TESTING GENERAL RELATIVITY WITH GRAVITATIONAL WAVES

by

Laura Mae Sampson

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

of

Doctor of Philosophy

in

Physics

MONTANA STATE UNIVERSITY
Bozeman, Montana

March 2014

ACKNOWLEDGEMENTS

Thank you to my advisor, Neil Cornish, to the rest of my committee members, in particular Nico Yunes, and to all of my professors in the physics department. A special thank you also goes to Margaret Jarrett, who made sure that all of the paperwork hoops involved in graduate school were as painless as possible.

TABLE OF CONTENTS

TABLE OF CONTENTS – CONTINUED

LIST OF TABLES

LIST OF TABLES – CONTINUED

LIST OF FIGURES

LIST OF FIGURES – CONTINUED

Figure Page

LIST OF FIGURES – CONTINUED

Figure Page

LIST OF FIGURES – CONTINUED

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

LIST OF FIGURES – CONTINUED

Figure Page

LIST OF FIGURES – CONTINUED

ABSTRACT

In the next few years, physicists will make the first direct detection of the phenomenon called gravitational waves. These waves, which are propagating perturbations in the metric that describes the geometry of space-time, are Einstein's last unconfirmed prediction. Among the most interesting science we will be able to accomplish with these observations is in the area of testing General Relativity. In this dissertation, I give a brief introduction to General Relativity and gravitational waves, and then spend the bulk of the document explaining how we can test General Relativity using gravitational waves. In particular, I focus on the data analysis techniques that will be necessary for performing such tests.

1. INTRODUCTION

General Relativity (GR) has been widely accepted for almost a century as the most accurate theory available to predict the outcome of gravitational interactions. Its description of gravity as a curvature in space-time rather than a force that can change instantaneously revolutionized our understanding of the universe we live in. In part because of the large philosophical departure in this theory from Newtonian gravity that came before it, GR has been tested extensively since the date of its publication. It has survived all of the tests unscathed.

It is important to remember, however, that to date all of the systems that have been used to test GR are either slow-moving (compared to the speed of light), not in a regime of large space-time curvature, or both. This means that there is still a possibility that GR is not the correct theory of gravity. In the near future, gravitational waves (GWs) will allow us to probe a gravitational regime that has not been accessible to any other tests. My thesis work has been focused on the implementation of these GW tests of GR.

The rest of this document is organized as follows. In Ch. 2, I give a brief introduction of linearized gravity within GR, GW detectors, and a few types of GW sources. In Ch. 3, I summarize the non-GW tests of GR that have been performed to date, and in Ch. 4 I introduce some of the non-GR theories of gravity that are still considered viable. Next, in Ch. 5, I derive the connections that can be drawn between the various known schemes for parameterizing deviations from GR. The rest of the thesis is focused on data analysis. In Ch. 6, I give a brief introduction of Markov chain Monte Carlo (MCMC) and Bayesian data analysis techniques, and, finally, in Ch. 7 I show how these techniques can be applied to testing GR using GW observations. Chapters 5 and 7 are based on my original work.

2. GRAVITATIONAL WAVES IN GENERAL RELATIVITY

In this chapter, I give a brief introduction to General Relativity and the description of gravitational waves within that theory.

2.1. Linearized Gravity

General Relativity (GR) is a theory of gravity, first published by Einstein in 1915 [57]. In this theory, the Newtonian concept of a gravitational 'force' between two massive objects that causes them to, for example, orbit each other is described instead as as a curvature in space-time. Massive objects create this curvature, and objects in free fall follow straight lines, or geodesics, in the curved space-time. The movement of these objects changes the curvature of space, which influences their motion, which changes the curvature, and on and on. This relationship was nicely summed up by the oft-quoted John Wheeler: 'Matter tells space how to curve; space tells matter how to move.'

Mathematically, this relationship can be summed up as

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{8\pi G}{c^4} T_{\mu\nu}. \tag{2.1}$$

On the left-hand-side of this expression, $g_{\mu\nu}$ is the space-time metric, $R_{\mu\nu}$ is the Ricci tensor, and R is the Ricci scalar. All of these quantities encode information about the curvature of space-time. On the other side of the equation, $T_{\mu\nu}$ is the energy-momentum tensor, which encodes the nature of the mass (or energy) that creates this curvature. G is the gravitational constant, and c is the speed of light. (For the rest of this document, we will use units in which $G = c = 1$. Additionally, we will use the convention that Greek letters refer to all four space-time indices, and Latin letters refer only to the three spatial indices.)

This equation is a deceptively simple way to write 10 coupled, partial differential equations. These equations are highly non-linear, and extremely difficult to solve in general. There are certain situations, however, in which exact or approximate solutions can be found. (The following discussion follows closely with Maggiore, Ch. 1 [94].)

For example, consider a region of space that is nearly flat - i.e. far from any large concentrations of mass or energy. In this situation, the metric can be approximated as a perturbation from the flat space-time, or Minkowski, metric, $\eta_{\mu\nu}$:

$$\eta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \tag{2.2}$$

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \qquad |h_{\mu\nu}| \ll 1 \tag{2.3}$$

where $h_{\mu\nu}$ is the small perturbation. By writing $|h_{\mu\nu}|$, we have implicitly chosen a coordinate system useful for the physical situation we are interested in. In this particular coordinate system, Eq. (2.3) holds in a sufficiently large region of space. We insert this definition into the definition of the Ricci tensor and Ricci scalar,

$$R_{\mu\nu} = R^{\alpha}_{\mu\alpha\nu}, \qquad R = g^{\mu\nu} R_{\mu\nu}, \tag{2.4}$$

where the right-hand-side of the first equation is a contraction of the Riemann tensor, which is calculated by taking derivatives of the metric. We calculate both of these quantities, but keep only terms that are linear in our perturbation, $h_{\mu\nu}$.

From this procedure, we arrive at a linearized version of Einstein's equation, which will tell us about possible solutions in this nearly-flat space. Before we write this linearized equation, we define the quantities $h = \eta^{\mu\nu} h_{\mu\nu}$, and

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2} h \eta_{\mu\nu}. \tag{2.5}$$

With these definitions in hand, the linearized Einstein equation takes the form

$$\Box \bar{h}_{\mu\nu} + \eta_{\mu\nu} \partial^{\rho} \partial^{\sigma} \bar{h}_{\rho\sigma} - \partial^{\rho} \partial_{\nu} h_{\mu\rho} - \partial^{\rho} \partial_{\mu} \bar{h}_{\nu\rho} = -16\pi T_{\mu\nu}. \tag{2.6}$$

We now use our freedom of gauge choice to impose

$$\partial^{\nu} \bar{h}_{\mu\nu} = 0, \tag{2.7}$$

known as the Lorentz condition. Under this condition, Eq. 2.6 becomes the very simple

$$\Box \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}, \tag{2.8}$$

which is a wave equation with the source term $T_{\mu\nu}$. This is essentially what we mean by the term gravitational wave (GW) - a wavelike solution to the linearized Einstein equations.

This result, Eq. (2.8), which has the source on the right hand side, is how we calculate the generation of GWs in the linearized theory. To learn about how GWs travel through space and interact with detectors, we are interested in the vacuum solution to this equation, i.e. the solution outside the source:

$$\Box\bar{h}_{\mu\nu} = 0. \tag{2.9}$$

We can simplify this equation even further. The imposition of the Lorentz gauge, Eq. (2.7), does not use up all of our degrees of freedom. It reduces the 10 independent components of $\bar{h}_{\mu\nu}$ to six independent components. We use our further gauge freedom to impose that the trace of $\bar{h}_{\mu\nu}$ is zero, $\bar{h} = 0$. Under this condition, $\bar{h}_{\mu\nu} = h_{\mu\nu}$. Next, we choose that $h^{0i} = 0$. These new requirements simplify the Lorentz condition to

$$\partial^0 h_{00} = 0. \tag{2.10}$$

That is, h_{00} is constant in time. The GW is defined to be the time-dependent portion of the metric, so for the purpose of calculating GWs, we have effectively set $h_{0\mu} = 0$, and we are left with only the spatial components. These choices collectively define the *transverse-tracelses gauge*, and are summarized as

$$h^{0\mu} = 0, \quad h^i_i = 0, \quad \partial^j h_{ij} = 0. \tag{2.11}$$

The transverse-traceless gauge will be represented in the following by superscript TT.

Eq. (2.9) admits solutions of the form $h^{TT}_{ij}(x) = e_{ij}(\vec{k})e^{ikx}$, which are plane waves with wave vector $k^\mu = (\omega, \vec{k})$ and $\omega = |\vec{k}|$. The tensor e_{ij} is defined as the polarization tensor. The gauge choices laid out in Eq. (2.11) imply that the non-zero components of h^{TT}_{ij} are all in the plane transverse to $\hat{n} = \vec{k}/|\vec{k}|$, the direction of propagation (hence the word 'transverse' in 'transverse-traceless'). For instance, if we define $\hat{n}$ to be along the z direction, we get, after taking the real part of the solution,

$$h_{ij}^{TT} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix} \cos[\omega(t-z)]. \tag{2.12}$$

Here, h_+ and $h_\times$ are called the 'plus' and 'cross' polarizations, respectively.

Next, we examine the polarization tensors, $e_{ij}^{+/\times}$. Their definitions are

$$e_{ij}^{+} = \hat{\mathbf{u}}_i\hat{\mathbf{u}}_j - \hat{\mathbf{v}}_i\hat{\mathbf{v}}_j, \qquad e_{ij}^{\times} = \hat{\mathbf{u}}_i\hat{\mathbf{v}}_j + \hat{\mathbf{v}}_i\hat{\mathbf{u}}_j, \tag{2.13}$$

where the plus and cross again label the two polarizations. In this definition, $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ are unit vectors that are orthogonal to the direction of propagation, $\hat{\mathbf{n}}$, and to each other. Again choosing the frame where $\hat{\mathbf{n}}$ is along the z direction, the polarization tensors can be written

$$e^{+} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad e^{\times} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{2.14}$$

To understand why the two polarizations are named in the way that they are, we examine the effect of a GW on a ring of test masses in free fall. Because they are in free fall, all of the masses will follow geodesics of the space-time. These geodesics will be perturbed as the wave passes, and the masses will follow the new geodesics. In curved space, geodesics that are initially close together can grow further apart with time. This is called geodesic deviation. In the rest-frame of the ring of particles, which we will call our detector, the equation for geodesic deviation, caused by the metric perturbation h_{ij}^{TT}, is

$$\ddot{\xi}^i = \frac{1}{2}\ddot{h}_{ij}^{TT}\xi^j, \tag{2.15}$$

where ξ^i is the separation of the two geodesics.

Finally, we can examine the effect of a GW passing through our detector. Again, let the wave be traveling in the z direction, and take the ring of particles to be oriented in the xy plane. Let us examine the effect of a wave that is polarized entirely in the plus mode. The wave can be written

$$h_{ab}^{TT} = h_+ \sin \omega t \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.16}$$

where h_+ is the amplitude of the wave. Eq. (2.15) then becomes

$$\delta \ddot{x} = -\frac{h_+}{2}(x_0 + \delta x)\omega^2 \sin \omega t, \tag{2.17}$$

$$\delta \ddot{y} = \frac{h_+}{2}(y_0 + \delta y)\omega^2 \sin \omega t. \tag{2.18}$$

At linear order, the terms δx and δy on the right hand side can be neglected, and these equations have the solutions

$$\delta x(t) = \frac{h_+}{2} x_0 \sin \omega t, \tag{2.19}$$

$$\delta y(t) = -\frac{h_+}{2} \sin \omega t. \tag{2.20}$$

This means that the separation of two geodesics in the y direction gets smaller as the separation in the x direction gets larger, and vice versa. This leads to the ring being 'stretched,' first along y, then along x, as illustrated in Fig. 1. Also shown in this figure is the effect of the cross mode of polarization, which can be calculated in direct analogy with the plus mode.

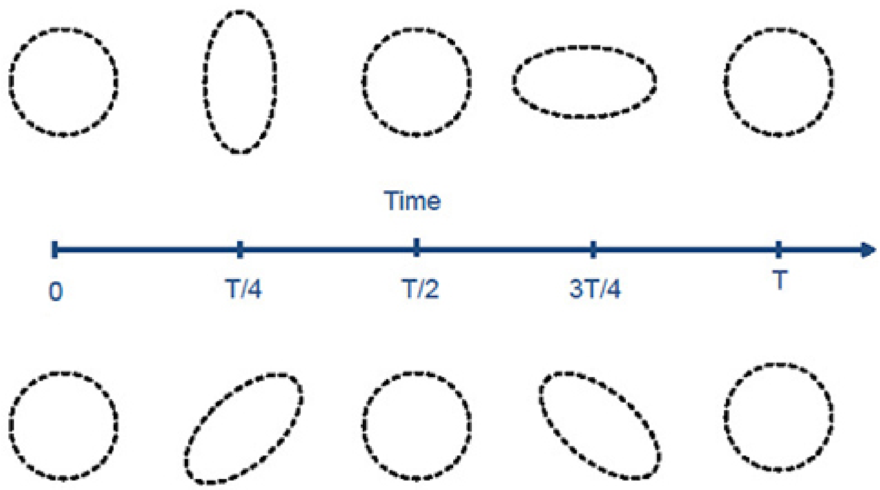

Figure 1: The effect of the two different polarizations of GWs on a ring of test particles. The top panel shows the plus polarization, and the bottom panel shows the cross. [75]

2.2. Gravitational Wave Detectors

Currently, there are two types of GW detectors in active operation or active development. These are interferometric detectors, notably the Laser Interferometer Gravitational Wave Observatory (LIGO) detectors in the US and the VIRGO detector in Italy [107], and pulsar timing arrays, which are the NanoGRAV [97] collaboration in the US, the European Pulsar Timing Array (EPTA) [62] in Europe, and the Parkes Pulsar Timing Array in Australia [96], which will soon begin collaborating on the International Pulsar Timing Array (IPTA) [95]. In addition, there is a concerted effort to launch a space-based interferometric detector similar to the Laser Interferometer Space Antenna (LISA) mission [8, 65] currently going forward in the European Union.

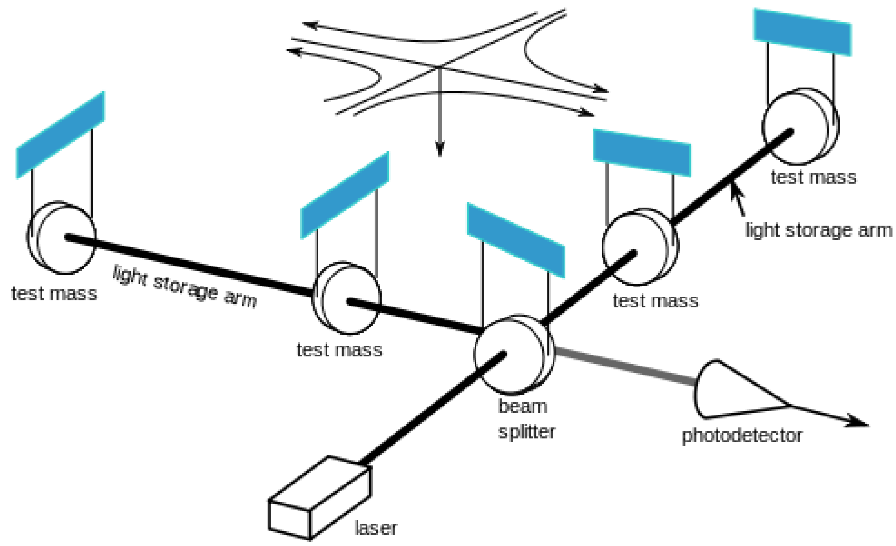

Figure 2: Cartoon of a Michelson interferometer, like the LIGO and VIRGO detectors. A beam of light is split in two and sent down both arms of the detector. Each half of the beam bounces off of the test mass in one arm, and the beam is recombined at the photodetector. Changes in the length of the two arms will results in changes of the light intensity at the photodetector, which allows us to use these interferometers to detect GWs. [36]

2.2.1. Interferometric Detectors

Interferometric detectors can be understood by examining a cartoon picture of the LIGO-type detectors, shown in Figure 2. The detectors consist of essentially a Michelson interferometer, with two long arms that are perpendicular to each other. At the end of each arm is a mirror. A laser is sent into the detector, split in half at the intersection of the two arms, reflected back at the end of each arm, and recombined in the middle. The interference pattern produced by the recombined beam can be used to measure the length of the two arms with high precision.

When a GW passes through the detector, the arms are stretched - first in one direction, then the other. Because of the quadrupolar nature of gravitational radiation, as one arm of the interferometer gets longer, the other gets shorter. Thus the interference pattern of the recombined laser beam shifts with a predictable pattern when a GW is incident on the detector.

That is essentially the idea behind interferometric detection of GWs, whether in space or on the ground. The reality is, of course, much more complicated, and there are excellent review articles on the subject [107, 64].

2.2.2. Pulsar Timing Array

The other type of detector that is currently being assembled is a pulsar timing array. Pulsars are quickly-rotating neutron stars, that have a beam of radiation shooting out along an axis that is not perfectly aligned with the spin axis of the star. This beam sweeps around the star, analogously to the beam of light from a lighthouse, and periodically hits the Earth, where it can be detected by radio telescopes. This periodicity is incredibly regular, especially in the case of pulsars with a spin timescale on the order of milliseconds, which means that millisecond pulsars can be thought of as some of the most precise clocks in the universe, and we can predict the arrival time of their pulses of electromagnetic radiation incredibly well with our timing models [25, 73].

The pulses travel to Earth from the pulsars at the speed of light, and so their arrival time at our detectors depends on the distance between Earth and the pulsar that emitted them. As a GW passes through the space between Earth and the pulsar, this distance is stretched and squeezed, exactly like the arms of an interferometer. This stretching and squeezing means that the pulse does not arrive in our detector at the exact time that it is predicted by our timing model, and this difference can be detected in order to measure the GW that caused it.

Pulsar timing arrays are called arrays because they consist of many pulsars, each being timed independently. These independent data sets can be used to increase our sensitivity to GWs, and also to increase the precision with which we can measure source parameters [52, 149].

2.3. Gravitational Wave Sources

The form of GWs that will be produced in nature can be calculated from Eq. (2.8), with T_μ determined by the particular source configuration. There are several main types of sources that we expect to detect in the coming years. These are inspiraling binary systems, composed of black holes, neutron stars, and white dwarfs; 'bursts' of gravitational radiation, for example from supernovae; and stochastic GW backgrounds, possibly from quantum gravity effects before the epoch of recombination.

Inspiraling binary systems are perhaps the most important type of signal we expect to detect - in particular binaries that do not merge within the frequency band of a particular detector. In LIGO, for example, we expect to be able to detect the inspiral of neutron-star binaries, but their merger will be at a frequency that is too high to be noticeable in these detectors. These inspiral signals are useful because their precise form, both amplitude and phase, can be calculated in the approximation scheme called the post-Newtonian expansion. In this scheme, discussed in more detail in Chapter 5, the equations of motion governing a binary system in GR are expanded in the small parameter, v/c, where v is the characteristic velocity of the system and c is the speed of light. The expected gravitational waveform can then be calculated, and the precise form depends sensitively on the physical parameters of the system.

For a non-spinning binary in a circular orbit, for instance, the first few terms of the gravitational wave (GW) phase are [24]

$$\Psi(f) = \frac{3}{128\eta}(\pi M f)^{-5/3}\left[1 + (\pi M f)^{2/3}\left(\frac{3715}{756} + \frac{55}{9}\eta\right) + \ldots\right] + 2\pi f t_c - \phi_c, \quad (2.21)$$

where t_c and ϕ_c are the coalescence time and phase, M is the total mass of the system, η is the symmetric mass ratio, $\eta = (m_1 m_2)^{3/5}/M^{1/5}$, and f is the frequency of the GW. Measuring the phase terms that depend on M and η will tell you the mass and mass ratio of the system that is producing the GWs. An important consequence of our ability to predict the precise form of these inspiral signals within GR is that we can use these signals as tests of GR. If we detect an inspiral signal that is significantly different from our predictions, this could point to GR being the incorrect theory of gravity. This concept is explored in detail in Chapter 7.

In contrast to binary signals, burst type GW signals, for instance those predicted to be caused by supernovae, are not well-modeled. Although they could not be used to measure source parameters such as masses and spins, the detection of such signals would nonetheless be significant. For instance, they could tell us how much energy from a typical supernova is emitted in GWs. Perhaps more excitingly, it is wholly possible that we will detect bursts of GWs from sources that we have not been able to dream up. These types of signals would, of course, be un-modeled, but would point us to previously unknown phenomena in our universe [118].

A final type of GW signal we may soon be able to detect is that of a stochastic background. This background will be due to both the superposition of many distinct sources, and to fundamental processes such as the Big Bang. In one detector over a short period, stochastic GW backgrounds are indistinguishable from stochastic noise. One way around this is to have one detector with a long observing time, paired with a signal that changes over, say, the course of a year due to the location of the Earth in its orbit. In this case, the background would change with time and the noise would not, and so the two could be separated. The other option is to have more than one detector, and to cross-correlate the signals in the two. In this case, the stochastic background signal in the two detects would be correlated and so add coherently, creating a larger than expected correlation in the data streams of the two detectors [118].

Both burst sources and stochastic backgrounds are interesting phenomena that have been studied extensively [118]. Burst sources, in particular, have even been proposed as a means for testing GR [74]. Our work, however, has concentrated on the use of inspiral signals to test GR with GWs, and so that is the type of signal that we will concentrate on for the rest of this document.

3. NON-GW TESTS OF GR

General relativity was a revolution in scientific thought. It fundamentally changed the way that we conceive of the force of gravity, and of how matter, energy, and space interact. Because of its large shift away from Newtonian physics, GR was met almost immediately with a bevy of experimental challenges, all of which it surmounted. As time went on and the theory proved hugely successful in predicting the behavior of gravitational fields, the motivation for testing it changed, and the precision of these tests grew. We are now in an era in which the predominant motivation for testing GR is not the belief of scientists that Newtonian gravity will be vindicated. It is rather in the hope that in some energy regime or on some length scale we will detect deviations from GR that will point us in a theoretical direction to help overcome the fundamental disagreement between GR and quantum mechanics.

In this chapter we present a brief summary of some of the non GW-based tests of GR, starting with some of the earliest. Each of these tests has served to enhance our confidence in GR as the correct theory of gravity that describes our universe. In the near future, GW tests will allow us to probe a regime with stronger fields and higher velocities than we have yet been able to access, and it is possible that at that point we will finally find results that are not in agreement with Einstein's theory.

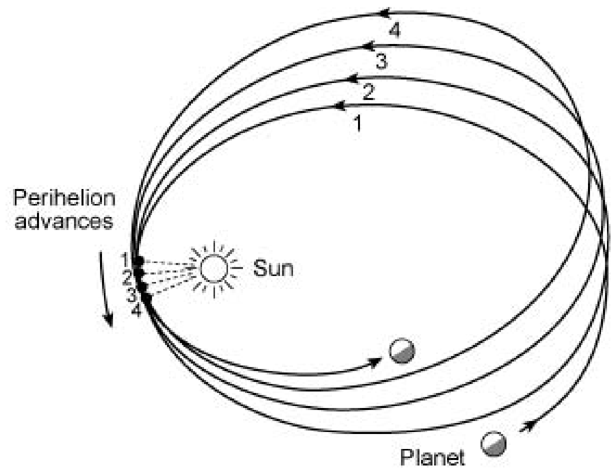

Figure 3: The perihelion precession of a planet around the Sun - greatly exaggerated so as to be visible. (Copyright 2010, Professor Kenneth R. Lang, Tufts University)

3.1. Perihelion Precession of Mercury

One of the firsts vindications of the theory of GR was its successful explanation of the anomalous perihelion precession of Mercury. Both Newtonian gravity and GR predict that objects should orbit the sun in elliptical orbits with the sun at one focus. In both theories, the ellipse of a planet's orbit does not stay static with time, but precesses around the sun, creating a lissajous pattern as shown in Figure 3. GR predicts a precession of this type independent of any external perturbations to the orbit. Newtonian gravity, on the other hand, predicts this precession as a consequence perturbations away from a perfect, ever-repeating ellipse - for instance due to the presence of more than one planet in the Solar System [122]. The rate of precession predicted from Newtonian gravity is smaller than the observed rate by 43 arc seconds per century.

This discrepancy was noted as early as 1859 by Urbain le Verrier, who re-examined observations of Mercury's transit across the sun collected from 1697 to 1848 [122]. He found that even after the perturbations of the other planets and the effects of the precession of the equinoxes were taken into account, there was still an anomalous precession of Mercury's orbit [144]. After the anomaly was discovered, many ad hoc solutions were proposed. For example, some physicists proposed that Newton's inverse square law for gravitational forces was incorrect, and that gravity was in fact described by a more complicated power law [144]. None of these explanations were successful, until finally in the early 20th century, when it was shown that GR correctly predicts the rate of Mercury's precession.

In GR, there is automatically a discrepancy between the angular and radial phase of an object orbiting in a Schwarzschild space-time such as the space-time around our Sun. This is because all objects which are moving under only gravitational influences will follow geodesics defined by

$$u^\alpha \nabla_\alpha u^\beta = 0, \tag{3.1}$$

where u is the four-velocity of the object in question, and ∇ represents the covariant derivative.

In Schwarzschild geometry, the geodesics obtained from this equation can be manipulated to produce orbital equations given by

$$\left(\frac{dr}{d\phi}\right)^2 = \frac{r^4}{b^2} - \left(1 - \frac{r_s}{r}\right)\left(\frac{r^4}{a^2} + r^2\right), \tag{3.2}$$

where r_s is the Schwarzschild radius, and a and b are constants of the motion that have been introduced for convenience [144].

