform models.

This is not the first time that parameter extraction is greatly improved when a more detailed model is used. Initial studies on the projected bounds of the graviton mass and the Brans-Dicke coupling parameter [189,229,230,232] assumed nonspinning signals. The introduction of aligned spins in the models brought along degeneracies that degraded the bounds by about an order of magnitude [41]. However, Stavridis and Will [204] and Yagi and Tanaka [238] showed that the inclusion of precessional effects in the models can bring the bounds back to almost their nonspinning values.

PROBING THE INTERNAL COMPOSITION OF NSS WITH GWS

GWs from NSNS inspirals contain information about the as yet unknown equation of state (EoS) of supranuclear matter. In the absence of definitive experimental evidence that determines the correct EoS, a number of diverse models that give the pressure inside a NS as function of its density have been constructed by nuclear physicists. These models differ not only in the approximations and techniques they employ to solve the many-body Schrödinger equation, but also in the internal NS composition they assume. We study whether GW observations of NSNS binaries in quasicircular inspirals up to contact will allow us to distinguish between EoSs of differing internal composition, thereby providing important information about the properties and behavior of extremely high density matter. We carry out a Bayesian model selection analysis, and find that second generation gravitational wave detectors can heavily constrain EoSs that contain only quark matter, but hybrid stars containing both normal and quark matter are typically harder to distinguish from normal matter stars. A gravitational wave detection with a SNR of 20 and masses around $1.4M_{\odot}$ would provide indications of the existence or absence of strange quark stars, while a SNR 30 detection could either detect or rule out strange quark stars with a 20 to 1 confidence. The presence of kaon condensates or hyperons in neutron star inner cores cannot be easily confirmed. For example, for the EoSs studied here, even a GW signal with a SNR as high as 60 would not allow us to claim a detection of kaon condensates or hyperons with confidence greater than 5 to 1. On the other hand, if kaon condensates and hyperons do not form in NSs, a GW signal with similar SNR would be able to constrain their existence with an 11 to 1 confidence for high-mass systems.

Introduction

With the main goal of the advanced detectors already achieved, the scientific community has started shifting focus to an even more interesting endeavor: What can we infer about Nature from GW observations?

One interesting possibility is linked to finite-size effects that NSs experience because they are extended bodies with structure. When objects with a finite size are subjected to the tidal field of another object their multipole moments are affected in a way that depends on the EoS – for barotropic fluids, a relation between pressure and density– of their matter. The densities encountered in NS interiors are extremely high; in the inner core they even exceed nuclear densities. In this high density regime, laboratory experiments and observations have still to provide a definitive answer on the correct EoS. We, therefore, study whether GWs can be used to answer the following question: Given the detection of the NSNS quasicircular inspiral, can we use finite-size effects to learn about the EoS of the extremely dense NS interior [78, 88, 89, 95, 103, 112, 134, 233]?

To leading order, finite-size effects cause the quadrupole moment tensor of a star Q_{ij} to be affected by the tidal field tensor of its companion $\mathcal{E}_{ij}$ through $Q_{ij} = -\lambda \mathcal{E}_{ij}$. The constant of proportionality λ is called the tidal deformability and it is a function of the mass and the EoS. This tidal interaction causes NSs to be distorted during the inspiral phase and torn apart before merger. By the 'inspiral phase' we define the evolution of the binary up to an orbital separation of six times the total mass or contact, whichever comes first. The plunge and final collision after this orbital separation is called the 'merger phase'. Since this merger is expected to happen at high frequencies of $\mathcal{O}(10^3)$Hz, where the detector noise is likely to dominate, we

here focus on the better modeled inspiral part[1]. The PN waveform, in fact, becomes less accurate near the merger [93, 225, 241]. For nonspinning NS binaries, a more accurate waveform is available using the effective-one-body (EOB) approach [40]. In this study, we consider precessing NS binaries as we will explain in more detail below. Since the EOB waveform for precessing NS binaries is currently unavailable, we use the double-precessing PN waveform of Chapter 4.

In the PN framework, the first finite-size effect enters the waveform at 2PN order in the GW phase through spin corrections to the quadrupole moment of the objects $Q_{1,2}$ [172], and here we also include the 1PN correction to this [235]. Then, at 5PN order and above, the tidal deformabilities $\lambda_{1,2}$ enter the phase directly [78, 93, 95, 112, 222, 241]. All these parameters are EoS-dependent, however, two of the authors showed that their interrelation is approximately EoS-independent [239, 240]. A lot of work has been put into understanding and extending this result [29, 56, 57, 81–83, 109, 123, 133, 145, 151, 152, 168, 170, 195, 237, 242, 243], in particular to include higher-order multipole moments and tidal deformabilities [65, 170, 205, 234, 236]. We can therefore use the Love-Q relation to choose $Q_{1,2}$ in favor of $\lambda_{1,2}$ or vice versa from the GW phase.

The problem of the detectability of finite size effects with gravitational waves has gathered a lot of attention in recent years. Initial studies, based on quantifying the differences between waveforms with different EoSs or on a Fisher information matrix analysis [78, 112, 127, 128, 153, 169, 179, 180] suggested that aLIGO has the potential of providing useful information on the NS EoS. However, due to the expected low signal-to-noise ratio (SNR) in aLIGO detections and due to the strong correlations

[1] A number of studies have examined the possibility of determining the EoS from the merger phase (see [128] for an example). However, in the absence of a full and accurate template bank of merger waveforms, it is not clear how one would perform a full data analysis study.

between the different GW parameters, the applicability of a Fisher study, and the conclusions derived from it, is limited [66, 162, 183, 215].

For this reason, several Markov-Chain Monte-Carlo (MCMC) [70, 140] studies have recently been carried out in order to address EoS detectability in a more robust way in the context of Bayesian inference. The first such study was performed by Del Pozzo et al. [79], which showed that a few tens of detections of moderate brightness can be combined to provide strong constraints on the EoS, though this result seems to be highly dependent on the mass distribution of the sources [8]. Wade et al. [225] studied the effects of systematic and statistical errors in EoS extraction, while Lackey and Wade [130] employed a more realistic parametrization of the EoS and agreed that a few bright sources can determine the EoS of NSs.

With the exception of [79] and [8], all previous work mentioned above consisted of *parameter estimation studies*, where the tidal deformability is treated as a system parameter and searched over with an MCMC analysis. Any EoS that predicts a value of the deformability within the recovered uncertainty is compatible with the results of the MCMC. Any EoS that does not fall in the deformability error bars can be ruled out. We choose a more direct approach here and compare the different EoS models directly, an approach known as *model selection*. In the latter, the EoS itself, instead of the tidal deformability is treated as an independent parameter of the system and the analysis allows the data to select which EoS is preferred. Our study is unique in that we use this tool to perform a comprehensive study on whether *we can extract important physical information about the composition of NSs, such as the existence of exotic species*.

The different EoS models proposed in the literature differ not only in the NS composition they assume, but also in the approximate schemes they employ to solve the many body Schrödinger equation. These approximate schemes include approaches

such as the variational method [167], Skyrme-Hartree-Fock (SHF) models [202], Brueckner-Hartree-Fock (BHF) models [86], and relativistic mean field (RMF) theory [197, 198]. The NS internal composition might be that the EoS is constructed solely with normal matter (neutrons n, protons p, electrons e, muons μ), or it may contain kaon condensates (K), hyperons (H), pion condensates (π), or quark matter (Q).

Determining that 2 EoSs with the *same* composition but *different* approximations are distinguishable will not result in any new information about Nature. On the other hand, determining that 2 EoSs with *different* compositions and *different* approximations are distinguishable must be treated with caution. Can we claim that the difference between the EoSs we detected is due to their actual physical differences, or due to their distinct mathematical approaches? In order to avoid this obstacle, we compare EoSs that employ the *same*, or as similar as possible, approximations and differ only in their internal composition.

For example, consider an EoS, which contains normal matter, and is constructed with the variational method[2], and another EoS, which is also constructed with the variational method, but includes both normal matter and hyperons. When we perform a comparison on these 2 EoSs, we are essentially comparing a hyperon EoS and a hyperonless EoS. The result of this comparison can directly be translated to physical information. If the hyperon EoS is preferred, we have detected hyperons in NS cores assuming one of the 2 EoSs is correct. If the hyperonless EoS is preferred, we have constrained hyperons in NS cores.

At this stage one may reasonably object: Could we confuse a hyperon EoS constructed with one method with a normal matter EoS constructed with a different

[2]Refer to Appendix A of [64] for a description of all EoSs used here along with their physical content and approximation schemes.

method? This concern can be alleviated by performing a large number of comparisons. Concluding that a hyperon EoS constructed with one method is preferred over a normal matter EoS with the same method is not enough. We need to find as many pairs of EoSs that have been constructed with a variety of different methods and compare all of them. This approach (i) ensures we have not confused the effects of internal composition and approximation schemes, and (ii) provides us with insight on how hyperons, or other particles, affect EoSs in general. Effectively, by comparing many pairs of EoSs, we 'average out' any effect coming from how each pair of EoSs is constructed and isolate the effect of the common difference between the 2 EoSs of all pairs i.e. the exotic matter.

Another aspect in which our analysis generalizes previous studies is that for the first time we include spin-precession both in our simulated signal and in the templates. We use the fully analytic double-precession model, derived under the assumption that the binary components have small spins. This turns out to be an excellent approximation for NSs in the LIGO band, as their dimensionless spin parameter (the spin angular momentum over the mass squared) is not expected to exceed ~ 0.1 [147].

The inclusion of spin-precession in the templates is crucial.In Chapter 7 we showed that allowing for the systems to precess around the orbital angular momentum can break degeneracies between the masses and the spins, improving mass extraction by about an order of magnitude. This improved mass determination is directly translated to better λ extraction, making EoSs easier to distinguish. This is because in the context of model selection, it is the EoS that is a GW parameter, and not λ, which is determined through a relation of the form $\lambda(m, \mathrm{EoS})$.

The main results of our analysis are summarized below.

We find that advanced detectors will be able to place strong constraints on the existence of quark stars comprised solely of quark matter. Given the EoSs available

today, a NSNS binary with masses in the $(1.2, 1.5)M_\odot$ range with SNR = 30 can effectively rule out strange quark stars, or make a positive detection of them. We, furthermore, argue that even for the plausible EoSs constructed in the future, there exists some mass $\in (1M_\odot, 1.8M_\odot)$ where quark stars are distinguishable from normal matter NSs.

The prospects of detecting or ruling out hybrid NSs including both normal and quark matter are worse. If the strong interactions between quarks are close to those predicted by a perturbative analysis [96] and the transition between nuclear and quark matter phases happens around twice the nuclear saturation density, the detection of a $(1.4, 1.35)M_\odot$ NS binary with SNR $30 - 40$ could provide significant evidence of whether quarks form in NS interiors. However, if the strong interactions are weaker, aLIGO will not be able to reach confident conclusions.

It is unlikely that aLIGO will be able to claim a detection of hyperons or kaons, since that would require high masses and SNR $\gtrsim 60$. The detection of hyperon or kaon condensates in NS interiors requires high mass stars, since it is only at these high masses that you encounter densities large enough for these condensates to form. This poses a significant problem; the importance of finite size effects is reduced with increasing mass since λ decreases with increasing mass. Moreover, most NSs are expected to have masses around $\sim 1.4M_\odot$, rarely reaching the $2M_\odot$ required for hyperons and kaons detection. We therefore conclude that it is unlikely that aLIGO will be able to positively identify hyperons or kaons in NSs. On the other hand, if hyperons and kaons are *not* formed in NS interiors, aLIGO could place constraints on their existence.

Our analysis suggests that aLIGO can distinguish between models that differ at low central densities, like normal matter EoSs and EoSs containing quark matter. In order to probe the high density regime we need SNRs higher than what aLIGO is

likely to achieve. This is due to the fact that high mass systems (i) present smaller finite size effects, and (ii) have masses close to the maximum mass allowed, which causes some interesting effects related to the prior boundary.

Among the various noise configurations aLIGO can be tuned to, the optimal for EoS determination is the default Zero-Detuned, High-Power (Zero-Det., High-P) one [201]. Tuned configurations include the NSNS Optimized (NSNS Opt.) and the High Frequency (High F.) ones; both perform in an inferior way when it comes to EoS extraction. In the case of NSNS Opt. this is due to its low sensitivity at frequencies above 600Hz, when it is exactly at these frequencies that finite size effects are more prominent. As far as High F. is concerned, its improved sensitivity is limited to a very narrow frequency range around 10^3Hz. This fact coupled to its worsened sensitivity at low frequencies makes High F. unsuitable for EoS studies. We conclude that in the high frequency regime we are interested in, the Zero-Det., High-P. noise curve has the overall higher sensitivity and more accumulated SNR.

Stacking, i.e. combining multiple low SNR sources, might improve the results obtained here, but it might also lead us to erroneous conclusions about the true EoS. We find that when the SNR is low our results can be dominated by prior information, rather than any new information we get from the GW data. This could lead to each individual low SNR binary system providing some confidence, but in favor of the wrong EoS. Stacking all these systems will inevitably lead to great confidence in favor of the wrong EoS. We emphasize that stacking must always be treated with caution.

As far as the spin of the bodies is concerned, we find that the magnitude of the injected spin angular momentum has a negligible effect on our results, provided that the templates allow for spin precession. This is in agreement with Ref. [60], where the order-of-magnitude improvement in mass extraction was achieved over spin-aligned

templates even for nonspinning systems, as long as spin-precessing templates were used in the recovery of the signal.

Model Selection

Model selection in the Bayesian framework requires an explicit statement of the models compared. In this section, we describe in detail the models we use, as well as our methodology when comparing them. We describe the *Bayes factor*, a quantity that assesses which model is preferred by the data, and give an overview of our simulated signals. We conclude this section with a brief discussion of the power of stacking signals versus detecting a single loud signal.

Bayesian Inference

In the context of Bayesian inference, the probability that a hypothesis H_1 is correct given some data d is [70, 140]

$$p(H_1|d) = \frac{p(H_1)p(d|H_1)}{p(d)}, \tag{8.1}$$

where $p(H_1|d)$ is the posterior belief in the hypothesis after the data has been analyzed, $p(H_1)$ is the prior belief based on all information we have before analyzing the data, and $p(d)$ is the probability of the data, an unimportant normalization constant in our case. The evidence $p(d|H_1)$ is given by an integral over the parameters of the model $\boldsymbol{\theta}$

$$p(d|H_1) = \int d\boldsymbol{\theta}\; p(\boldsymbol{\theta}|H_1)\; p(d|\boldsymbol{\theta}H_1), \tag{8.2}$$

where $p(\boldsymbol{\theta}|H_1)$ is the prior information on the model parameters, and $p(d|\boldsymbol{\theta}H_1)$ is the likelihood, where $\ln p(d|\boldsymbol{\theta}H_1) = -1/2\,(\,s-h|\,s-h)$ in Gaussian noise, with s the signal, h the template model and $(\cdot|\,\cdot)$ the noise-weighted inner-product [70, 140].

When we have to select between 2 competing hypotheses, we compare their posterior beliefs through the *odds ratio* (OR), defined by

$$\mathcal{O} = \frac{p(H_1|d)}{p(H_2|d)} = \frac{p(H_1)p(d|H_1)}{p(H_2)p(d|H_2)}. \tag{8.3}$$

The OR is the 'betting odds' of H_1 compared to H_2 and includes the both prior belief in each hypothesis and any new information that is extracted from the data. These two contributions can be separated by defining the *Bayes factor*

$$\mathrm{BF} = \frac{p(d|H_1)}{p(d|H_2)}, \tag{8.4}$$

which includes only the data contribution to the OR. A BF > 1 means that H_1 is supported better by the data, while a BF < 1 means that H_2 is preferred. In the case of uninformative priors, i.e. $p(H_1) = p(H_2)$, the OR equals the BF. For this study, we choose to work with the BF instead of the OR because we are interested in whether the data lends more support to some hypothesis over another, irrespective of our prior beliefs in them. Of course, once we can confidently quantify our prior belief in a hypothesis, we can trivially go from the BF to the OR.

Working with the BF, however, has one major problem: we cannot draw the same conclusions about two different pairs of hypotheses that have the same BF. For example, consider the problem of whether a given GW signal is better described by GR or by a modified gravity theory. In this case, we have a strong prior belief in favor of GR, given the many successes of experimental relativity [231,246]. We would therefore require very large BFs in favor of the modified gravity hypothesis to claim a detection of a deviation from GR. On the other hand, if we are interested in whether the signal is better described by one of two competing EoSs, we do not have such

strong prior beliefs in favor of any of the models. That means that we would not need such large BFs to claim that we have identified the correct EoS of Nature, provided we remain agnostic about the two EoSs.

For this reason, it is crucial that we explicitly define what we mean by "a BF that is large enough" on a problem by problem basis. As we will discuss later in this section, we work with models for which we do not have strong experimental prior knowledge. Therefore, we adopt the *Jeffreys scale of interpretation of BFs* [117] to define how significant a BF is. When BF < 1 it is *negative*, for $1 <$ BF < 3 it is *barely worth mentioning*, for $3 <$ BF < 10 it is *strong*, for $10 <$ BF < 100 it is *very strong*, and finally for BF > 100 it is *decisive*. This is in contrast to the analysis of Refs. [69, 185–187], which dealt with tests of GR, where BFs around 100 were considered strong and not decisive, given the strong prior in favor of GR.

In order to calculate the BF between two models, a number of different techniques can be used: thermodynamic integration [140, 220], nested sampling [220] and reverse jump, Markov chain-Monte Carlo [70]. Here, we employ the third technique which requires promoting the model to a parameter of the Markov chains. Then, the BF is given by

$$\mathrm{BF} = \frac{\text{time the chains spend in model 1}}{\text{time the chains spend in model 2}}, \tag{8.5}$$

with error bars calculated with the technique suggested in [71].

Models

Our models represent GWs emitted in the late inspiral of NS binaries assuming GR is correct from the time they enter the LIGO band until the NSs come into contact, or their separation becomes six times the total mass, whichever comes first. The

difference in the models will only be in the finite size effects they include. Comparison between these models can be viewed as traditional parameter estimation, only now the model itself is an extra discrete parameter. In parameter estimation, the waveform of a NS binary inspiral with a circular orbit depends on the following parameters[3]:

$$\boldsymbol{\theta}_{\rm PE} = \{m_1, m_2, \theta_N, \phi_N, D_L, \theta_L, \phi_L, t_c, \phi_c, \boldsymbol{S}_1, \boldsymbol{S}_2, \bar{\lambda}_1, \bar{\lambda}_2\}, \tag{8.6}$$

while in a model selection study the template depends on the parameters

$$\boldsymbol{\theta}_{\rm MS} = \{m_1, m_2, \theta_N, \phi_N, D_L, \theta_L, \phi_L, t_c, \phi_c, \boldsymbol{S}_1, \boldsymbol{S}_2, {\rm EoS}\}, \tag{8.7}$$

where m_i are the component masses, (θ_N, ϕ_N) give the sky location of the source, D_L is the luminosity distance, (θ_L, ϕ_L) give the direction of the orbital angular momentum, (t_c, ϕ_c) are the time and phase of coalescence respectively, $\boldsymbol{S}_i \equiv \chi_i m_i^2 (\sin\theta_i \cos\phi_i, \sin\theta_i \sin\phi_i, \cos\theta_i)$ are the spin angular momentum vectors of each binary component, with $\chi_i := |\boldsymbol{S}_i|/m_i^2$ the dimensionless spin parameter, and $\bar{\lambda}_i \equiv \lambda_i/m_i^5$ are the dimensionless tidal deformabilities. The two parameter sets are equivalent, since knowledge of $(m_1, m_2, {\rm EoS})$ from $\boldsymbol{\theta}_{\rm MS}$ can be used to construct the quantities $\bar{\lambda}_1(m_1, {\rm EoS})$ and $\bar{\lambda}_2(m_2, {\rm EoS})$ in $\boldsymbol{\theta}_{\rm PE}$.

The EoS is what *defines* a model. For some EoS the model predicts a GW with $\boldsymbol{\theta}_{\rm PE}$ and $\bar{\lambda}_{1,2}$ as a function of $(m_{1,2}, {\rm EoS})$ for $(m_1, m_2) \leq M_{\rm max}({\rm EoS})$, and no GW otherwise. The quantity $M_{\rm max}({\rm EoS})$ is the maximum NS mass that can be stably supported, given an EoS. For the GW template itself, we will use the small-spins, double-precessing waveform. We use an extended version of these templates by

[3] In this argument, we encode the EoS dependence through the dimensionless tidal deformabilities $\bar{\lambda}_{1,2}$; the argument goes through if one uses any other quantity to parametrize tidal deformations.

EoS	Composition
AP4 [12], SV [138], SGI [138], SkI4 [181] DBHF$^{(2)}$(A) [120], MPa [244], G4 [129] GA-FSU2.1 [105], Shen [197, 198]	n, p, e, μ
SGI-YBZ6-SΛΛ3 [137], NlY5KK* [120], SkI4-YBZ6-SΛΛ3 [137] MPaH [244], H4 [129]	n, p, e, μ, H
SGI178 [138] SV222 [138] GA-FSU2.1-180 [105]	n, p, e, μ, K
ALF4 [13], ALF5	n, p, e, μ, π, Q
SQM3 [176]	Q (u, d, s)

Table 8.1: Classification of EoSs with respect to internal composition. The first cluster corresponds to EoSs with normal matter. The second and third clusters include hyperons and kaon condensates respectively. The last two rows list EoSs that include quark matter.

adding finite-size effects at 2 and 3PN order (due to the quadrupole moment) [172,235] and at 5 PN order and higher (due to tidal deformability effects) [78, 95, 112, 222]. More specifically, we include tidal terms that depend on the $\ell = 2$ electric tidal deformability λ up to 7.5PN order given in [78]. We also include the contribution from the $\ell = 2$ magnetic tidal deformability σ_2 at 6PN order and the $\ell = 3$ electric tidal deformability λ_3 at 7PN order [234]. We further take into account the leading correction to the adiabatic approximation entering at 8PN order that depends on the $\ell = 2$ f-mode frequency f_2 of a NS [95, 112]. σ_2, λ_3 and f_2 can approximately be expressed in terms of λ thanks to the universal relations found in [58, 234]. In addition, we use the Love-Q relations [239, 240] to rewrite the quadrupole moment in terms of the tidal deformability.

Which EoS models should we allow the data to select from? There is a great number of EoSs available in the literature, varying both in the type of matter they consider (quarks, hyperons, kaons, muons, pions, neutrons, protons, electrons), and in the approximation schemes used to construct the EoS. We can classify the EoSs by the type of matter they employ; the subset of EoSs discussed in this study are presented in Table 8.1. The EoSs in each of these categories differ in the approximation schemes used to solve the many-body Schrödinger equation.

Given the great variety of EoSs, how should we carry out a model selection study? Could we select a characteristic set of models, perhaps chosen by looking at how the $\bar{\lambda} - m$ relations behave for a set of EoSs? Figure 8.1 shows these relations for a few of the EoS listed above. Notice how the normal matter EoSs (AP4, Shen) form a band in the $\bar{\lambda} - m$ space that contains both SQM3 (consists of pure quark matter) and H4 (contains hyperons). This *lack of clustering* with respect to the internal composition classification in the $\bar{\lambda} - m$ space indicates that comparing arbitrary EoSs will not produce physically meaningful results. For example, say we concluded that AP4 (red solid line) is distinguishable from H4 (black, dot-dashed line). Can we claim that hyperons are detectable from GW observations? The answer is no; there is another normal matter EoS (Shen, turquoise dotted line) much closer to H4 than AP4. Unless we compare Shen to H4 too, we cannot claim detectability of hyperons.

Another possibility would be to compare each model against every other model *simultaneously*, but (i) the computational cost would be prohibitive, and (ii) it is not clear what physical question such a comparison would address. Comparing models within each category that differ only in the approximation scheme used to calculate the EoS is analogous to comparing post-Newtonian models for the GWs emitted in the inspiral of compact objects. In the end, all of such models are only approximations to the exact solution of Nature. Instead, we are interested in specific physical questions

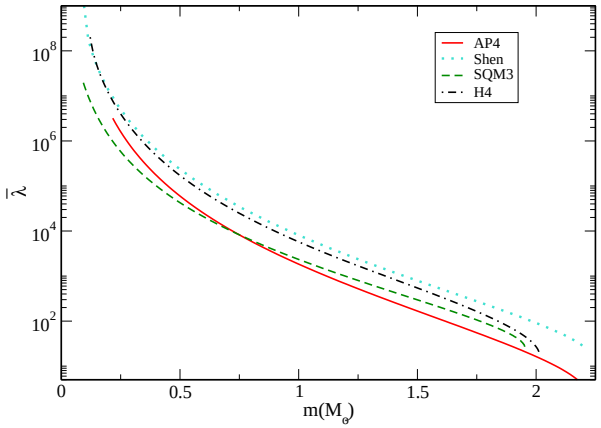

Figure 8.1: Dimensionless tidal deformability as a function of mass for a number of EoSs with very different physical contents. The lack of clustering in the $\bar{\lambda} - m$ space shows that we cannot perform model selection with all EoSs simultaneously.

regarding the internal composition of NSs. To do so, we will choose pairs of EoSs that are as similar as possible (using the same approximation schemes), but differ only in the matter degrees of freedom. In this way, we can directly translate *distinguishability between models* into *distinguishability between physical scenarios*.

Signal Injections

Whether the data can distinguish a given EoS model over another depends strongly on the particular signal detected. The parameters that affect EoS distinguishability the most are the two masses, the distance to the source, which effectively controls the SNR, and the EoS itself. We therefore simulate different signals with various values for these parameters. For the remaining parameters we select the following injection values: $(\cos\theta_N, \phi_N) = (-0.105, 3.705)$, $(\cos\theta_L, \phi_L) = (0.801, 3.216)$, $(t_c, \phi_c) = (1,024\,\mathrm{s}, 4.461)$, $(\cos\theta_1, \phi_1) = (0.774, 2.248)$, $(\cos\theta_2, \phi_2) = (0.968, 5.311)$, and $(\chi_1, \chi_2) = (0.04, 0.04)$. All parameters have been randomly chosen so that they do not lead to any 'special' orientation of the binary (face on, edge on

etc.). We have also performed simulations with other dimensionless spin magnitudes and found that the spin has a very small effect on EoS distinguishability.

The initial misalignment of the spin and the orbital angular momenta means that the system will undergo precession [21]. Indeed, with these choices of parameters, the angle between the orbital angular momentum and the total spin angular momentum is $\sim 30^\circ$. Spin precession is modeled through the small-spins, double-precession approach of Chapter 4, which has been shown to be highly accurate for modeling NS binaries in Chapter 6. This approach is valid in the inspiral phase only, since it is based on a post-Newtonian expansion. The signal is thus modeled with the same approach as the templates that define the EoS models.

The approximations used to describe the orbital motion in the small-spins, double-precession approach are only valid up to a given frequency. We thus carry out our analysis up to $\min(f_{\rm ISCO}, f_{\rm c}(\rm EoS))$, where $f_{\rm ISCO}$ is the Keplerian frequency at $r = 6(m_1 + m_2)$ [35, 213] and $f_{\rm c}$ is the *contact frequency*, that at which the separation of the two bodies is equal to the sum of their radii. Although Mandel et al. [146] showed that terminating the waveforms at a certain frequency can affect the results through the addition of artificial information, our cutoffs are at such high frequencies that they are not expected to affect the results.

Even though the particular noise realization in the detector at the time the GW passes through will have an effect on parameter extraction, we do not inject noise in our analysis. Given that it is impossible to predict the noise instance, the best we can do is average our results over multiple noise realizations. However, Nissanke et al. [160] showed that such averaging is equivalent to zero injected noise in the signal. Sampson et al. [185] showed that a given noise realization causes the likelihood to shift as a whole, without significantly changing its shape (Fig. 4 of [185]), suggesting that the integral of the likelihood (the evidence) is minimally affected by noise fluctuations.

This picture, however, does not hold for poorly constrained parameters, like the ones studied here, where the posterior extends over a large fraction of the prior volume [215, 225].

Of course, the noise curve of the detector does play a very important role in the calculation of BFs, through the noise-weighted inner product in the likelihood [70]. We will mostly adopt the zero-detuned, high-power noise spectral density of the advanced detectors [201], though we will explore other choices later. The specific form of the likelihood we use assumes that the noise in the detectors is stationary and Gaussian [see below Eq. (8.2)], neither of which is strictly true. Cornish and Littenberg, however, have shown how to model the non-Gaussian features [71], and deal with the nonstationarity of the noise [141], leaving us with only the modeling of the stationary, Gaussian noise component.

When recovering the parameters, we use a uniform prior in $(0.1, 3)M_{\odot}$ for the masses, a uniform prior on the sphere for all directions, a uniform prior in $(0, 1)$ for the dimensionless spin magnitudes, and a uniform prior in the log of the distance. All prior ranges are selected such that they are wide enough to not affect our results.

Stacking vs High SNR

In reality aLIGO will probably reach physically interesting conclusions by combining information from multiple detections, rather than by waiting for a very loud one. Along those lines, one would argue it makes more sense to stack a sufficient number of moderate SNR sources rather than study the BF as a function of the SNR. Our results, however, suggest that stacking should be performed with caution.

For low SNR detections we find 2 rather counterintuitive effects: (i) it is possible for the wrong model to be preferred over the correct one, and (ii) it is possible for the correct model to be preferred less and less as the SNR increases (see Appendix D).

The first effect is not new; it has already been encountered in the context of comparing models with different dimensionality [69], where it takes the form of an *Occam Penalty* on the more complicated model. The second effect is perhaps less familiar: why would it be that, as the signal strength increases, the data fails to increasingly support the correct model? The answer to this question and the root cause of these effects can be traced down to sharp cutoffs of the prior distribution.

When computing BFs, we must compare models with different maximum masses. In fact, most of the time the difference is rather large, with one model allowing NSs up to $\sim 2M_{\odot}$, while the competing model going up to $\sim 2.5M_{\odot}$. When the injected mass is close to the maximum allowed mass and the SNR is sufficiently low, the posterior distribution is affected by this cutoff. In order to understand and visualize this effect, we construct a simple $1-$D model in Appendix E. In the context of this simple model we explain both the effect on the BFs that favor the wrong model, and the BFs that decrease with increasing SNR.

These results show that stacking many weak sources is *not* necessarily equivalent to a single bright source. When each new observation is informative, i.e. the likelihood dominates over the prior, then the data will prefer the correct model and BF > 1. In that case, adding the extracted information from this observation will push the analysis in the right direction, and eventually, we will recover the same results as from a single loud event. On the other hand, if the observation is not informative and the result is prior dominated rather than likelihood dominated, the wrong model might be preferred and BF < 1. In that case, adding this observation in the stack will lead the analysis in the wrong direction: a large number of weak observations that favor the wrong model will build confidence in the wrong conclusion.

The above results suggest that stacking must be treated with caution. For sufficiently high SNR events, the posterior will be narrow enough that it will not

be affected by the maximum mass cutoff. In this case, one recovers the expected result: the correct model is preferred and it is preferred more and more as the signal strength increases. However, if the signals are of lower SNRs, the observations may not be informative, and then, the final cumulative result may be largely influenced by the prior and not by the new information contained in the signals.

Comparing EoSs

Our goal is to study whether GW detections of inspiraling NSs can be used to learn about their interior composition and in particular, whether they contain kaon condensates, hyperons, and quark matter. To do so, we need to isolate their respective effects in the EoSs. We accomplish this by comparing pairs of EoS models that are as similar as possible, but differ only in the inclusion of one of these particles. Appendix A of [64] provides a comprehensive classification of all the EoSs we use in this section.

Kaon Condensates

We first address the question of whether NSs with kaon condensates in their inner cores leave an observable signature on inspiral GWs. To do so, we choose 3 pairs of EoSs constructed with (i) the Skyrme-Hartree-Fock (SHF) scheme, (ii) relativistic mean field (RMF) theory, and (iii) the SHF scheme including three-nucleon interactions (TNI). Each pair consists of one EoS with a kaon condensate and one without.

Kaon condensates could emerge in stars at high central densities and, therefore, kaon models differ from kaonless ones only for NSs with sufficiently high masses. This makes the extraction of physical information from these systems more difficult than that for low-mass systems for 3 reasons. From a data analysis point of view, systems with masses close to the maximum allowed mass will suffer from the edge

effects described in Appendix D, making it more difficult to get likelihood-dominated results. From a physical point of view, NSs with high masses have smaller values of $\bar{\lambda}$, as seen in Fig. 8.1, making finite-size effects less relevant in the GW phase. From an astrophysical point of view, NSs with masses around $2M_{\odot}$ are expected be rare.

As we will show below, we find that if kaon condensates indeed form in the inner core of NSs, they will be hard to detect with the signals expected from aLIGO. On the other hand, if condensates do not form in NSs, then loud aLIGO signals may be able to establish this with a high mass detection.

Models The 3 pairs of models we compare are constructed with different approximations.

1. *SHF models* [138]: SV(kaonless)/ SV222 (with kaons).

2. *RMF theory* [105]: GA-FSU2.1 (kaonless)/ GA-FSU2.1-180 (with kaons).

3. *SHF models + TNI* [138]: SGI (kaonless)/ SGI178 (with kaons).

The results of comparison (3) are very similar to the results of comparison (2), so we will only present pairs (1) and (2). Figure 8.2 shows the mass-radius relation for the EoSs we present. This figure suggests that kaons affect the properties of only the most massive NSs. The pairs of circles, triangles, and squares indicate the values of the masses in the injected signals of our analysis.

Bayes Factor Figure 8.3 shows the BF in favor of the kaon model of each pair (left panel) and the kaonless model of each pair (right panel) as a function of the SNR of the signal. The different lines correspond to the different injected masses indicated by the blue symbols in Fig. 8.2. We always plot the BF in favor of the correct (injected) model. Any BF > 1 in the left panel means that we correctly detected

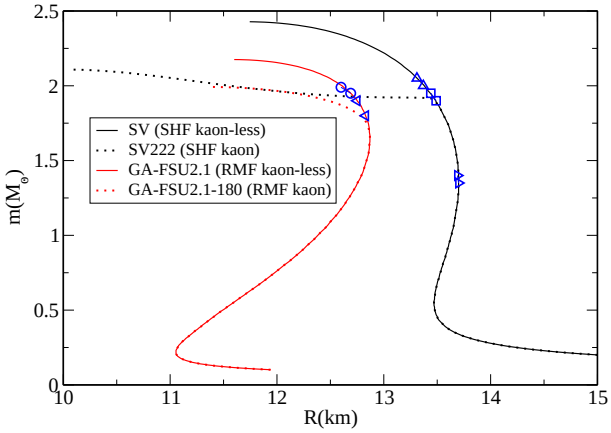

Figure 8.2: m-R relation for the EoS pairs that test kaons: The black lines correspond to SV (solid) and SV222 (dotted), which are constructed through SHF models. The red lines are for GA-FSU2.1 (solid) and GA-FSU2.1-180 (dotted), which are constructed through RMF theory. The presence of kaons causes the kaon model of each pair to differ from the normal matter model for high masses. The pairs of similar blue symbols indicate the component masses we use in our analyses.