In analogy to analysis in Newtonian gravity, we can introduce an effective potential for the Schwarzschild geometry, given by

$$V(r) = -\frac{Mm}{r} + \frac{L^2}{2\mu r^2} - \frac{(M+m)L^2}{\mu r^3}, \tag{3.3}$$

where M is the mass of the sun, m is the mass of the planet, r is their separation distance, L is the orbital angular momentum of the Sun-planet system, and $\mu = mM/(m+M)$. The first two terms are familiar from Newtonian gravity. The third term is not part of the Newtonian equation, but is present in GR. It leads to many interesting effects, in particular to the mismatch between the angular and radial phases of the orbit.

To understand where this mismatch comes from, we consider a circular orbit that has been perturbed by a small amount. It will oscillate with radial frequency

$$\omega_r^2 = \omega_\phi^2 \sqrt{1 - \frac{3r_s^2}{a}}, \tag{3.4}$$

where ω_ϕ is the Newtonian orbital frequency for a circular orbit, given by

$$\omega_\phi = \frac{r_s^4}{16a^6}. \tag{3.5}$$

The precession of perihelion is due to the difference between these two frequencies. It is given by

$$\Delta\phi = T(\omega_\phi - \omega_r), \tag{3.6}$$

where T is the orbital period. When appropriate numbers for Mercury are used in this calculation, the predicted rate of perihelion precession is $\Delta\phi = 42.98 \pm 0.04$ arc seconds per century [144] - precisely in line with the observed value.

3.2. Gravitational Lensing

Another early test of GR was the confirmation of the GR prediction that massive objects will 'bend' light rays - an effect known as gravitational lensing. In Newtonian gravity, only particles with mass are affected by gravitational forces. Before GR was written down, scientists had already seen evidence that at some times light acted as if it had mass and so calculated the amount that a light ray would be bent by the Sun assuming that the Sun pulled on it as it does every massive particle. In GR, on the other hand, every particle, including massless photons, will travel along a geodesic if it is not perturbed by external forces. Thus GR has its own prediction for the bending of light by the Sun. For a light ray traveling near to our Sun, the Newtonian prediction for the deflection angle ends up being about half of the GR prediction.

In GR, a light ray passing near the Sun will be bent from its straight-line path by an angle

$$\delta\theta = \frac{1}{2}(1+\gamma)\frac{4M_{\odot}}{d}\frac{1+\cos\Phi}{2}. \tag{3.7}$$

A more useful quantity is the change in relative angular separation between two sources when the light rays from one of them passes very close to the Sun. This

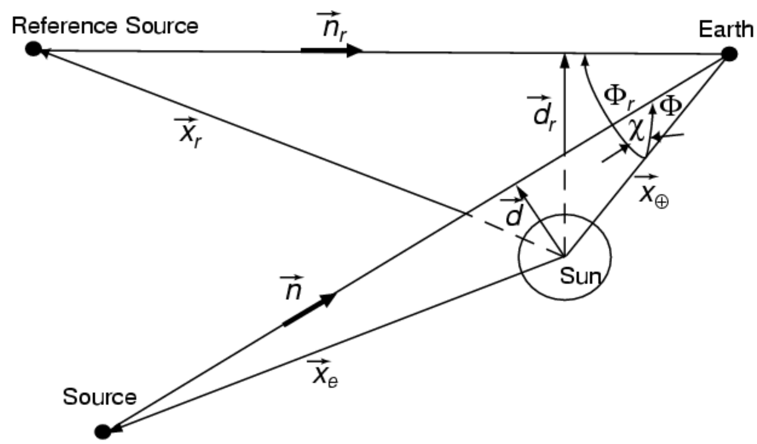

Figure 4: Light from a distant star being bent by the Sun as it travels to Earth. [144]

separation is given by

$$\delta\vartheta = \frac{1}{2}(1+\gamma)\left[-\frac{4M_\odot}{d}\cos\chi + \frac{4\text{M}_\odot}{d_r}\frac{1+\cos\Phi}{2}\right], \tag{3.8}$$

where the geometric variables are as in Figure 4 [144], and γ is the same parameter from the previous section, and is equal to 1 in GR.

In 1919 Sir Arthur Eddington led an expedition to the West coast of Africa so that he could make observations of the position of stars near the sun during a total solar eclipse. Because the mass of the Sun is not sufficient to produce a large effect on the path of light rays, they must pass very close to the Sun for the deflection to be observable at all. For this reason the expedition team set out to measure the apparent location of stars very nearby the sun relative to reference stars whose light does not pass near enough to the Sun to be detectably bent. Eddington's measurements agreed with the GR prediction, however we know now that they were only accurate to around 30%. Modern techniques such as radio interferometry have since been used to measure the angular separations to better than 100 μarcseconds, and have yielded a constraint of

$$\frac{\gamma+1}{2} = 0.99992 \pm 0.00023, \tag{3.9}$$

which is in perfect agreement with the GR prediction.

3.3. Shapiro Time Delay

Another classic Solar System test of GR is the Shapiro time delay. This effect is observed when light is bounced off of a massive object and returned to the emitter. In GR, this trip takes slightly longer than in classical theories of gravity, because light is effected by the gravitational well of the massive body. In classical gravity, the effect of the gravitational pull on the photon would be to speed it up on the way towards the body, and slow it down on the way back. These two effects would precisely cancel, and the photon could be treated as traveling at a uniform speed throughout. In GR, the photon really *does* travel at a uniform speed, c, throughout the trip, but the space-time it is traveling through is distorted away from flat space by the presence of the massive body. This means that the distance the photon must travel is larger than it would be if the body were not there, resulting in a delay.

The potential for this effect was first noticed in 1964 by Irwin J. Shapiro, for whom it was named. Shapiro calculated that a beam shot from Earth, traveling close to the sun, and reflected off of Venus should be delayed by $\sim 200\mu s$ due to GR effects. This was well within the ability of 1960's technology to measure, and so the first tests were performed in 1966 and 1967 with the MIT Haystack radar antenna [144, 120]. Shapiro and his team used the 120 ft. antennae, which were donated to MIT by the US military, to measure the delay of light bounced off of Venus. Their measurements agreed with predictions.

The simplicity of this prediction makes it an elegant test for the theory of GR, and as such many more experiments have been performed. Various objects have been used as targets, including Venus and Mercury as passive reflectors, and artificial satellites, including Mariners 6 and 7, Viking 2, and the Viking Mars orbiters and landers as active retransmitters.

The Shapiro delay can also be thought of as an effect of time dilation caused by gravitation. Although the speed of light is constant, the curvature of space-time in the area near a massive body causes a discrepancy between the proper time of a photon traveling through that curved space and that of an inertial observer far away. For a ray of light that passes close to the Sun, the discrepancy is approximately equal to

$$\delta t = \frac{1}{2}(1+\gamma)(240 - 20\log\frac{d^2}{r})\mu s, \tag{3.10}$$

where d is the distance of closest approach to the Sun, and r is the distance from the Sun to the source. Finally, γ is a parameter that is equal to 1 in GR. Tests of GR via the Shapiro delay are usually quoted in the constraints they can put on the quantity $1-\gamma$.

When performing an experiment to measure this effect, there is no Newtonian baseline against which to compare the measured value. That is, we can predict how much longer it would take for light to gravel from Earth to Venus and back in both Newtonian gravity and GR, but we are only able to measure one of these numbers, not the difference between the two. Thus to claim that the data supports GR, it is necessary to take multiple measurements in order to observe the logarithmic nature of Eq. (3.10). In order for these measurements to be meaningful, the motion of the Earth relative to the object being observed must be taken into account. This is accomplished by using long-range radar to estimate the ephemeris of the object when it is far from superior conjunction, and using this measurement to predict the position of the object when it is near superior conjunction (i.e. when it is on the opposite side of the Sun from Earth). Then, as the object approaches and passes through superior conjunction, successive measurements of the light travel time can be made, and these delays can be fit to Eq. (3.10) using a least-squares regression which returns the best-fit value for $\frac{1}{2}(\gamma + 1)$.

3.4. Testing the Equivalence Principle

In GR, there are several versions of 'the equivalence principle,' each of which stems from the equivalence between inertial and gravitational mass. The weak equivalence principle (WEP) of GR states that all test particles should behave the same in a gravitational field, regardless of their makeup. In Newtonian mechanics, this is the statement of the fact that a bowling ball and a feather should accelerate at the same rate in the same gravitational field. This principle has been tested in the laboratory using torsion balances [136]. In these experiments, test bodies of the same mass but different compositions are suspended at opposite ends of a rod, which in turn is suspended by a fiber. A difference in the gravitational force felt by the different bodies will result in a torque on the rod, which can be measured. To this point, there have been no violations of the WEP detected.

The strong equivalence principle (SEP) is similar to the WEP, but states that all bodies should behave the same in a gravitational field regardless of their size. That is, bodies of the same mass but different densities, meaning they posses different 'gravitational self energies,' should respond in the same way to the same gravitational field. As with the other tests of GR, a deviation from the SEP can be parameterized as [144]

$$\frac{m_G}{m_I} = 1 - \eta \frac{E_G}{M}, \tag{3.11}$$

where m_G is the gravitational mass, m_I is the inertial mass, M is the total mass energy, and E_G is the gravitational self energy. In GR, $\eta = 0$.

The value of η has been experimentally constrained using Lunar Laser Ranging (LLR) data [98]. Barring any contribution from the WEP, a difference between the acceleration of the Moon towards the Sun and the Earth towards the Sun would indicate a violation of the SEP. This difference would lead to a 'polarization' of the Moon's orbit towards the Sun as seen from Earth. That is, the Moon's orbit would be slightly stretched in the direction of the Sun. This polarization can be described as a perturbation in the distance between the Earth and the Moon:

$$\delta r = -2.9427 \times 10^{10} \left[\left(\frac{m_G}{m_I} \right)_{\oplus} - \left(\frac{m_G}{m_I} \right)_{☾} \right] \cos D[m], \tag{3.12}$$

weigh D the angle between the mean longitude of the Moon and the mean longitude of the Sun as observed from Earth. This orbital effect is known as the Nordtvedt effect, named after Ken Nordtvedt [103, 145], who noticed the possibility of its existence in various non-GR metric theories of gravity. If it exists, it can be detected by very accurately measuring the distance between the Earth and the Moon throughout their orbit. This is what is accomplished in LLR experiments.

The first measurement of laser light reflected off of the retroreflectors left on the Moon by the Apollo 11 astronauts was made in 1969. Since then, the LLR experiment has made regular measurements of the roundtrip travel time between the Earth and Moon, currently accurate at the 50 ps ($\sim$ 1cm) level [144]. These measurements are then fit to a theoretical model of the Moon's motion, which takes into account tidal interactions, perturbations due to the Sun and other planets, and GR effects as decried by the post-Newtonian formalism. Because of the sensitivity necessary to detect the Nordtvedt effect, researchers must also take into account librations of the moon and atmospheric effects when evaluating the travel time of the laser light. So far, with several teams performing data analysis, no evidence has been found for the existence of this effect.

3.5. Constraints from Binary Pulsars

The most relativistic systems that we have yet been able to access are stellar binaries consisting of two neutron stars. These systems reach velocities of $v \sim 2 \times 10^{-3}$, and have gravitational compactnesses of $\frac{M}{r} \sim 6 \times 10^{-6}$, where r is the orbital separation. We can measure the dynamics of these systems because many of the stars are pulsars - neutron stars with a powerful beam of radiation emitted along their magnetic field, which is not aligned with the spin axis of the star. These beams of radiation sweep past the Earth as the neutron star rotates, acting like incredibly precise light houses. By measuring the time of arrival of these pulses of radiation, we can extrapolate the orbital dynamics of the binary system, and thus determine if the gravitational physics is consistent with the predictions of GR.

The most famous of these systems is the Hulse-Taylor binary [128]. In 1974, Russell Hulse and Joseph Taylor discovered this system using the radio Arecibo telescope. The pulse period was 59 milliseconds, but there were periodic changes in the arrival time of the pulses. With a predictable period of 7.75 hours, the pulses would arrive either a little early or a little late, which showed that the pulsar was a member of a binary system. If GR is correct, the orbit of the binary should decay slowly due to the emission of gravitational waves. By measuring the period of the orbit over many years, Hulse and Taylor were able to show that the rate of orbital decay is 0.997 ± 0.002 times that predicted by GR [138]. The measured data are plotted on top of the GR prediction in Figure 5.

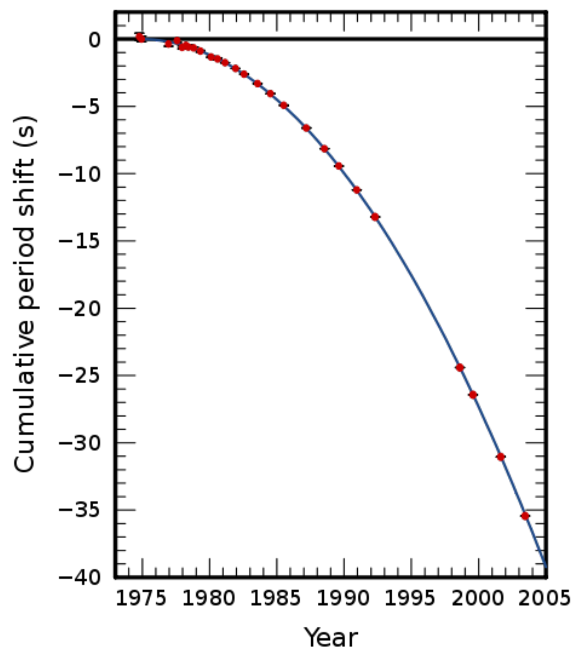

Figure 5: The measured orbital period for the Hulse-Taylor pulsar (points) plotted with the GR prediction (line). The agreement between the measured values and the theoretical predictions is very precise. [138]

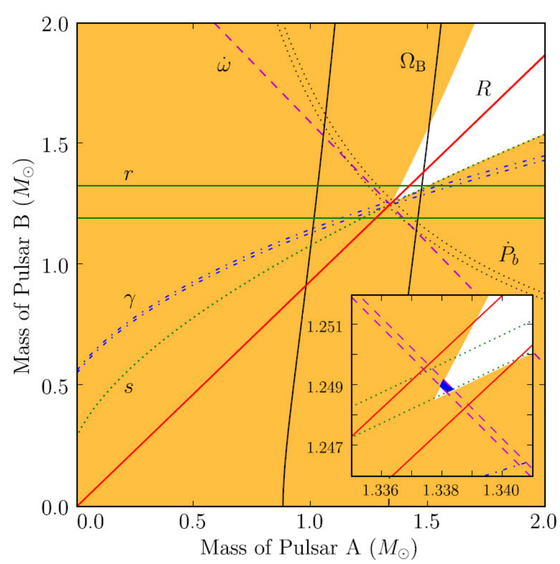

Figure 6: The measured value of parameters that characterize possible deviations from GR in binary pulsar orbits. Each parameter is a function of the two masses. If all of the measured curves overlap in the same region of the mass1- mass2 plane, the data is consistent with GR. As we can see, the curves do, in fact, overlap. [85]

Another very important pulsar system is the double pulsar binary, PSR J0737-3039 [29]. This is the only known system in which both components are pulsars - i.e., the beam of radiation from *both* neutron stars passes over the Earth. This has allowed for a very precise measurement of the masses of both components of the system, and for very precise tests of GR.

In Ch. 5 we will discuss a system for parameterizing deviations from GR that was developed for these pulsar systems. For now, suffice it to say that the time of arrival of pulses from neutron star binaries can be predicted very precisely, taking into account many effects that are predicted by GR. Using the double pulsar binary system, Kramer et. al [85] have been able to measure many of these effects, and show that they are in agreement with GR. A plot summarizing their results is shown in Figure 6.

4. NON-GR THEORIES OF GRAVITY

As discussed in Ch. 3, GR has been tested extensively since its publication, and every experimental test of the theory has vindicated its predictions. Because of this, GR remains the accepted theory of gravity to describe our universe, but it is important to remember that none of these tests have been in the strong-field, dynamical regime that will be probed by GWs. It is because of this, amongst other reasons, that developing non-GR theories of gravity is still an active field of research. In this chapter, we present a few of these alternative theories that have been explored in the context of GWs. This chapter follows closely with Section II of [165].

4.1. Theoretical Properties

There are many theoretical alternatives to GR, but here we will only consider a handful in detail. In selecting which theories to discuss, we are driven by a few theoretical constraints. These are:

- **Weak-field agreement with GR**. This is necessary because GR has proven to be so accurate in all of the weak-field tests that we have so far been able to perform. Another way of stating this point is that there must be some limit in which the theory reproduces the predictions of GR to within experimental precision, or it is already ruled out.
- **Motivation from fundamental physics**. The alternative theory is derived from some fundamental underlying theory. This fundamental theory should

be proposed to solve some physical problem, such as the conflict between GR and quantum mechanics. For example, many alternative theories of gravity are low-energy effective theories derived from string theory.

- **Well-posedness**. There must be a unique solution to the modified field equations that is determined by initial data.
- **Strong-field deviation from GR**. In order for the theory to be detectable as different from GR using GW observations, the theory must predict a deviation from GR in the strong-field, dynamical regime in which GWs will be produced. This means that we are not interested theories of gravity that produce, e.g. changes from GR at cosmological length scales, for instance BLAH, but fail to produce large deviations for inspiraling binaries.

With these properties in mind, we will now discuss a couple of different classes of theory that have been well explored in the context of GWs.

4.2. Scalar-Tensor Theories

Scalar-Tensor theories of gravity add a scalar field to the action that couples to the metric. This results in an action of the form [60, 59]

$$S = \frac{1}{16\pi}\int d^4x\sqrt{-g}[R - 2(\partial^\mu\varphi)(\partial_\mu\varphi) - V(\varphi)] + S_m[\psi_m, A^2(\varphi)g_{\mu\nu}]. \tag{4.1}$$

This expression is in the so-called Einstein frame, and in it, $g_{\mu\nu}$ is not the physical metric. In this expression, φ represents the scalar field, $V(\varphi)$ is a potential function, ψ_m represents the matter degrees of freedom, and S_m is the action for matter.

In order to place this theory in the more intuitive frame in which the metric governs space-time separations, the Jordan frame, it can be re-written by making the conformal transformation $\bar{g}_{\mu\nu} = A^2(\varphi)g_{\mu\nu}$. In this new frame, $\bar{g}_{\mu\nu}$ is the physical metric, and the action takes the form

$$S = \frac{1}{16\pi}\int d^4x\sqrt{-\bar{g}}\left[\phi R - \frac{\omega(\phi)}{\phi}(\partial^\mu\phi)(\partial_\mu\phi) - \phi^2 V\right] + S_m. \tag{4.2}$$

The new scalar field is represented by ϕ, where $\phi \equiv A^{-2}$, and $\omega(\phi)$ is called the coupling field. When this field is constant, i.e. $\omega(\phi) = \omega_{\rm BD}$, this theory reduces to massless Jordan-Fierz-Brans-Dicke (Brans-Dicke) theory.

In massless Brans-Dicke theory weak-field agreement with GR is recovered if $\omega_{\rm BD} < 1/40000$. This bound is set by tracking data from the Cassini spacecraft [22]. More generally, scalar-tensor theories of this form, barring a certain class that are described by homogeneous solutions to the scalar field evolution equations, recover GR in the limit that $\omega \to \infty$, and so our first property is fulfilled. The second property is also fulfilled - scalar-tensor theories can be derived as the low-energy limit of certain string theories [61, 66]. Scalar-tensor theories can also be shown to be well-posed, i.e. in [113, 87], which satisfies our third condition. The final property, that there must be strong-field deviations from GR, is not actually met in general by this type of scalar tensor theories, as they do not lead to corrections at high order in the curvature. Rather, the primary modification of this theory in the GW sector is to introduce dipolar gravitational radiation, which dominates in the weak-field regime.

A more general class of scalar-tensor theories, introduced by Damour and Esposito-Farese [45], are also defined by the action in Eq. (4.1), but defined by a different conformal factor, $A(\psi) = e^{\beta\psi^2/2}$, where β is some constant. In these theories, for $\beta < -4$, it is possible for neutron stars that are initially non-scalarized to acquire a strong scalar field when the gravitational energy is large enough. For instance, if two neutron stars with initially weak scalar fields are inspiraling towards each other, the scalar fields can suddenly become much larger when the separation between the two stars reaches a critical distance. This process is referred to as *spontaneous scalarization*, and results in a distinctive signature in the GW signal from the inspiral, satisfying our final condition. GW signals of this type are explored in Ch. 7, Sec. 7.8.

4.3. Massive Graviton Theories

Einstein's theory predicts that gravity propagates at the speed of light, which in the language of field theory implies that the force-mediating boson associated with gravity, the graviton, is massless. Much work has gone into developing theories in which the graviton is *not* massless, and thus in which gravitational signals travel slower than electromagnetic ones.

These theories possess the second property that we desire from alternative theories - they are well motivated by fundamental physics. In loop quantum cosmology, which is the cosmological extension of loop quantum gravity, gravitons acquire a mass in the loop quantization process [26]. In Dvali's string theory-inspired effective theory [56], gravitons have massive modes. Massive gravitons also arise in theories such as Rosen's bimetric theory, in which they couple to a different metric than photons [110]. There are more examples, but these suffice to show that massive gravitons arise in many well-motivated theories.

The first property, however, is not trivially met in massive graviton theories - in the limit $m_g \rightarrow 0$, GR is not in general recovered. The mass of the graviton leads to a scalar mode that is not present in GR, even as that mass goes to zero. This problem can be avoided by invoking certain nonlinearities. It has been shown in [133, 84, 50] that near enough to any mass, M, linear theories are no longer accurate. Because our particular examples of massive graviton theories are typically only understood at linear order, it is possible that they would give the proper GR limit if non-linear corrections were included. There is much ongoing work in this area.

Because we have no expression of an action in a massive graviton theory that produces the correct GR limit, we cannot examine the precise properties of such a theory. There are, however, some general effects of this type of theory that can be described phenomenologically. In particular, theories of gravity that include massive gravitons will lead to changes in the Newtonian gravitational potential, and to the propagation of GWs.

The modifications in the Newtonian potential correspond to a change of $V = (M/r) \rightarrow (M/r)\exp(-r/\lambda_g)$. That is, the potential is replaced by a Yukawa type potential, where r is the distance from the body, and λ_g is the wavelength of the graviton. This type of potential can be tested in many ways, including tidal interactions between galaxies [69] and weak gravitational lensing [35].

The propagation of GWs is modified in massive graviton theories due to the introduction of a dispersion relation. This dispersion relation takes the form

$$\frac{v_g^2}{c^2} = 1 - \frac{m_g^2 c^4}{E^2}, \tag{4.3}$$

where E is the energy of the graviton, calculated via $E = hf$. This type of dispersion relation indicates that GWs of different energies will travel at different speeds, and that all GWs will travel slower than the speed of light. The difference in speed for different energies will lead to phase accumulations in GWs that differ from the GR prediction, as GWs emitted at higher frequencies catch up with those emitted at lower frequencies. The slowness of GWs relative to c will lead to a delay of the GW signal from a given event relative to the electromagnetic signal.

4.4. Modified Quadratic Gravity

Modified quadratic gravity theories are a class of theories that add terms to the Einstein-Hilbert action that are quadratic in curvature. The general action for this type of theory is [166, 150].

$$\begin{aligned} S = \int d^4x \sqrt{-g} \Big\{ &\kappa R + \alpha_1 f_1(\theta) R^2 + \alpha_2 f_2(\theta) R_{\mu\nu} R^{\mu\nu} + \alpha_3 f_3(\theta) R_{\mu\nu\sigma\alpha} R^{\mu\nu\sigma\alpha} \\ &+ \alpha_4 f_4(\theta) R^*_{\mu\nu\sigma\alpha} R^{\mu\nu\sigma\alpha} - \frac{\beta}{2}[\nabla_\mu \theta \nabla^\mu \theta + 2V(\theta)] + \mathcal{L}_{\text{mat}} \Big\}. \end{aligned} \tag{4.4}$$

In this expression, $\mathcal{L}_{\rm mat}$ is the matter Langrangian, $f_i(\theta)$ are functionals of the field θ, and α_i and β are coupling constants. This action assumes that all quadratic terms couple to the same field, θ, which is a common restriction but not necessary. It is well motivated by particular examples of non-GR gravitational theories, such as Dynamical Chern-Simons, which is recovered when $\alpha_4 = -\alpha_{\rm CS}/4$, and all other $\alpha_i = 0$. Dynamical Chern-Simons theory is the low-energy expansion of certain types of string theories [6].

Because these theories smoothly recover the limit of GR as $\alpha_i \to 0$, these theories satisfy our first condition, that the GR limit is recovered to experimental precision, provided that the coupling constants are small enough. They also satisfy the final requirement, that deviations from GR appear in the strong-field, because the deviations are high-order in curvature. This guarantees that they will become more important as the curvature grows, i.e., in the strong-field regime. It is not clear that the requirement of well-posedness is met by this type of theory. It is important to remember, however, that quadratic gravity is not in general thought of as an exact solution, but rather as an effective theory that is only valid in certain energy regimes. Thus the well-posedness of the underlying theory is what is actually of interest, and not of the particular quadratic gravity theory that emerges.

There are a handful of other theories that have been explored in their relationship to gravitational waves, such as non-commutative geometries and theories with gravitational parity violation. And there are many other possible alternatives to GR, such as Einstein-Aether theory and Born-Infeld gravity, that are not ruled out by current experiment but have not been explored in detail in the context of gravitational waves. This large array of non-GR theories serves as the motivation for the next chapter, Ch. 5, which describes methods for simultaneously parameterizing many alternative gravity theories. These parameterizations allow us to constrain many theories with a single measurement, instead of undertaking the infinite task of ruling out every possible theory of gravity individually.