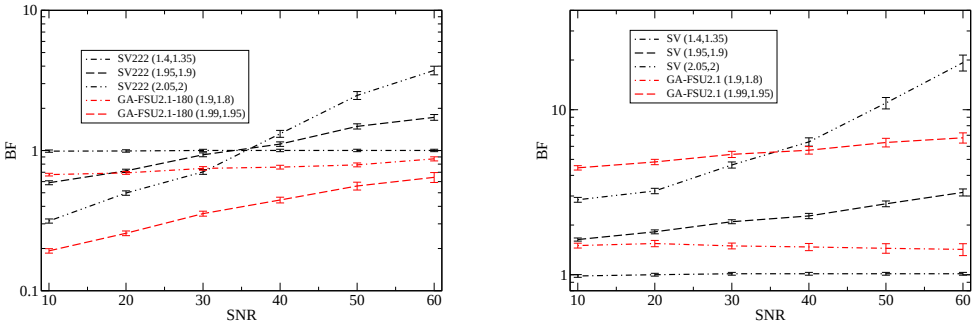

Figure 8.3: BF in favor of the kaon models (left) and the kaonless models (right) as a function of the SNR for different injected masses given in the brackets. BFs are always quoted in favor of the correct (injected) model. For example, the black dotted-dashed curve labeled "SV222 (1.4,1.35)" in the left panel means that the injected model is the SV222 EoS with the NS mass $(1.4, 1.35)M_\odot$, and $\mathcal{H}_1 =$ SV222 and $\mathcal{H}_2 =$ SV in Eq. (8.4). BF > 1 in the left panel means that we correctly detected the presence of kaons inside the NSs, while that in the right panel means that we correctly ruled out such a presence of kaons. The kaonless models give higher BFs than the kaon models, making it easier to establish that NSs do not have kaon condensates than the opposite.

the presence of kaons in the interior of the NSs, while any BF > 1 in the right panel means that we correctly concluded that there are no kaons in their interiors.

As first suggested by the mass-radius plot, Fig. 8.3 confirms that in order to detect the presence of kaons on NS EoS we need a high mass system. The lowest mass system ($m_1 = 1.4M_\odot, m_2 = 1.35M_\odot$) gives BFs that are consistent with 1, in agreement with the m$-$R relation. For large SNRs, as we increase the injected masses we recover BFs that correctly favor the kaon model (left panel) and correctly favor the kaonless models (right panel).

At low SNRs, however, we encounter BFs that incorrectly disfavor the kaon models on the left panel of Fig. 8.3. This is because the injected masses are very close to the maximum mass allowed by the correct model. When an injected parameter is closer to the edge of the prior than the width of the posterior, the posterior has to be cut off (see Appendices D and E). In our case, the injected masses are close enough to the maximum mass of the kaon models that these models "lose" some chain points because the models simply cannot produce such high mass systems. What the kaon models are in fact trying to do when the mass is above their maximum allowed mass is match the GW signal with pure Gaussian noise. How this affects the BF is clear from Eq. (8.5). Recall that the BF in favor of SV222, for example, is equal to the number of times the chains visited this model, divided by the number of times they visited the competing model SV. When SV222 has an abrupt cutoff close to the injected masses, a lot of the chain points will be disfavored because the model fails to produce a signal for these values of the masses. This leads to the chains visiting the SV222 model less, and in the end a BF that does not favor the correct model.

This indicates that when we have BFs favoring the wrong model, we are in a regime where it is mostly the prior that dominates our results. Obviously, the results in this regime do not offer new physical information and cannot be used to claim that

we constrained the EoS. For example, imagine we detected a $(1.99, 1.95)M_\odot$ system at SNR 30 and we want to claim something about the presence or absence of kaons in the system. If we choose to test whether GA-FSU2.1 or GA-FSU2.1-180 fit the data better, we will recover BFs in favor of the kaonless model GA-FSU2.1 regardless of whether kaons are actually present or not. The data coming from such a system can clearly not be trusted to give the correct result.

For the case studied here, the kaon model comparison starts to become likelihood-dominated (BF > 1) when the SNR $\gtrsim 40$ in the SV/SV222 case, and when the SNR is somewhere above 60 in the GA-FSU2.1/GA-FSU2.1-180 case. A detection of a $(2.05, 2)M_\odot$ system at SNR $= 40$ and above will lead to a strong BF in the Jeffreys scale in favor of the kaonless model for a kaonless injection. On the other hand, for a kaon injection, the SNR needs to reach 60 and above in order to obtain strong BFs. For lower mass systems, the BFs are lower, and barely worth mentioning in the Jeffreys scale.

Hyperons

Moving on to the study of the detectability of hyperons in the inner cores of NSs, we select 5 EoSs constructed with (i) RMF theory, (ii) a nonrelativistic Brueckner-Hartree-Fock (BHF) approach, (iii) a relativistic BHF approach, (iv) and (v) a SHF approach. Each pair consists of one EoS with a hyperon in the inner core and one EoS without.

Hyperons, much like kaons, affect the EoS of only the most massive NSs, since they form at the very highest central densities. For that reason, the results of this section are very similar to the previous one on kaons: we find that aLIGO can constrain the existence of hyperons in NSs, but detecting them will be much harder.

Models The models we use to determine hyperon detectability are the following.

1. *RMF theory* [129]: G4 (hyperonless)/ H4 (with hyperons).

2. *Nonrelativistic BHF models* [244]: MPa (hyperonless)/ MPaH (with hyperons).

3. *relativistic BHF models* [120]: DBHF$^{(2)}$(A) (hyperonless)/ NlY5KK* (with hyperons).

4. *SHF model*: SGI [138] (hyperonless)/ SGI-YBZ6-SΛΛ3 [137] (with hyperons).

5. *SHF models*: SkI4 [181] (hyperonless)/ SkI4-YBZ6-SΛΛ3 [137] (with hyperons).

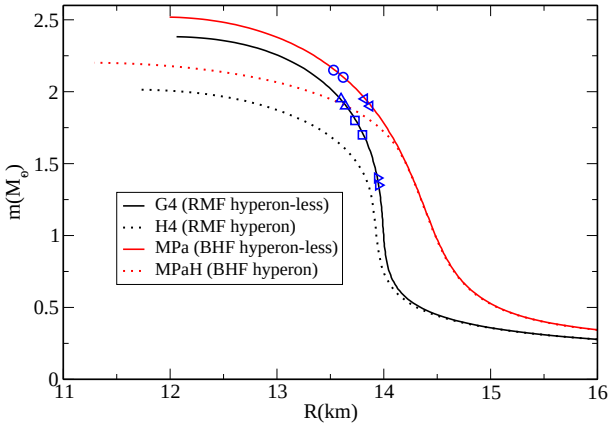

Figure 8.4: m-R relation for the EoS pairs that test hyperons: The black lines correspond to G4 (solid) and H4 (dotted), which are constructed through RMF theory. The red lines are for Mpa (solid) and MPaH (dotted), which are constructed through the nonrelativistic BHF method. The presence of hyperons causes the hyperon model of each pair to differ from the normal matter model for high masses. The pairs of similar symbols indicate the component masses we use in our analyses.

The results of comparisons (3), (4) and (5) are very similar to the results of the MPa and MPaH comparison, so we will not present them here. Figure 8.4 shows the mass-radius relation for the EoSs for which we will present comparisons. Clearly, hyperons affect the EoS of only the most massive NSs, exactly like kaons. The pairs of symbols indicate the masses we inject.

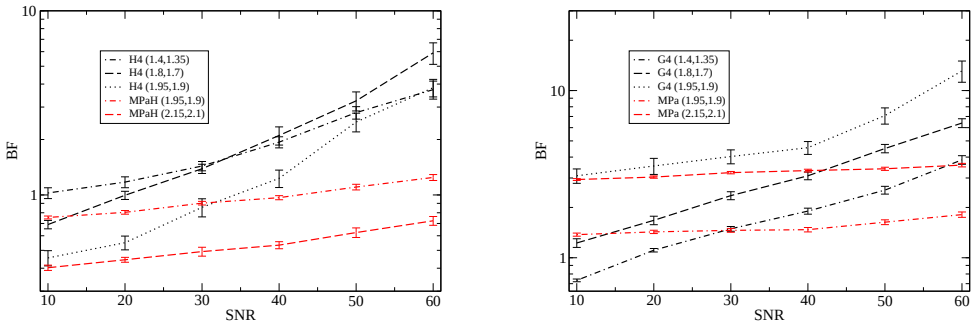

Figure 8.5: BF in favor of the hyperon models (Left) and the hyperonless models (Right) as a function of the SNR for different injected masses. BFs are always quoted in favor of the correct (injected) model. We conclude that it is easier to constrain hyperons than detect them. A similar result was reached in the case of kaons as shown in Fig. 8.3.

Bayes Factors The results of our analysis are presented in Fig. 8.5, which shows the BF in favor of the hyperon models (left) and the hyperonless models (right) as a function of the SNR of the injected signal. The message of this plot is clear if we take into account the reasoning presented in the previous section. In the G4-H4 comparison for the hyperonless comparison at low SNR (below 40), we encounter the effect of the wrong model being preferred over the correct one. On the other hand, in the MPa-MPaH comparison, this continues to hold until the SNR = 60. This indicates that it is the prior (and more specifically the maximum mass of the hyperon model) that dominates our results and not the likelihood. When the SNR exceeds 40, we start extracting interesting information from the comparison. Assuming hyperons do not form in NS cores, a $(1.95, 1.9)M_{\odot}$ detection with SNR = 40 and 60 will give a strong and very strong indication in the Jeffreys classification of the nonexistence of hyperons respectively. On the other hand, if hyperons do form in NS cores, then a signal with SNR = 60 would only provide *strong* evidence, as the BFs do not exceed 7.

Quark Matter

Unlike kaon condensates and hyperons that can only exist in combination with normal matter, quark matter can form both with and without normal matter. In the first case we have hybrid NSs with EoSs of the ALF [13] type that have quarks formed after a certain transition density. The second case results in quark stars [176] comprised solely of quark matter.

The EoSs of pure quark stars differ from normal matter EoSs even at low densities, making it possible for aLIGO to detect or rule them out. On the other hand, hybrid normal/quark matter EoSs are constructed by stitching the nuclear matter EoSs in the low density regime to quark matter ones in the high density regime, with appropriate phase transitions in between. Therefore, they reduce to normal matter EoSs at low densities. The transition density n_c and the strength of the strong interactions c determine how much the hybrid EoS differs from the normal matter EoS it is stitched to. The constant c is predicted to be $c \sim 0.37$ [96] by a perturbative calculation. When c is close to this value and the transition from a nuclear matter to quark matter happens at around twice the saturation density, we find that hybrid EoSs might be distinguishable from normal matter EoSs for signals with SNR $\sim 30 - 40$.

Hybrid EoSs In order to study hybrid normal/quark matter stars we select EoSs of the ALF family [13], the normal matter part of which are based on AP4 [12]. All EoSs include ordinary matter, while the ALF EoSs also include pions and quarks in the inner core[4]. The left panel of Fig. 8.6 shows the mass–radius relation for

[4]In principle, we should treat pions as a separate particle, like kaons and hyperons. However, we are not aware of any EoS model that includes only pion condensates and predicts a maximum NS mass above $2M_\odot$.

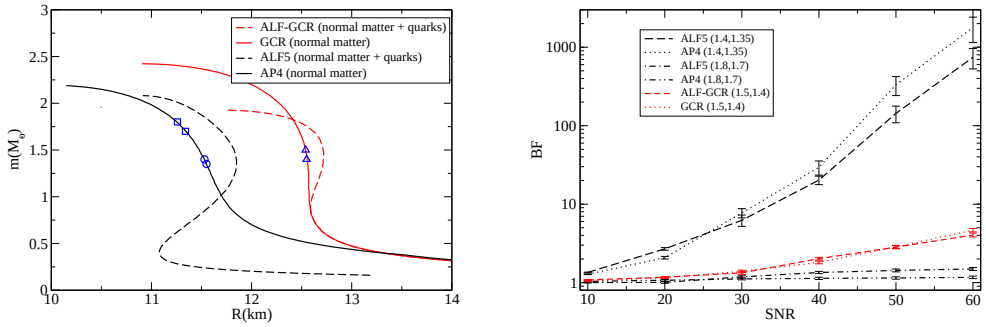

Figure 8.6: (Left) m-R relation for the hybrid EoSs and the nuclear matter EoSs they are based on and compared to: AP4 (solid black), ALF5 (dashed black), GCR (solid red), and ALF-GCR (dashed red). The pairs of similar symbols denote the injected masses. The larger the value of the strong interactions, the larger the deviation between the normal matter and the hybrid EoS. (Right) BF for the ALF5/AP4 comparison (black lines) and the ALF-GCR/GCR comparison (red lines) as a function of the SNR for different injected masses. BFs are given in favor of the correct model denoted in the legend. When the strong interactions between the quarks are close to the value predicted from perturbative calculations, it is possible to distinguish between normal and hybrid NSs.

ALF5 and AP4. ALF5 has $(n_c, c) = (2n_0, 0.4)$, where $n_0 = 0.16$ fm^{-3} is the nuclear saturation density.

We also constructed new hybrid star EoSs (GCR-ALF) by stitching the nuclear matter GCR EoS constructed in [98] to the same quark matter EoS as the ALF family. The m-R relation for a GCR-ALF EoS with $(n_c, c) = (2n_0, 0.35)$ is shown again in Fig. 8.6, together with the corresponding nuclear matter GCR EoS with the symmetry energy of $E_{\rm sym} = 33.8$MeV. We have also chosen different values for n_0, c and $E_{\rm sym}$ and found that the difference between the nuclear matter and hybrid EoSs are typically even smaller than the one in Fig. 8.6.

The right panel of Fig. 8.6 shows the BFs we recover for the ALF5/AP4 (black) and the ALF-GCR/GCR comparison (black) as a function of the SNR for different values of the masses. The only case we find where the hybrid EoS could be distinguishable from the normal matter one is for masses around $1.4M_\odot$. A SNR ~ 30

detection with such masses can provide strong BFs in the Jeffreys scale if we compare ALF5 to AP4. However, if we compare ALF-GCR to GCR we recover smaller BFs. We also find that we recover similar results when comparing hybrid stars and a normal NS regardless of which one is the correct star of Nature. This is different from the kaon and hyperon cases studied before, where we obtain more conclusive BFs when kaons or hyperons are not present in NSs. Of course, if the true hybrid EoS of Nature contains weaker strong interactions between the quarks, the prospects of detecting a hybrid star reduce even further.

Quark Stars Our final study case is SQM3 [176], an EoS that contains no normal matter at all, but rather it is constructed solely with quark matter. Comparing it to a normal matter EoS amounts to comparing normal NSs to strange quark stars. Although the latter have already been heavily constrained [14], it is still interesting to study the bounds aLIGO could place on them. However, it is not clear what EoS we should compare SQM3 to. Throughout our analysis, we compare EoSs that contain exotic matter to the EoS that we obtain if we remove the exotic matter, but change nothing else in how it is constructed; this is precisely how we defined EoS pairs. In the case of SQM3, if we remove the quarks there is no matter left, so there is no normal matter counterpart that can form a pair with a strange quark star.

For that reason and in order to arrive at conservative conclusions, we will compare SQM3 to AP4 [12], a reliable EoS that has both not yet been ruled out by observations and leads to NSs that are the most similar to SQM3 quark stars in the set of EoSs we considered. The left panel of Fig. 8.7 shows the mass–radius relation for SQM3 and AP4, along with symbols that indicate injected masses. Being the softest EoS in the set of non-ruled out EoSs we considered, the AP4 mass-radius relation is the farthest to the left in Fig. 8.7 and thus the closest to the SQM3 one. Other

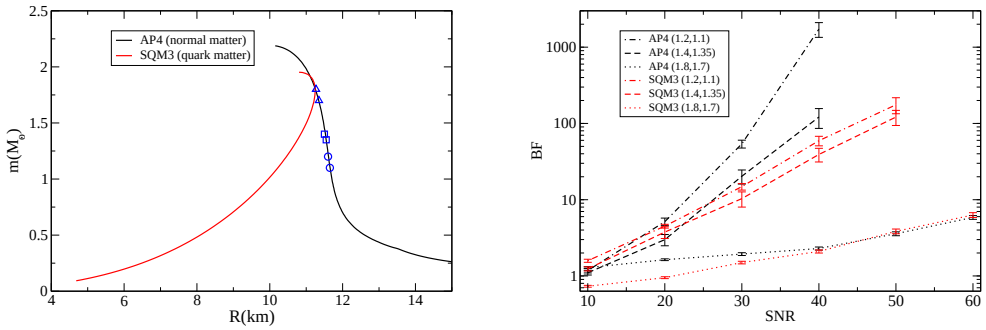

Figure 8.7: (Left) m-R relation for SQM3 and AP4. The pairs of similar symbols denote the injected masses. (Right) BF in favor of AP4 (black) and SQM3 (red) as a function of the SNR for different injected masses. BFs are given in favor of the correct model. aLIGO will be able to place strong constraints on the existence of strange quark stars, both detecting them if present in Nature or strongly disfavoring their existence if not.

normal matter EoSs are stiffer than AP4, predicting larger radii for the same mass, which leads to m − R curves even farther away from SQM3 than AP4. As we show below, AP4 is already clearly distinguishable from SQM3, and thus, distinguishing between SQM3 and other stiffer normal-matter EoSs would be even easier.

Clearly, SQM3 is very different from all other EoSs studied here; its m − R relation differs qualitatively from normal matter EoSs, and we thus expect it to be distinguishable. Observe that SQM3 fails to produce a $2M_\odot$ NS, though it is still consistent with the current observational bound on the maximum NS mass [17] within a 2-σ statistical error. Of course, it is possible that the star observed in [17] is simply a NS, and not a quark star, without necessarily ruling out the existence of the latter. We have thus decided to study whether aLIGO can distinguish between strange quark stars and normal NSs. If all compact stars that aLIGO sees are NSs or BHs, then this would build confidence for the nonexistence of quark stars.

The right panel of Fig. 8.7 shows the BF in favor of AP4 (black) and SQM3 (red) as a function of the SNR of the signal for different mass combinations. The

really high BFs we recover indicate that aLIGO will be able to both detect strange quark stars if they exist or provide very strong evidence for their nonexistence. For example, the detection of a $(1.4, 1.35)M_{\odot}$ system with SNR = 20 results in strong BFs in the Jeffreys scale in favor of the correct model. For even brighter sources, or less massive systems, we recover very strong or even decisive evidence in favor of the correct model of Nature, be it strange quark stars or normal NSs.

We have argued that the results given above represent the worst case scenario when comparing quarks stars to normal NSs, given the reliable normal matter EoSs available today and not yet ruled out by observations. But what if nuclear physicists construct a viable EoS that is softer than AP4? Even in this scenario, we can make some claims based on the unique shape of SQM3. Revisit the left panel of Fig. 8.7 and notice that between $1.8M_{\odot}$ and $1M_{\odot}$ the radius SQM3 predicts increases by about 2km. A normal matter EoS that is softer than AP4 will, roughly speaking, have a similar shape to AP4, but it will be shifted to the left in the m − R plane. Even in this case, there will exist some mass between $1M_{\odot}$ and $1.8M_{\odot}$ where the difference between the radius of a quark star and a normal NS is at least 1km. Our results indicate that typically a radius difference around 0.5km is about enough for distinguishability for this type of study, provided the mass is not close to the maximum mass allowed. We can, therefore, claim that even if a softer than AP4 EoS is constructed, it will still be distinguishable from SQM3 for some value of the mass between 1 and $1.8M_{\odot}$, given an SNR around 30, depending on its exact shape.

Edge Effects

So far, we have injected signals with masses for which both models can produce a NS (or a hybrid star or a quark star). In the case of hyperons and kaons, we saw that the large value of the masses required to tell EoS models apart are close to the

maximum mass of the exotic matter model. Low strength signals are greatly affected by this maximum mass cutoff. Here, we examine the case where one of the two objects can only be supported by the model that predicts the higher maximum mass.

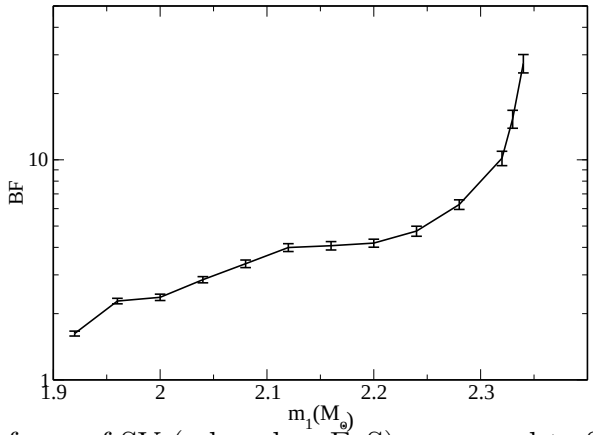

Figure 8.8: BF in favor of SV (a kaonless EoS) compared to SV222 (a kaon EoS) as a function of m_1 for $m_2 = 1.9M_\odot$ with SNR= 30.

We revisit the pair SV-SV222 and fix the SNR to 30 and the mass of the smaller NS in the binary to $1.9M_\odot$. We inject signals constructed with SV and in Fig. 8.8 we plot BF in favor of the correct model as a function of the NS mass of the larger star m_1. The maximum mass supported by SV222 is $\sim 2.11M_\odot$. When both simulated masses can be supported by the wrong model, then we find BFs that are only barely worth mentioning in the Jeffreys scale in favor of the correct model. However, when $m_1 > 2.11M_\odot$ the BF in favor of SV starts increasing and it diverges around $m_1 = 2.35M_\odot$. At this point, the chirp mass of the system is so large that it cannot be matched by a system with both masses below $2.11M_\odot$. Beyond this point, SV222 cannot produce any systems with the correct chirp mass of the injected signal.

To illustrate this transition we plot the 2−D scatter plot of the chain points in the $m_1 - m_2$ plane in Fig. 8.9 for $m_1 = 1.92M_\odot$ and $m_1 = 2.34M_\odot$. The red

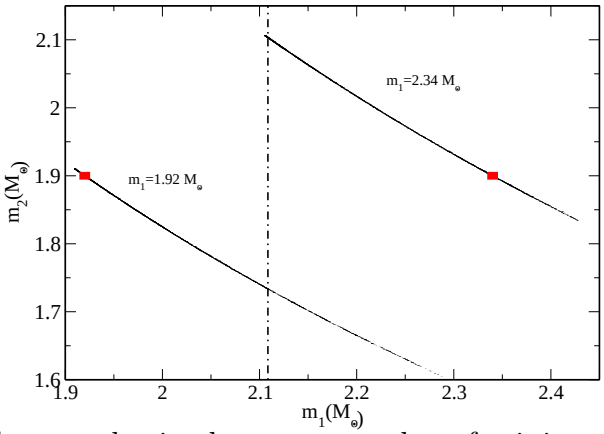

Figure 8.9: 2D Scatter plot in the $m_1 - m_2$ plane for injected masses $m_2 = 1.9M_\odot$ and $m_1 = 1.92M_\odot$ and $m_1 = 2.34M_\odot$. The red box indicates the injected masses and the vertical line gives the maximum mass SV222 can support.

boxes mark the injected masses and the vertical line is the maximum mass SV222 can support. The scatter plots have support only along constant chirp mass lines. For $m_1 = 1.92M_\odot$ most points fall on the left of the vertical line, and SV222 can provide a good match for them. However, as we increase m_1, more and more points move beyond SV222's maximum mass, resulting in BFs that favor it less and less (recall that the BF, as given in Eq. (8.5), is the ratio of the points in SV222 over the points in SV). When $m_1 = 2.34M_\odot$ only a small number of points can be supported by SV222, and we recover very strong BFs in favor of SV. If we increase the mass even more, no points fall on the left of the vertical line, and the BF goes to infinity.

Noise Curves

Apart from the Zero-Det., High-P. noise configuration, there are a few other tuned noise curves for aLIGO [201]. Among them, the NSNS Optimized configuration gives the optimal SNR for a NSNS coalescence, while the high frequency one achieves the best sensitivity around 1000Hz. Figure 8.10 shows these 3 noise curves. The NSNS

Opt. noise curve is tuned to give the highest SNR by reducing the noise levels around $(60-500)$Hz, however this comes at the expense of higher noise in the kHz region. The finite-size effects that we are looking for manifest themselves at frequencies above 400 Hz, suggesting that NSNS Opt. is suboptimal for EoS determination. On the other hand, the High F. configuration has the lowest noise in a small window around 10^3Hz, but it is clearly inferior for lower frequencies. This loss of SNR at low frequencies makes its suitability for EoS studies questionable.

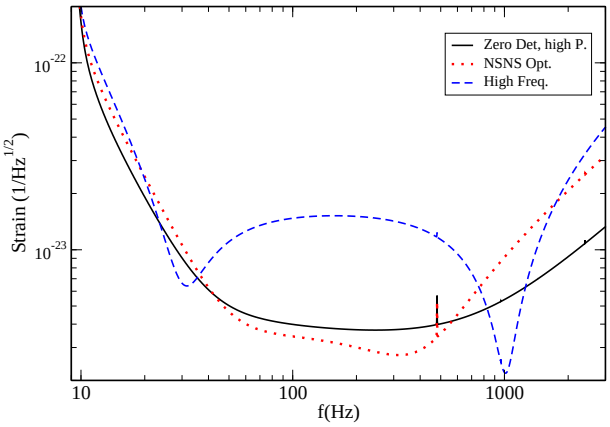

Figure 8.10: Sensitivity curves of various aLIGO configurations. The High-Freq. optimized curve (blue dashed line) is very sensitive in a small window around 1000 Hz, but it has much higher noise at lower frequencies. The NSNS Opt. noise curve (red dotted line) has slightly lower noise at frequencies below 600 Hz, but much higher noise above this. EoS effects become important at frequencies above 400 Hz.

Figure 8.11 demonstrates how the standard High-P., Zero-Det. configuration is the optimal noise curve for measuring the EoS both for systems of constant SNR (top panel) and for systems at the same distance (bottom panel). We plot the BF for a system with the same parameters (apart from the distance in the top panel) achieved with the 3 noise configurations. In the constant distance case, NSNS Opt. achieves the best SNR value while High F. has the worst, as expected. In both cases, the lower high-frequency noise of the Zero-Det., High-P. configuration together with its

low-enough noise at lower frequencies outperforms both tuned configurations in EoS extraction. This is another manifestation of the effect discussed in [187]; when we are looking for an effect that appears only at certain frequencies, it is not the total SNR that matters, but the SNR accumulated in those frequencies.

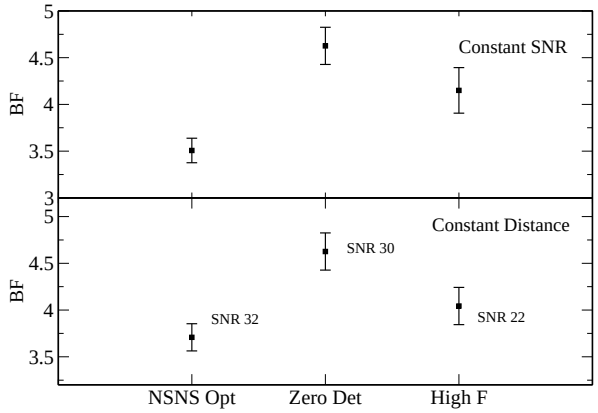

Figure 8.11: BF for the same system with the different aLIGO noise curves. In the top panel we keep the SNR constant, while at the bottom panel we keep the luminosity distance constant. The SNR to which this distance corresponds with each noise curve is indicated in the plot. In both cases, the Zero-Det., High-P. configuration gives the highest BFs at these distances and SNRs, making it the optimal noise curve for EoS studies.

Conclusions

In this chapter we studied whether future GW data from the advanced, ground-based detectors will allow us to learn about the internal composition of NSs. We find that aLIGO can efficiently distinguish between NSs with EoSs that differ at relatively low to moderate central densities. This is the case, for example, for EoSs that model just pure quark matter. If the NS is a hybrid of normal and quark matter, higher SNR values are required, as well as relatively strong interactions between the quarks. On the other hand, aLIGO will not be able to efficiently distinguish between NSs with

EoS that differ only at high central densities. This is the case, for example, for EoSs that model normal matter and those that include a hyperon or a kaon condensate in the inner core. In this case, aLIGO would require a very loud detection to be able to discern between normal and exotic matter NSs.

Our results could be improved if an accurate description of the merger, where the finite size effects are more prominent, is included in the GW models. For example, the effect of hyperons and strange quark matter on GWs from merging NS binaries have been studied e.g. in [115] and [32] respectively. However, NS mergers always occur at kHZ frequencies[5], where the detector sensitivity is lower and it is the inspiral phase that falls in the most sensitive frequency band. Moreover, to this day, no complete template bank of NSNS mergers exists. It would be interesting to study how much information can be extracted from the merger phase, and how this could aid EoS determination by carrying a full inspiral-merger analysis.

As a side note, we find that the aLIGO noise configuration that maximizes the gain of relevant physical information about NS EoS is the standard Zero-Detuned, High-Power one. This is perhaps contrary to the belief that a high-frequency tuned noise configuration would do best. The reason why the Zero-Detuned, High-Power configuration does better is that the finite-size effects that depend on the interior composition of the stars accumulate from roughly a few hundred Hz all the way up to merger, and not in a narrow band in the kHz range. Of course, these effects are much smaller at hundreds of Hz than at contact, however, the noise configuration that accumulates the most SNR when finite-size effects are non-negligible is the Zero-Detuned, High-Power one.

[5]The NS merger frequency cannot be pushed arbitrarily low since the NS mass cannot exceed roughly $3M_{\odot}$ by causality.

The above picture could, again, be altered if the merger was included. Numerical relativity studies [30–34,38,39,67,114–116,206–208] have shown that the EoS leaves an important imprint in the merger and post-merger phases in the form of resonant-like features in the Fourier GW amplitude due to oscillations of hypermassive NSs. For such features that spike in a very narrow band, it may be that a high-frequency tuned noise curve would be best. As of today, it is unclear how much physical information could be extracted from such very high frequency features.

In our study we attempted to focus on physical questions of model selection as opposed to analyzing all possible EoSs constructed. There may exist other interesting physical features of NS interiors, other than the ones studied here, that might be worth investigating. To discover what other features would be interesting to measure, a stronger synergy between the GW and the nuclear physics community should be encouraged.

SUMMARY

The era of GW astronomy has just begun and the possibilities for interesting physical inferences seem endless. In this dissertation we studied one such interesting possibility – that of detecting GWs from spin-precessing compact binaries. Interactions between the spins of the binary components cause the orbital plane to precess in space inducing modulation in the amplitude and the phase of the emitted GW. These modulations add rich structure to signals seen by second generation ground-based detectors and carry a lot of information about the underlying physics that governs the GW source.

Our first task was to construct GW models that could be used to detect and characterize such spin-precessing binaries. We constructed two such models: one that is valid for binaries with slowly spinning components (Chapter 4), and one that captures precessional effects from generic binaries (Chapter 5). Both models were shown to be accurate within their region of validity inducing small systematic errors when used for parameter estimation.

Precessing GW models require that we first solve the spin-precession equations, a system of 10 coupled ordinary differential equations that govern the evolution of all angular momenta of a compact binary system. In the absence of gravitational radiation the spin-precession equations are reduced to 9 coupled ordinary differential equations. We solved this simplified system first as an expansion in small spin magnitudes and then generically. We then perturbatively accounted for radiation reaction employing elements of MSA due to radiation reaction evolving on a longer timescale than precession. This resulted in solutions to the spin-precession equations that were valid for the early inspiral, where this separation of timescales is more robust.

GW data analysis is significantly simplified in the frequency domain, hence Fourier-transforming waveforms is an essential step of GW modeling. In the case of our small-spin GW model, the traditional SPA is valid, allowing us to approximate the Fourier integrand by its value at the stationary point – the point were its phase is not varying rapidly. The situation is different for our generically precessing GW model; the SPA has long been known to not be valid for highly precessing systems. In this case, the SPA is generalized to the SUA, which approximates the Fourier integral by a series of time-shifted waveforms approximated in the SPA.

With accurate GW models that are not plagued by systematic errors at hand, our second task was to study physical inference from spin-precessing sources. We first assessed the performance of our small-spin model in the detection of GWs. We found that for slowly spinning sources, as expected for NSs, non precessing waveforms could detect the majority of the signals. This picture changes for sources displaying significant precession, as expected for binaries containing BHs; failure to account for any type of precession in the models used for detection would lead to the loss of a large fraction of sources (see Chapter 6).

The need for precessing models is even more pronounced if we want to measure the parameters of the source. Our spin-precessing models only suffer from small systematic errors in their region of validity, while all non precessing models are inadequate for parameter estimation of precessing sources. The need for precessing models is not only due to the large systematic errors of the non precessing ones. We have found that the inclusion of spin-precession in GW models leads to a significant reduction in the statistical errors of the parameters as well (see Chapter 7).

The degeneracy between the mass ratio and the spin magnitudes has long been known to deteriorate measurement of both. Using our new spin-precessing model we show that precession largely breaks this degeneracy leading to order-of-magnitude

improvements in both mass and spin extraction. The improvement is such that it allows us to correctly characterize the GW source as consisting of NSs or BHs based on mass estimates only and without any information from matter effects (see Chapter 7).

The final study we carried out for this thesis concerns these matter effects that are present when the GW source contains at least one NS. The specific way the GW is altered by matter effects depends on the still-unknown EoS of extremely dense matter. Measurement of such a GW could provide nuclear physicists with invaluable information. In Chapter 8 we studied how much information we can extract about the nature of the constituents of NS matter. We conclude that second generation ground-based detectors can provide information about the presence or absence of particles formed at relatively low densities/energies. Higher density effects are harder to detect with aLIGO-type detectors, however partial information can still be extracted in this regime.

In general, this thesis is divided into two interconnected but logically distinct parts; quality of extracted data from spin-precessing binaries and physical information extracted from the data. In the first part, we construct accurate models for extraction of minimally biased physical parameters from GW data produced by spin-precessing sources. These parameters could be of unprecedented importance for physical inference. The second part of this thesis deals with this issue. We have selected interesting questions from astrophysics and nuclear physics and assessed the extend to which these questions can be answered by spin-precessing GWs. The hopefully not-so-distant detection of a spin-precessing GW source will allow us to apply the results of this thesis.

REFERENCES CITED

[1] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. *Phys.Rev.*, D88:062001, 2013.

[2] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. 2013.

[3] J. Abadie et al. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *Class.Quant.Grav.*, 27:173001, 2010.

[4] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, 1972.

[5] F. Acernese et al. Status of Virgo detector. *Class. Quant. Grav.*, 24:S381–S388, 2007.

[6] R. Adhikari. private communication.

[7] M. Agathos, W. Del Pozzo, T. G. F. Li, C. V. D. Broeck, J. Veitch, et al. TIGER: A data analysis pipeline for testing the strong-field dynamics of general relativity with gravitational wave signals from coalescing compact binaries. 2013.

[8] M. Agathos, J. Meidam, W. Del Pozzo, T. G. F. Li, M. Tompitak, et al. Constraining the neutron star equation of state with gravitational wave signals from coalescing binary neutron stars. 2015.

[9] P. Ajith. Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins. *Phys.Rev.*, D84:084037, 2011.

[10] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, et al. Phenomenological template family for black-hole coalescence waveforms. *Class.Quant.Grav.*, 24:S689–S700, 2007.