5. PARAMETERIZING ALTERNATIVES TO GENERAL RELATIVITY

Since its publication in 1915, GR has been tested extensively [144] - some of these tests are discussed in Chapter 3. To-date, the theory has survived every challenge. Beginning with the early tests of very weak-field gravity in our solar system, through to the most extreme laboratories of gravitational physics that we can yet access - binary pulsars - GR has had its predictions vindicated with extreme precision. There is, so far, no experimental evidence for GR being incorrect on small scales (i.e. solar system and compact binary systems), though flat galaxy rotation curves [100, 99] and the accelerated expansion of the Universe [37] have been put forward as evidence for deviation on much larger scales. On small scales, and for weak fields, the constraints on deviations from GR are so strong that we can say that alternative theories of gravity have effectively been ruled out in this regime.

But it is not true that we have effectively ruled out all alternative theories of gravity in all regimes. Aside from being an impossible endeavor, given our continued ingenuity in developing new theories of gravity, definitively ruling out all challengers to GR will require tests in the strong-field, dynamical gravitational regime. This is a regime that we have not yet been able to access.

The gravitational compactness of a system is defined as [165]

$$\mathcal{C} = \frac{G}{c^2}\frac{M}{R}, \tag{5.1}$$

where M is the characteristic mass of the system, and R is the characteristic length scale associated with gravitational radiation. The characteristic velocity of a system, $\mathcal{V} \sim \frac{v}{c}$, is a measure of the rate of change of the gravitational field. The weak field, then, is defined as the regime in which both $\mathcal{C}$ and $\mathcal{V}$ are very small, i.e. $\mathcal{C} \ll 1$ *and* $\mathcal{V} \ll 1$. The strong-field, dynamical regime is the regime in which *neither* of these conditions are met and a lowest-order perturbative analysis of the gravitational field equations does not suffice.

We have not yet been able to measure systems in the strong-field. Even for the double pulsar binary system (PSR J0737-3039) [86] , which boasts the strongest gravitational fields we have yet been able to directly probe, the compactness of the system is only $\mathcal{C} \sim 6 \times 10^{-6}$, with characteristic velocity of $\mathcal{V} \sim 2 \times 10^{-3}$. In contrast, compact binary coalescences will reach both compactness and velocity close to unity. Thus while binary pulsar data can teach us about strong gravitational fields, they are not dynamical, strong-field systems by our definition. Our lack of data in this regime means that there is still room left for alternative theories of gravity that do not predict outcomes strongly different from GR except in very strong-field, dynamical systems.

Most alternative gravity theories, including those discussed in Chapter 4, involve additional parameters in their field equations, such as coupling constants, or the mass of extra fields. These parameters are inherent to the theory, but their values are unknown and must be measured. Clearly, the parameters for each theory could be constrained via experiment, for instance by predicting the outcome of a given experiment for a given alternative theory of gravity. Although this is an obvious way forward, the work involved in making testable predictions for every known theory of gravity is enormous - even for the limited subset of theories that we have so far discovered. It is therefore important to develop methods for testing multiple theories of gravity simultaneously. It is with this goal in mind that the various generically parameterized models of non-GR gravity have been developed.

The parameterized post-Newtonian (ppN) formalism, developed by Will and Nordtvedt [103, 145, 104, 140], parameterizes the space-time metric in a way that captures many different types of gravitational theories. It is built on an expansion about Minkowski space, and is accurate in the weak-field regime, and on scales that are sufficiently small to ensure the accuracy of a linear perturbation. The ppN framework was designed for, and is thus ideally suited to, tests we can perform in our own solar system.

In contrast, the parameterized post-Keplerian (ppK) formalism was designed to allow us to perform tests of GR with pulsar timing data [63, 112, 78]. Developed in its current form primarily by Damour, Deruelle, and Taylor [49, 44], it gives a timing formula for the arrival of pulses at Earth that builds on an expansion on Kepler's laws. The timing formula phenomenologically takes into account many different gravitational effects, both in the binary motion of the components of a pulsar system, and on the travel of electromagnetic (EM) radiation away from such a system towards Earth. Measuring these parameterized effects allows us to test the nature of the underlying theory of gravity that describes our universe in regimes where the gravitational field is as strong as that induced by binary pulsars.

Adding to our knowledge from solar system experiments and measurements of the pulses from binary pulsar systems, in the near future we will have yet another tool for testing the fundamental nature of gravity - gravitational waves (GWs). The Fourier amplitude and phase of GWs from inspiraling compact objects have been predicted to extreme precision within GR [23], but not within alternative theories. To solve this problem, and also to allow a means of detecting or constraining un-modeled deviations from GR, Yunes and Pretorious [163] developed the parameterized post-Einsteinian (ppE) family of waveform templates. The ppE formalism has since been extended by Chatziioannou, Yunes, and Cornish [34] to describe the multiple polarization states that generically appear in extensions to GR. The parameters in the ppE model can be constrained by GW detections, and these constraints in turn will either build our confidence that GR is the correct theory of gravity, or indicate that we must deepen our knowledge of other possible theories in order to explain our observations.

The ppN and ppK formalisms have already been used successfully to constrain GR in relatively weak fields. New tests of GR using GWs will allow us to probe gravitational fields that are much stronger and more dynamical than anything we have had access to before. Yet we know that the underlying theory of gravity describing these types of systems is the same as that which governs the physics of our solar system and of pulsars. It would therefore be sub-optimal to proceed as if we have no knowledge of the constraints already placed on alternative theories of gravity when analyzing new data.

The purpose of this chapter is to draw connections between the ppE, ppN, and ppK formalisms, so that measurements that constrain parameters in one can be used to enhance our knowledge of constraints in the other two. Knowing the connection between the ppN, ppK, and ppE parameters will allow us to use known constraints from the solar system and binary pulsars to inform our search for deviations from GR in GWs. For instance, they can be used to construct a well-informed prior when conducting a Bayesian model-selection analysis [38, 115]. Going in the other direction, constraining ppE parameters using GWs could allow us to place stronger constraints on ppN and ppK parameters than we have been able to achieve using solar system and binary pulsar tests.

The main result of this chapter is the mathematical mapping between the three different parameterizations. First, we find the mapping between ppN and ppE parameters. In doing so, we discover that ppN modifications to the binding energy generically lead to 1 post-Newtonian (PN) corrections to the GW phase, if and only if the ppN parameters $\beta_{\rm ppN}, \alpha_1^{\rm ppN}$, or $\alpha_2^{\rm ppN}$ are modified. Most alternative theories studied thus far do not modify these parameters, which is perhaps why leading-order GR deviations that enter at 1PN order have not yet been found. Second, we find the mapping between ppK and ppN parameters. We find that although the Shapiro shape parameter is not modified by ppN corrections, the range parameter, the redshift parameter and the pericenter rate of change parameter are modified and directly related. Third, we find the mapping between ppK and ppE parameters by investigating ppE modifications to the conservative dynamics.

With these mappings at hand, we can then investigate how current bounds on ppN and ppK parameters already constrain ppE parameters. We find that binary pulsar constraints on ppK parameters already stringently bound ppE parameters in a certain weak-field regime of ppE parameter space, ie. when the ppE phase exponent parameter $b < 4$. In this regime, these bounds are stronger than future projected bounds with GWs [38]. We also find that Solar System constraints on the ppN parameters already stringently bound ppE parameters at 1PN order. Although this constraint holds only at a single point of ppE parameter space, they beat projected bounds from GWs by a factor of a few. Multiple GW observations, or one especially bright signal, will however be able to provide a more stringent bound than that inferred from bounds on ppN parameters from Solar System observations.

The rest of this chapter is arranged as follows. In the first few sections, we give a brief introduction of the ppN, ppK, and ppE parameterizations. In Sec 5.4, we find the correspondence between the ppN metric parameters and the ppE GW parameters. In Sec 5.5, we find the correspondence between ppN and the pulsar timing parameters of ppK. Next, in Sec 5.6, we show the connection between the orbital decay parameter of ppK, and the phase parameters of ppE. Finally, in Sec 5.7, we discuss current constraints on the various parameters, and finally we conclude and point to future research.

5.1. Parameterized post-Newtonian

Post-Newtonian (PN) formalism expands Einstein's equations beginning with the lowest-order deviations from Newtonian gravity. The small expansion parameter is typically the characteristic velocity of the system, which must be small in comparison to the speed of light. The formalism is valid only for weak fields, in systems composed of objects that travel slowly when compared to the speed of light - for instance, in our own solar system.

The ppN formalism, due primarily to Nordtvedt and Will [103, 145, 140], is a first-order PN framework to parameterize a large class of gravitational theories. It is useful for testing GR in the weak-field, and was developed with solar-system experiments in mind. The following outline of the ppN formalism closely follows Chapter 4 of [141].

In order for alternative theories of gravity to be studied in the ppN framework, they must be metric theories - i.e. theories in which there exists a symmetric tensor called the metric that governs proper distances and proper times, and in which matter and fields respond when acted upon by gravity via the equation

$$\nabla_\nu T^{\mu\nu} = 0, \tag{5.2}$$

where $T^{\mu\nu}$ is the stress-energy tensor for all matter and non-gravitational fields, and ∇_ν is the covariant derivative with respect to the metric. This requirement is equivalent to the statement that all theories of gravity in the ppN formalism must satisfy the Einstein equivalence principle, which states that, in a freely falling reference frame, all physical laws behave as if gravity were absent. In other words, there are no local experiments that one can perform to differentiate between a uniformly accelerated reference frame, and one that is freely falling with respect to the local gravitational field [141].

To generate the ppN formalism, we must begin with some book-keeping that allows us to keep track of the various "smallness" parameters used in the expansion. The first assumption is that the gravitational fields are *weak*. In this context, this means that the classical, Newtonian potential of all rest masses in the system is small:

$$U(t, \mathbf{x}) \equiv \int \frac{\rho(t, \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{x}' \ll 1, \tag{5.3}$$

where $\rho(t, \mathbf{x}')$ is the rest-mass density at $(t, \mathbf{x}')$.

In Newtonian gravity, the velocities of bodies in orbit around each other are related via the virial theorem to U through $v^2 \lesssim U$. These quantities are both considered to be of second-order, whereas quantities with a single power of velocity are first-order. For example, vU is a perturbative quantity of $\mathcal{O}(3)$.

With this book-keeping in hand, and recalling that in this formalism all gravitational theories are governed by a space-time metric, we can proceed to construct a general, ppN metric. As the ppN formalism is based on a perturbation of Newtonian gravity, we know that the ppN metric must reduce to the Minkowski metric at lowest order. Thus the most general metric one could construct would begin with the Minkowski metric, and then add PN metric terms from all possible functionals of matter variables, each multiplied by an arbitrary coefficient that could be set by matching to cosmological conditions. There are, however, an infinite number of these functionals, and so to produce a workable formalism, one typically adopts some restrictions:

- The metric coefficients should be of Newtonian or (first) PN order, and no higher;
- Perturbations to the Minkowski metric should go to zero at spatial infinity, so that the metric is asymptotically flat;
- The metric should be dimensionless;
- The metric should contain no explicit reference to the spatial origin or the initial moment of time;
- The metric components g_{00}, g_{0j}, and g_{jk} should transform as a scalar, vector, and tensor respectively;
- The functionals should depend on the rest-mass, energy, pressure, and velocity - not on gradients of these quantities;
- The functionals should be "simple."

We assume that the theory of gravity of interest can be described by a least-action principle, in which the Lagrangian is defined as

$$L = \left(-g_{\mu\nu} \frac{dx^\mu}{dt} \frac{dx^\nu}{dt} \right)^{1/2} = \left(g_{00} - 2 g_{0j} v^j - -g_{jk} v^j v^k \right)^{1/2} . \tag{5.4}$$

Because the velocity, v, is a first-order parameter, in order to keep each of these terms at the same order, we need to know g_{00} to $\mathcal{O}(4)$, g_{0j} to $\mathcal{O}(3)$, and g_{jk} to $\mathcal{O}(2)$.

Parameter	Value in GR	Value in semi-cons. theories	What does it measure?
γ_{ppN}	1	γ_{ppN}	How much space-time curvature is produced by a unit rest mass?
β_{ppN}	1	β_{ppN}	How much "nonlinearity" is there in the superposition law for gravity
ξ	0	ξ	Are there preferred location effects?
$\alpha_1^{\text{ppN}}, \alpha_2^{\text{ppN}}, \alpha_3^{\text{ppN}}$	0	$\alpha_1^{\text{ppN}}, \alpha_2^{\text{ppN}}, 0$	Are there preferred frame effects?
$\zeta_1^{\text{ppN}}, \zeta_2^{\text{ppN}}, \zeta_3^{\text{ppN}}, \zeta_4^{\text{ppN}}$	0	0	Violation of mom. conservation?

Table 1: The ten ppN parameters, as well as their physical significance, and their value in GR and in semi-conservative theories, in which energy and momentum are conserved.

As an example, let us consider g_{jk}. We have determined that this metric element must transform as a tensor, and contain functionals of the rest-mass, energy, pressure, and velocity that are no higher than second order in our expansion parameters. The only terms that can appear in g_{jk} that satisfy these restrictions, as well as the full list of restrictions above, are $U\delta_{jk}$ and U_{jk} where U_{jk} is given by

$$U_{jk} \equiv \int \frac{\rho(\mathbf{x}', t)(x - x')_j (x - x')_k}{|\mathbf{x} - \mathbf{x}'|} d^3x'. \tag{5.5}$$

The other metric components can be written in terms of similar functionals, each meeting the requirements described above. There are ten such functionals in total, after many have been eliminated for failing the final, and rather subjective requirement of "simplicity." After the metric is written in terms of these *metric potentials*, we make a choice of gauge and coordinate system that results in the final ppN parameterization.

The metric components are then written, in this coordinate system, as a collection of constants multiplying the metric potentials. For instance, the g_{jk} component becomes

$$g_{jk} = (1 + 2\gamma_{\text{ppN}} U)\delta_{jk}. \tag{5.6}$$

Here, γ_{ppN} is one of the aforementioned constants. These constants, of which there are ten, are called "ppN parameters." They are listed in Table 1 with their physical significance. In different theories of gravity, these ten parameters take on different values. Measurements of or constraints on these parameters can then either constitute evidence for a non-GR theory of gravity or place bounds on GR deviations.

These ppN parameters have been constrained by many experiments within our solar system, such as lunar laser ranging [148], gravitational redshift experiments [68, 53], and the measurement of Earth's tides [155]. The current limits on the ppN parameters, as well as the sources of those limits, are listed in Table 2.

Parameter	Effect	Limit
$\bar{\gamma}_{\rm ppN}$	time delay	2.3×10^{-5}
$\bar{\beta}_{\rm ppN}$	perihelion shift	3.0×10^{-3}
$\alpha_1^{\rm ppN}$	orbital polarization	10^{-4}
$\alpha_2^{\rm ppN}$	spin precession	4×10^{-7}

Table 2: The current experimental constraints on the four ppN parameters that we will consider in this thesis, along with the effect used to measure that constraint [144]. We use the definitions $\bar{\beta}_{\rm ppN} = \beta_{\rm ppN} - 1, \bar{\gamma}_{\rm ppN} = \gamma_{\rm ppN} - 1$

As stated, all theories of gravity consistent with the ppN formalism obey the Einstein equivalence principle. In addition to these restrictions, one can impose the conservation of total momentum, that is, both momentum and energy. Theories with this conservation law are called "semi-conservative", and have only five non-zero ppN parameters. These are $\gamma_{\rm ppN}, \beta_{\rm ppN}, \alpha_1^{\rm ppN}, \alpha_2^{\rm ppN}$, and $\xi_{\rm ppN}$. In this paper, in order to be able to compare the ppN results to the post-Keplerian and post-Einsteinian, we restrict ourselves to considering these semi-conservative theories.

Finally, in order to best understand the bounds that have been placed on this particular type of ppE term, we change variables in many expressions to a system in which all of the ppN parameters are equal to zero in GR. That is, $\{\beta_{\rm ppN}, \gamma_{\rm ppN}\} \to \{\bar{\beta}_{\rm ppN}, \bar{\gamma}_{\rm ppN}\}$, where $\bar{\beta}_{\rm ppN} = \beta_{\rm ppN} - 1, \bar{\gamma}_{\rm ppN} = \gamma_{\rm ppN} - 1$.

5.2. Parameterized post-Keplerian

We have not only been able to test GR with experiments in our Solar System, but also by analyzing the timing data from pulsars in binary systems [125]. Pulsars are some of the best clocks in the universe - the arrival time at Earth of their EM pulses can be predicted with extreme precision by fitting the measured arrival times to a timing formula. This formula must include gravitational effects on the emission and travel time of the pulses, as well as non-gravitational effects intrinsic to the pulsar itself. For example, aberration due to the fact that the beam of EM radiation comes from a concentrated point on the star, and not the entire star. The precision and complexity of these systems make them excellent laboratories for testing GR - the simultaneous measurement of several of these effects allows for consistency checks on the theory of gravity used to generate the timing formula.

Blandford and Teukolsky [25] derived a timing model that assumed that the components of the binary behaved according to Kepler's laws of planetary motion. The five "Keplerian" parameters which enter this timing model are the orbital period, P_b, the epoch of periastron passage, T_0, the eccentricity, e, the longitude of periastron, ω, and the projected semi major axis of the orbit, $x = a \sin\iota/e$, where a is the semi-major axis of the orbit, and ι is the inclination of the binary, measured from the line of sight to the binary. In addition to the Keplerian parameters, their model allowed for secular drifts of these parameters, as well as one extra parameter, $\gamma_{\rm ppK}$, to account for special-relativistic time dilation effects. Although Blandford and Teukolsky had intended their model only as a way to measure parameters in GR, the phenomenological approach that they took allowed it to fit timing predictions from other theories as well. The parameters and structure of the timing formulas was theory-independent

- it was the functional relationship between the parameters and the masses of the pulsar and its companion that were determined by a given theory of gravity.

Later, Epstein [58] and Haugan [73] attempted to include the 1PN corrections to the timing formula, which come from the Shapiro time delay and the gravitational redshift due to the mass of the companion, as well as post-Keplerian effects on the orbital motion. The resulting formula was very complicated, and moreover was not theory independent, as it had been calculated within GR.

Damour and Deruelle [44] showed that all 1PN corrections to the timing formula could be captured in a simple way that was applicable to many theories of gravity. This work led to their ppK formalism for a pulsar timing formula. This formula includes eight separately measurable post-Keplerian parameters, as well as four post-Keplerian parameters that are not separately measurable. The timing formula including these eight parameters can be written as [48]

$$t_b - T_0 = F\left[\tau; \{p^K\}; \{p^{pK}\}; \{q^{pK}\}\right]. \tag{5.7}$$

Here, t_b is the time at which a signal from the pulsar would be detected at the solar-system barycenter if there were no GR effects on its propagation, and τ is the proper time of the pulsar.

In addition to proper time, the right-hand-side (RHS) of Eq. (5.7) depends on

$$\{p^K\} = \{P_b, T_0, e_0, \omega_0, x_0\}, \tag{5.8}$$

the set of Keplerian parameters that describe an elliptical orbit, with ω_0 the initial position of periastron and x_0 the initial semi-major axis;

$$\{p^{\rm ppK}\} = \{k, \gamma, \dot{P}_b, r, s, \delta_\theta, \dot{e}, \dot{x}\}_{\rm ppK}, \tag{5.9}$$

the set of separately measurable post-Keplerian parameters; and finally

$$\{q^{\rm ppK}\} = \{\delta_r, A, B, D\}_{\rm ppK}, \tag{5.10}$$

the set of not separately measurable post-Keplerian parameters. For a detailed discussion of what exactly each of these parameters is, and how they fit into the timing model, see [48].

The right hand side of Eq. (5.7) can be written schematically as

$$F(\tau) = D^{-1}\left[\tau + \Delta_R(\tau) + \Delta_E(\tau) + \Delta_S(\tau) + \Delta_A(\tau)\right], \tag{5.11}$$

where each term is due to different effects. In this expression, Δ_A is called the "aberration" delay, and is due to the fact that the pulsar is not simply a radial pulsation, but a rotating beacon. Δ_R is the modulation in arrival time due to the motion of the Earth about the Sun, as well as the orbital motion of the pulsar and its companion, known as the Roemer time delay. Δ_E is the Einstein time delay, or gravitational redshift, caused by the pulsar's binary companion. Finally, Δ_S is the Shapiro time delay, also due to the pulsar's binary companion.

Each of these effects can be written in terms of a combination of Keplerian and ppK parameters.

$$\Delta_R = x \sin\omega \left[\cos u - e(1+\delta_r)\right] + x\left[1 - e^2(1+\delta_\theta)^2\right]^{1/2} \cos\omega \sin u \tag{5.12}$$

$$\Delta_E = \gamma_{\mathrm{ppK}} \sin u \tag{5.13}$$

$$\Delta_S = -2r_{\mathrm{ppK}} \ln\Big\{1 - e\cos u - s_{\mathrm{ppK}}\big[\sin\omega(\cos u - e) + (1-e^2)^{1/2}\cos\omega\sin u\big]\Big\} \tag{5.14}$$

$$\Delta_A = A\Big\{\sin\left[\omega + A_e(u)\right] + e\sin\omega\Big\} + B\Big\{\cos\left[\omega + A_e(u)\right] + e\cos\omega\Big\} \tag{5.15}$$

Here, A_e and ω are functions of u, described by

$$A_e(u) = 2\arctan\left(\left[\frac{1+e}{1-e}\right]^{1/2}\tan\frac{u}{2}\right), \qquad \omega = \omega_0 + k_{\mathrm{ppK}} A_e(u), \tag{5.16}$$

and finally u is eccentric anomaly, and is a function of proper time, τ. It is defined by solving Kepler's equation

$$u - e\sin u = \frac{2\pi}{P_b}(\tau - T_0), \tag{5.17}$$

Later, Damour and Taylor [48] extended the model to include fully nineteen separately measurable ppK parameters, but measurable in theory does not always mean measurable in practice. Of these nineteen parameters, five have been measured using available pulsar data. These are the perihelion precession, $\dot{\omega}_{\mathrm{ppK}}$ (related to k_{ppK}), the gravitational redshift due to the pulsar's companion, $\gamma_{\rm ppK}$, the range and shape of the Shapiro time delay, $r_{\rm ppK}$ and $s_{\rm ppK}$, and the rate of decay of the orbital period, $\dot{P}_b^{\rm ppK}$. We will restrict ourselves in this chapter to considering these five parameters.

Our ability to test GR using these ppK parameters depends on the precision with which we can measure them. The current uncertainties in the measured values of these parameters are listed in Table 3.

Parameter	Effect	Measured value
$\dot{P}_b^{\rm ppK}$	orbital decay	$-1.252(17) \times 10^{-12}$
$r_{\rm ppK}$	range of Shapiro delay	$6.21(33)(\mu s)$
$s_{\rm ppK}$	shape of Shapiro delay	$0.99974(+16, -39)$
$\dot{\omega}_{\rm ppK}$	periastron precession	$016.89947(68)(^\circ/\mathrm{yr})$
$\gamma_{\rm ppK}$	gravitational red-shift	0.3856(26) (ms)

Table 3: Uncertainty in measured values for PSR J0737-3039A. [86]

5.3. Parameterized post-Einsteinian

In the near future, the detection of GWs will open a new window for testing GR. There are essentially two approaches for doing so. One method, a *top-down* approach, demands that we have a particular alternative gravity theory that we wish to test. In this method, we could calculate what GWs would look like in this theory, and develop non-GR GW templates for data analysis. For a given signal, one could then calculate the Bayesian odds ratio[1] between this specific theory and GR, and in this way decide which theory is better supported by the data. The advantage to this

[1]The Bayesian odds ratio, or *Bayes factor*, between two models, A and B, is the betting odds that model A is supported by the data better than model B. A Bayes factor of 3 in favor of model A would indicate that the data shows a 3 to 1 preference for model A.

method is that we would have the full equations of motion for the model, and be able to answer theoretical questions such as well-posedness, in addition to being able to predict many observables. On the other hand, there is no particularly compelling alternative to GR presently known, and the effort involved in fleshing out all possible contenders is highly non-negligible.

In contrast, the second method for testing GR is a *bottom-up* [165] approach . In this approach, one uses experimental data to learn about a possible gravitational theory. If there is an indication of deviation from GR, these data then motivate the development of an alternative to GR. In order to use GWs for this type of analysis, we need a set of waveform templates that do not assume that GR is the correct theory of gravity.

With these templates as their aim, Yunes and Pretorius [163] developed the ppE family of waveform templates. They focused on creating templates for the Fourier transform of the quadrupole GW strain signal from a system of two inspiraling, non-spinning, compact objects in quasi-circular orbits. In the future, the restriction to circular, non spinning systems can be relaxed, building on work done in Ref. [156] for eccentric systems in GR and Ref. [33] for spinning systems in GR. The restriction to the quadruple mode has already been lifted by Chatziioannou et al in [34].

The full waveform from two coalescing bodies is typically split into three phases - inspiral, merger, and ringdown. ppE templates have been developed for all three phases, but here we restrict ourselves to the inspiral only. The inspiral is the part of the waveform that is generated while the two bodies are still widely separated, and thus slowly spiraling towards each other due to the emission of GWs. The definition of the end of inspiral is somewhat arbitrary, but we follow typical convention and define the transition from inspiral to merger as occurring at the innermost stable circular orbit of the system in center of mass (COM) coordinates. The simplest, quadrupole ppE inspiral templates have the form:

$$\tilde{h}(f) = \tilde{h}^{GR} \cdot (1 + \alpha_{\text{ppE}} u^a) e^{i\beta_{\text{ppE}} u^b}, \qquad u = (\pi \mathcal{M} f)^{1/3}, \tag{5.18}$$

where $\tilde{h}^{GR}$ is the GW waveform in GR, $\mathcal{M}$ is the chirpmass of the system, $\mathcal{M} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$, and f is the GW frequency. These simple ppE modified waveforms consist of an additional amplitude term, $\alpha_{\text{ppE}} u^a$, and an additional phase term, $\beta_{\text{ppE}} u^b$, relative to GR. We refer to α_{ppE} and β_{ppE} as the *strength* parameters of the ppE deviations.

The ppE templates are constructed by introducing parameterized modifications to both the binding energy and the energy balance equations of GR. Both types of modifications lead to changes in the GW phase, which leads to a degeneracy if we are only sensitive to the phase. This means that if a deviation from GR is detected in the phase, it is impossible to determine from only this measurement whether the deviation is from the conservative or dissipative sector. The combination of GW measurements with other experiments, as well as more sensitive measurements of GW amplitude, could possibly lift this degeneracy.

Theory	a	$\alpha_{\rm ppE}$	b	$\beta_{\rm ppE}$
Variable G(t)	-8	$\alpha'_{\rm ppE}$	-13	$\beta'_{\rm ppE}$
Brans-Dicke	-2	$\alpha'_{\rm ppE}$	-7	$\beta'_{\rm ppE}$
Dynamical Chern-Simons	a	0	-1	$\beta'_{\rm ppE}$

Table 4: The values that the ppE parameters take on in various non-GR theories [165].

The ppE waveforms cover all known inspiral waveforms from specific alternative theories of gravity [38] that are analytic in the frequency evolution of the GWs. Some specific examples are listed in Table 4.

5.4. ppN-ppE Correspondence

The first correspondence we calculate is that between the ppE and ppN parameters. In particular, we are interested in how the ppE phase parameters can be related to the ppN metric. The ppE amplitude parameters can also be related to ppN corrections to the metric. These corrections, like the phase ones, enter at first PN order. Because we are much more sensitive to the GW phase in analyzing GW data, we typically do not include any PN corrections to the amplitude, and so we do not show this calculation.

To calculate the GW phase, we use

$$\phi(t) = 2\pi \int_{t_0}^{t} dt' \int_{t_0'}^{t'} \dot{f}(t'')dt''. \tag{5.19}$$

We can calculate $\dot{f}(t)$ from

$$\frac{df}{dt} = \frac{df}{dE}\frac{dE}{dt}, \tag{5.20}$$

where E is the binding energy of the binary, and $\dot{E}$ is the GW luminosity. This is the same technique used by Yunes and Pretorius in developing the ppE framework [163]. We find the integrand in Eq. (5.20) by calculating the binding energy as a function of velocity, using the ppN modified Kepler's law to change this to a function of frequency, and then inverting this expression so we have the frequency as a function of energy. We then use the standard GR expression for the GW luminosity, coupled with our non-GR expression for df/dE, and we can calculate the phase of the gravitational waveform.

Each of these steps needs to be carried out to consistent PN order. In our case, we are only interested in the 1PN correction to GR, as this is the order to which the ppN framework is valid. In the COM frame of the two-body system, the binding energy in the ppN formalism, correct to 1PN order, is [141]

$$E_{COM} = -\frac{M\mu}{r_{12}} + \frac{1}{2}\mu v^2 + \frac{3}{8}\frac{\mu}{M}v^4(M-3\mu) + \frac{M\mu}{2r_{12}} \times \Big[(2\gamma_{\rm ppN}+1)v^2 + \frac{M}{r_{12}}(2\beta_{\rm ppN}-1) + v^2\frac{\mu}{M}\left(1+\alpha_1^{\rm ppN}-\alpha_2^{\rm ppN}\right)\Big]. \tag{5.21}$$

Here M is the total mass, μ is the reduced mass, r_{12} is the separation distance between the two bodies, and v is the magnitude of the relative velocity between the two bodies in the COM frame.

Next, and again from [141], the magnitude of the acceleration between the two bodies, in the COM frame, in the ppN formalism, correct to 1PN order, is

$$a = r_{12}\omega^2 - \frac{M}{r_{12}^2}\left\{1 - \frac{M}{r_{12}}\left[2\beta_{\rm ppN} + \gamma_{\rm ppN} - \eta\left(1 + \frac{\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN}}{2}\right)\right]\right\}. \quad (5.22)$$

Note that the ppN parameter, ξ, having to do with preferred location effects, does not appear in either the binding energy or the acceleration. Due to symmetry, the system being considered would need to consist of at least three bodies for these effects to be non-zero. This means that the set of ppN parameters relevant to our problem has been reduced from 10 to 4, which are $\beta_{\rm ppN}, \gamma_{\rm ppN}, \alpha_1^{\rm ppN}$, and $\alpha_2^{\rm ppN}$.

To proceed, we must re-write Eq. (5.21) as a function of the GW frequency, f. To do so, we use the fact that, in a circular orbit, correct to 2.5 PN order, $v = r_{12}\omega$ and $a = r_{12}\omega^2$. We can thus us Eq. (5.22) to find a relation between v and r_{12}, which we can then use to re-write Eq. (5.21). We then have an expression for the binding energy in terms of GW frequency:

$$E = -\frac{1}{2}\mu(2\pi M f)^{2/3}\Bigg(1 + (2\pi M f)^{2/3} \times \left\{-\frac{3}{4} + \frac{2}{3}(\beta_{\rm ppN} - \gamma_{\rm ppN}) - \eta\left[\frac{1}{12} + \frac{5}{3}(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\right]\right\}\Bigg). \quad (5.23)$$

We then differentiate Eq. (5.23) with respect to f, and invert the result to find df/dE

$$\frac{df}{dE} = -\frac{\mu}{3}(2\pi M)\Bigg((2\pi M f)^{-1/3} + 2(2\pi M f)^{1/3} \times \left\{-\frac{3}{4} + \frac{2}{3}(\beta_{\rm ppN} - \gamma_{\rm ppN}) - \eta\left[\frac{1}{12} + \frac{5}{3}(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\right]\right\}\Bigg). \quad (5.24)$$

where we have expanded in $Mf \ll 1$.

Lastly, we need the energy carried away from the system as GW luminosity. At this point, we make the assumption that the luminosity in GWs for alternative theories can be calculated from the mass and current multipoles of the system in the same way as in GR. Put another way, we are relating the changes in the conservative, as opposed to dissipative, sector of gravity to effects seen in the waveforms. This is one instance in which it is clear that the ppE-ppN mapping is not perfect, because the ppE formalism includes changes to the dissipative sector as well as to the binding energy, whereas the ppN formalism is concerned only with the conservative sector.

In GR, and to 1PN order, the GW luminosity is given by [23]

$$L_{GW} = \frac{1}{5} M_{ij}^{(3)} M_{ij}^{(3)} + \frac{1}{189} M_{ijk}^{(4)} M_{ijk}^{(4)} + \frac{16}{45} S_{ij}^{(3)} S_{ij}^{(3)}, \tag{5.25}$$

where the mass quadrupole moment is M_{ij}, the mass octupole is M_{ijk}, and the current quadrupole is S_{ij}; the superscript (n) implies n time derivatives; and the source multipole moments of the binary are calculated in the standard way [23].

Theory	Un-known functions or constants	$\bar{\gamma}_{\rm ppN}$	$\bar{\beta}_{\rm ppN}$	$\alpha_1^{\rm ppN}$	$\alpha_2^{\rm ppN}$
Brans-Dicke	ω_{BD}	$\frac{1}{\omega_{BD}}$	0	0	0
General scalar-tensor	$A(\varphi), V(\varphi)$	$\frac{1}{\omega}$	$\frac{1}{4\omega^3} \frac{d\omega}{d\phi}\Big\rvert_{\phi_0}$	0	0
Einstein-Aether	c_1, c_2, c_3, c_4	0	0	α_1'	α_2'

Table 5: Values of ppN parameters for a selection of alternative theories. These expressions are in the large ω limit. [144, 28, 79, 60]

Using this assumption, we arrive at an expression for the GW luminosity,

$$\frac{dE}{dt} = \frac{32}{5}\eta^2 M^2 r^4 (2\pi f)^6 \left[1 + \frac{M}{r_{12}}\left(-\frac{10}{7} + \frac{16}{7}\eta\right)\right]. \tag{5.26}$$

We then use the mapping between v and r_{12} to re-write this equation as:

$$\begin{aligned}\frac{dE}{dt} &= \frac{32}{5}\eta^2 (\pi M f)^{10/3}\Bigg(1 + (\pi M f)^{2/3} \\ &\times \left\{\frac{97}{336} - \frac{8}{3}\beta_{\rm ppN} - \frac{4}{3}\gamma_{\rm ppN} + \eta\Big[-\frac{35}{12} + \frac{2}{3}(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\Big]\right\}\Bigg). \end{aligned} \tag{5.27}$$

Multiplying Eq. (5.27) and Eq. (5.24), and again expanding in $Mf \ll 1$, we arrive at

$$\begin{aligned}\frac{df}{dE}\frac{dE}{dt} &= \frac{96}{5}\pi^{8/3} f^{11/3}\mathcal{M}^{5/3}\Bigg(1 + (\pi M f)^{2/3} \\ &\times \left\{\frac{601}{336} - 4\beta_{\rm ppN} - \eta\left[\frac{11}{4} - 4(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\right]\right\}\Bigg). \end{aligned} \tag{5.28}$$

We next integrate Eq. (5.28) once, and then invert perturbatively to find

$$\begin{aligned}t(f) &= t_c - 5(8\pi f)^{-8/3}\mathcal{M}^{-5/3}\Bigg\{1 + (\pi M f)^{2/3} \\ &\times \left(-\frac{601}{252} + \frac{16}{3}\beta_{\rm ppN} + \eta\left[\frac{11}{3} - \frac{16}{3}(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\right]\right)\Bigg\}. \end{aligned} \tag{5.29}$$

Here t_c is the time of coalescence. We integrate Eq. (5.29) again to find $\phi(t)$, and then replace t with Eq. (5.29) to arrive at

$$\begin{aligned}\phi(f) &= \phi_c - 2(8\pi \mathcal{M} f)^{-5/3}\Bigg(1 + (\pi M f)^{2/3} \\ &\times \left\{-\frac{3005}{1008} + \frac{20}{3}\beta_{\rm ppN} + \eta\Big[\frac{110}{24} - \frac{20}{3}(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN})\Big]\right\}\Bigg), \end{aligned} \tag{5.30}$$

where ϕ_c is the phase at coalescence.

Finally, we can use the stationary phase approximation (SPA) to find the phase of the GW in the Fourier domain. To calculate this, we use $\Psi_{\rm SPA} = 2\pi f t(f) - \phi(f) - \pi/4$, which gives

$$\begin{aligned}\Psi_{\rm SPA}(f) = 2\pi f t_c - \phi_c - \frac{\pi}{4} \\ + \frac{3}{128} u^{-5} \left\{ 1 + u^2 \left[\frac{20}{9\eta^{2/5}} \left(\frac{743}{336} + \frac{11}{4}\eta \right) \right] \right\} \\ + u^{-3} \frac{5}{24\eta^{2/5}} \left[\bar{\beta}_{\rm ppN} - \eta(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN}) \right]. \quad (5.31)\end{aligned}$$

where recall that $u = (\pi \mathcal{M} f)^{1/3}$.

The last term in Eq. (5.31) represents a ppE correction with $b = -3$, and strength parameter

$$\beta_{\rm ppE}^{b=-3} = \frac{5}{24\eta^{2/5}} \left\{ \bar{\beta}_{\rm ppN} - \eta(\alpha_1^{\rm ppN} - \alpha_2^{\rm ppN}) \right\}. \quad (5.32)$$

Note that in the test-particle limit, when $\eta \to 0$, the $\beta_{\rm ppE}$ and $\bar{\beta}_{\rm ppN}$ parameters are simple rescalings of each other. It is also interesting to note that the parameter $\gamma_{\rm ppN}$ does not appear in the final expression for the phase. This tells us that within alternative theories of gravity that do not alter the value of $\beta_{\rm ppN}, \alpha_1^{\rm ppN}$, or $\alpha_2^{\rm ppN}$, there will be no phase corrections at first PN order.

Table 5 lists the values of the ppN parameters for a selection of alternative theories of gravity. None of these theories have $\bar{\beta}_{\rm ppN}$ that is strongly different from the GR value of $\bar{\beta}_{\rm ppN} = 0$. This tells us that if a significant departure from GR is detected at 1 PN order, it must come from the dissipative sector of these alternative theories, or from a different theory altogether.

5.5. ppN-ppK Correspondence

Now we derive the connections between the ppN metric parameters and the ppK timing parameters. The arrival time of a pulse of EM radiation at the Earth is expressed schematically in Eq. (5.11). The Roemer time delay for a binary pulsar, as stated above, includes not only the (non-GR) modulation due to the motion of the Earth around the Sun, but also modulations due to the orbital motion of the binary itself. This orbital motion includes GR corrections at 1PN order, which lead to perihelion precession, $\dot{\omega}$, one of the ppK parameters. This parameter can be related to the ppN parameters by calculating the equations of motion for the pulsar using the ppN metric. This calculation is done by Will [141], and, in the semi-conservative theories of gravity that we are considering, is equal to

$$< \dot{\omega} >= \frac{(2M)^{2/3}\pi^{5/3}}{(e^2-1)P_b^{5/3}}\Big[2(1+\bar{\beta}_{\rm ppN}-2\bar{\gamma}_{\rm ppN})-\eta(2\alpha_1^{\rm ppN}-\alpha_2^{\rm ppN})\Big]. \tag{5.33}$$

where the angled brackets indicate integration over one orbit.

In order to be consistent with our other analyses, this result, and all results in this section, neglects any contributions from the self-gravity or structure of the pulsars. These types of effects have been explored in [139].

The Shapiro time delay is calculated from the formula for null geodesics. In GR, to the appropriate order, this is simply

$$dt = \big[1-2\phi(x)\big]dx, \tag{5.34}$$

where $\phi(x)$ is the Newtonian gravitational potential due to the pulsar's companion, m_c/r_{12}, with m_c the mass of the companion, and r_{12} the separation between them[94]. In the ppN formalism, this equation becomes

$$dt = \Big[1 - (1 + \gamma_{\rm ppN})\phi(x)\Big]dx. \tag{5.35}$$

Integrating this over the path of the photon from the pulsar, past its companion, and to the solar system barycenter, we get

$$\begin{aligned}\Delta_S = -2r_{\rm ppK} \log\Big\{[1 - e\cos(u)] - s\Big[\sin(\omega)(\cos(u) - e) \\ + \sqrt{1-e^2}\cos(\omega)\sin(u)\Big]\Big\},\end{aligned} \tag{5.36}$$

which is the same as in GR, except now

$$r_{\rm ppK} = m_c\left(1 + \frac{\bar{\gamma}_{\rm ppN}}{2}\right), \tag{5.37}$$

is the range of the Shapiro delay. The shape, s, is still equal to $\sin(\iota)$, as in GR, but it is generally written in terms of the two masses in the binary, m_p and m_c, using Kepler's law. Kepler's law is altered at 1PN order from the GR expression, but this alteration should not be considered, as $r_{\rm ppK}$ is already a 1PN correction. Thus, the formula for s in terms of the masses of the system is unaltered from GR:

$$s = \frac{x}{(m_c + m_p)^{1/3}}\left(\frac{P_b}{2\pi}\right)^{-2/3}, \tag{5.38}$$

where recall that x is the projected semi-major axis.

Finally, we calculate the Einstein time delay, which is really just the gravitational redshift. This redshift arises from the time dilation experienced by a photon as it travels out of a gravitational potential well. It is expressed as the rate of change of proper time with respect to coordinate time

ppK Parameter	ppN expression
$\gamma_{\rm ppK}$	$\frac{P_b}{2\pi}\frac{e}{2a(m_c+m_p)}\left\{3m_c^2+m_cm_p+(\bar{\gamma}_{\rm ppN}+1)(m_c^2+m_cm_p)\right\}$
$r_{\rm ppK}$	$\frac{1}{2}(\bar{\gamma}_{\rm ppN}+2)m_c$
$< \dot{\omega} >$	$\frac{(2M)^{2/3}\pi^{5/3}}{(e^2-1)P_b^{5/3}}\left\{2(1+\bar{\beta}_{\rm ppN}-2\bar{\gamma}_{\rm ppN})-\eta(2\alpha_1^{\rm ppN}-\alpha_2^{\rm ppN})\right\}$

Table 6: ppK parameters expressed as combinations of the ppN parameters.

$$d\tau^2 = [1+(1+\gamma_{\rm ppN})\phi(x)]\,dt^2 - [1-(1+\gamma_{\rm ppN})\phi(x)]\,dx^2, \tag{5.39}$$

which leads, in the weak-field approximation, to

$$\frac{d\tau}{dt} = 1 + \frac{1+\gamma_{\rm ppN}}{2}\phi(x) - \frac{1}{2}v_p^2, \tag{5.40}$$

where $v_p = dx/dt$ is the velocity of the pulsar. We can next use the virial theorem to replace v_p^2.

$$\frac{1}{2}v_{12}^2 = \frac{m_p+m_c}{r_{12}} - \frac{m_p+m_c}{2a}, \quad v_p = \frac{m_c}{m_p+m_c}v_{12}, \tag{5.41}$$

with the separation for eccentric orbits given by $r_{12} = a(1-e\cos u)$, and v_{12} the relative velocity between the two bodies. Just as with Kepler's law, there are 1PN corrections, but, in order to keep all terms to the proper order, we use the Newtonian approximation. We use the relationships in Eq. (5.41) to replace v_p in Eq. (5.40) with an expression in terms of v_{12}.

Next, we want to change variables from the eccentric anomaly, u, to P_b. Differentiating Kepler's equation, Eq. (5.17), and neglecting any terms that go as $\dot{P}_b$, we arrive at

$$\frac{du}{dt} = \frac{2\pi}{P_b}\frac{1}{1 - e\cos u}. \tag{5.42}$$

We change coordinates using Eq. (5.42) and the fact that $d\tau/dt = (d\tau/du)(du/dt)$, and, finally, Eq. (5.40) becomes

$$\begin{aligned}\frac{2\pi}{P_b}\frac{d\tau}{dt} = &\left\{1 - \frac{m_c[2m_c + m_p + \gamma_{\rm ppN}(m_c + m_p)]}{2a(m_c + m_p)}\right\} \\ &\times \left\{1 - e\cos u\left[1 + \frac{3m_c^2 + m_c m_p + \gamma_{\rm ppN}(m_c^2 + m_c m_p)}{2a(m_c + m_p)}\right]\right\}.\end{aligned} \tag{5.43}$$

The part of this expression that is unmodulated with u is not detectable. We can absorb it into a rescaling of the proper time

$$\tau \to \tau \times \left\{1 - \frac{m_c[2m_c + m_p + \gamma_{\rm ppN}(m_c + m_p)]}{2a(m_c + m_p)}\right\}. \tag{5.44}$$

This leaves us with the formula for the gravitational redshift in the standard form

$$\frac{d\tau}{dt} = \frac{P_b}{2\pi}(1 - e\cos u) - \gamma_{\rm ppK}\cos u, \tag{5.45}$$

where we find that $\gamma_{\rm ppK}$ is related to ppN parameters via

$$\gamma_{\rm ppK} = \frac{P_b}{2\pi}e\frac{3m_c^2 + m_c m_p + \gamma_{\rm ppN}(m_c^2 + m_c m_p)}{2a(m_c + m_p)}. \tag{5.46}$$

In summary, we have a correspondence between the ppN parameters and four of the ppK parameters, in semi-conservative theories of gravity in a reference frame at rest with respect to any universal reference frame, and neglecting effects due to self-gravity and structure of the pulsars. This correspondence is summarized in Table 6. Because of the combinations of ppN parameters that appear in these expressions, it is not possible to use the results from Sec. 5.4 and re-write them entirely in terms of ppE parameters.

5.6. ppE-ppK Correspondence

The final piece missing from our correspondence puzzle is the relationship between the decay of the orbital period of a binary system, the ppK parameter $\dot{P}_b$, and either ppN or ppE parameters. The correspondence with ppE parameters was tackled by Yunes and Hughes [159]. We here carry out a similar calculation.

Assuming that the binding energy of a system is the same as in GR, the decay rate of a binary system can be calculated via

$$\frac{\dot{P}_b}{P_b} = -\frac{3}{2}\frac{\dot{E}}{E_b}, \tag{5.47}$$

where $\dot{E}$ is the energy carried away by GWs, and E_b is the binding energy of the system.

Yunes and Hughes assumed in their calculation that only the dissipative sector is affected by the ppE parameters. This is because the functional form of the ppE phase corrections does not depend on whether the binding energy, the GW luminosity, or both are modified from GR. This degeneracy makes it impossible to determine whether the ppE changes in the GW phase arise from the dissipative or conservative sector. Thus, we first take the presence of ppE parameters in the phase of the GW to come only from the expression for $\dot{E}$ [159]:

$$\dot{E} = \dot{E}_{\rm GR}\left(1 + \pi^2\mathcal{M}^2 u^{-1}\frac{d^2\Psi_{\rm GR}}{df^2}\right), \tag{5.48}$$

where Ψ_{GR} is the phase of the GW in GR.

The ppE corrections have only been calculated assuming circular binaries. However, all known pulsars are in eccentric orbits. The corrections to $\dot{E}$ from eccentricity are known in GR and have the form

$$\dot{E}_{\text{GR}} = -\frac{32}{5}\eta^2\frac{M^5}{a^5}(1-e^2)^{-7/2}\left(1+\frac{73}{24}e^2+\frac{37}{96}e^4\right). \tag{5.49}$$

The same type of corrections, however are not known in the ppE framework, and will necessarily involve the introduction of new ppE parameters. Fortunately, the eccentricity of some pulsar systems is small enough, that we can accurately model the ppE parameters as expansions in small e, where we keep only the lowest order term in our calculations.

With this assumption, we can use Eq. (5.47) to show that the expression for the orbital decay rate becomes [159]

$$\left(\frac{\dot{P}_b}{P_b}\right)_{Phase} = \left(\frac{\dot{P}_b}{P_b}\right)_{\text{GR}}\left[1+\frac{48}{40}\beta_{\text{ppE}}b(b-1)u^{b+5}\right], \tag{5.50}$$

when corrected with phase ppE parameters. The term $(\dot{P}/P)_{\text{GR}}$ stands for the orbital decay rate in GR for an eccentric inspiral. Because the observed value of $\dot{P}/P$ is very close to the GR prediction, we can write $(\dot{P}_b/P_b)_{obs} = (\dot{P}_b/P_b)_{\text{GR}}(1+\delta)$, where δ, the observational error, is small.

Because the periods of binary pulsars have been observed to decay at the GR rate, the ppE strength parameter must satisfy [159]

$$|\beta_{\text{ppE}}| \leqslant \frac{40}{48|b||b-1|}\frac{\delta}{u^{b+5}}\ . \tag{5.51}$$

As stated, the preceding calculation was done with the assumption that only the GW luminosity of the pulsar system was different from GR. The binding energy was assumed to be the same as in GR. This is the opposite of what we assumed in Sec. 5.4, in which only the binding energy was altered. We now calculate the relation between $\beta_{\rm ppE}$ and $\dot{P}_b$, but this time assuming that it is the conservative sector that is altered from GR.

Following [34], we parameterize the binding energy of a binary system as

$$E = E_{\rm GR}\left[1 + A\left(\frac{M}{r_{12}}\right)^q\right], \tag{5.52}$$

where A is small, and therefore E differs from $E_{\rm GR}$ by only a small perturbation. This binding energy leads to a modified Kepler's law:

$$\omega^2 = \frac{M}{r_{12}^3}\left[1 + \frac{1}{2}Aq\left(\frac{m}{r}\right)^q\right], \tag{5.53}$$

which we can use to re-write the energy in terms of the orbital period of the pulsar system:

$$E = -\frac{1}{2}\eta^{-2/5}\left(\frac{2\pi\mathcal{M}}{P_b}\right)^{2/3} \times \left[1 - \frac{1}{3}A(5q-6)\eta^{-2q/5}\left(\frac{2\pi\mathcal{M}}{P_b}\right)^{2q/3}\right]. \tag{5.54}$$

We can find an expression for $\dot{P}_b/P_b$ by differentiating Eq. (5.54) with respect to time. This will give us an expression in terms of the GW luminosity, $\dot{E}$. Although we are not explicitly changing the GR expression for this luminosity, $\dot{E}$ is modified when we use Eq. (5.53) to relate r_{12} to P_b.

$$\dot{E} = \frac{32}{5}\left(\frac{2\pi\mathcal{M}}{P_b}\right)^{10/3} \times \left[1 - \frac{1}{3}Aq\eta^{-2q/5}\left(\frac{2\pi\mathcal{M}}{P_b}\right)^{2q/3}\right]. \tag{5.55}$$

Differentiate Eq. (5.54) with respect to time, and replace $\dot{E}$ with Eq. (5.55) and arrive at

$$\frac{\dot{P}_b}{P_b} = \left(\frac{\dot{P}_b}{P_b}\right)_{\rm GR}\left[1 - \frac{1}{3}A\eta^{-2q/5}(5q^2 - 2q - 6)u^{2q}\right], \tag{5.56}$$

where we have expanded in $A \ll 1$.

The final step is to relate q and A to the more standard ppE parameters, b and $\beta_{\rm ppE}$. The relations are [34]

$$2q = b + 5, \tag{5.57}$$

$$-A(5q^2 - 2q - 6)\eta^{-2q/5} = \frac{32}{5}\beta_{\rm ppE}(4 - q)(5 - 2q). \tag{5.58}$$

With these replacements, we finally arrive at

$$\frac{\dot{P}_b}{P_b} = \left(\frac{\dot{P}_b}{P_b}\right)_{\rm GR}\left[1 + \frac{16}{15}\beta_{\rm ppE}b(b - 3)u^{b+5}\right], \tag{5.59}$$

which leads to the constraint on $\beta_{\rm ppE}$:

$$|\beta_{\rm ppE}| \leqslant \frac{15}{16}\frac{\delta}{|b||b - 3|u^{b+5}}. \tag{5.60}$$

Both Eq. (5.51) and Eq. (5.60) are constraints on the ppE strength parameter associated with the phase of a GW. In the first case, this constraint comes from the assumption that the GW luminosity of the system is not as described by GR, while in the second case it is the binding energy that is changed from the GR expression. The fact that both approaches lead to constraints on the GW phase illustrates a degeneracy in the ppE formalism. Both changes to the binding energy and changes to the luminosity of a system lead to the same type of non-GR terms in the GW phase.