[11] P. Ajith and S. Bose. Estimating the parameters of non-spinning binary black holes using ground-based gravitational-wave detectors: Statistical errors. *Phys.Rev.*, D79:084032, 2009.

[12] A. Akmal, V. Pandharipande, and D. Ravenhall. The Equation of state of nucleon matter and neutron star structure. *Phys.Rev.*, C58:1804–1828, 1998.

[13] M. Alford, M. Braby, M. Paris, and S. Reddy. Hybrid stars that masquerade as neutron stars. *Astrophys.J.*, 629:969–978, 2005.

[14] M. G. Alford, G. Burgio, S. Han, G. Taranto, and D. Zappal. Constraining and applying a generic high-density equation of state. 2015.

[15] M. Anderson, E. W. Hirschmann, L. Lehner, S. L. Liebling, P. M. Motl, et al. Simulating binary neutron stars: Dynamics and gravitational waves. *Phys.Rev.*, D77:024006, 2008.

[16] W. Anderson, P. R. Brady, D. Chin, J. D. E. Creighton, K. Riles, and J. T. Whelan. *Beam pattern response functions and times of arrival for earthbound interferometer.* LIGO Document T010110, 2001.

[17] J. Antoniadis, P. C. Freire, N. Wex, T. M. Tauris, R. S. Lynch, et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340:6131, 2013.

[18] T. A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev. D*, 52:605–620, Jul 1995.

[19] T. A. Apostolatos. Construction of a template family for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 54:2421–2437, Aug 1996.

[20] T. A. Apostolatos. Influence of spin-spin coupling on inspiraling compact binaries with $M_1 = M_2$ and $S_1 = S_2$. *Phys. Rev. D*, 54:2438–2441, Aug 1996.

[21] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys.Rev.*, D49:6274–6297, 1994.

[22] K. Arun, L. Blanchet, B. R. Iyer, and M. S. Qusailah. The 2.5PN gravitational wave polarisations from inspiralling compact binaries in circular orbits. *Class.Quant.Grav.*, 21:3771–3802, 2004.

[23] K. Arun, B. R. Iyer, B. Sathyaprakash, and P. A. Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-Newtonian gravitational wave phasing: The Non-spinning case. *Phys.Rev.*, D71:084008, 2005.

[24] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev.*, D79:104023, 2009.

[25] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy. Degeneracy between mass and spin in black-hole-binary waveforms. *Phys.Rev.*, D87:024035, 2013.

[26] R. Balasubramanian and S. Dhurandhar. Estimation of parameters of gravitational wave signal from coalescing binaries. *Phys.Rev.*, D57:3408–3422, 1998.

[27] R. Balasubramanian, B. Sathyaprakash, and S. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters. *Phys.Rev.*, D53:3033–3055, 1996.

[28] B. Barker and R. O'Connell. The gravitational interaction: Spin, rotation, and quantum effects-a review. *Gen. Relativ. and Gravit.*, 11:149–175, 1979.

[29] M. Baubock, E. Berti, D. Psaltis, and F. Ozel. Relations Between Neutron-Star Parameters in the Hartle-Thorne Approximation. *Astrophys.J.*, 777:68, 2013.

[30] A. Bauswein, H. Janka, K. Hebeler, and A. Schwenk. Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Phys.Rev.*, D86:063001, 2012.

[31] A. Bauswein and H.-T. Janka. Measuring neutron-star properties via gravitational waves from binary mergers. *Phys.Rev.Lett.*, 108:011101, 2012.

[32] A. Bauswein, R. Oechslin, and H.-T. Janka. Discriminating Strange Star Mergers from Neutron Star Mergers by Gravitational-Wave Measurements. *Phys.Rev.*, D81:024012, 2010.

[33] A. Bauswein and N. Stergioulas. Unified picture of the post-merger dynamics and gravitational wave emission in neutron star mergers. *Phys.Rev.*, D91(12):124056, 2015.

[34] A. Bauswein, N. Stergioulas, and H.-T. Janka. Revealing the high-density equation of state through binary neutron star mergers. *Phys.Rev.*, D90(2):023002, 2014.

[35] M. Bejger, D. Gondek-Rosinska, E. Gourgoulhon, P. Haensel, K. Taniguchi, et al. Impact of the nuclear equation of state on the last orbits of binary neutron stars. *Astron.Astrophys.*, 431:297, 2005.

[36] K. Belczynski, G. Wiktorowicz, C. Fryer, D. Holz, and V. Kalogera. Missing Black Holes Unveil The Supernova Explosion Mechanism. *Astrophys.J.*, 757:91, 2012.

[37] C. M. Bender and S. A. Orszag. *Advanced mathematical methods for scientists and engineers 1, Asymptotic methods and perturbation theory.* Springer, New York, 1999.

[38] S. Bernuzzi, T. Dietrich, and A. Nagar. Towards a description of the complete gravitational wave spectrum of neutron star mergers. 2015.

[39] S. Bernuzzi, T. Dietrich, W. Tichy, and B. Brugmann. Mergers of binary neutron stars with realistic spin. *Phys.Rev.*, D89(10):104021, 2014.

[40] S. Bernuzzi, A. Nagar, T. Dietrich, and T. Damour. Modeling the Dynamics of Tidally Interacting Binary Neutron Stars up to the Merger. *Phys. Rev. Lett.*, 114(16):161103, 2015.

[41] E. Berti, A. Buonanno, and C. M. Will. Estimating spinning binary parameters and testing alternative theories of gravity with LISA. *Phys.Rev.*, D71:084025, 2005.

[42] T. Binnington and E. Poisson. Relativistic theory of tidal love numbers. *Phys. Rev. D*, 80:084018, Oct 2009.

[43] L. Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 17:2, 2014.

[44] L. Blanchet, A. Buonanno, and G. Faye. Tail-induced spin-orbit effect in the gravitational radiation of compact binaries. *Phys. Rev. D*, 84:064041, Sep 2011.

[45] T. Bogdanovic, C. S. Reynolds, and M. C. Miller. Alignment of the spins of supermassive black holes prior to merger. *Astrophys. J. Lett.*, 661:L147, 2007.

[46] A. Bohe, G. Faye, S. Marsat, and E. K. Porter. Quadratic-in-spin effects in the orbital dynamics and gravitational-wave energy flux of compact binaries at the 3PN order. 2015.

[47] A. Bohe, S. Marsat, and L. Blanchet. Next-to-next-to-leading order spin-orbit effects in the gravitational wave flux and orbital phasing of compact binaries. 2013.

[48] A. Bohe, S. Marsat, G. Faye, and L. Blanchet. Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries. *Class.Quant.Grav.*, 30:075017, 2013.

[49] M. Boyle, R. Owen, and H. P. Pfeiffer. Geometric approach to the precession of compact binaries. *Phys. Rev. D*, 84:124011, Dec 2011.

[50] R. P. Breton, V. M. Kaspi, M. Kramer, M. A. McLaughlin, M. Lyutikov, S. M. Ransom, I. H. Stairs, R. D. Ferdman, F. Camilo, and A. Possenti. Relativistic Spin Precession in the Double Pulsar. *Science*, 321:104–, July 2008.

[51] A. J. Brizard. A primer on elliptic functions with applications in classical mechanics. *ArXiv e-prints*, Nov. 2007.

[52] A. Buonanno, Y. Chen, and T. Damour. Transition from inspiral to plunge in precessing binaries of spinning black holes. *Phys.Rev.*, D74:104005, 2006.

[53] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two- body dynamics. *Phys. Rev.*, D59:084006, 1999.

[54] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys.Rev.*, D80:084043, 2009.

[55] J. R. Cash and A. H. Karp. A variable order runge-kutta method for initial value problems with rapidly varying right-hand sides. *ACM Trans. Math. Softw.*, 16(3):201–222, Sept. 1990.

[56] S. Chakrabarti, T. Delsate, N. Gurlebeck, and J. Steinhoff. I-Q relation for rapidly rotating neutron stars. *Phys.Rev.Lett.*, 112:201102, 2014.

[57] T. Chan, A. P. Chan, and P. Leung. I-Love relations for incompressible stars and realistic stars. *Phys.Rev.*, D91(4):044017, 2015.

[58] T. Chan, Y. H. Sham, P. Leung, and L. M. Lin. Multipolar universal relations between f-mode frequency and tidal deformability of compact stars. *Phys.Rev.*, D90(12):124023, 2014.

[59] K. Chatziioannou, N. Cornish, A. Klein, and N. Yunes. Detection and Parameter Estimation of Gravitational Waves from Compact Binary Inspirals with Analytical Double-Precessing Templates. *Phys.Rev.*, D89:104023, 2014.

[60] K. Chatziioannou, N. Cornish, A. Klein, and N. Yunes. Spin-precession: Breaking the black hole-neutron star degeneracy. *The Astrophysical Journal Letters*, 798(1):L17, 2015.

[61] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish. in preparation.

[62] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish. Gravitational Waveforms for Precessing, Quasicircular Compact Binaries with Multiple Scale Analysis: Small Spin Expansion. *Phys.Rev.*, D88(6):063011, 2013.

[63] K. Chatziioannou, E. Poisson, and N. Yunes. Tidal heating and torquing of a Kerr black hole to next-to-leading order in the tidal coupling. *Phys. Rev.*, D87(4):044022, 2013.

[64] K. Chatziioannou, K. Yagi, A. Klein, N. Cornish, and N. Yunes. Probing the Internal Composition of Neutron Stars with Gravitational Waves. *Phys. Rev.*, D92(10):104008, 2015.

[65] K. Chatziioannou, K. Yagi, and N. Yunes. Toward realistic and practical no-hair relations for neutron stars in the nonrelativistic limit. *Phys.Rev.*, D90(6):064030, 2014.

[66] H.-S. Cho, E. Ochsner, R. O'Shaughnessy, C. Kim, and C.-H. Lee. Gravitational waves from BH-NS binaries: Effective Fisher matrices and parameter estimation using higher harmonics. *Phys.Rev.*, D87:024004, 2013.

[67] J. Clark, A. Bauswein, L. Cadonati, H.-T. Janka, C. Pankow, et al. Prospects For High Frequency Burst Searches Following Binary Neutron Star Coalescence With Advanced Gravitational Wave Detectors. *Phys.Rev.*, D90(6):062004, 2014.

[68] L. S. Collaboration. LIGO: The Laser Interferometer Gravitational-Wave Observatory. 2007.

[69] N. Cornish, L. Sampson, N. Yunes, and F. Pretorius. Gravitational Wave Tests of General Relativity with the Parameterized Post-Einsteinian Framework. *Phys.Rev.*, D84:062003, 2011.

[70] N. J. Cornish and T. B. Littenberg. Tests of Bayesian Model Selection Techniques for Gravitational Wave Astronomy. *Phys.Rev.*, D76:083006, 2007.

[71] N. J. Cornish and T. B. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Class.Quant.Grav.*, 32(13):135012, 2015.

[72] C. Cutler, T. A. Apostolatos, L. Bildsten, L. S. Finn, E. E. Flanagan, et al. The Last three minutes: issues in gravitational wave measurements of coalescing compact binaries. *Phys.Rev.Lett.*, 70:2984–2987, 1993.

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

[74] C. Cutler and M. Vallisneri. LISA detections of massive black hole inspirals: parameter extraction errors due to inaccurate template waveforms. *Phys. Rev.*, D76:104018, 2007.

[75] T. Damour. Coalescence of two spinning black holes: An effective one- body approach. *Phys. Rev.*, D64:124013, 2001.

[76] T. Damour, B. R. Iyer, and B. S. Sathyaprakash. Improved filters for gravitational waves from inspiralling compact binaries. *Phys. Rev.*, D57:885–907, 1998.

[77] T. Damour and A. Nagar. Relativistic tidal properties of neutron stars. *Phys. Rev. D*, 80:084035, Oct 2009.

[78] T. Damour, A. Nagar, and L. Villain. Measurability of the tidal polarizability of neutron stars in late-inspiral gravitational-wave signals. *Phys. Rev. D*, 85:123007, Jun 2012.

[79] W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale. Demonstrating the feasibility of probing the neutron star equation of state with second-generation gravitational wave detectors. *Phys.Rev.Lett.*, 111(7):071101, 2013.

[80] W. Del Pozzo, J. Veitch, and A. Vecchio. Testing General Relativity using Bayesian model selection: Applications to observations of gravitational waves from compact binary systems. *Phys.Rev.*, D83:082002, 2011.

[81] T. Delsate. I-Love irrotationally. 2015.

[82] D. D. Doneva, S. S. Yazadjiev, K. V. Staykov, and K. D. Kokkotas. Universal I-Q relations for rapidly rotating neutron and strange stars in scalar-tensor theories. *Phys.Rev.*, D90(10):104021, 2014.

[83] D. D. Doneva, S. S. Yazadjiev, N. Stergioulas, and K. D. Kokkotas. Breakdown of I-Love-Q universality in rapidly rotating relativistic stars. *Astrophys.J.*, 781:L6, 2013.

[84] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen. Gravitational waves from inspiraling compact binaries: Validity of the stationary phase approximation to the Fourier transform. *Phys.Rev.*, D59:124016, 1999.

[85] M. D. Duez. Numerical relativity confronts compact neutron star binaries: a review and status report. *Class.Quant.Grav.*, 27:114002, 2010.

[86] L. Engvik, E. Osnes, M. Hjorth-Jensen, G. Bao, and E. Ostgaard. Asymmetric Nuclear Matter and Neutron Star Properties. *Astrophys. J.*, 469:794, Oct. 1996.

[87] F. Douchin and P. Haensel. A unified equation of state of dense matter and neutron star structure. *A&A*, 380(1):151–167, 2001.

[88] J. A. Faber, P. Grandclement, and F. A. Rasio. Mergers of irrotational neutron star binaries in conformally flat gravity. *Phys.Rev.*, D69:124036, 2004.

[89] J. A. Faber, P. Grandclement, F. A. Rasio, and K. Taniguchi. Measuring neutron star radii with gravitational wave detectors. *Phys.Rev.Lett.*, 89:231102, 2002.

[90] J. A. Faber and F. A. Rasio. Binary Neutron Star Mergers. *Living Rev.Rel.*, 15:8, 2012.

[91] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, et al. The Mass Distribution of Stellar-Mass Black Holes. *Astrophys.J.*, 741:103, 2011.

[92] M. Favata. Post-newtonian corrections to the gravitational-wave memory for quasicircular, inspiralling compact binaries. *Phys. Rev. D*, 80:024002, Jul 2009.

[93] M. Favata. Systematic parameter errors in inspiraling neutron star binaries. *Phys.Rev.Lett.*, 112:101101, 2014.

[94] G. Faye, L. Blanchet, and A. Buonanno. Higher-order spin effects in the dynamics of compact binaries. I. Equations of motion. *Phys. Rev.*, D74:104033, 2006.

[95] E. E. Flanagan and T. Hinderer. Constraining neutron star tidal Love numbers with gravitational wave detectors. *Phys.Rev.*, D77:021502, 2008.

[96] E. S. Fraga, R. D. Pisarski, and J. Schaffner-Bielich. Small, dense quark stars from perturbative QCD. *Phys.Rev.*, D63:121702, 2001.

[97] R. Fujita. Gravitational Waves from a Particle in Circular Orbits around a Schwarzschild Black Hole to the 22nd Post-Newtonian Order. *Prog.Theor.Phys.*, 128:971–992, 2012.

[98] S. Gandolfi, J. Carlson, and S. Reddy. The maximum mass and radius of neutron stars and the nuclear symmetry energy. *Phys.Rev.*, C85:032801, 2012.

[99] D. Gerosa, M. Kesden, R. O'Shaughnessy, A. Klein, E. Berti, U. Sperhake, and D. Trifiro. Precessional instability in binary black holes with aligned spins. *Phys. Rev. Lett.*, 115:141102, 2015.

[100] D. Gerosa, M. Kesden, U. Sperhake, E. Berti, and R. O'Shaughnessy. Multi-timescale analysis of phase transitions in precessing black-hole binaries. *Phys. Rev.*, D92:064016, 2015.

[101] D. Gerosa, R. O'Shaughnessy, M. Kesden, E. Berti, and U. Sperhake. Distinguishing black-hole spin-orbit resonances by their gravitational-wave signatures. *Phys. Rev.*, D89(12):124025, 2014.

[102] B. Giacomazzo and R. Perna. Formation of Stable Magnetars from Binary Neutron Star Mergers. 2013.

[103] D. Gondek-Rosiska, M. Bejger, T. Bulik, E. Gourgoulhon, P. Haensel, F. Limousin, K. Taniguchi, and L. Zdunik. The final phase of inspiral of neutron stars: Realistic equations of state. *Advances in Space Research*, 39(2):271 – 274, 2007.

[104] S. Gossan, J. Veitch, and B. Sathyaprakash. Bayesian model selection for testing the no-hair theorem with black hole ringdowns. *Phys.Rev.*, D85:124056, 2012.

[105] N. Gupta and P. Arumugam. Antikaons and higher order couplings in the relativistic mean-field study of neutron stars. *Journal of Physics G Nuclear Physics*, 40(2):025203, Feb. 2013.

[106] M. Hannam, D. A. Brown, S. Fairhurst, C. L. Fryer, and I. W. Harry. When can gravitational-wave observations distinguish between black holes and neutron stars? *Astrophys.J.*, 766:L14, 2013.

[107] M. Hannam, P. Schmidt, A. Boh, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Prrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014.

[108] G. M. Harry. Advanced LIGO: The next generation of gravitational wave detectors. *Class.Quant.Grav.*, 27:084006, 2010.

[109] B. Haskell, R. Ciolfi, F. Pannarale, and L. Rezzolla. On the universality of I-Love-Q relations in magnetized neutron stars. *Mon. Not. Roy. Astron. Soc.*, 438:L71, 2013.