It is impossible to tell if a ppE term in a GW signal arises from the conservative or dissipative sector of a gravitational theory, or some combination of the two, if one is sensitive only to the phase.

5.7. Current Constraints

With expressions for the correspondences between parameters in the different systems in hand, we can calculate current constraints on the ppE parameters from known constraints on ppN and ppK parameters. Sec. 5.4, Eq. (5.32) gives the bound on $\beta_{\rm ppE}$ from the constraints on ppN parameters. Using the best constraints on the ppN parameters from Table 2, in Fig. 7 we plot this limit as a function of mass ratio. We see that $\beta_{\rm ppE} < 0.008$ from Solar System tests.

Similar bounds on $\beta_{\rm ppE}$ from GW detections (with signal-to-noise ratio of 20) were calculated in [38]. For both a 2 : 1 and a 3 : 1 mass ratio, the authors found a limit of $\beta_{\rm ppE} \lesssim 0.003$, which is comparable to current bounds from the Solar System.However, with multiple detections, the GW bounds should eventually surpass the latter.

Source of constraint	Constraint
Pulsar - $\dot{E}$ non-GR	$\beta_{\rm ppE} \leqslant 215$
Pulsar - E_b non-GR	$\beta_{\rm ppE} \leqslant 182$
Solar-System tests	$\beta_{\rm ppE} \lesssim 0.001$
Anticipated GW detections	$\beta_{\rm ppE} \lesssim 0.008$

Table 7: The current constraints that can be placed on $\beta_{\rm ppE}$ from Solar System and binary pulsar tests. These values are only for $b = -3$.

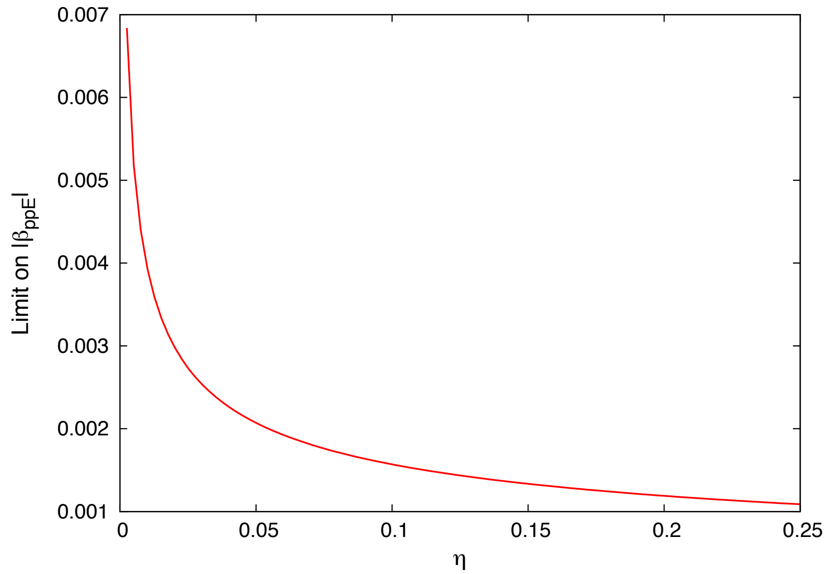

Figure 7: Limits that can be placed on the ppE strength parameter, $\beta_{\rm ppE}$, using the known limits on ppN parameters, and the results from Sec. 5.4. Regions above the curve are ruled out. The limits are a function of mass ratio, and the ppE b parameter is set to $b = -3$ in this plot.

The observed values for the ppK parameters from PSR J0737-3039 can similarly be used to place constraints on the ppE parameters, this time using results from Sec. 5.6. For this pulsar system, the uncertainty in the measurement of $\dot{P}_b/P_b$ is $\delta = 0.017 \times 10^{-12}/(1.242 \times 10^{-12}) \simeq 10^{-2}$. The chirp mass is $\mathcal{M} \simeq 5.5399 \times 10^{-6}$ s, and the GW frequency is $f \simeq 2.263842976 \times 10^{-4}$ Hz. Using these values in Eq. (5.51) and Eq. (5.60), we find that a GW measurement can set better constraints than the pulsar measurements starting around $\sim b = -4$. This means that there are regimes in which measurements of ppK parameters can help to constrain ppE parameters, and also regimes in which the opposite is true. This is consistent with the conclusions of [38].

The current constraints for β_{ppE}, with $b = -3$, from Solar System and pulsar experiments are listed in Table 7, as well as the anticipated constraint from future GW detections. The bounds from pulsar data do not depend strongly on whether it is the conservative or dissipative sector that is changed from GR. Both of these constraints are weaker than the Solar System and GW bounds, which differ from each other by a factor of ~ 3.

5.8. Conclusion

The many parameterizations of gravitational theory that have been developed over the years are designed for, and thus ideally suited to, testing the nature of gravity in quite different situations. We know, though, that the underlying theory of gravity that describes the universe we live in is the same in our solar system as it is in binary pulsar systems and colliding black holes. We should therefore be able to learn about the ppE parameters from our knowledge of ppN and ppK parameters, and vice versa.

In this chapter, we have found correspondences between the ppE, ppN, and ppK parameters that allows us to apply constraints from one formalism to the parameters in the others. In addition to finding the connections between the parameters in the different formalisms, in this work we have found that alternative theories of gravity that do not alter the β_{ppN}, α_1^{ppN}, or α_2^{ppN} parameters do not result in 1PN corrections to the GW phase. We also found that the bounds we will be able to place on deviations from GR at the 1PN level using GWs will be comparable to those already known from solar system tests.

The correspondences that we have calculated are not perfect. The ppN-ppE correspondence assumes semi-conservative theories of gravity, in a reference frame at rest with respect to any universal preferred frames. The ppN-ppK correspondence makes the same assumptions. And the ppE-ppK correspondence is only perfectly accurate for circular binaries. Finally, both the ppN-ppE and ppE-ppK correspondences assume that the generation of GWs, used to calculate the luminosity, is the same as in GR.

Future work could focus on relaxing some of the assumptions we used in this analysis. For instance, by allowing for changes in Eq. (5.25) by including source or current multipoles that are not present in GR. It may also be possible to introduce eccentricity into the ppE formalism, which could improve the accuracy of the ppE-ppK correspondence.

6. BAYESIAN PROBABILITY THEORY AND MARKOV METHODS

In this chapter, I give an introduction to the methods of Bayesian model selection, in particular to the use of Markov chain Monte Carlo algorithms

6.1. Bayesian Inference

Questions of model selection and parameter biases can be addressed very naturally in the framework of Bayesian inference. Within this framework, we are interested in comparing the hypothesis $\mathcal{H}_0$ that gravity is described by GR with the hypothesis $\mathcal{H}_1$ that gravity is described by an alternative theory belonging to the ppE class. Here we are dealing with nested hypotheses, as the ppE models include GR as a limiting case. When new data d comes available, our prior belief $p(\mathcal{H})$ in hypothesis $\mathcal{H}$ is updated to give the posterior belief $p(\mathcal{H}|d)$. Bayes' theorem tells us that

$$p(\mathcal{H}|d) = \frac{p(d|\mathcal{H})p(\mathcal{H})}{p(d)}\,, \tag{6.1}$$

where $p(d|\mathcal{H})$ is the (marginal) likelihood of observing the data d if the hypothesis holds, and $p(d)$ is a normalization constant. For hypotheses described by models with continuous parameters, the likelihood $p(d|\mathcal{H})$ is found by marginalizing the likelihood $p(d|\vec{\theta},\mathcal{H})$ of observing data d for model parameters $\vec{\theta}$:

$$p(d|\mathcal{H}) = \int d\vec{\theta}\; p(\vec{\theta},\mathcal{H})p(d|\vec{\theta},\mathcal{H})\,, \tag{6.2}$$

where $p(\vec{\theta},\mathcal{H})$ is the prior distribution of the parameters. The marginal likelihood, $p(d|\mathcal{H})$, is also known as the evidence for a given model. Hypotheses are compared by computing the odds ratio, or Bayes factor:

$$BF = \mathcal{O}_{1,0} \equiv \frac{p(\mathcal{H}_1|d)}{p(\mathcal{H}_0|d)} = \frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)} \frac{p(d|\mathcal{H}_1)}{p(d|\mathcal{H}_0)} \,, \tag{6.3}$$

which gives the 'betting odds' of $\mathcal{H}_1$ being a better description of Nature than $\mathcal{H}_0$. The normalization constant $p(d)$ cancels in the odds-ratio. The prior odds ratio $p(\mathcal{H}_1)/p(\mathcal{H}_0)$ gets updated by the likelihood ratio, $p(d|\mathcal{H}_1)/p(d|\mathcal{H}_0)$, which is also known as the evidence ratio. In Bayesian analysis 'today's posterior is tomorrow's prior' [132], and $p(\mathcal{H}|d)$ is used in place of $p(\mathcal{H})$ in subsequent analyses. While a single black hole inspiral event may not yield strong evidence for a departure from GR, several such observations can be combined to make a more compelling case.

In addition to simply detecting deviations from GR, we are also interested in studying how departures from GR might affect parameter estimation. This can be assessed by looking at the posterior distribution function $p(\vec{\theta}|d,\mathcal{H})$, which describes the probability distribution for parameters $\vec{\theta}$ under the assumption that the signals are described by model $\mathcal{H}$ given data d. The posterior distribution is given by the product of the prior and the likelihood, normalized by the evidence:

$$p(\vec{\theta}|d,\mathcal{H}) = \frac{p(\vec{\theta},\mathcal{H})p(d|\vec{\theta},\mathcal{H})}{p(d|\mathcal{H})} \,. \tag{6.4}$$

Once the prior distribution and the likelihood function have been specified we are left with the purely mechanical task of computing the posterior distributions and odds ratio for competing hypotheses. For a thorough introduction to Bayesian model selection in the context of gravitational waves, see Ref. [93].

6.2. Computational Techniques

In order to calculate Bayes factors and posterior distributions, we use the technique of Markov Chain Monte Carlo (MCMC). MCMC algorithms are ways of searching a parameter space in such a way that the probability distribution of a given function is fully explored. In Bayesian analysis, this function is the likelihood. The 'Monte Carlo' part of MCMC implies a random element to the technique, and this implication is born out. In an MCMC analysis, a 'chain' is just a collection of points in parameter space where a function has been evaluated. The chain is started at a given point in the space, and tries other points at random to determine if these points are at larger or smaller values of the function being evaluated. If the randomly chosen point is 'better,' (either larger or smaller, depending on the problem being solved), then the chain will 'jump' there, and then propose a new point chosen from that spot. If the new point is 'worse,' the chain might still go there, or it might not. The proposal of points and the possibility to go to a point that has a worse fit are both random aspects of MCMC. The beauty of the technique is that it allows the posterior space of a given likelihood function to be explored in a way such that areas of high posterior weight are explored more thoroughly than areas of low weight.

All of this can, of course, be quantified. In order for a Markov chain to accurately reproduce the weight of different parts of the posterior, that is, in order for the distribution of points in parameter space contained in a Markov chain to reflect the actual posterior distribution of a function, the proposal of new points for a chain to try needs to satisfy the detailed balance condition:

$$\pi(x_i)p(x_{i+1}|x_i) = \pi(x_{i+1})p(x_i|x_{i+1}) \tag{6.5}$$

Here π represents the target distribution function, and $p(x_i|x_{i+1})$ is the probability of moving from x_i to x_{i+1}, if this move is proposed. Because we don't know *a priori* what the probability of jumping between two given points is, we write the probability as:

$$p(x_{i+1}|x_i) = q(x_{i+1}|x_i)\kappa(x_{i+1}|x_i) \tag{6.6}$$

$q(x_{i+1}|x_i)$ is called the 'proposal distribution.' It could be, for instance, an n-dimensional Gaussian centered at the current point in parameter space. More simply, it could be a random draw from the full prior range of a given parameter. In order for proposals from this distribution to accurately recover the desired posterior distribution, we need the other factor, $\kappa(x_{i+1}|x_i)$. This term gives the probability of accepting a given move. In order to satisfy detailed balance, κ must take the form:

$$\kappa(x_{i+1}|x_i) = \frac{\pi(x_{i+1})q(x_i|x_{i+1})}{\pi(x_i)q(x_{i+1}|x_i)} \tag{6.7}$$

This quantity is referred to as the Hastings Ratio, H. In order that κ is a probability, we actually define it as $min[1, H]$. This ensures that if a point in parameter space is proposed that is at a higher value of the function we are investigating than the point we are currently at, the chain will always move to that point. If a proposed point has a lower value, there is still a finite probability of moving to that point. This feature keeps Markov chains from becoming 'stuck' at a local maximum in a likelihood function, and allows them to fully map out the parameter space.

The ppE waveforms introduce a number of complications that make parameter estimation and model selection challenging. These complications can be seen when using the quadratic Fisher matrix approximation $\Gamma_{ij} = -\partial_i\partial_j\langle \ln p(\vec{\theta}|d)\rangle$ to estimate the parameter correlation matrix $C^{ij} = \langle \Delta\theta^i \Delta\theta^j \rangle \approx \Gamma^{-1}_{ij}$. When evaluated at the GR limit point $(\alpha, \beta) = (0, 0)$, the quadratic approximation to the Fisher matrix is singular, and it is necessary to include higher order derivatives to obtain a finite covariance matrix. The situation is worse when $a = 0$, as then α is fully degenerate with D_L, and when $b = 0$, as then β is fully degenerate with Φ_c. Partial degeneracies also exist whenever the a or b exponents match the exponents found in the post-Newtonian expansion of GR.

The various degeneracies and parameter correlations do not constitute a fundamental problem with the ppE formalism, but they do demand that we use very effective MCMC samplers that are able to fully explore the parameter space. The algorithm described in Ref. [93] uses parallel tempering with multiple, coupled chains, with each chain exploring a tempered likelihood surface $p(d|\vec{\theta})^{1/T}$. The high temperature chains explore more widely, and can communicate this information via parameter exchange to the $T = 1$ chain that is used for parameter estimation. Parallel tempering helps the Markov chains explore complicated posterior distributions, but convergence can still be slow if the proposal distributions are not well chosen.

We use the proposals that are outlined in this reference, which include

- **Jiggle proposals** : $\vec{y}$ is proposed by changing all of the parameters of $\vec{x}$ by a very small amount.
- **Uniform proposals**: $\vec{y}$ is proposed by drawing all of the chain parameters from their full prior range.
- **Fisher proposals**: $\vec{y}$ is proposed by calculating the local Fisher matrix, defined below, and jumping along an eigenvector of this matrix.

The jiggle and uniform proposals are quite self-explanatory. The Fisher proposal, as stated, depends on the calculation of the local Fisher matrix, Γ, defined by

$$\Gamma_{ij} = (h_{,i} | h_{,j}), \tag{6.8}$$

where i and j run through all of the system parameters, and commas denote partial differentiation, and here $(a|b)$ indicates the noise weighted inner product between a and b. The Fisher matrix is the inverse of the correlation matrix, and its eigenvectors point along directions in which the parameters are highly correlated. Thus jumping along these directions is usually a well-accepted proposal.

In addition to these three techniques, we have implemented two types of proposals that are not used in Ref. [93], but that are very helpful in exploring highly correlated parameter spaces. These are the techniques of differential evolution and Langevin diffusion.

6.2.1. Differential Evolution

The ultimate proposal distribution is the posterior distribution itself, but since that is unavailable in advance, we have to make do with approximations to this ideal. Differential Evolution (DE) provides an approximation to the posterior distribution based on the past history of the chains. In its original formulation DE [129] was designed to work with a population of N parallel chains (all with temperature $T = 1$). The idea is very simple and can be coded in a few lines: Chain i is updated by randomly selecting chains j and k with $j \neq k \neq i$, forming the difference vector $\vec{\theta}_j - \vec{\theta}_k$ and proposing the move

$$\vec{y}_i = \vec{\theta}_i + \gamma(\vec{\theta}_j - \vec{\theta}) \,. \tag{6.9}$$

For D-dimensional multivariate normal distributions, the optimal choice for the scaling is $\gamma = 2.38/\sqrt{2D}$. Since the difference vector points along the D-dimensional error ellipse, the jumps are usually "in the right direction." It is a good idea to occasionally (e.g. 10% of the time) propose jumps with $\gamma = 1$, which act as mode-hopping jumps when the samples (j, k) come from separate modes of the posterior.

The original formulation of DE is not very practical since it requires $N > 2D$ parallel chains for each rung on the temperature ladder. A more economical approach is to use samples from the past history of each chain [130]. It can be shown that this approach is asymptotically Markovian in the limit as one uses the full past history of the chain. We have implemented a variant of the DE algorithm as follows:

- Create a history array for each parallel chain. Initialize a counter M. Store every 10^{th} sample in the history array and add to the counter each time a sample is added. DE moves are more effective if points during the burn-in phase of the search are discarded from the history array.

- Draw two samples from the history array: $j \in [1, M]$, $k \in [1, M]$ and repeat if $k = j$. Propose the move to

$$\vec{y} = \vec{\theta} + \gamma(\vec{\theta}_j - \vec{\theta}_k)\,. \tag{6.10}$$

Here we draw γ from a Gaussian of width $2.38/\sqrt{2D}$ for 90% of the DE updates and set $\gamma = 1$ for the rest.

Figure 8 illustrates this technique.

The standard DE proposal seeks to update all the parameters at once, but it is often more effective to update smaller sub-blocks of highly correlated parameters.

6.2.2. Langevin Proposals

The Metropolis Adjusted Langevin Algorithm (MALA) is a technique for proposing jumps that takes into account the curvature of the posterior distribution. In it, a new point, $\vec{y}$, is proposed from the current point, $\vec{x}$, via [67]

$$\vec{y} = \vec{x} + \frac{\epsilon^2}{2}\mathbb{C}(\vec{x})\nabla \log \pi(\vec{x}) + \epsilon\sqrt{\mathbb{C}(\vec{x})}\vec{z}, \tag{6.11}$$

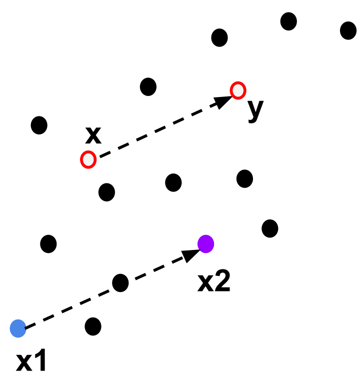

Figure 8: An illustration of the differential evolution technique. Points x1 and x2 are drawn from the history array of the chain, and the vector connecting them is calculated. We then propose to jump along this vector from $\vec{x}$, but multiply by a randomly selected scale. This produces point $\vec{y}$.

where $\vec{z}$ is drawn from a multivariate normal distribution, π is the posterior distribution,$\mathbb{C}$ is the covariance matrix, which is the inverse of the local Fisher matrix, and $\sqrt{\mathbb{C}}$ is the square root of the determinant of $\mathbb{C}$.

In theory, the gradient could be calculated along any direction. In practice, we have found it useful to propose jumps along eigendirections of the Fisher matrix. Thus the gradient is a directional derivative along whichever eigendirection is chosen. Furthermore, ϵ is a parameter that can be set to any value that works well. We have found that $\epsilon = 1.0$ is a successful choice. With these assumptions, the MALA proposals take the form

$$\vec{y} = \vec{x} + \hat{e}_{(i)} \frac{\hat{e}_{(i)} \cdot \nabla \log \pi}{\lambda_{(i)}} + \frac{\delta}{\lambda_{(i)}} \hat{e}_{(i)}, \tag{6.12}$$

when moving along the i^{th} eigenvector. Here, δ is drawn from a zero mean, unit variance Normal distribution, $\mathcal{N}(0,1)$. Note that if the term involving the gradient is dropped from the expression, this is just the same as jumping along Fisher eigendirections.

Thus the MALA proposal distribution is

$$q(\vec{y}|\vec{x}) = \mathcal{N}\left(\vec{x}|\mu(\vec{x}), \mathbb{C}(\vec{x})\right). \tag{6.13}$$

That is, it is a normal distribution, but one that depends on the current location. In particular, $q(\vec{x}|\vec{y}) \neq q(\vec{y}|\vec{x})$, and so one must be careful when using this proposal that the Hastings ratio is calculated correctly.

Figure 9 illustrates how useful this technique is for exploring the posterior distributions of highly correlated parameters. In both panels, we plot the first 100 points of a chain (in blue), showing the values of total mass and the ppE strength parameter. On the left is a chain that is only using Fisher proposals, and on the right is a chain that is only using Langevin proposals. Both chains were started slightly off of the peak likelihood, and both find the peak (i.e. the injected parameters), quite quickly. Also plotted on both panels, in red, is a chain that has been allowed to run for much longer, using both proposals, and thus has mapped out the full posterior distribution. It is clear that the Langevin proposal is much more efficient at exploring the full parameter space.

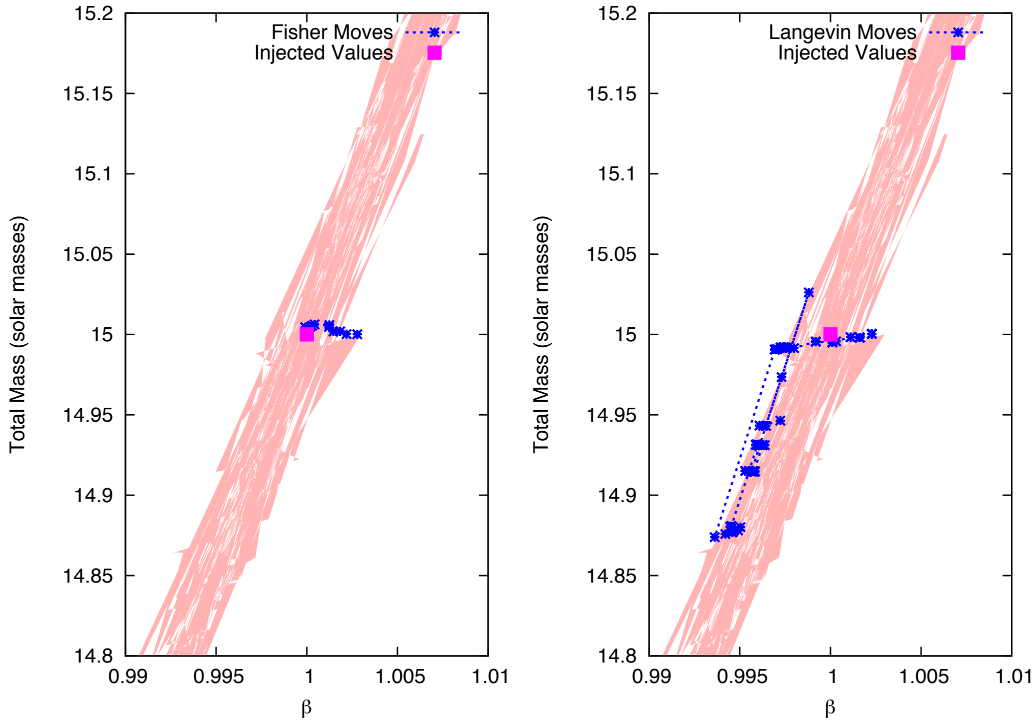

Figure 9: The first 100 steps of a chain (blue) that is using only Fisher proposals (left panel) or only Langevin proposals (right panel). Also plotted (red) is a chain that has run for 20000 steps, using both types of proposals, and has mapped out the full posterior, as well as the injected total mass and β parameter. It is clear that the Langevin proposal is much more efficient at exploring these highly correlated parameters.

6.3. Calculating Bayes Factors

The calculation of Bayes factors is easily the most expensive part of any MCMC analysis. Because the BF between two models is the ratio of the evidences for each of the models, calculating it requires the evaluation of large, multidimensional integrals. There are numerous techniques for doing this, a few of which we will discuss in this section. Some of these techniques are much cheaper computationally than others, but not all are applicable in all situations. We have used all of these techniques at different points in the research presented in this thesis.

6.3.1. Thermodynamic Integration

If we had the time and computing power to evaluate our function at every point in parameter space, the problem of calculating the evidence, and thus of evaluating the Bayes factor, would be moot. Because we are limited, however, it is important that we ensure that the parameter space has been sufficiently sampled that our technique for estimating the integrand is sufficiently accurate. With a single Markov chain exploring a given likelihood surface, it can often take a prohibitively large number of iterations for the full space to be well sampled. This is particularly a problem when the likelihood function involves multiple peaks of differing sizes. Although eventually the chain will find all of the peaks, it is possible for it to become stuck at a local maximum for a large number of iterations. One technique for avoiding this trap, and for efficiently calculating the evidence, is called parallel tempering. All information in this paper about parallel tempering is adapted from [4].

Parallel tempering is parallel because it involves running several Markov chains simultaneously, and allowing them the communicate information to each other. The 'tempering' piece refers to the fact that each chain sees a likelihood surface of a different 'temperature', according to:

$$p(\vec{\theta}|d,\beta) \propto p(\vec{\theta})p(d|\vec{\theta})^{\beta} \tag{6.14}$$

where β is the inverse of the temperature. Each chain is assigned a different β, and so each sees a likelihood surface that is slightly different. The effect of raising the likelihood to the power of $1/T$ is that the likelihood becomes progressively smoother for each higher temperature chain. Peaks become lower, valleys become higher, and it becomes more difficult for a chain to get stuck on a local maximum. This effect is illustrated in Figure 1, which shows the recovered posterior distribution of a Gaussian for chains of successively higher heat.

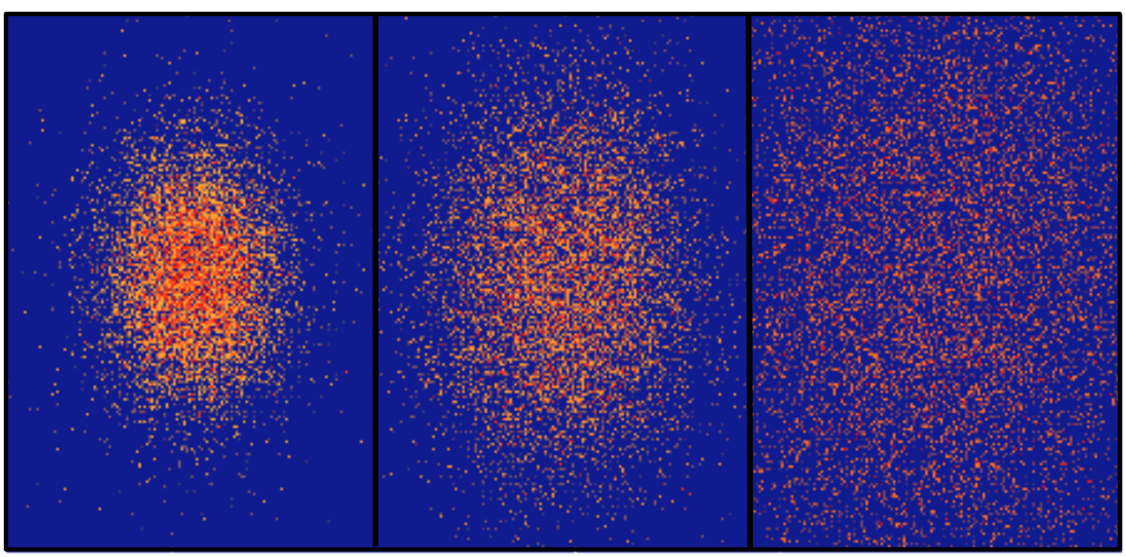

Figure 10: The recovered PDF's for three chains. The cold chain is on the left, and the hottest chain is on the right. All chains were searching the same Gaussian surface, but it is clear that the hottest chain saw a significantly smoother surface than the cold chain.

We then allow the chains to 'talk' to each other - that is, they compare the points in parameter space that each is exploring. They then have the option of switching places with each other according to the Hastings ratio:

$$H_{i\leftrightarrow j} = \frac{p(d|\vec{\theta}_i, \beta_j)p(d|\vec{\theta}_j, \beta_i)}{p(d|\vec{\theta}_i, \beta_i)p(d|\vec{\theta}_j, \beta_i)} \tag{6.15}$$

Using several chains in parallel running at different temperatures allows for a very fast exploration of even complicated parameter spaces. They are also helpful for accurately calculating the evidence. The trick in using parallel tempering in this way is called 'thermodynamic integration,' (TI) so named because it involves treating the posterior distribution function as a partition function, $Z(\beta)$, such that:

$$\begin{aligned} Z(\beta) &\equiv \int d\vec{\theta} p(d|\vec{\theta}, \beta)p(\vec{\theta}) \\ &= \int d\vec{\theta} (d|\vec{\theta})^{\beta} p(\vec{\theta})) \end{aligned} \tag{6.16}$$

Because the second function under the integral, the prior, is independent of β, we can write:

$$\frac{d}{d\beta} \ln Z(\beta) = \langle \ln p(d|\vec{\theta}) \rangle_{\beta} \tag{6.17}$$

where on the right we have the expectation value of the likelihood for a chain with inverse temperature $1/\beta$. We can then calculate the evidence by integrating over β.

$$\ln p(d|\mathcal{M}) = \int_0^1 d\beta \langle p(d|\vec{\theta}, \mathcal{M}\rangle_{\beta}) \tag{6.18}$$

Here we write $p(d|\mathcal{M})$ to explicitly draw attention to the fact that this is the evidence for a particular model, $\mathcal{M}$. This equation tells us that in order to calculate the evidence for a given model, or, in the context of this paper, the integral of whatever function we're interested in, all we need to do is have a sufficient number of parallel chains exploring the space well. We can then combine the average likelihoods for each of these chains in a particular way, weighted by the inverse of their heat, in order to find the answer we are looking for. This is much easier in practice than getting a single chain to explore a complicated parameter space well enough to confidently calculate the integral, because the hottest chains will very quickly explore the entire parameter range.

6.3.2. Volume Tessellation Algorithm

The volume tessellation algorithm, or VTA, relies on Markov chain sample that has accurately reproduced the posterior. Given this sample, the idea of the VTA is to partition the points in the chain into volumes of parameter space, then calculate the integral directly by adding together all of these volumes. Because Markov chains spend little time in areas of the posterior distribution with low likelihood, these volume elements will be poorly sampled, but they add very little to the integral anyway. There are many ways that one could go about partitioning the volume of parameter space covered by the Markov chain sample. In this particular implementation of the algorithm, a k-d tree decomposition has been used. This algorithm was developed by Martin Weinberg, as in [137].

k-d tree decompositions are a particular way to split up multi-dimensional spaces. Each 'node' in a k-d (short for k-dimensional) tree effectively splits the space in half with a hyperplane. Then half of the points in the distribution are on one side of the node, and the other half are on the other side. In the implementation of the VTA discussed in this paper, the nodes of the k-d tree are at the median of a given parameter. So the first cut is made along the median of each parameter for the entire distribution. Then a median is determined for each of these halves of the distribution, and another cut is made there. This procedure is continued until there are a set number of points in each volume element delineated by the k-d nodes. The number of points in each box, the 'boxing number,' is a parameter that can be changed. Once the points are sorted into their particular volume elements, the average value of the integrand is calculated in each element, and multiplied by dV, in what is essentially a direct calculation of the integral. This algorithm is illustrated in Figure 2.

There is of course a tradeoff between high and low boxing numbers. If the number is higher, the estimate for the average integrand in a given volume is better, but of course the volume itself is larger and so small features in the distribution are easier to miss. And, as always with Markov methods, the more samples in the Markov chain you are working with, the more accurate the estimate for the integral will be.

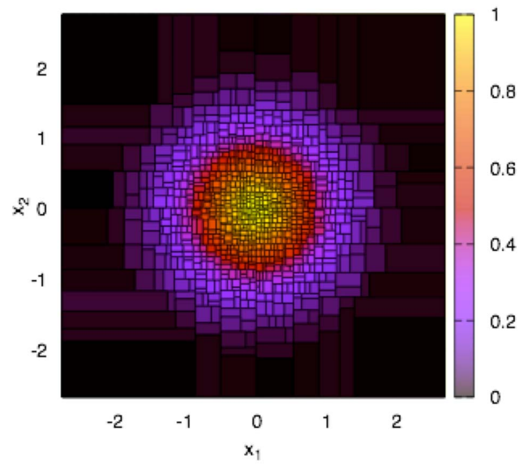

Figure 11: An illustration of the kd tree decomposition of a two-dimensional Gaussian. The color-coding corresponds to areas of high posterior density. It is easy to see that the peak of the distribution is sampled most densely, and thus has its volume calculated most accurately. This figure is reproduced from [137]

6.3.3. Reversible Jump MCMC

Both thermodynamic integration and the VTA are methods for directly calculating the evidence for a given model. The Bayes factor between two models is then simply the ratio of the evidences for the two models.

In Reversible Jump MCMC (RJMCMC) [71, 93, 92], the Markov chains are allowed to transition between competing models during the run. Effectively, this makes the model itself one of the search parameters. The chain spends more time exploring the model that is better supported by the data, and thus the Bayes factor between, say model $\mathcal{M}_1$ and model $\mathcal{M}_2$ is just the ratio of the number of iterations spent in $\mathcal{M}_1$ to the number in $\mathcal{M}_2$. That is:

$$B_{12} = \frac{\#\text{ of iterations in }1}{\#\text{ of iterations in }2} \tag{6.19}$$

In order for the RJMCMC algorithm to converge to the correct posterior distribution, the Hastings ratio for jumping between two models is calculated via [93]

$$H_{\mathcal{M}_i \to \mathcal{M}_j} = \frac{p(d|\vec{\theta}_j, \mathcal{M}_j) p(\vec{\theta}_j, \mathcal{M}_j) q(\vec{u}_i, \mathcal{M}_i)}{p(d|\vec{\theta}_i, \mathcal{M}_i) p(\vec{\theta}_i, \mathcal{M}_i) q(\vec{u}_j, \mathcal{M}_j)} |J_{ij}|. \tag{6.20}$$

In this expression, $\vec{\theta}_j$ represents the set of parameters in model j, and $\vec{u}_j$ represents a set of random numbers used to generate the parameters in $\mathcal{M}_j$, and the Jacobian factor [92]

$$|J_{ij}| = \left| \frac{\partial(\vec{\theta}_j, \vec{u}_i)}{\partial(\vec{\theta}_i, \vec{u}_j)} \right| \tag{6.21}$$

accounts for any difference in dimensionality between the two models.

For example, consider two models, one of dimension 1 and one of dimension 2. Let model $\mathcal{M}_1$ has free parameter, θ, and model $\mathcal{M}_2$ has free parameters, θ and ϕ. When proposing a jump from $\mathcal{M}_1$ to $\mathcal{M}_2$, we generate ϕ using a random number, u_ϕ, via $\phi = u_\phi \Delta\phi$, where Δ_ϕ is the prior range of ϕ. This Jacobian factor is then equal to

$$|J_{12}| = \begin{vmatrix} \frac{\partial\theta}{\partial\theta} & \frac{\partial\theta}{\partial u_\phi} \\ \frac{\partial\phi}{\partial\theta} & \frac{\partial\phi}{\partial u_\phi} \end{vmatrix} = \Delta\phi. \tag{6.22}$$

6.3.4. Savage-Dickey Density Ratio

As we have noted, the hypothesis that is GR correct is nested within the hypothesis that ppE is correct. That is, if all of the ppE strength parameters are equal to zero, $\beta_b = 0 \ \forall \ b$, the resulting waveform is the GR waveform. Because of this nested nature, we can use a technique called the Savage-Dickey (SD) density ratio to calculate the Bayes factor between ppE and GR.

The SD technique works as follows. Say we have a model M_0, described by $p(y|\omega, \vec{\theta}, M_0)$, which depends on the parameters ω and $\vec{\theta}$. Given another model, M_1, that is a restricted version of M_0, i.e.

$$p(y|\vec{\theta}, M_1) = p(y|\omega = \omega_0, \vec{\theta}, M_0), \tag{6.23}$$

the Bayes factor between M_1 and M_0 can be calculated via

$$B_{01} = \frac{\pi(\omega = \omega_0|y, M_0)}{\pi(\omega = \omega_0|M_0)}. \tag{6.24}$$

In this expression, the numerator is the posterior probability that $\omega = \omega_0$, and the denominator is the prior probability.

This concept is illustrated in Figure 12. In this figure, the prior probability on the parameter ω is uniform between = 0.5 and 0.5 - giving it a prior density of 1.0 in M_0. The value of ω in M_1 is taken to be $\omega = 0$. For the top panel, the posterior probability density at $\omega = 0$ is much larger than the prior probability density, resulting in a Bayes factor that favors the lower-dimensional model. In the middle panel, the opposite is the case. In the bottom panel, the prior and posterior probabilities are equal, and neither model is favored.

To use the Savage-Dickey density ratio in our studies, we simply use MCMC techniques to map out the posteriors of the ppE strength parameters, and compare the posterior probability at $\beta = 0$ to the prior probability at that same point.

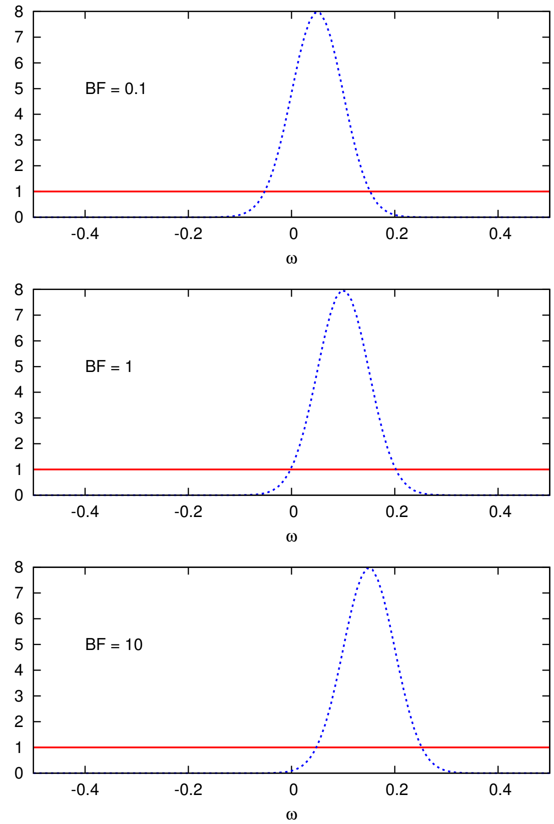

Figure 12: An example of calculating Bayes factors from the Savage-Dickey density ratio. The red lines show the prior probability for parameter ω, and the blue lines show possible posterior distributions for this parameter. Here, $\omega = 0$ corresponds to model M_0. The Bayes factor for M_0 vs. M_1 is equal to the ratio of the prior to the posterior. In the top panel, M_0 is favored. In the middle panel, M_1 is favored. In the bottom panel, neither model is better supported by the data.

6.3.5. Comparison of Techniques

TI is the most expensive of these four techniques for calculating the BF. This is because it requires many chains to be run in parallel - many more than are useful for promoting exploration of the parameter space. Furthermore, in order for the results to be robust, the temperature spacing of the chains needs to be carefully adjusted such, often necessitating lengthy pilot runs in order to determine the ideal spacing. On the other hand, TI is applicable in all situations.

The VTA is also applicable in all situations, but the results can be quite sensitive to tunable parameters in the algorithm, such as the boxing number. We have done a small study comparing TI and VTA applied to calculating the integral of a multivariate Gaussian, integrated over all dimensions from -3σ to $+3\sigma$. We chose $\sigma = 2$ for all standard deviations for ease of programming and calculation. For both thermodynamic integration and the VTA, we examined how the integration performed for different numbers of iterations in the Markov chain, as well as over integrals of different dimensionality.

The results for thermodynamic integration are shown in Figures **??** and 14. Figure**??**shows the value recovered for the integral of a four dimensional Gaussian, using both different numbers of iterations and different numbers of parallel chains. The value returned from TI is $-2\log(\text{integral})$, where in the plots $E = \text{integral}$. We find that for all numbers of chains, the returned value quickly loses its dependence on number of iterations, but it is absolutely necessary to have enough chains to get an accurate value for the integral. In this case, 20 chains were necessary. The error bars in this plot reflect the standard usage that thermodynamic integration on this type of function returns an answer that is good to ± 0.5.

Figure 14 explores how TI performs on integrals of different dimensionality. Because four was the lowest number of dimensions we were looking at, I chose to use

20 and 25 chains to explore the parameter spaces. We ran all searches for 150000 iterations, as this was more than enough for all of the cases in Figure **??** to fully converge. As expected, both numbers of chains are able to accurately recover the integral for lower dimensionality, but as the parameter space grows, so does the error. This could be rectified with longer Markov chain runs. The values returned in this plot are also good to ± 0.5, but for ease of reading we have not included error bars in this plot.

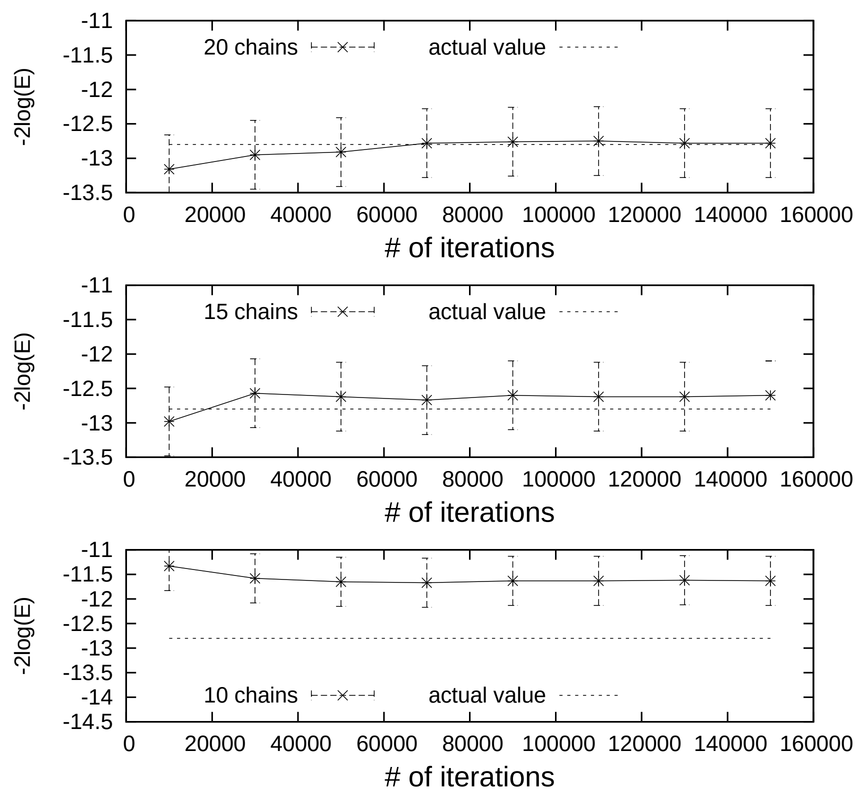

Figure 13: This figure shows the recovered value for a the integral of a four-dimensional Gaussian using thermodynamic integration. The different sets of points use different numbers of chains. The y axis shows ln(integral), and the x axis shows number of points in the Markov chain. For all numbers of chains, the returned answer becomes very stable for high iterations.

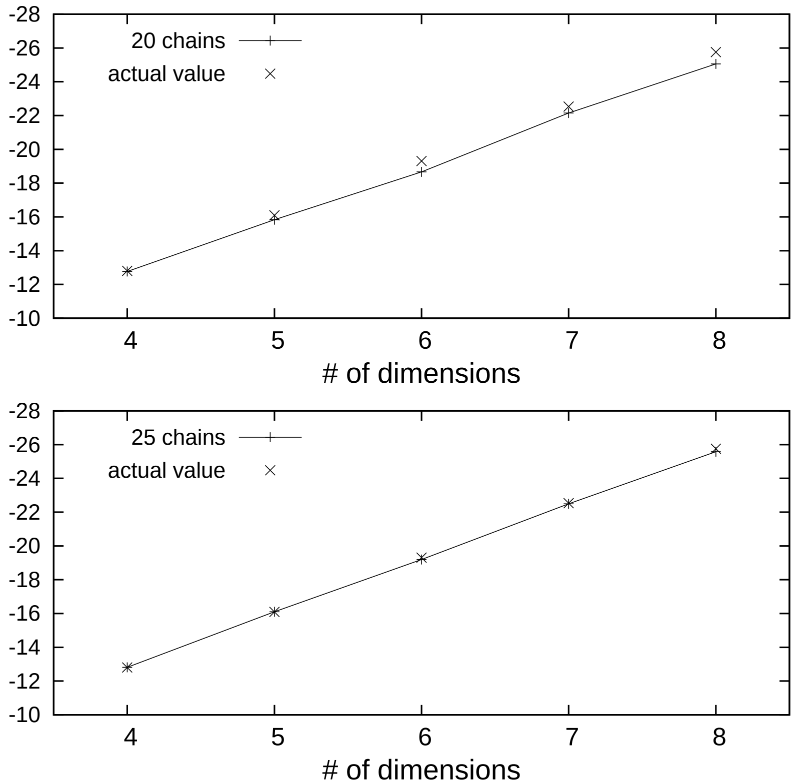

Figure 14: This figure shows the performance of thermodynamic integration for integrals of different dimensionality. We plot the recovered value for TI using 20 chains and 25 chains. All runs were for 150000 points. For higher dimensional problems, more chains or more iterations are required to recover the correct value.

Next we looked at the VTA algorithm. Similarly to thermodynamic integration, there are two parameters in VTA that can be adjusted. These are the number of samples in the chain and the boxing size, i.e. the number of samples that are put into each volume element by the decomposition tree. Figure 15 shows how the returned integrand for a four-dimensional Gaussian depended on boxing size and number of iterations. In general, the higher boxing numbers led to higher values for the integral, although this effect became smaller for chains with larger numbers of iterations. The chain with the smallest number of iterations showed the greatest variance of answers for different boxing sizes. The VTA returns a value for $4\log(\text{integral})$. The error bars in this figure show an error of ± 1, which was calculated in [137].

We expect that boxing size should have a somewhat large effect on the returned value for the integral, as if the boxing size is too large the average value of the integrand (larger than the actual value of the integral) will be taken to apply to a large volume of parameter space. Conversely, if the boxing size is too small, there will be many parts of parameter space that have far too few samples in them to accurately calculate the average of the function being integrated. The more points in the Markov chain, the smaller this sensitivity on boxing size becomes.

Finally, in Figure 16 we show an analysis of how the VTA performs on integrals of different dimensions. We find the boxing sizes of 16 and 32 perform about as well as each other, and, as expected, the performance drops off as the dimensionality increases. Each of these calculations was done using a Markov chain containing 150000 points, and a search that used eight chains.

RJMCMC is a wonderful tool for choosing between models of different dimensionality. Its primary drawback is that it can be difficult to get the chains to transition between models, unless a tailored move is made for this purpose. When we have used this method for calculating BFs, we have first run a pilot analysis in each of the models we are interested in differentiating between. In this initial run, we keep the model fixed, and explore the posterior distribution to find a region of high likelihood. We then calculate the local Fisher matrix in this region, and use that matrix (different in

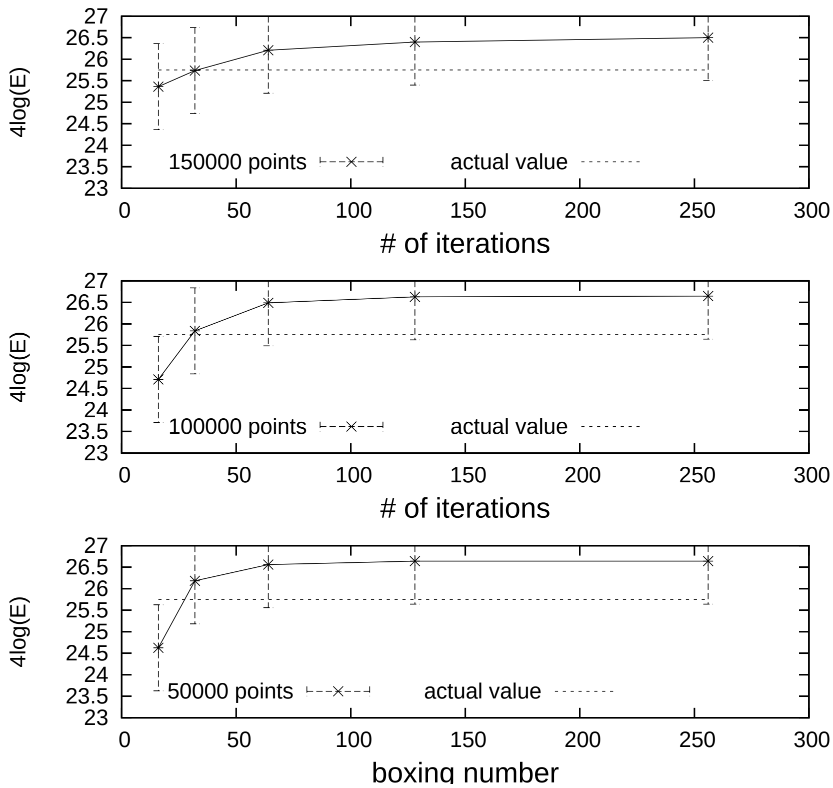

Figure 15: This figure shows the recovered value for a the integral of a four-dimensional Gaussian using the VTA. The different lines are different numbers of points, and the x axis shows different boxing sizes. Large boxing sizes tend to overestimate the value of the integral, as they assign a large number for the average value of the integrand to a large volume of space.

each model) to propose jumps between models. This method is extremely effective, but it is also quite expensive.

By far the least expensive method for calculating BF's is the Savage-Dickey density ratio. It is also quite robust to different choices that can be made in the MCMC run, such as the number of chains, the different types of jumps, etc. It is, however,

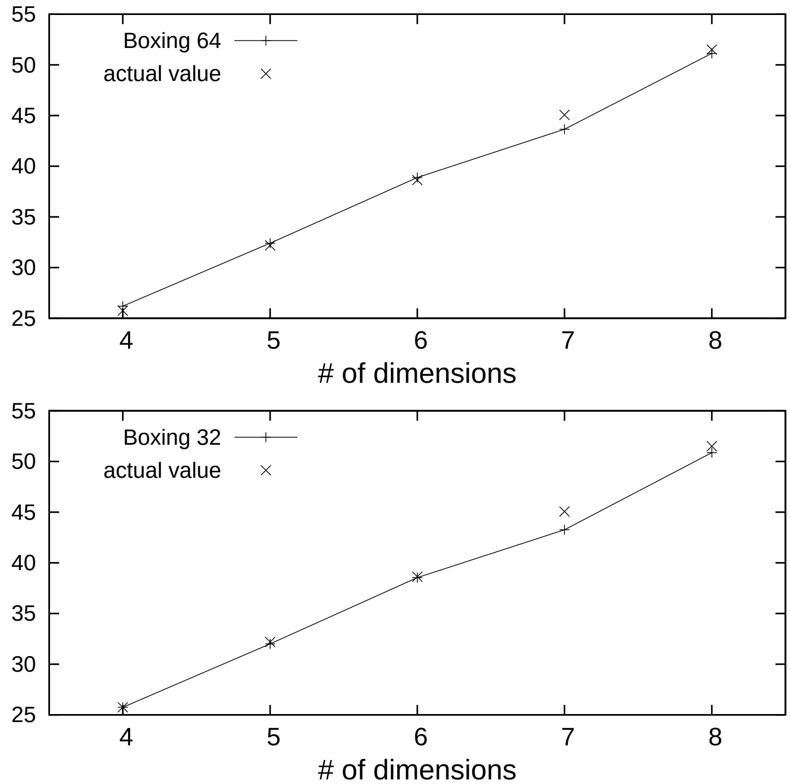

Figure 16: This figure shows the recovered value for a the integral of Gaussians of different dimensionality, using two boxing sizes. All runs used 150000 points. Although both boxing sizes recovered the actual value with good accuracy, the accuracy gets worse and worse with higher dimensionality.

only useful for calculating the BF between nested models. Luckily, most of our analyses involve nested models, and so it is this technique that we have employed most often.