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

[111] T. Hinderer and E. E. Flanagan. Two timescale analysis of extreme mass ratio inspirals in Kerr. I. Orbital Motion. *Phys.Rev.*, D78:064028, 2008.

[112] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read. Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral. *Phys.Rev.*, D81:123016, 2010.

[113] A. W. Hotan, M. Bailes, and S. M. Ord. Geodetic Precession in PSR J1141-6545. *Astrophys. J.*, 624:906–913, May 2005.

[114] K. Hotokezaka, K. Kyutoku, H. Okawa, and M. Shibata. Exploring tidal effects of coalescing binary neutron stars in numerical relativity. II. Long-term simulations. *Phys.Rev.*, D91(6):064060, 2015.

[115] K. Hotokezaka, K. Kyutoku, H. Okawa, M. Shibata, and K. Kiuchi. Binary Neutron Star Mergers: Dependence on the Nuclear Equation of State. *Phys.Rev.*, D83:124008, 2011.

[116] K. Hotokezaka, K. Kyutoku, and M. Shibata. Exploring tidal effects of coalescing binary neutron stars in numerical relativity. *Phys.Rev.*, D87(4):044001, 2013.

[117] H. Jeffreys. *The Theory of Probability.* Oxford, 1961.

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

[119] D. L. Kaplan, S. Chatterjee, B. M. Gaensler, and J. Anderson. A Precise Proper Motion for the Crab Pulsar, and the Difficulty of Testing Spin-Kick Alignment for Young Neutron Stars. *Astrophys. J.*, 677:1201, 2008.

[120] T. Katayama and K. Saito. Neutron stars with Hyperons in Dirac-Brueckner-Hartree-Fock approach. *ArXiv e-prints*, Oct. 2014.

[121] M. Kesden, D. Gerosa, R. O'Shaughnessy, E. Berti, and U. Sperhake. Effective potentials and morphological transitions for binary black-hole spin precession. 2014.

[122] K. Kiuchi, Y. Sekiguchi, K. Kyutoku, and M. Shibata. Gravitational waves, neutrino emissions, and effects of hyperons in binary neutron star mergers. *Class.Quant.Grav.*, 29:124003, 2012.

[123] B. Kleihaus, J. Kunz, and S. Mojica. Quadrupole Moments of Rapidly Rotating Compact Objects in Dilatonic Einstein-Gauss-Bonnet Theory. *Phys.Rev.*, D90:061501, 2014.

[124] A. Klein, N. Cornish, and N. Yunes. Gravitational Waveforms for Precessing, Quasi-circular Binaries via Multiple Scale Analysis and Uniform Asymptotics: The Near Spin Alignment Case. 2013.

[125] A. Klein, N. Cornish, and N. Yunes. Fast Frequency-domain Waveforms for Spin-Precessing Binary Inspirals. *Phys. Rev.*, D90:124029, 2014.

[126] A. Klein, P. Jetzer, and M. Sereno. Parameter estimation for coalescing massive binary black holes with lisa using the full 2-post-newtonian gravitational waveform and spin-orbit precession. *Phys. Rev. D*, 80:064027, Sep 2009.

[127] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman. Extracting equation of state parameters from black hole-neutron star mergers: Nonspinning black holes. *Phys. Rev. D*, 85:044061, Feb 2012.

[128] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman. Extracting equation of state parameters from black hole-neutron star mergers: aligned-spin black holes and a preliminary waveform model. *Phys.Rev.*, D89(4):043009, 2014.

[129] B. D. Lackey, M. Nayyar, and B. J. Owen. Observational constraints on hyperons in neutron stars. *Phys.Rev.*, D73:024021, 2006.

[130] B. D. Lackey and L. Wade. Reconstructing the neutron-star equation of state with gravitational-wave detectors from a realistic population of inspiralling binary neutron stars. *Phys.Rev.*, D91(4):043002, 2015.

[131] R. N. Lang and S. A. Hughes. Measuring coalescing massive binary black holes with gravitational waves: The Impact of spin-induced precession. *Phys.Rev.*, D74:122001, 2006.

[132] R. N. Lang, S. A. Hughes, and N. J. Cornish. Measuring parameters of massive black hole binaries with partially aligned spins. *Phys.Rev.*, D84:022002, 2011.

[133] J. M. Lattimer and Y. Lim. Constraining the Symmetry Parameters of the Nuclear Interaction. *ApJ. 771,*, 51, 2013.

[134] J. M. Lattimer and M. Prakash. Neutron Star Observations: Prognosis for Equation of State Constraints. *Phys.Rept.*, 442:109–165, 2007.

[135] J. M. Lattimer and F. D. Swesty. A generalized equation of state for hot, dense matter. *Nuclear Physics A*, 535(2):331 – 376, 1991.

[136] T. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, et al. Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence. 2011. 26 pages, 23 figures, Accepted by PRD.

[137] Y. Lim, C. H. Hyun, K. Kwak, and C.-H. Lee. $\Lambda \Lambda$ Interaction and Neutron Stars. *ArXiv e-prints*, Dec. 2014.

[138] Y. Lim, K. Kwak, C. H. Hyun, and C.-H. Lee. Kaon Condensation in Neutron Stars with Skyrme-Hartree-Fock Models. *Phys.Rev.*, C89(5):055804, 2014.

[139] L. Lindblom, B. J. Owen, and D. A. Brown. Model Waveform Accuracy Standards for Gravitational Wave Data Analysis. *Phys. Rev.*, D78:124020, 2008.

[140] T. B. Littenberg and N. J. Cornish. A Bayesian Approach to the Detection Problem in Gravitational Wave Astronomy. *Phys.Rev.*, D80:063007, 2009.

[141] T. B. Littenberg and N. J. Cornish. Bayesian inference for spectral estimation of gravitational wave detector noise. *Phys.Rev.*, D91(8):084034, 2015.

[142] Y. T. Liu, S. L. Shapiro, Z. B. Etienne, and K. Taniguchi. General relativistic simulations of magnetized binary neutron star mergers. *Phys.Rev.*, D78:024012, 2008.

[143] M. Luna and A. M. Sintes. Parameter estimation of compact binaries using the inspiral and ringdown waveforms. *Class.Quant.Grav.*, 23:3763–3782, 2006.

[144] A. Lundgren and R. O'Shaughnessy. A single-spin precessing gravitational wave in closed form. 2013.

[145] B. Majumder, K. Yagi, and N. Yunes. Improved Universality in the Neutron Star Three-Hair Relations. 2015.

[146] I. Mandel, C. P. Berry, F. Ohme, S. Fairhurst, and W. M. Farr. Parameter estimation on compact binary coalescences with abruptly terminating gravitational waveforms. *Class.Quant.Grav.*, 31:155005, 2014.

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

[148] C. Markakis, J. S. Read, M. Shibata, K. Uryu, J. D. Creighton, et al. Inferring the neutron star equation of state from binary inspiral waveforms. pages 743–745, 2010.

[149] S. Marsat. Cubic order spin effects in the dynamics and gravitational wave energy flux of compact object binaries. *Class. Quant. Grav.*, 32(8):085008, 2015.

[150] S. Marsat, A. Bohe, L. Blanchet, and A. Buonanno. Next-to-leading tail-induced spin-orbit effects in the gravitational radiation flux of compact binaries. *Class.Quant.Grav.*, 31:025023, 2014.

[151] G. Martinon, A. Maselli, L. Gualtieri, and V. Ferrari. Rotating protoneutron stars: Spin evolution, maximum mass, and I-Love-Q relations. *Phys.Rev.*, D90(6):064026, 2014.

[152] A. Maselli, V. Cardoso, V. Ferrari, L. Gualtieri, and P. Pani. Equation-of-state-independent relations in neutron stars. *Phys.Rev.*, D88:023007, 2013.

[153] A. Maselli, L. Gualtieri, and V. Ferrari. Constraining the equation of state of nuclear matter with gravitational wave observations: Tidal deformability and tidal disruption. *Phys.Rev.*, D88(10):104040, 2013.

[154] D. McKechan, C. Robinson, and B. Sathyaprakash. A tapering window for time-domain templates and simulated signals in the detection of gravitational waves from coalescing compact binaries. *Class.Quant.Grav.*, 27:084020, 2010.

[155] B. Metzger and E. Berger. What is the Most Promising Electromagnetic Counterpart of a Neutron Star Binary Merger? *Astrophys.J.*, 746:48, 2012.

[156] B. Mikoczi, B. Kocsis, P. Forgacs, and M. Vasuth. Parameter estimation for inspiraling eccentric compact binaries including pericenter precession. *Phys.Rev.*, D86:104027, 2012.

[157] B. Mikóczi, M. Vasúth, and L. A. Gergely. Self-interaction spin effects in inspiralling compact binaries. *Phys. Rev. D*, 71:124043, Jun 2005.

[158] Y. Mino et al. Black hole perturbation. *Prog. Theor. Phys. Suppl.*, 128:1–121, 1997.

[159] D. Nicholson and A. Vecchio. Bayesian bounds on parameter estimation accuracy for compact coalescing binary gravitational wave signals. *Phys.Rev.*, D57:4588–4599, 1998.

[160] S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers. Exploring Short Gamma-ray Bursts as Gravitational-wave Standard Sirens. *Astrophys. J.*, 725:496–514, Dec. 2010.

[161] R. O'Shaughnessy, B. Farr, E. Ochsner, H. Cho, V. Raymond, et al. Parameter Estimation of Gravitational Waves from Precessing BH-NS Inspirals with higher harmonics. 2014.

[162] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, C. Kim, et al. Parameter Estimation of Gravitational Waves from Nonprecessing BH-NS Inspirals with higher harmonics: Comparing MCMC posteriors to an Effective Fisher Matrix. *Phys.Rev.*, D89:064048, 2014.

[163] F. Ozel, D. Psaltis, R. Narayan, and J. E. McClintock. The Black Hole Mass Distribution in the Galaxy. *Astrophys.J.*, 725:1918–1927, 2010.

[164] F. Ozel, D. Psaltis, R. Narayan, and A. S. Villarreal. On the Mass Distribution and Birth Masses of Neutron Stars. *Astrophys.J.*, 757:55, 2012.

[165] Y. Pan, A. Buonanno, J. G. Baker, J. Centrella, B. J. Kelly, et al. A Data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. *Phys.Rev.*, D77:024014, 2008.

[166] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mrou, H. P. Pfeiffer, M. A. Scheel, and B. Szilgyi. Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism. *Phys. Rev.*, D89(8):084006, 2014.

[167] V. Pandharipande. Hyperonic matter. *Nucl.Phys.*, A178:123–144, 1971.

[168] P. Pani and E. Berti. I-Love-Q, Spontaneously: Slowly Rotating Neutron Stars in Scalar-Tensor Theories. *Phys.Rev.*, D90:024025, 2014.

[169] F. Pannarale, L. Rezzolla, F. Ohme, and J. S. Read. Will black hole-neutron star binary inspirals tell us about the neutron star equation of state? *Phys.Rev.*, D84:104017, 2011.

[170] G. Pappas and T. A. Apostolatos. Effectively universal behavior of rotating neutron stars in general relativity makes them even simpler than their Newtonian counterparts. *Phys.Rev.Lett.*, 112:121101, 2014.

[171] P. C. Peters. Gravitational radiation and the motion of two point masses. *Phys. Rev.*, 136:B1224–B1232, Nov 1964.

[172] E. Poisson. Gravitational waves from inspiraling compact binaries: The quadrupole-moment term. *Phys. Rev. D*, 57:5287–5290, Apr 1998.

[173] E. Poisson. Absorption of mass and angular momentum by a black hole: Time-domain formalisms for gravitational perturbations, and the small-hole / slow-motion approximation. *Phys. Rev.*, D70:084044, 2004.

[174] E. Poisson and C. Will. *Gravity: Newtonian, Post-Newtonian, Relativistic*. Cambridge University Press, 2014.

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

[176] M. Prakash, J. Cooke, and J. Lattimer. Quark - hadron phase transition in protoneutron stars. *Phys.Rev.*, D52:661–665, 1995.

[177] E. Racine. Analysis of spin precession in binary black hole systems including quadrupole-monopole interaction. *Phys.Rev.*, D78:044021, 2008.

[178] E. Racine, A. Buonanno, and L. Kidder. Recoil velocity at second post-newtonian order for spinning black hole binaries. *Phys. Rev. D*, 80:044010, Aug 2009.

[179] J. S. Read, L. Baiotti, J. D. E. Creighton, J. L. Friedman, B. Giacomazzo, et al. Matter effects on binary neutron star waveforms. *Phys.Rev.*, D88:044042, 2013.

[180] J. S. Read, C. Markakis, M. Shibata, K. Uryu, J. D. Creighton, et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys.Rev.*, D79:124033, 2009.

[181] P.-G. Reinhard and H. Flocard. Nuclear effective forces and isotope shifts. *Nuclear Physics A*, 584:467–488, Feb. 1995.

[182] J. Rhoades, Clifford E. and R. Ruffini. Maximum mass of a neutron star. *Phys.Rev.Lett.*, 32:324–327, 1974.

[183] C. L. Rodriguez, B. Farr, W. M. Farr, and I. Mandel. Inadequacies of the Fisher Information Matrix in gravitational-wave parameter estimation. *Phys.Rev.*, D88:084013, 2013.

[184] C. L. Rodriguez, B. Farr, V. Raymond, W. M. Farr, T. Littenberg, et al. Basic Parameter Estimation of Binary Neutron Star Systems by the Advanced LIGO/Virgo Network. 2013.

[185] L. Sampson, N. Cornish, and N. Yunes. Gravitational Wave Tests of Strong Field General Relativity with Binary Inspirals: Realistic Injections and Optimal Model Selection. *Phys.Rev.*, D87(10):102001, 2013.

[186] L. Sampson, N. Cornish, and N. Yunes. Mismodeling in gravitational-wave astronomy: The trouble with templates. *Phys.Rev.*, D89(6):064037, 2014.

[187] L. Sampson, N. Yunes, N. Cornish, M. Ponce, E. Barausse, et al. Projected Constraints on Scalarization with Gravitational Waves from Neutron Star Binaries. *Phys.Rev.*, D90(12):124091, 2014.

[188] L. Santamaria, F. Ohme, P. Ajith, B. Bruegmann, N. Dorband, et al. Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries. *Phys.Rev.*, D82:064016, 2010.

[189] P. D. Scharre and C. M. Will. Testing scalar tensor gravity using space gravitational wave interferometers. *Phys.Rev.*, D65:042002, 2002.

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

[191] P. Schmidt, M. Hannam, S. Husa, and P. Ajith. Tracking the precession of compact binaries from their gravitational-wave signal. *Phys.Rev.*, D84:024046, 2011.

[192] P. Schmidt, F. Ohme, and M. Hannam. Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter. 2014.

[193] J. D. Schnittman. Spin-orbit resonance and the evolution of compact binary systems. *Phys. Rev.*, D70:124020, 2004.

[194] S. Shah, M. van der Sluys, and G. Nelemans. Using electromagnetic observations to aid gravitational-wave parameter estimation of compact binaries observed with LISA. 2012.

[195] Y. H. Sham, L. M. Lin, and P. Leung. Testing universal relations of neutron stars with a nonlinear matter-gravity coupling theory. *Astrophys.J.*, 781:66, 2014.

[196] H. Shen, H. Toki, K. Oyamatsu, and K. Sumiyoshi. Relativistic equation of state of nuclear matter for supernova and neutron star. *Nuclear Physics A*, 637(3):435 – 450, 1998.

[197] H. Shen, H. Toki, K. Oyamatsu, and K. Sumiyoshi. Relativistic equation of state of nuclear matter for supernova and neutron star. *Nuclear Physics A*, 637:435–450, July 1998.

[198] H. Shen, H. Toki, K. Oyamatsu, and K. Sumiyoshi. Relativistic Equation of State of Nuclear Matter for Supernova Explosion. *Progress of Theoretical Physics*, 100:1013–1031, Nov. 1998.

[199] M. Shibata, M. Sasaki, H. Tagoshi, and T. Tanaka. Gravitational waves from a particle orbiting around a rotating black hole: PostNewtonian expansion. *Phys. Rev.*, D51:1646–1663, 1995.

[200] M. Shibata, K. Taniguchi, and K. b. o. Uryū. Merger of binary neutron stars with realistic equations of state in full general relativity. *Phys. Rev. D*, 71:084021, Apr 2005.

[201] D. Shoemaker. *Advanced LIGO anticipated sensitivity curves*. Tech. Rep. LIGO-T0900288-v3, 2010.

[202] T. H. R. Skyrme. CVII. The nuclear surface. *Philosophical Magazine*, 1:1043–1054, Nov. 1956.

[203] H. Spruit and E. S. Phinney. Birth kicks as the origin of pulsar rotation. *Nature*, 393:139–141, 1998.

[204] A. Stavridis and C. M. Will. Bounding the mass of the graviton with gravitational waves: Effect of spin precessions in massive black hole binaries. *Phys.Rev.*, D80:044002, 2009.

[205] L. C. Stein, K. Yagi, and N. Yunes. Three-Hair Relations for Rotating Stars: Nonrelativistic Limit. *Astrophys. J.*, 788:15, 2014.

[206] N. Stergioulas, A. Bauswein, K. Zagkouris, and H.-T. Janka. Gravitational waves and nonaxisymmetric oscillation modes in mergers of compact object binaries. *Mon.Not.Roy.Astron.Soc.*, 418:427, 2011.

[207] K. Takami, L. Rezzolla, and L. Baiotti. Constraining the Equation of State of Neutron Stars from Binary Mergers. *Phys.Rev.Lett.*, 113(9):091104, 2014.

[208] K. Takami, L. Rezzolla, and L. Baiotti. Spectral properties of the post-merger gravitational-wave signal from binary neutron stars. *Phys.Rev.*, D91(6):064001, 2015.

[209] T. Tanaka, H. Tagoshi, and M. Sasaki. Gravitational waves by a particle in circular orbits around a Schwarzschild black hole: 5.5 post-Newtonian formula. *Prog. Theor. Phys.*, 96:1087–1101, 1996.

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

[211] K. S. Thorne and J. B. Hartle. Laws of motion and precession for black holes and other bodies. *Phys. Rev.*, D31:1815–1837, 1984.

[212] R. Trotta. Applications of Bayesian model selection to cosmological parameters. *Mon.Not.Roy.Astron.Soc.*, 378:72–82, 2007.

[213] K. Uryu, M. Shibata, and Y. Eriguchi. Properties of general relativistic, irrotational binary neutron stars in close quasiequilibrium orbits: Polytropic equations of state. *Phys.Rev.*, D62:104015, 2000.

[214] M. Vallisneri. Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Phys.Rev.*, D77:042001, 2008.

[215] M. Vallisneri. Beyond Fisher: exact sampling distributions of the maximum-likelihood estimator in gravitational-wave parameter estimation. *Phys. Rev. Lett.*, 107:191104, 2011.

[216] M. van der Sluys, I. Mandel, V. Raymond, V. Kalogera, C. Rover, et al. Parameter estimation for signals from compact binary inspirals injected into LIGO data. *Class.Quant.Grav.*, 26:204010, 2009.

[217] M. van der Sluys, V. Raymond, I. Mandel, C. Rover, N. Christensen, et al. Parameter estimation of spinning binary inspirals using Markov-chain Monte Carlo. *Class.Quant.Grav.*, 25:184011, 2008.

[218] A. Vecchio. LISA observations of rapidly spinning massive black hole binary systems. *Phys.Rev.*, D70:042001, 2004.

[219] J. Veitch, I. Mandel, B. Aylott, B. Farr, V. Raymond, et al. Estimating parameters of coalescing compact binaries with proposed advanced detector networks. *Phys.Rev.*, D85:104045, 2012.

[220] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys.Rev.*, D91(4):042003, 2015.

[221] I. Verdinelli and L. Wasserman. Computing bayes factors using a generalization of the savage-dickey density ratio. *Journal of the American Statistical Association*, 90(430):pp. 614–618, 1995.

[222] J. Vines, E. E. Flanagan, and T. Hinderer. Post-1-Newtonian tidal effects in the gravitational waveform from binary inspirals. *Phys.Rev.*, D83:084051, 2011.

[223] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani. Measuring the spin of black holes in binary systems using gravitational waves. 2014.

[224] S. Vitale and M. Zanolin. Parameter estimation from Gravitational waves generated by non-spinning binary black holes with laser interferometers: beyond the Fisher information. *Phys.Rev.*, D82:124065, 2010.

[225] L. Wade, J. D. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, et al. Systematic and statistical errors in a bayesian approach to the estimation of the neutron-star equation of state using advanced gravitational wave detectors. *Phys.Rev.*, D89(10):103012, 2014.

[226] C. Wang, D. Lai, and J. Han. Neutron star kicks in isolated and binary pulsars: observational constraints and implications for kick mechanisms. *Astrophys. J.*, 639:1007–1017, 2006.

[227] C. Wang, D. Lai, and J. Han. Spin-Kick Correlation in Neutron Stars: Alignment Conditions and Implications. *Astrophys. J.*, 656:399–407, 2007.

[228] J. M. Weisberg and J. H. Taylor. General Relativistic Geodetic Spin Precession in Binary Pulsar B1913+16: Mapping the Emission Beam in Two Dimensions. *Astrophys. J.*, 576:942–949, Sept. 2002.

[229] C. M. Will. Testing scalar - tensor gravity with gravitational wave observations of inspiraling compact binaries. *Phys.Rev.*, D50:6058–6067, 1994.

[230] C. M. Will. Bounding the mass of the graviton using gravitational wave observations of inspiralling compact binaries. *Phys.Rev.*, D57:2061–2068, 1998.

[231] C. M. Will. The Confrontation between General Relativity and Experiment. *Living Rev.Rel.*, 17:4, 2014.

[232] C. M. Will and N. Yunes. Testing alternative theories of gravity using LISA. *Class.Quant.Grav.*, 21:4367, 2004.

[233] Z.-G. Xing, J. M. Centrella, and S. L. McMillan. Gravitational radiation from the coalescence of binary neutron stars: Effects due to the equation of state, spin, and mass ratio. *Phys.Rev.*, D54:7261–7277, 1996.

[234] K. Yagi. Multipole Love Relations. *Phys.Rev.*, D89:043011, 2014.

[235] K. Yagi, K. Chatziioannou, and N. Yunes, in preparation.

[236] K. Yagi, K. Kyutoku, G. Pappas, N. Yunes, and T. A. Apostolatos. No-Hair Relations for Neutron Stars and Quark Stars: Relativistic Results. 2014.

[237] K. Yagi, L. C. Stein, G. Pappas, N. Yunes, and T. A. Apostolatos. Why I-Love-Q: Explaining why universality emerges in compact objects. *Phys.Rev.*, D90(6):063010, 2014.

[238] K. Yagi and T. Tanaka. Constraining alternative theories of gravity by gravitational waves from precessing eccentric compact binaries with LISA. *Phys.Rev.*, D81:064008, 2010.

[239] K. Yagi and N. Yunes. I-Love-Q Relations in Neutron Stars and their Applications to Astrophysics, Gravitational Waves and Fundamental Physics. 2013.

[240] K. Yagi and N. Yunes. I-Love-Q: Unexpected Universal Relations for Neutron Stars and Quark Stars. *Science*, 341:365, 2013.

[241] K. Yagi and N. Yunes. Love can be Tough to Measure. *Phys.Rev.*, D89:021303, 2014.

[242] K. Yagi and N. Yunes. I-Love-Q Anisotropically. *Phys. Rev.*, D91:123008, 2015.

[243] K. Yagi and N. Yunes. Relating follicly-challenged compact stars to bald black holes: A link between two no-hair properties. *Phys.Rev.*, D91(10):103003, 2015.

[244] Y. Yamamoto, T. Furumoto, N. Yasutake, and T. A. Rijken. Hyperon mixing and universal many-body repulsion in neutron stars. *Phys. Rev. C*, 90(4):045805, Oct. 2014.

[245] N. Yunes, K. G. Arun, E. Berti, and C. M. Will. Post-Circular Expansion of Eccentric Binary Inspirals: Fourier-Domain Waveforms in the Stationary Phase Approximation. 2009.

[246] N. Yunes and X. Siemens. Gravitational-Wave Tests of General Relativity with Ground-Based Detectors and Pulsar Timing-Arrays. *Living Rev.Rel.*, 16:9, 2013.

APPENDICES

APPENDIX A

JUSTIFICATION OF THE ϕ_Z CALCULATION

The precession-averaged $\langle \dot{\phi}_z \rangle_{\rm pr}$ given in Eqs. (5.61)-(5.63) is exact. In principle, we could calculate D_2 and D_4 as functions of v, substitute them in Eq. (5.61), and carry out a PN expansion and integration to obtain $\langle \phi_z \rangle_{\rm pr} = \phi_{z,-1}$.

Though this approach should work, in practice we run into 2 considerable problems. Firstly, the resulting $\phi_{z,-1}$ is ill-behaved in the small mass ratio limit, despite never having assumed comparable masses. Secondly, $\phi_{z,-1}$ diverges when, at any point in the evolution of a precessing system, the total spin angular momentum is (anti)aligned with the orbital angular momentum. We stress that this does *not* mean that $\boldsymbol{S}$ is approximately (anti)aligned with $\boldsymbol{L}$ all the time; a brief moment of (anti)alignment suffices.

Both issues are *not* caused by real physical divergences in Ω_z. Firstly, at no point did we assume comparable masses. The second issue is more subtle. It might be true that the denominator of Ω_z vanishes if $\boldsymbol{S}$ and $\boldsymbol{L}$ are (anti)aligned. However, the binary (and ϕ_z) is well behaved at the moment of (anti)alignment since the numerator of Ω_z vanishes too, leading to a 0/0 type situation[1].

We argue that even though Eq. (5.61) is well behaved in both the small mass ratio and the (anti)alignment between $\boldsymbol{S}$ and $\boldsymbol{L}$ limit, the same need not be true for its PN expansion. Consider the following function[2]

$$h(x; h_2, h_1, h_0) = \sqrt{\frac{h_2^2}{x^2} + \frac{h_1}{x} + h_0}, \tag{A.1}$$

and its expansion around $x = 0$

$$h^{\rm exp}(x; h_2, h_1, h_0) = \frac{h_2}{x} + \frac{h_1}{2h_2} - \frac{h_1^2 - 4h_0 h_4}{8h_2^3} x + \mathcal{O}(x^2). \tag{A.2}$$

[1] We have verified that this is the case both analytically and numerically.

[2] The similarity between our toy function and J given in Eq. (5.41) is not accidental.

Clearly, $h(x)$ is finite as $h_2 \to 0$. However, $h^{\rm exp}(x)$ is not, and the $h_2 \to 0$ limit is worse and worse as we keep more terms in the x expansion.

This is exactly the situation we encounter with Eq. (5.61) both in the small mass ratio limit, and in the approximate $\boldsymbol{S}$ and $\boldsymbol{L}$ (anti)alignment limit. Fixing the small mass ratio limit is straightforward: we identify the problem as originating from expanding the J multiplying the entire right hand side of Eq. (5.61), and factor it out. This is the reason behind the form Eq. (5.64) has.

The second problem is more complicated. We can still identify the terms that, when expanded, cause the limit when $\boldsymbol{S}$ and $\boldsymbol{L}$ are (anti)aligned to be problematic. However, if we do not expand them, we can no longer perform the integral of Eq. (5.64). Using this fully expanded $\langle \phi_z \rangle_{\rm pr}$ causes 5% of the systems studied here to have faithfulnesses below threshold.

In light of this, we tried a number of alternative, approximate methods for calculating $\langle \phi_z \rangle_{\rm pr}$. We discovered that if we keep the terms D_2 and D_4 in Eqs. (5.62) and (5.63), our results are greatly improved by about an order of magnitude: only 0.8% of the systems are below the faithfullness threshold. We examined a number of different definitions for D_2 and D_4, from using their initial value as given directly from Eqs. (5.62) and (5.63) to retaining different orders in a PN expansion, but evaluated at the initial time. We found out that these methods give comparable results, so we choose, for simplicity, to set D_2 and D_4 equal to their leading PN order:

$$D_2 \to \frac{c_p - \sqrt{c_p c_m}}{R_m \eta^2}, \tag{A.3}$$

$$D_4 \to \frac{cp(c_p - \sqrt{c_p c_m})}{R_m^2 \eta^4} - \frac{\sqrt{c_p c_m}}{2 R_m \eta^2}, \tag{A.4}$$

where c_p, c_m, R_m are defined in Appendix D of [61].

We expect this problem to be solved if we consistently PN expand both (D_2, D_4) and the roots S^2_+, S^2_- (see Appendix D of [61]). The complexity of the roots' expansion poses some serious problems in this calculation and we here opt for the approach described above and the partial resummation of the roots explained in Appendix D of [61]. This approach yields satisfactory results for the waveform precision required for aLIGO (see Fig. 5.6), but can be improved if need be through expansions appropriate for these systems, like a small misalignment between $\boldsymbol{S}$ and $\boldsymbol{L}$ expansion.

APPENDIX B

SYSTEMATIC AND STATISTICAL ERROR

In practice no waveform model will be a perfect match to the signals produced by Nature. The mismatch between the true signal and the theoretical model leads to a *systematic* error in parameter recovery. On top of that, parameter extraction also suffers from *statistical* errors arising from noise in the detectors. The systematic errors correspond to a deterministic, while the statistical errors correspond to a random displacement away from the signal manifold. To leading order, for nearby waveform families, the two types of error are independent. Let us assume we have two waveform models h_1 and h_2 that are qualitatively the same, so that we can write

$$h_2 = h_1 + \delta h, \tag{B.1}$$

where δh is small in the sense $(\delta h|h) \ll (h|h)$ and $(\delta h|\delta h) \ll (h|h)$.

The statistical error on the extracted parameters for either model can be approximated through the Fisher information matrix

$$\Gamma_{ij} = (h_{,i}|h_{,j})\,. \tag{B.2}$$

To leading order in the waveform mismatch we have

$$\Gamma_{2ij} = \Gamma_{1ij} + (h_{1,i}|\delta h_{,j}) + (\delta h_{,i}|h_{1,j})\,. \tag{B.3}$$

The correction terms will be small if the derivatives of δh are also of order δh. To see this is indeed the case, we let h_1 be the most accurate PN model we can construct (the numerical PN one) and h_2 be the double-precessing model of [62]. The difference between the two models is that the latter misses corrections of $\mathcal{O}(\chi^2)$, and thus, δh

is of 1.5PN order. To see this, write the waveforms as

$$h_1 = Ae^{i\Phi_1}, \quad h_2 = Ae^{i\Phi_2}, \tag{B.4}$$

assuming the amplitudes are the same, since they do not have a large impact from a data analysis point of view. Then,

$$\delta h = A(e^{i\Phi_1} - e^{i\Phi_2}) = A(e^{i\Phi_1} - e^{i\Phi_1 + i\delta\Phi}) \sim -Ae^{i\Phi_1} i\delta\Phi, \tag{B.5}$$

and its derivative is

$$\delta h_{,i} \sim -Ae^{i\Phi_1} i\delta\Phi_{,i} + Ae^{i\Phi_1}\Phi_{1,i}\delta\Phi. \tag{B.6}$$

Since all the derivatives are with respect to the parameters, they do not change the PN order of the terms. Therefore, $\delta h_{,i} \sim \delta h$ and $\Gamma_{2ij} \simeq \Gamma_{1ij}$.

We also present a numerical demonstration of the above result by computing the entries of the Fisher matrix for the two waveform models (the numerical PN and the analytical double-precessing one), evaluated at the same parameters. We do so for the eight physical parameters of interest: the chirp mass $\mathcal{M}$, the total mass m, and the six spin parameters $(\theta_1, \phi_1, \chi_1, \theta_2, \phi_2, \chi_2)$ and a system with $(m_1, m_2, \chi_1, \chi_2) = (7.4M_\odot, 6M_\odot, 0.5, 0.5)$. The mismatch between the numerical PN and analytic waveform, $1 - F$, for this system was 0.043. The ratio of the entries for

the two matrices $\Gamma^{an}_{ij}/\Gamma^{num}_{ij}$ is

$$\begin{pmatrix} 1.01 & 1.02 & 1.05 & 1.06 & 1.01 & 1.01 & 0.97 & 1.04 \\ 1.02 & 1.03 & 1.06 & 1.08 & 1.02 & 1.02 & 0.99 & 1.05 \\ 1.05 & 1.06 & 1.09 & 1.11 & 1.05 & 1.04 & 1.01 & 1.08 \\ 1.06 & 1.08 & 1.11 & 1.12 & 1.07 & 1.06 & 1.03 & 1.09 \\ 1.01 & 1.02 & 1.05 & 1.07 & 1.01 & 1.01 & 0.98 & 1.04 \\ 1.01 & 1.02 & 1.04 & 1.06 & 1.01 & 1.00 & 0.96 & 1.03 \\ 0.97 & 0.99 & 1.01 & 1.03 & 0.98 & 0.98 & 0.94 & 1.01 \\ 1.04 & 1.05 & 1.08 & 1.09 & 1.04 & 1.03 & 1.01 & 1.07 \end{pmatrix}. \tag{B.7}$$

As expected, the fractional difference in the Fisher matrix entries is of order the mismatch. The similarity of the corresponding entries of the Fisher matrix demonstrates that the mismatch between the numerical PN and the analytical waveform has little effect on the statistical errors calculated here.

APPENDIX C

HIGHER-D SAVAGE DICKEY DENSITY RATIO

Let M_1 be a simple model nested in the more complex M_2. Suppose that the two nested models have a common set of parameters $\vec{\lambda}$. M_2 has additional amplitudes $\vec{A}$ and angular parameters $\vec{\theta}$. The likelihood for the more complex model $p(d|\vec{A},\vec{\theta},\vec{\lambda},M_2)$ reduces to the likelihood of the simpler model $p(d|\vec{\lambda},M_1)$ when $\vec{A}=\vec{0}$. In other words, the angular parameters only affect the likelihood when the amplitude parameters are nonzero. The evidence for the simpler model is

$$p(d|M_1)=\int p(d|\vec{\lambda},M_1)p(\vec{\lambda}|M_1)d\vec{\lambda}, \tag{C.1}$$

and the evidence for the more complex model is

$$p(d|M_2)=\int p(d|\vec{A},\vec{\theta},\vec{\lambda},M_2)p(\vec{\lambda},\vec{A},\vec{\theta}|M_2)d\vec{\lambda}d\vec{A}d\vec{\theta}. \tag{C.2}$$

The posterior density for model M_2 is then

$$p(\vec{\lambda},\vec{A},\vec{\theta}|d,M_2)=\frac{p(d|\vec{A},\vec{\theta},\vec{\lambda},M_2)p(\vec{\lambda},\vec{A},\vec{\theta}|M_2)}{p(d|M_2)}. \tag{C.3}$$

Now consider the situation where $p(\vec{\lambda},\vec{A},\vec{\theta}|M_2)=p(\vec{\lambda})p(\vec{A})p(\vec{\theta})$ and $p(\vec{\lambda}|M_1)=p(\vec{\lambda})$. The marginal posterior density for model M_2 over $\vec{\lambda}$ is then

$$p(\vec{A},\vec{\theta}|d,M_2)=\int p(\vec{\lambda},\vec{A},\vec{\theta}|d,M_2)\,d\vec{\lambda}, \tag{C.4}$$

and

$$p(\vec{A}=\vec{0},\vec{\theta}|d,M_2)=\frac{p(d|M_1)p(\vec{A}=\vec{0})p(\vec{\theta})}{p(d|M_2)}, \tag{C.5}$$

where we have used $p(d|\vec{A}=\vec{0},\vec{\theta},\vec{\lambda},M_2)=p(d|\vec{\lambda},M_1)$. Going a step further, we can marginalize over the angular parameters to arrive at $p(\vec{A}=\vec{0}|d,M_2)=\int p(\vec{A}=$

$\vec{0}, \vec{\theta}|d, M_2)d\vec{\theta}$ which yields

$$p(\vec{A} = \vec{0}|d, M_2) = \frac{p(d|M_1)p(\vec{A} = \vec{0})}{p(d|M_2)}. \tag{C.6}$$

We then see that the BF between models 2 and 1, defined as $B_{21} = p(d|M_2)/p(d|M_1)$, is given by the Savage-Dicke density ratio

$$B_{21} = \frac{p(\vec{A} = \vec{0})}{p(\vec{A} = \vec{0}|d, M_2)}. \tag{C.7}$$

APPENDIX D

BAYES FACTORS AND PRIOR CUTOFFS

The NSNS Optimized noise curve has higher noise levels at high frequencies than the Zero-Detuned, High-Power one. Therefore, as far as EoS determination is concerned, the likelihood calculated with a NSNS Opt. noise curve does not contain a lot of information in the high-frequency part of the waveform where finite-size effects become important. Calculating BFs in a regime where it is not the likelihood, but the prior that dominates the results can lead to an interesting and rather counterintuitive effect: BFs that initially decrease with increasing SNR. In other words, as the signal strength increases, the correct model is preferred less and less.