6.4. Fitting Factor

In addition to the Bayes factor, another quantity that can be calculated via MCMC techniques is the fitting factor, which measures how well one template family can recover an alternative template family. To define the fitting factor, we must first define the match between two templates h and h' as

$$\mathcal{M} = \frac{(h|h')}{\sqrt{(h|h)}\sqrt{(h'|h')}} \,. \tag{6.25}$$

The match is related to the metric distance between templates [105] by $\mathcal{M} = 1 - \frac{1}{2} g_{ij} \Delta x^i \Delta x^j$, where the metric is evaluated with the higher-dimensional model (appropriate when dealing with nested models). The fitting factor FF is then defined as the best match that can be achieved by varying the parameters of the h' template family to match the template belonging to the the other family, h.

Another interpretation for the fitting factor is as the fraction of the true signal-to-noise ratio $\mathrm{SNR} = \sqrt{(h|h)}$ that is recovered by the frequentist statistic $\rho = \max[(h|h')/\sqrt{(h'|h')}]$. The imperfect fit leaves behind a residual $(h-h')$ with $\mathrm{SNR}^2_{\mathrm{res}} = \chi^2$, which can be minimized by adjusting the amplitude of h' to yield

$$\mathrm{SNR}^2_{\mathrm{res}} = (1 - \mathrm{FF}^2)\mathrm{SNR}^2 \,. \tag{6.26}$$

Assuming that a residual with SNR_* is detectable, and working in the limit where $\mathrm{FF} \sim 1$, we have

$$1 - \mathrm{FF} \simeq \frac{\mathrm{SNR}^2_*}{2\,\mathrm{SNR}^2} \,. \tag{6.27}$$

We see then that the ability to detect departures from GR scales inversely with the square of the SNR, as given by Eq. (6.27). On the other hand, the detectable difference between the parameters in the two theories will scale inversely with a single power of the SNR. This is because this detectable difference is proportional to the square-root of the minimized match function and

$$\sqrt{\min(g_{ij}\Delta x^i \Delta x^j)} \simeq \frac{\mathrm{SNR}_*}{\mathrm{SNR}}\,, \tag{6.28}$$

and the metric is independent of SNR. This reasoning applies to both the additional model parameters of the alternative theory, *e.g.* $\Delta x^i = (\alpha, \beta)$, and the physical source parameters such as the masses and distance. We then expect both the bounds on the ppE model parameters and the biases caused by using the wrong template family to scale inversely with SNR. This scaling is in keeping with the usual scaling of parameter estimation errors that follows from a Fisher matrix analysis where $\langle \Delta x^i \Delta x^j \rangle \simeq (h_{,i}|h_{,j})^{-1} \sim \mathrm{SNR}^{-2}$. Figure 17 shows that the errors in the recovery of the ppE parameters follows the expected scaling with SNR.

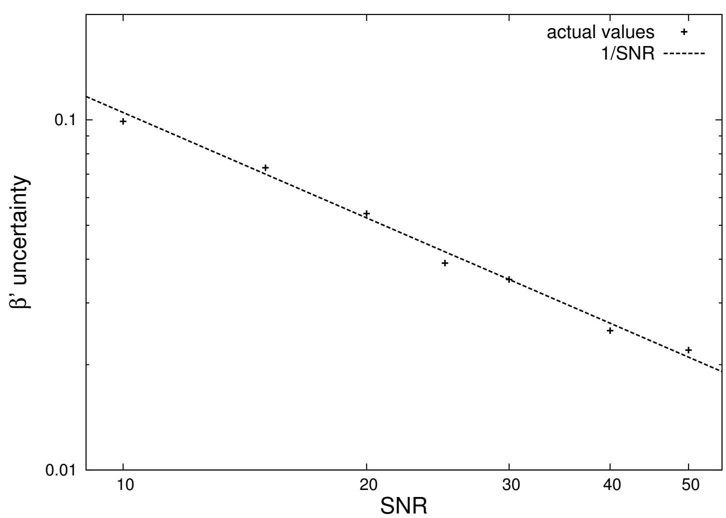

Figure 17: The scaling of the parameter estimation error in the ppE parameter β for an aLIGO simulation with ppE parameters $(a, \alpha, b, \beta) = (0, 0, -1.25, 0.1)$. The parameter errors follow the usual 1/SNR scaling.

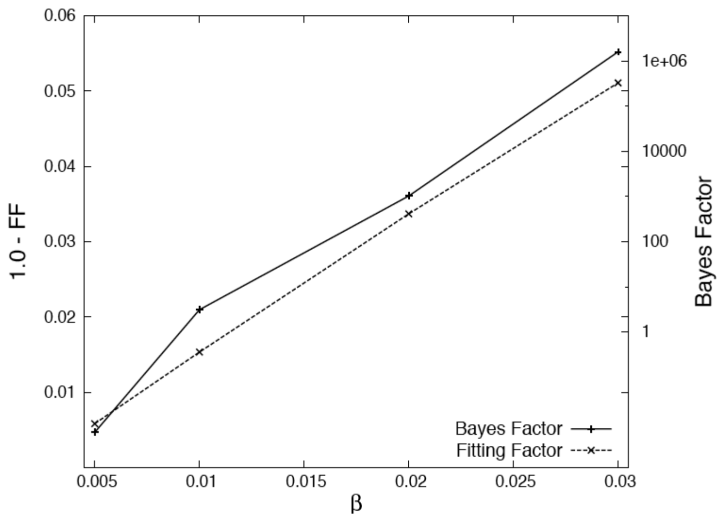

Figure 18: The log Bayes factors and $(1 - \mathrm{FF})$ plotted as a function of β for a ppE injection with parameters $(a, \alpha, b, \beta) = (0, 0, -3.75, \beta)$. The predicted link between the fitting factor and Bayes factor is clearly apparent.

Alternative models that are not well-fitted by GR will be more easily distinguished than models that can be well-fitted. This suggests that there should be a correlation between the fitting factor and the Bayes factor. The relationship can be established using the Laplace approximation to the evidence [Eq. (7.9)], from which it follows that the log Bayes factor is equal to

$$\begin{aligned}
\log \mathrm{BF} &= \log \frac{e^{-\chi^2(\mathcal{H}_1)/2}}{e^{-\chi^2(\mathcal{H}_0)/2}} \frac{\mathcal{O}_1}{\mathcal{O}_0} \\
&= \frac{\chi^2_{\min}}{2} + \Delta \log \mathcal{O} \\
&= (1 - \mathrm{FF}^2) \frac{\mathrm{SNR}^2}{2} + \Delta \log \mathcal{O} \,, \qquad (6.29)
\end{aligned}$$

where $\mathcal{O}$ is the Occam factor, defined in the discussion following [Eq. (7.9)]. Thus, up to the difference in the log Occam factors, the log Bayes factor should scale as $2(1 - \mathrm{FF})$ when $\mathrm{FF} \sim 1$. This link is confirmed in Figure 18.

Different non-GR theories of gravity lead to modifications to the phase and amplitude of GWs at different post-Newtonian (PN) orders. For theories that lead to low PN-order[1] deviations from GR, e.g. Brans-Dicke gravity leads to changes in the phase at -1PN order relative to the leading-order GR prediction [163], there are already strong constraints on the possible size of the deviation from current Solar System and pulsar timing experiments. The current bounds placed on ppE parameters by pulsar timing experiments are shown in Fig. 19 [158]. As stated, for large, negative values of b, i.e. low-PN order terms, current data from pulsar binary systems places very tight bounds on possible deviations in the phase of GW signals from the GR expectation.

On the other hand, for theories that lead to high PN-order deviations, the existing constraints are very weak. This is also shown in Fig. 19, which shows that at high-PN order, for less negative values of b, the bounds become very weak. Therefore, it is possible that GW signals will differ greatly from the GR predictions, but only once the characteristic velocity of the system becomes quite high, thus avoiding current experimental bounds. In this section, we examine the possible effects of using GR templates to search for non-GR signals that contain large deviations from GR at high PN order. Although 'large' in the sense that they would be easily differentiable from GR, these deviations are not currently ruled out by any experimental evidence.

[1]A modification at NPN order is one which is proportional to $(v/c)^{2N}$ relative to the leading-order term, where v is the characteristic orbital velocity of the binary and c is the speed of light.

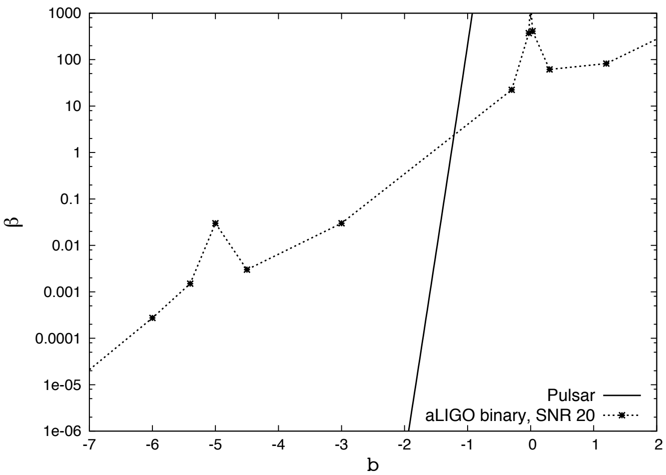

Figure 19: Bounds that can be placed on the ppE strength parameter, β, for various values of the ppE exponent, b, using GW measurements [38] and measurements from binary pulsars [158]. The regions above the pulsar line are already ruled out by experiment.

To test what effect a signal containing these high-order PN deviations could have on our ability to detect and characterize GWs using GR templates, we inject non-spinning ppE inspiral signals of the form of Eq. (5.18), containing two phase corrections, with $b = -1$ and $b = 1$. This corresponds to adding both a 2PN and a 3PN order correction to the GW phase. The β_{-1} parameter is chosen to be a deviation from GR that is not ruled out by current experimental bounds, and β_1 is chosen to be larger by a factor of 1000. This factor is a conservative one - based on the bounds on β_b shown in Fig. 19, it is clear that an even larger ratio would be well within experimental limits. These system studied is a $1.4M.$ neutron star and a $3.5M_\odot$ black hole in a quasi-circular orbit.

The fractional loss of events due to mis-modeling errors in the templates scales as $1 - \mathrm{FF}^3$, so to achieve 90% efficiency we must have FF > 0.97. FFs below 0.5 imply

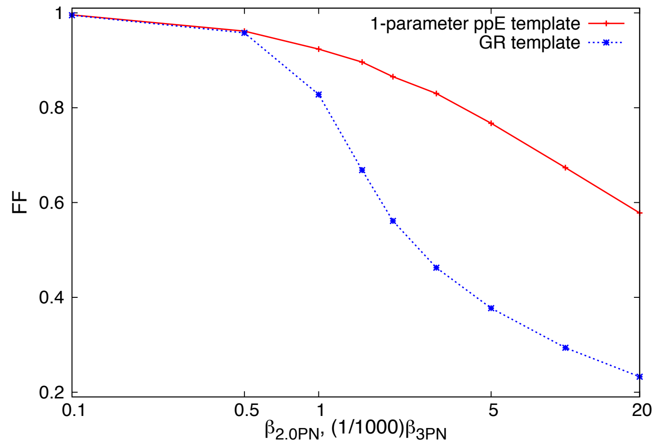

Figure 20: FF between both a GR (dashed/blue line) and a one-parameter ppE template with $b = -1$ (solid/red line) and the injected, non-GR signal as a function of β_{2PN}. The FF drops well below the desired value of 0.97 for values of β_{2PN} that are fully consistent with existing experimental bounds.

a 90% loss of signals, while FFs below 0.2 imply a 99% loss of signals. Thus, if the FF of GR templates with plausible non-GR signals is significantly lower than 0.97, we could miss GW signals that are present in our data streams.

Figure 20 shows the FFs between our injected, non-GR signals and non-spinning, circular GR templates. For small values of β_{-2}, as expected, the GR templates can achieve a near-perfect FF. Quite quickly, however, as β_{-2} increases, the FF drops below the desired level of FF $= 0.97$. Well before the injected signals are ruled out by current pulsar constraints, the FF drops to 0.2, and it is likely that all such signals would be missed. The decrease in FF is accompanied by an increasing bias in the recovered chirp mass, shown in Fig. 21.

Figure 20 also shows the FF and recovered chirp mass for a simple ppE template that contains only one strength parameter, β_{-1}, with exponent $b = -1$ (recall that the

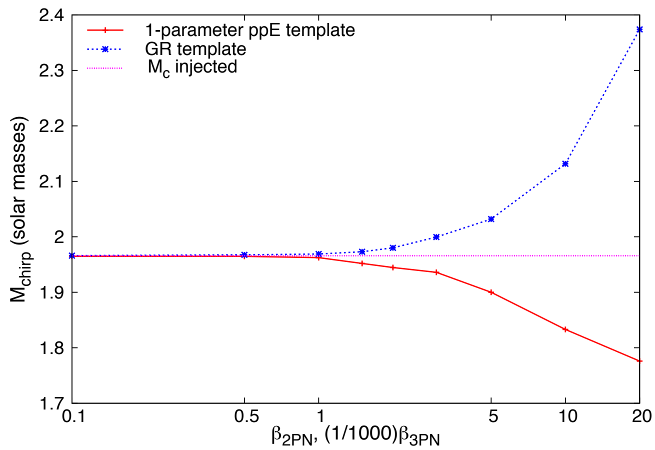

Figure 21: Recovered value of $\mathcal{M}$ from both a GR (dashed/blue line) and a one-parameter ppE template with $b = -1$ (solid/red line) used to fit a non-GR signal, as a function of $\beta_{1.5PN}$. The true value of $\mathcal{M}$ is $\mathcal{M} = 10M_{\odot}$. Using GR templates to fit non-GR signals leads to large biases in the recovered parameters.

injected signals have two ppE strength parameters -β_{-2} and β_{-1}). These templates perform much better than the GR ones at detecting the signal, but suffer from similar issues in biased recovery of the chirp mass.

This outcome is not surprising – we know that templates are only effective in detecting signals that are at least somewhat similar to them. It is important, though, to be aware that there are non-GR signals, completely consistent with current experiment, that would be entirely missed using GR templates.

7. TESTING GR USING PPE TEMPLATES

Convincing alternative models to GR are hard to find because none of the currently proposed alternatives, several of which are discussed in Ch. 4, can satisfy key criteria that physicists would like to require. On the observational front, one wishes that any GR alternative passes all Solar System and binary pulsar tests with flying colors, only predicting deviations from GR in the strong-field regime, where tests are currently lacking. Many theories, such as Brans-Dicke theory [142, 119, 147, 19, 152], are heavily constrained by this requirement [144]. On the theoretical front, one would wish viable GR alternatives to lead to well-posed theories, with a postive definite Hamiltonian and free of instabilities. All perturbative string theory and loop quantum gravity low-energy effective theories [6, 166] currently lead to higher-derivative theories, which might violate this theoretical criteria.

The paucity of concrete alternative models to GR [123] has impacted testing grounds beyond simply GWs. For example - those based on solar system observations, or binary pulsar systems. In those instances the standard approach has been to develop models that parameterize a wide class of possible departures from GR - the parameterized post-Newtonian formalism [103, 140, 145, 104] and the parameterized post-Keplerian formalism [49]. It is natural to adopt the same strategy when analyzing gravitational wave data, which leads to the parameterized post-Einsteinian (ppE) formalism derived in Ref. [163]. These parameterizations are discussed in detail in Ch. 5. In this chapter, we are interested only in the ppE parameterization, the simplest form of which is

$$\tilde{h}(f) = \tilde{h}_{\rm GR}(f)\,[1 + \alpha\; u^a]\,e^{i\beta\,u^b}, \tag{7.1}$$

where (α, a) are amplitude ppE parameters and (β, b) are phase ppE parameters. As noted previously, both α and β can depend on the spin angular momenta and mass difference of the two bodies, as well as the symmetric mass ratio of the system.

In one study, Arun *et.al.* [10, 11, 101] considered what can now be interpreted as a restricted version of the ppE formalism, in which the exponents, a_i and b_i, are required to match those found in GR. This amounts to asking how well the standard PN expansion coefficients could be recovered from gravitational wave observations. They also developed internal self-consistency checks based on the observation that each coefficient $\psi_k(\eta)$ provides an independent estimate of the mass ratio η. While interesting, these tests are limited in scope as few of the well known alternative theories of gravity (Brans-Dicke [142, 119, 147, 19, 152], Massive Graviton [143, 147, 19, 126, 12, 81, 152], Chern-Simons [5, 161, 162, 124, 6], Variable G [164], TeVeS [17]*etc.*) have corrections with exponents a_i and b_i that match those of GR [163]. The full ppE formalism allows us to look for a much wider set of possible departures from GR.

Our goal here is to study how the ppE formalism can be used to search for waveform deviations from GR using data from the next generation of ground based interferometers (aLIGO/aVirgo) and future space based interferometers (*e.g.* LISA). Bayesian model selection is used to determine the level at which departures from GR can be detected (See Ref.[51] for a related study that uses Bayesian inference to study constraints on Massive Graviton theories). Advanced MCMC techniques are used to map out the posterior distributions for the models under consideration. From these distributions, we are able to quantify the degree of fundamental bias in parameter extraction, and in particular, if the fundamental bias can be significant in situations where there is no clear indication that there are departures from GR.

Recently, Pozzo et.al. [37] performed a similar study that applied Bayesian model selection to estimate the bounds that could be placed on massive graviton theory. As such, their work is a sub-case of the ppE framework, *i.e.* a particular choice of (b, β). Their implementation differed from ours in that they used Nested Sampling while we used MCMC techniques, but as we will show, our results are in agreement with theirs for the relevant sub-case.

We find that gravitational wave observations will allow us to extend the existing bounds derived from pulsar orbital decay [159] into the region of parameter space that covers strong field departures from GR ($a_i, b_i > 0$) (see Fig. 22–23 in Sec. **??**). As expected, we find that the strength of the bounds on the ppE parameters are inversely proportional to the signal-to-noise ratio (SNR), and the extent to which deviations between GR templates and non-GR signals can be detected (the departure of the "fitting factor" from unity) scales as $1/\mathrm{SNR}^2$. The logarithm of the odds ratio used to decide if a signal is described by GR or some alternative theory follows the same $1/\mathrm{SNR}^2$ scaling. A more surprising result is the possibility of "stealth bias" whereby the parameters recovered using GR templates can be significantly biased even when the odds ratio shows no clear preference for adopting an alternative theory of gravity.

Next, we revisit the ppE framework and carry out a more realistic data analysis study. First, we examine the effect of more realistic non-GR injections that include modifications to several terms in the PN GR phase, instead of a single one. Generic deviations from GR will be characterized by an infinite number of phase corrections. Ground-based detectors will not be sensitive to all of them, just as they are not sensitive to GR signals to arbitrarily high PN order. The presence of the first few higher-order terms can affect our ability to test GR. We find that the presence of multiple phase modifications will improve our chances of detecting departures from GR if they are of the same sign. However, if the phase modifications are of alternating sign, they can cancel out to some degree, and make a non-GR signal appear to be well described by GR.

As something of an aside, we consider the issue of adding explicit noise realizations to the simulated signals, especially for low SNR signals. This is done because some concerns have been voiced about the conclusion of the Cornish *et al* [38] work due to the relatively high SNR of the signals used, and their technique of accounting for the noise solely through the weighting of the likelihood function. We analytically and numerically show that the conclusions of [38] remain unaffected when adding an explicit noise realization. We also show that these results scale linearly with SNR down to values close to the detection threshold, which for this source was SNR ~ 7.5.

We then tackle the problem of determining the *optimal* ppE model for detecting departures from GR. On the one hand, including additional phase terms will improve the fit and increase the likelihood. On the other hand, adding additional parameters to the model incurs an "Occam penalty". We find, on balance, that in almost all cases, templates with only one ppE parameter are preferred over those with multiple parameters. These suggests that the simple one-parameter ppE model may well be the ideal one to search for GR deviations in early data from advanced detectors.

7.1. Analysis Framework

Here we discuss the instrument response, noise curves, and likelihood function that were used in our analysis.

7.1.1. Instrument Response

The aLIGO/aVirgo analysis was performed using simulated data from the 4 km Hanford and Livingston detectors and the 3 km Virgo detector. The time delays between the sites and the antenna beam patterns were computed using the expression quoted in Ref. [9]. Since the detectors barely move relative to the source during the time the signal is in-band, the antenna patterns can be treated as fixed and the time delays Δt between the sites can be inserted as phase shifts of the form $2\pi f \Delta t$. For the instrument noise spectral density, we assumed all three instruments were operating in a wide-band configuration with

$$S_n(f) = 10^{-49}\left(x^{-4.14} - 5x^{-2} + 111\frac{(2 - 2x^2 + x^4)}{2 + x^2}\right), \tag{7.2}$$

and $x = (f/215\text{Hz})$.

The space based (LISA) analysis was performed using the A and E Time Delay Interferometry channels [109] in the low frequency approximation [41, 111]. It is known that this approximation can lead to biases in some of the recovered parameters, such as polarization and inclination angles. This, however, is an example of a modeling bias introduced by inaccurate physical assumptions, and not of a fundamental bias resulting from incomplete knowledge of the theory describing gravity. We therefore do not include its effects in our current analysis. In contrast to the ground based detectors, the signals seen by LISA are in-band for an extended period of time, and the motion of the detector needs to be taken into account. The time dependent phase delay between the detector and the barycenter and the time dependent antenna pattern functions are put into a form that can be used with the stationary phase approximation waveforms by mapping between time and frequency using $t(f) = (d\Phi/df)/2\pi$. Details of this procedure can be found in Ref. [42]. The noise spectral density model includes instrument noise and an estimate of the foreground confusion noise from unresolved galactic binaries, matching those quoted in Ref. [82].

7.1.2. Likelihood Function

Under the assumption that the noise is Gaussian, the likelihood that the data d would arise from a signal with parameters $\vec{\theta}$ is given by

$$p(d|\vec{\theta}) = Ce^{-\chi^2(\vec{\theta})/2}\,, \tag{7.3}$$

where C is a constant that depends on the noise level. Here

$$\chi^2(\vec{\theta}) = (d - h(\vec{\theta})|d - h(\vec{\theta}))\,, \tag{7.4}$$

and the brackets denote the noise weighted inner product

$$(a|b) = 2\int \frac{\tilde{a}(f)\tilde{b}^*(f) + \tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}\,df\,. \tag{7.5}$$

For a theoretical study that assumes the noise is Gaussian and has a known spectrum, there is no need to add simulated noise to the data - the appropriate spread in the parameter values and overall topography of the likelihood surface follow from the functional form of the signal and the noise weighting in Eq. (7.5). Thus, we may write $d = h(\vec{\theta}')$ where $\vec{\theta}'$ are the true source parameters.

Many alternative theories of gravity predict the existence of polarization states beyond the usual "plus" and "cross" polarizations of GR that complicate the treatment of the instrument response, whose Fourier transform is

$$\begin{aligned} \tilde{h}_{inst} &= F_+h_+ + F_\times h_\times + F_S h_S \\ &\quad + F_L h_L + F_{V1}h_{V1} + F_{V2}h_{V2}\,, \end{aligned} \tag{7.6}$$

Here $h_{+\times}$ are the usual plus and cross-polarization states, h_S is a scalar (breathing) mode, h_L is a scalar longitudinal model and $h_{V1,V2}$ are two vectorial modes [146], while the F's are the detector antenna patterns which depend on the sky location (θ, ϕ) and polarization angle ψ of the signal. These are given in [141]

To simplify the analysis we assume the usual polarization content for a circular binary viewed at inclination angle ι and neglect the other contributions:

$$\begin{aligned}
\tilde{h}_{+} &= (1+\cos^2\iota)\Re(\tilde{h}) + 2\cos\iota\,\Im(\tilde{h})\,, \\
\tilde{h}_{\times} &= (1+\cos^2\iota)\Im(\tilde{h}) - 2\cos\iota\,\Re(\tilde{h})\,. \qquad (7.7)
\end{aligned}$$

In other words, we have assumed that the signal in the detector has the form $\tilde{s}(f) = F(\theta, \phi, \psi, \iota)\ \tilde{h}(f)$ with the function $F(\theta, \phi, \psi, \iota)$ given by the usual GR expression. If additional polarization states were present, this assumption would result in a reduction in detection efficiency and biases in the recovery of the extrinsic parameters $(\theta, \phi, \psi, \iota)$.

The justification for making this simplification is that we are primarily interested in how well the intrinsic parameters (α, a, β, b) can be constrained, and we expect these parameters to be only weakly correlated with the extrinsic parameters. The presence of additional polarization states will provide an additional handle on detecting departures to GR [76, 88, 131], and we plan to explore this possibility in the context of the ppE formalism in future work.

Defining $A_+ = |F^+\tilde{h}_+(f;\vec{\theta})|$ and $A_\times = |F^\times\tilde{h}_\times(f;\vec{\theta})|$, and similarly for $\vec{\theta}'$, the chi-squared goodness of fit of Eq. (7.4) can be re-expressed as

$$\begin{aligned}\chi^2(\vec{\theta}) &= 4\int \frac{df}{S_n(f)}\Big[A_+^2 + A_\times^2 + {A'}_+^2 + {A'}_\times^2 \\ &- 2(A_+{A'}_+ + A_\times {A'}_\times)\cos\Delta\Psi \\ &- 2(A_\times {A'}_+ - A_+{A'}_\times)\sin\Delta\Psi\Big]\,, \end{aligned} \tag{7.8}$$

where $\Delta\Psi = \Psi(\vec{\theta}) - \Psi(\vec{\theta}')$. As noted in Ref. [39], in the regime of interested where χ^2 is small, all the terms in the above integrand are slowly varying functions of frequency, so it is possible to compute the likelihood very cheaply using an adaptive integrator.