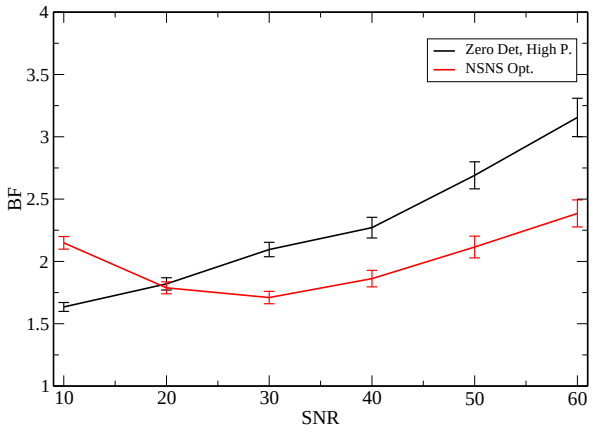

Figure D.1: BF in favor of SV compared to SV222 for the system with masses $(1.95, 1.9)$ with the Zero-Detuned, High-Power sensitivity curve (black) and the NSNS Optimized curve (red). For high SNR values, where the likelihood dominates, the detuning curve gives higher BFs. For lower SNR values with the NSNS Opt. curve, we encounter the counterintuitive effect of decreasing BFs with increasing SNR values.

Figure D.1 shows the BF in favor of SV compared to SV222 for the $(1.95, 1.9)M_{\odot}$ system calculated with the Zero-Det., High-P. curve (black) and the NSNS Opt. curve (red). The black line is the same as the black dashed line of Fig. 8.3. The red line presents some rather interesting behavior. At low SNR, the BF decreases with the

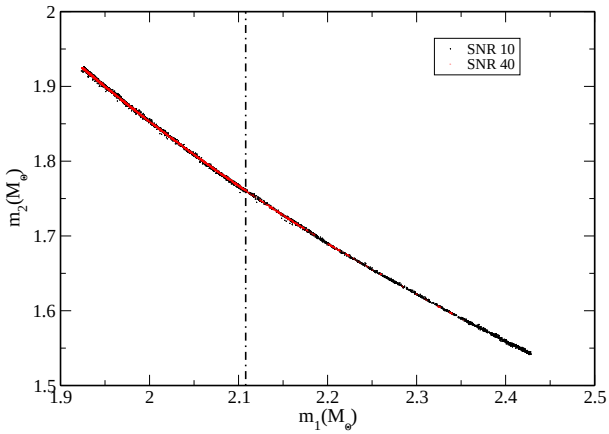

Figure D.2: 2D Scatter plot in the $m_1 - m_2$ plane for SNR 10 (black) and 40 (red) with the NSNS Opt. sensitivity curve and with the same system as in Fig. D.1. The vertical line denotes the maximum SV222 mass. Any points on the right of this line necessarily correspond to the SV model.

SNR, while after SNR = 40 it starts increasing, like one would expect. In order to understand this effect, consider Fig. D.2, where we plot the 2−D scatter plot of the chain points in the $m_1 - m_2$ plane for the system used in Fig. D.1, analyzed with NSNS Opt., and SNR 10 (black dots) and 40 (red dots). The vertical line indicates the maximum mass the competing model SV222 can support. This maximum mass is a hard cutoff on the prior mass of SV222; no system with masses higher than the maximum mass can be produced by SV222, and the model reduces to noise. This effectively means that SV222 has no posterior weight in that region, and all these points correspond to SV.

As the SNR increases the posterior width decreases. At SNR = 40 there are very few points above the mass cutoff. That means that with increasing SNR SV222 suffers less and less from this cutoff in its mass prior. This explains why the BF in

favor of SV decreases as we increase the SNR. Of course, at some values of the SNR the differences between the models will start dominating over the prior cutoff, and the BF will again start increasing with the SNR. Figure D.1 shows that in this case this is true for SNR > 40.

APPENDIX E

TOY MODEL

To demonstrate the rather counterintuitive effect of BFs that favor the wrong model, or decrease with the SNR, we construct a simple toy problem, thus obtaining a more robust explanation of these effects than the 2-D scatter plots presented in the previous appendix. Imagine we receive N data points from a very simple signal that obeys $d(f) = f$ and we try to match it with two competing $1-$D models $h_1(f) = af$ and $h_2(f) = af^{1.5}$, where a is the parameter of the models. The likelihood for model i is

$$L_i = \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{ -\sum^{N} \frac{[d(f) - h_i(f)]^2}{2\sigma^2} \right\}. \tag{E.1}$$

The parameter σ is the standard deviation of the data; in GW language, it is 1/SNR. Clearly h_1 is the correct model and L_1 is maximized when $a = 1$. On the other hand, h_2 cannot fit the signal perfectly; L_2 is maximized for $a = 1.12$ with a residual that depends on σ.

Now imagine that the two models have different prior ranges for the parameter a. In the case of h_1 we have $a \in (0, 2)$, while h_2 allows $a \in (0, 2\kappa)$, where κ is an arbitrary number. The evidence for each model is proportional to its likelihood integrated over the range of parameter a, while the BF in favor of h_1 (the correct model) is the ratio of the two evidences. Figure E.1 shows the BF in favor of h_1 as a function of the SNR ($= 1/\sigma$) for three different values of κ. The case $\kappa = 1$ corresponds to 2 models with the same parameter prior range, while $\kappa = 0.7$ and $\kappa = 1.4$ are similar to comparing EoS models with different maximum allowed mass. More specifically, $\kappa = 0.7$ corresponds to the case where the wrong model has the smaller prior range (the right panels of Figs. 8.3, 8.5, and Fig. D.1), while $\kappa = 1.4$ corresponds to the case where it is the correct model that has the smaller prior range (the left panels of Figs. 8.3 and 8.5).

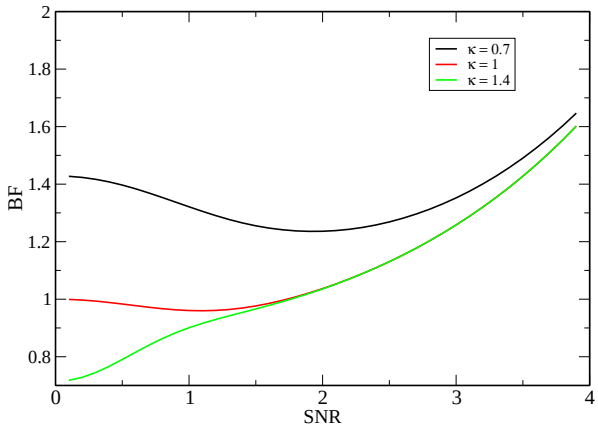

Figure E.1: BF in favor of h_1 as a function of the SNR for $\kappa = 1$ (red), $\kappa = 0.7$ (black), and $\kappa = 1.4$ (green). The effect of the BF decreasing with increasing SNR is present in both $\kappa = 1$ and $\kappa = 0.7$ cases, but it is more pronounced in the case where the two models have a different parameter range. As the SNR increases all three lines tend to coincide. This is because the likelihood becomes more and more peaked, and at some value of the SNR the limits of integration stop affecting the integral of the likelihood.

The $\kappa = 1.4$ case is easier to understand. The likelihood of the wrong model h_2 is integrated over a larger parameter region $a \in (0, 2.8)$ than the correct model where $a \in (0, 2)$. When the SNR is low and the models are not very different from each other, this can lead to a larger evidence for the wrong model. As the signal strength increases the differences between the two models will start dominating the evidence and the correct model will end up being preferred. Indeed we find that for SNR $\gtrsim 2$ the BF favors the correct model h_1.

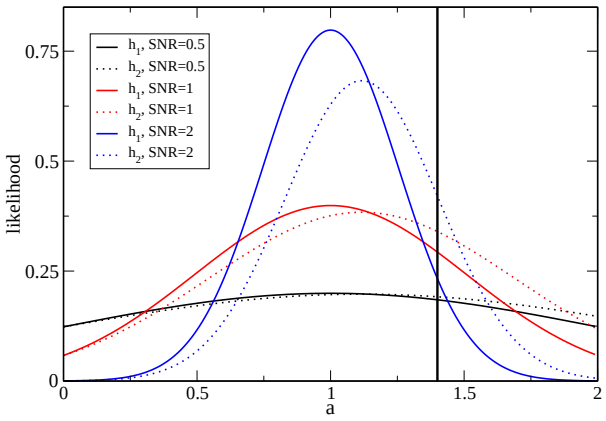

Figure E.2: Likelihood for the correct model (solid lines) and the wrong model (dotted lines) for different values of the SNR. The black vertical line shows the cutoff in the parameter prior range when $\kappa = 0.7$. The likelihood for the wrong model does not peak at the injected value of $a = 1$ and it is not a Gaussian.

Both $\kappa = 0.7$ and $\kappa = 1$ lead to BFs that decrease with increasing SNR, however, this effect is more pronounced in the $\kappa = 0.7$ case and it extends to higher values of SNR making it easier to identify. When $\kappa = 1$, one might expect the Laplace approximation to the evidence to be reliable since the posterior width on a is smaller than the prior range. However, this is not the case: the likelihood for h_2 is not a

Gaussian, and it is this small deviation from Gaussianity that we see as a decreasing BF. To visualize this effect, in Fig. E.2 we plot the likelihood for h_1 and h_2 for different values of SNR. For low SNR values the likelihoods are essentially the same and BF = 1. However, as the SNR increases the area under the red dotted line is larger that the area under the red solid line, leading to a small decrease in the BF in favor of h_1. Clearly at sufficiently high SNR the correct model will prevail and the BF will start increasing in favor of h_1.

When $\kappa = 0.7$ the effect of BFs that decrease with SNR is stronger and persists for higher values of SNR. Revisit Fig. E.2 and keep in mind now that the evidence for the wrong model is obtained by integrating the likelihood up to the black vertical line. As the SNR increases, the likelihood for h_2 is more peaked on the left of the vertical line, which means that the area permissible by the prior cutoff increases. This leads to an increase of the evidence of h_2 and a BF in favor of h_1 that decreases. Clearly at some point the differences between the models will overcome this effect, and the correct model will prevail. For our toy model this happens when SNR $\gtrsim 2$.

APPENDIX F

FAITHFULNESS REQUIREMENT

The agreement between two waveforms h, $\bar{h}$ with parameters $\vec{\lambda}$ is measured in terms of the faithfulness F

$$F(h,\bar{h}) = \frac{(h|\bar{h})}{\sqrt{(h|h)(\bar{h}|\bar{h})}}\,. \tag{F.1}$$

In the high SNR regime, a typical waveform sample from the posterior distribution function [220] is given by

$$\bar{h} = h + h_{,i}\Delta\lambda^i + \frac{1}{2}h_{,ij}\Delta\lambda^i\Delta\lambda^j + \ldots \tag{F.2}$$

where the $\Delta\vec{\lambda} = \vec{\lambda} - \vec{\lambda}_0$ are described by the multivariate normal distribution

$$p(\Delta\vec{\lambda}) = \sqrt{\det(\Gamma/2\pi)}e^{-\Gamma_{ij}\Delta\lambda^i\Delta\lambda^j/2}\,, \tag{F.3}$$

with $\Gamma_{ij} = (h_{,i}|h_{,j})$ and $\vec{\lambda}_0$ are the true parameters. Treating the $\Delta\lambda^i$ as small and expanding we get

$$F = 1 - \frac{1}{2}g_{ij}\Delta\lambda^i\Delta\lambda^j + \ldots \tag{F.4}$$

where

$$g_{ij} = \frac{(h_i|h_j)}{(h|h)} - \frac{(h|h_{,i})(h|h_{,j})}{(h|h)^2}\,. \tag{F.5}$$

Using $E[\Delta\lambda^i\Delta\lambda^j] = C^{ij} \simeq \Gamma^{-1}_{ij}$, we find

$$E[F] \simeq 1 - \frac{(D-1)}{2\,\mathrm{SNR}^2}\,, \tag{F.6}$$

for the expectation value of the faithfulness, where D is the dimension of $\vec{\lambda}$. The factor of D comes from $C^{ij}\Gamma_{ij} \simeq \delta^i_i = D$ and the factor of -1 from the $(h|h_{,i})(h|h_{,j})/(h|h)^2$

removing the dependence on the overall amplitude of the waveform, thus reducing the dimensions count by one.

The expected value of the faithfulness in Eq. (F.6) describes the impact of statistical errors. In deciding how accurate a waveform model needs to be, we should at a minimum demand that the systematic errors from mis-modeling are smaller than the statistical errors. If we wish to model spin-precessing binaries with $D = 8$ intrinsic parameters for systems with SNRs up to 50 then the modeling unfaithfulness should be below $8/5000 = 0.0016$ (there is no -1 for just intrinsic parameters, the amplitude is extrinsic). For and SNR of 25 we obtain the faithfulness requirement of 0.994 that we used in Ch. 5.

To calculate the variance, it is easier to work with the unfaithfulness, $1 - F$. The expectation of the square is given by

$$\begin{aligned} E[(1-F)^2] &= \frac{1}{4} g_{ij} g_{kl} E[\Delta\lambda^i \Delta\lambda^j \Delta\lambda^k \Delta\lambda^l] \\ &= \frac{1}{4} g_{ij} g_{kl} \left(C^{ij} C^{kl} + C^{ik} C^{jl} + C^{il} C^{jk} \right) \\ &\simeq \frac{3(D-1)^2}{4\,\mathrm{SNR}^4} \end{aligned} \tag{F.7}$$

Thus

$$\mathrm{var}[1-F] = \frac{2(D-1)^2}{4\,\mathrm{SNR}^4} \tag{F.8}$$

This shows that the average fathfulness is slightly less than 1-σ from a perfect faithfulness ($\sigma/\sqrt{2}$ to be precise). This agrees with what we see when computing the distribution of the match from MCMC waveform samples. The distribution is not Gaussian, and has a larger tail toward small values of the match.

An alternative derivation of the faithfulness requirement makes direct use of the posterior distribution function in the case of uniform priors

$$p(\vec{\lambda}) \sim e^{-\frac{(d-h|d-h)}{2}}\,, \tag{F.9}$$

where d is the data. The peak of the posterior, evaluated at the best-fit parameters is

$$\begin{aligned} p(\vec{\lambda}_{bf}) &\sim e^{-\frac{(d-h_{bf}|d-h_{bf})}{2}} \sim e^{-\frac{(d|d)+(h_{bf}|h_{bf})-2(d-h_{bf})}{2}} \\ &\sim e^{-\frac{\mathrm{SNR}^2+\mathrm{SNR}^2-2\mathrm{SNR}^2\,\mathrm{FF}}{2}} \sim e^{-\mathrm{SNR}^2(1-\mathrm{FF})}, \end{aligned} \tag{F.10}$$

where FF is the fitting factor, or the faithfulness maximized over all model parameters. The posterior on the true parameters is

$$p(\vec{\lambda}_0) \sim e^{-\frac{(d-h_0|d-h_0)}{2}} \sim e^{-\mathrm{SNR}^2(1-\mathrm{F})}, \tag{F.11}$$

From Eq. (F.3) we can calculate the value of the multidimensional posterior $1-\sigma$ away from the best-fit parameters

$$p(\vec{\lambda}_{1-\sigma}) \sim e^{-\frac{\Gamma_{ij}\Delta\lambda^i\Delta\lambda^j}{2}} \sim e^{-\frac{\Gamma_{ij}C^{ij}}{2}} \sim e^{-\frac{D}{2}}\,, \tag{F.12}$$

Assuming that the model can fit the data perfectly for some parameters (an assumption that will lead to a conservative faithfulness threshold) we set FF=1 and requiring that the true parameters are less than $1-\sigma$ away from the best-fit ones we find

$$1-\mathrm{F} < -\frac{D}{2\mathrm{SNR}^2}, \tag{F.13}$$

where D is the number of parameters whose measurability is affected by the model inaccuracy. For spin-precessing models with 8 intrinsic parameters, $D = 8$.

This derivation translates the results of [139] that were written in terms of requirements on the GW amplitude and phase to requirement on the faithfulness.