7.2. Approximate Bounds and Comparison with Pulsar Bounds

First, we seek to determine how well the four ppE parameters (α, a, β, b), from the simplest ppE parameterization, can be constrained. One approach to answering this question is to examine how a search using ppE templates would look when used to characterize a signal that is consistent with GR. That is, if the signal observed is described by GR to the given level of accuracy of our detectors, what values for the ppE parameters will be recovered from a search with ppE templates? Because we know that in GR the values of α and β should be 0 for all values of a and b, we wish to determine the typical spread in the recovered value of (α, β), centered at zero. The standard deviation in this spread then gives us a constraint on the magnitude of the deviation that is still consistent with observations, ie. deviations that are 'inside our observational error bars.'

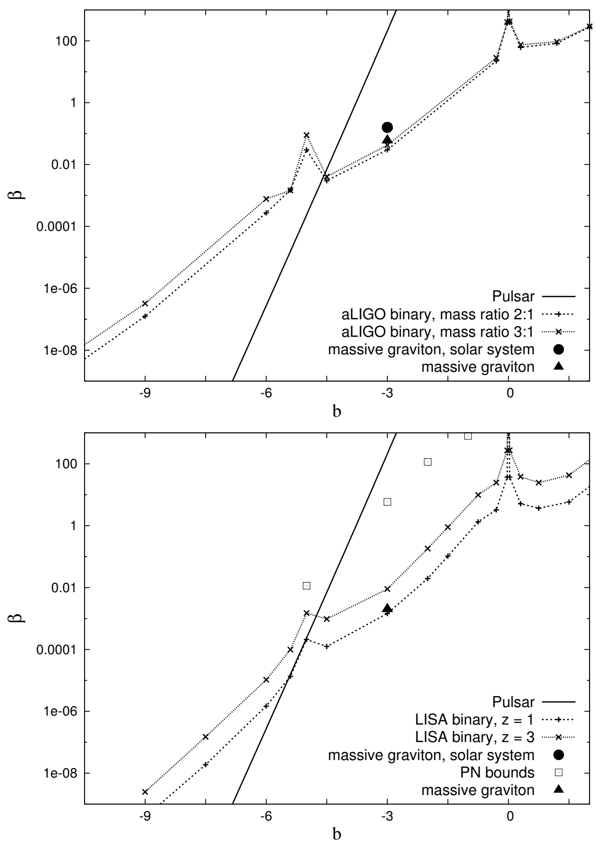

Figure 22: UPPER PANEL: Bounds on β for different values of b for a single SNR = 20 aLIGO/aVirgo detection. Plotted here is a (3σ) constraint, where σ is the standard deviation of the β parameter derived from the Markov chains. The two sources had different mass ratios, total masses, and sky locations, but were both at redshift $z = 0.1$ ($D_L = 462Mpc$). Also included is the bound on β derived from the golden pulsar (PSR J0737-3039) data, as well as bounds found from solar system experiments and other aLIGO analyses for massive graviton theory.
LOWER PANEL: Bounds on β for different values of b found using two LISA sources at redshift $z = 1$ and $z = 3$. The pulsar bound is shown for comparison, as well as bounds found from solar system experiments and other LISA analyses for massive graviton theory. These other bounds are scaled to a system with $z = 1$.

Approximate constraints will be defined as the (3σ)-bound on the posterior distribution of ppE parameters α or β, while keeping a or b fixed and marginalizing over all other system parameters. These bounds are 'approximate' because we do not have to re-run a search with pure GR templates and then compute the evidence, via integration of the posterior, to compute the Bayes factor (the latter is particularly computationally expensive). These approximate bounds are similar to constraints studied by looking at the (α, α) or (β, β) elements of the variance-covariance matrix. Our approximate constraints, however, are 3σ ones, while all other constraints studied to date with the variance-covariance matrix are 1σ.

Figures 22 and 23 show these approximate constraints on the ppE amplitude parameters as a function of the exponents a and b for a variety of aLIGO/aVirgo and LISA detections. To generate these plots, we injected GR signals and then searched on them with ppE templates. For each search, either a or b was held fixed at a specific value, while the other three ppE parameters (and all other system parameters) were allowed to vary. We then calculated the standard deviation of the posterior distribution of the relevant amplitude parameter α or β, and used three times this value as the approximate bound shown on the plots.

(It may seem that the natural course of action should be to marginalize over a and b as well, instead of keeping them fixed, and calculate constraints on α and β this way. Looking at Figures 22 - 23, however, shows why this analysis would not be particularly helpful. The uncertainty in α and β is so much higher at the positive ends of the prior ranges on a and b than at the negative ends, the Markov chains would spend almost all of their iterations exploring this area of parameter space if a and b were allowed to change. Thus, to get any knowledge about the uncertainties in α and β for negative values of a and b, we need to fix a and b.)

The aLIGO systems were chosen to have network SNR = 20 ($D_L = 462$ Mpc), but different masses and sky locations. One system had masses $m_1 = 6M_\odot$, $m_2 = 18M_\odot$ ($\eta = 0.1875$), while the other had $m_1 = 6M_\odot$, $m_2 = 12M_\odot$ ($\eta = 0.2222$). The LISA sources were at different redshifts and had different masses and SNRs. The system at redshift $z = 1$ had $m_1 = 1 \times 10^6 M_\odot$, $m_2 = 3 \times 10^6 M_\odot$ ($\eta = 0.1875$) and SNR = 879, while the system at redshift $z = 3$ had $m_1 = 2 \times 10^6 M_\odot$, $m_2 = 3 \times 10^6 M_\odot$ ($\eta = 0.24$) and SNR = 280.

Figures 22-23 are 'exclusion' plots, showing the region (above the curves) which could be excluded with a 99.73% confidence. These figures also plot the bound on the ppE parameters that have already been achieved through analysis of the 'golden pulsar' system, PSR J0737-3039 [159]. Observe that for the amplitude parameter α, the pulsar bounds beat the aLIGO bounds through almost the entire range of a; LISA can improve upon the pulsar bounds for $a > 0$. For the phase parameter β, however, both aLIGO and LISA do better than the pulsar analysis through a significant portion of the range. As expected, gravitational wave observations tend to do better in the strong field regime ($b > -5$ and $a > 0$), while the reverse is true for binary pulsar observations.

Vertical lines in Figs. 22-23 can be mapped to bounds on specific alternative theories, which we can then compare to current Solar System constraints. For example, consider the following cases:

- Brans-Dicke $[(\alpha, b, \beta_{\rm BD}) = (0, -7, \beta_{\rm BD})]$: The tracking of the Cassini spacecraft [22] has constrained $\omega_{\rm BD} > \bar{\omega}_{\rm BD} \equiv 4 \times 10^3$, which then forces $\beta_{\rm BD} < (5/3584) 4^{-2/5} (s_1 - s_2)^2/\bar{\omega}_{\rm BD}$, where $s_{1,2}$ are the sensitivities of the binary components (for BHs $s_{\rm BH} = 1/2$, and for NSs $s_{\rm NS} \approx 0.2 - 0.3$).
- Massive Graviton $[(\alpha, b, \beta_{\rm MG}) = (0, -3, \beta_{\rm MG})]$: Observations of Solar system dynamics [127] have constrained $\lambda_{\rm MG} > \bar{\lambda}_{\rm MG} \equiv 2.8 \times 10^{12}$ km, which then forces $\beta_{\rm MG} < \pi^2 (D/\bar{\lambda}_{\rm MG}) \mathcal{M} (1+z)^{-1}$ km^{-2}, where D is a distance measure to the source [143].

The Solar System constraint on β_{MG} is shown in Fig. 22 with a black circle[1]. Observe that the constraints we could place with aLIGO and particularly LISA can be orders of magnitude stronger than Solar System constraints (below the black circle). This is more easily seen by mapping our projected constraints on $\beta_{\rm MG}$ to constraints on $\lambda_{\rm MG}$; with the aLIGO source, we find $\lambda_{\rm MG} \lesssim 8.8 \times 10^{12}$ km, while for the LISA source, we find $\lambda_{\rm MG} \lesssim 3.763 \times 10^{16}$ km. This is consistent with results from previous Fisher [143, 147, 19, 126, 12, 81, 152, 142, 119] and Bayesian studies [51]. Plotted for comparison are the bounds from Pozzo et al. [51] on the upper panel of 22 and from Stavridis and Will [126] on the lower panel of 22 both labeled as "massive graviton."

[1]We don't show similar constraints for Brans-Dicke theory, as here we consider binary BH inspirals, for which the Brans-Dicke correction would vanish due to the no-hair theorem.

We find that our bound on β for $b = -3$ is quite comparable to those found in these previous studies. Finally, shown on the lower panel of 22 are the bounds found in the study by Arun et. al. [12], which allowed the PN coefficients themselves to vary as parameters. Their bounds on β are somewhat weaker than those we found in our analysis, but this is an expected effect of the covariance between the PN coefficients.

For all comparisons with previous studies, we took into account differences in SNR between the systems we analyzed and those we were comparing to. We also chose systems with the same or very similar total masses and mass ratios as those explored in previous papers. For the LISA systems, we compare the results from previous papers to our results for redshift $z = 1$.

These plots show several other features that deserve further discussion. First, observe that all results show very little dependence on the choice of system parameters. This is quantitatively true for the aLIGO sources, shown in the upper panels of Figs. 22 and 23, as these signals have the same SNR. The LISA sources, shown in the lower panels of Figures 22 and 23, show a factor of ~ 9 offset, since these curves correspond to signals with different SNRs. The SNR difference is a factor of ~ 3, which is a bit surprising as one would expect the spread on a parameter to scale with the SNR, and not the square of the SNR. However, we are working here in a region where the quadratic approximation to the Fisher matrix is singular, so the usual scaling does not hold. The more rigorous bounds derived in the next section do

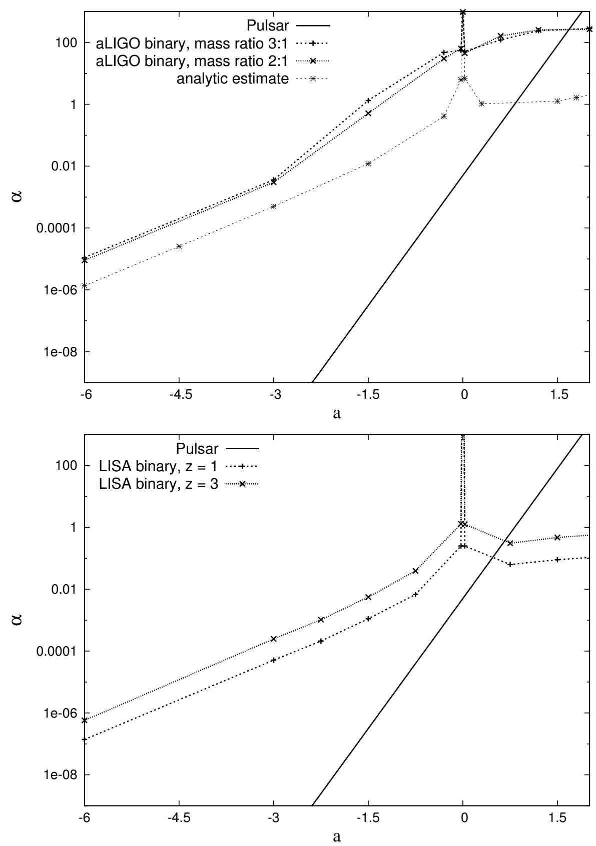

Figure 23: UPPER PANEL:Bounds on α for different values of a, found using two different aLIGO sources. The pulsar bound is shown for comparison. The sources injected had the same parameters as those from the upper panel in Figure 22 .
LOWER PANEL: Bounds on α for different values of a, found using two LISA sources at redshift $z = 1$ and $z = 3$. The pulsar bound is shown for comparison. The sources injected had the same parameters as those from the lower panel in Figure 22 .

follow a linear scaling with SNR, which is reasonable since they use ppE injections and have non-singular Fisher matrix elements for the ppE parameters.

Another interesting feature in these plots are the spikes at certain values of a and b. These spikes say that for those values of a and b, gravitational wave observations can say little about the magnitude of GR deviations. The reason for such spikes is that for those values of a and b, α and β become completely or partially degenerate with other parameters. For instance, when $a = 0$, α is fully degenerate with the luminosity distance, and when $b = 0$, β is fully degenerate with the initial orbital phase ϕ_c.

One can also develop 'approximate' bounds that use ppE instead of GR signal injections. For instance, one could start with injections with a range of values for α and β, and then look to see when the posterior distributions for these parameters no longer show significant support at the GR values of $\alpha = \beta = 0$. These two types of approximate bounds are illustrated in Fig. 24. Given an observation of a non-zero α, a approximate bound calculation as described in this section (solid curve) would indicate a value $|\alpha| < 1.5$ is still consistent with GR. A similar study with ppE injections, however, which produced the dashed-curve posterior distribution for α would indicate a preference for the ppE model over the GR model with a detection of $\alpha > 0.75$. Thus the technique used in this section, which is a variance-covariance study, answers an inherently different question from a model selection study. In the next section, we explore model selection in detail.

7.3. Noise Modeling and Signal Strength

Our results thus far have been from studies conducted on signals that do not have a noise realization explicitly added to the signal injection, although all analyses incorporate the noise spectrum of the detectors in the likelihood calculation. We chose not to include an explicit noise realization in order to expedite the calculation of the

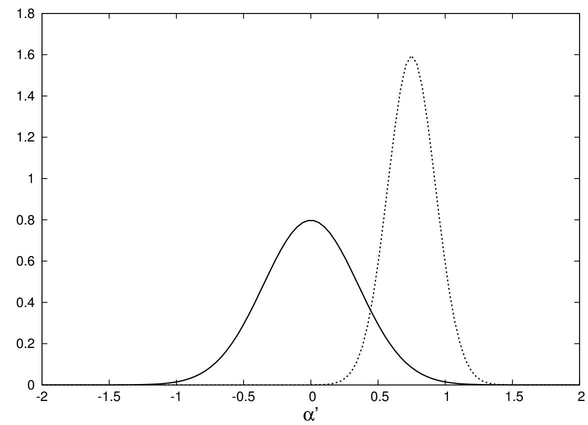

Figure 24: An illustration of the two approaches for calculating approximate bounds on the ppE amplitude parameters. The solid curve illustrates the bound that can be derived by looking at the spread in the amplitude α when applying the ppE search to GR signals. In this example, values of $|\alpha| > 1.5$ would be taken as indicating a departure from GR. The dashed curve shows the bound that can be derived by starting with ppE signals and determining how large the ppE amplitude needs to be for the posterior distribution to have little weight at the GR value of $\alpha = 0$. In this example, theories with $\alpha > 0.75$ would be considered distinguishable from GR.

likelihood [39], which then allows us to produce long Markov chains that fully explore the high dimensional parameter spaces. Unfortunately, our use of this technique has led some to question the reliability of our results [89, 134]. Here we show that those concerns are unfounded.

The inclusion of noise in our signals has little effect on the conclusions we drew in the previous section, as can be seen in Fig. 25. In this figure, we plot the (3σ)-bounds that we could place on the ppE phase parameters, if one has detected a NS-NS inspiral with SNR 15 that has no GR deviation. To calculate these bounds, we inject a GR signal and try to recover it using a single parameter ppE template, ie. Eq. (7.12) with a single β. For any given value of b, we integrate out over all other parameters and take the standard deviation of the β PDF as a 1σ bound. In other

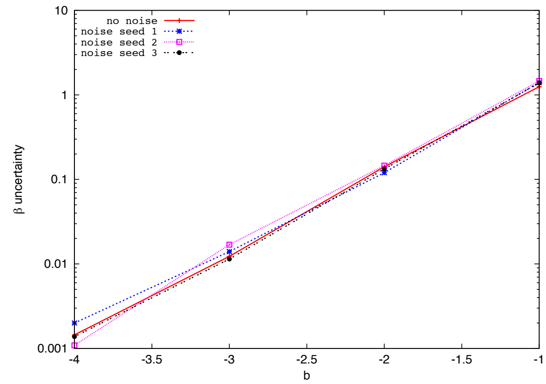

Figure 25: (3σ)-bounds on β that can be inferred for different values of b, calculated from the PDF's of β generated by recovering a GR signal with a ppE template. This plot shows the bounds for both a signal with no noise, and three that include Gaussian noise, generated from three different random seeds. The results are essentially identical. The signal parameters for this injection are in Table 8.

words, the curves show the upper limit of the magnitude a ppE parameter could be found to have, and still have the signal be consistent with GR. This plot shows that the bounds placed on the ppE parameters from a signal that includes an explicit noise realization are consistent with those found when no noise is added to the signal. That is, including an explicit noise realization does not affect the conclusions derived from a approximate-bound calculation with noise accounted for only through the detectors' noise spectrum in the likelihood.

To understand this result, it is useful to look at Figure 26, which shows the recovered PDF's for the β parameter from three different runs, each including noise generated with a different random seed. Since the injected signal was a GR NS-NS inspiral waveform with SNR 15, we would expect the β PDF's to peak at zero. It is clear from this figure that, although the peak of the PDF is shifted by the inclusion of noise, the uncertainty in the recovery of this parameter, i.e. the spread of the distribution, is not affected. This concept has been explored before, in [102] and [134]. In [102], the authors argue that when discussing our ability to measure system parameters in general, and not for a particular case, what we really want to do is examine the *noise-averaged* uncertainties in these parameters. That is, we are interested in how well we can measure parameters when averaged over many specific realizations of the noise. The authors show that the noise-averaged uncertainties are the same as the uncertainties calculated with zero noise injected into the signal. In [134] it is argued that the specific noise realization will affect our parameter estimation, and while this is technically true, we have shown in this section that the overall effect is minimal. In any case, for the type of analysis that we want to do in the rest of this chapter, the reasoning of [102] applies, and so we do not inject an explicit noise realization for any of our analyses in the other sections. It has also been claimed in [89] that simulated data that only includes a signal injection, ie. that does not include a noise realizations, will necessarily lead to posterior distributions for the system parameters that are Gaussian. This is patently false, as can easily be demonstrated by analytically calculating the posterior distribution for a signal of the form $(d_0/d)\cos(2\pi f t)$, which leads to a highly non-Gaussian distribution in the distance d.

Obviously, signals with high SNR will be better for testing GR, as they are better for any type of GW data analysis. When discussing how well GR can be tested using GW detections, the highest-SNR events are the ones that will lead to the strongest

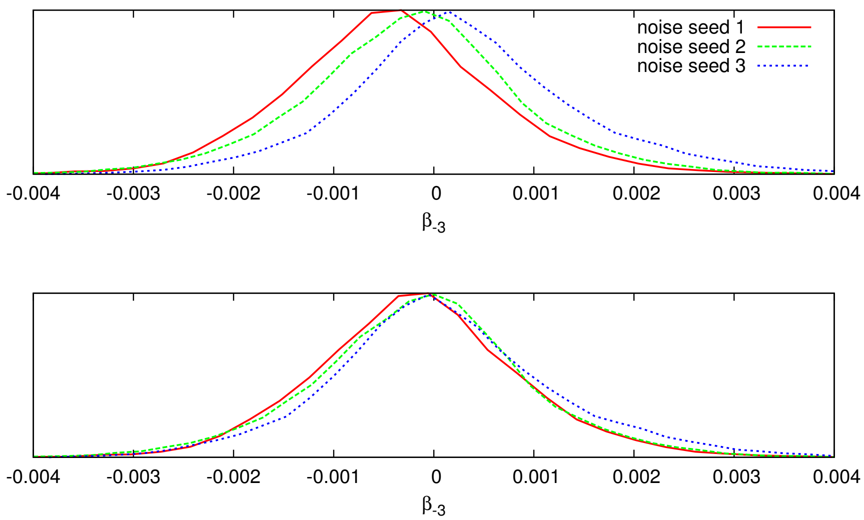

Figure 26: The top panel shows posterior distributions of β recovered from three ppE injections, including noise in the injection. Each of the three signals was generated using a different random seed for the noise, but the same system parameters. The lower panel shows the same distributions, now with the best-fit value of β subtracted. This illustrates that, although noise affects the peak of the posterior distribution for a given parameter, it does not affect the uncertainty in that parameter. Thus the *approximate bounds* of [38] are unaffected by the inclusion of noise.

constraints. In previous sections, we analyzed signals with SNR ~ 20, which would be considered a high SNR detection by the LIGO detectors. It is irrelevant, however, that most signals will probably have SNRs in the low 10s. There will always be one signal with highest SNR, and this is likely to be above 15. It is therefore still useful to study GR tests assuming detections with SNRs ~ 20, as it is not a hopeless proposition that we will have this type of event in our GW catalog. Throughout the rest of this section, however, we have taken a more pessimistic tack, and restricted ourselves to analyzing signals with SNR $\sim 10-12$. The results follow the theoretical linear scaling with SNR [39] down to values of the SNR that are close to the detection threshold. This scaling is shown in Figure 27.

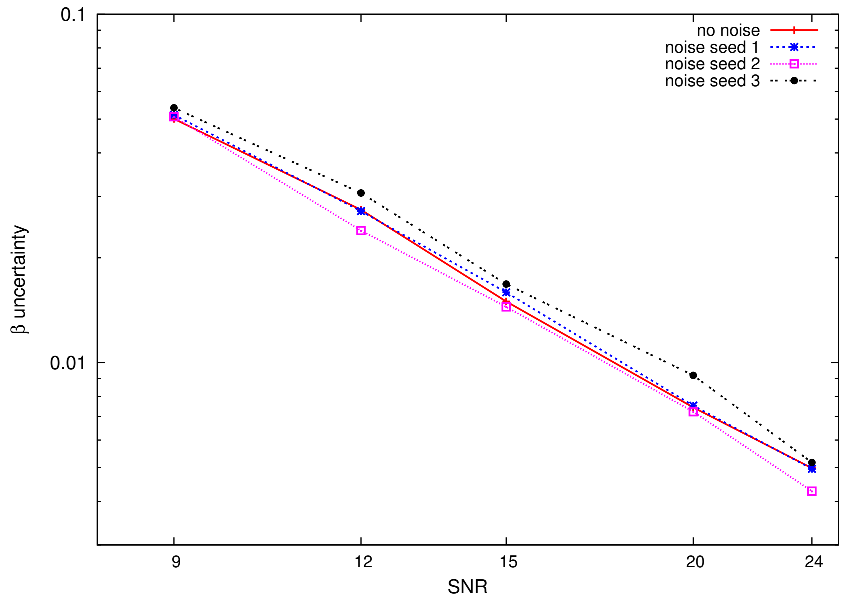

Figure 27: (3σ)-bounds on β for $b = -1.0$, calculated from the PDF's of β generated by recovering a GR signal with a ppE template. This plot shows the linear relationship between the bounds on β and the SNR of the signal. There are four lines shown - one for a signal that had no noise injected, and three for signals that had noise injected, each with a different random seed. The results are essentially identical. The signal parameters for this injection are in Table 8.

7.4. Full Bayesian Bounds and Model Selection

In order to see how accurate the approximate bounds found in Section SECTION are, we performed a full Bayesian model selection analysis on several different signals. We injected a signal with a given set of ppE parameters and ran a search using both GR and ppE templates. We then calculated the Bayesian evidence for each model and from this the Bayes factor. To compare these results to the approximate bounds, we ran the analysis on several different ppE signals, each with the same injected value of a or b, but with progressively larger values of α or β. This then allows us to determine the values of ppE amplitudes α or β where the evidence for the ppE hypothesis exceeds that of the GR hypothesis by some large factor, which we took to be Bayes factors in excess of 100 (in the Jeffery's classification [80], Bayes factors in excess of 100 represent *decisive* evidence in favor of that model).

We do not expect the approximate bounds to agree precisely with the more rigorous model selection bounds as they are based on quite different reasoning. The approximate bounds simulate what we would find if GR was consistent with observations, and establishes the spread in the ppE amplitude parameters that would remain consistent. If we were to analyze some data and find ppE amplitude parameters outside of this range, it would give us motivation to search more rigorously for departures from GR. In the more expensive model selection bounds, we start with non-GR signals and seek to determine how large the departures from GR have to be for the ppE hypothesis to be preferred. In the first case the distribution of α and β is known to be centered around zero, but in the second case it is explicitly not, and so the two analyses should not be expected to agree precisely.

One can derive a more detailed connection between the alternative form of the approximate bounds derived using ppE injections (discussed at the end of the previous section) and the more rigorous Bayesian evidence calculations using the Savage-Dickey density ratio [135]. The latter states that for nested hypotheses with separable priors, the Bayes factor is equal to the ratio of the posterior and prior densities evaluated at the parameter values that correspond to the lower dimensional model. If the posterior distribution was a Gaussian with width σ centered at $\alpha = n\sigma$, and we were using a uniform prior with width $N\sigma$, then the Bayes factor would equal BF $= Ne^{-n^2/2}/\sqrt{2\pi}$, where this Bayes factor shows the odds of the lower dimensional model being correct. For example, with $N = 100$ and $n = 4$ we get a Bayes factor of BF $= 0.013$, showing strong support for the higher dimensional model. While the approximate bounds that can be derived using ppE signal injections will be stronger than the approximate bounds that can be derived from GR signal injections, the computational cost is higher as multiple simulations have to be run to find the transition point, and this approach is only moderately cheaper than performing the full Bayesian model selection.

Examples of the full model selection procedure are shown in Fig. 28 for aLIGO/aVirgo detections with SNR = 20. Each panel shows Bayes factors for two types of ppE search, one with a or b held fixed at the injected value, and one in which all four ppE values were allowed to vary. The Bayes factor, defined in Eq. (7), is here the odds ratio between the ppE model and the GR model. A larger Bayes factor indicates a stronger preference for the ppE model. The search in which a or b was fixed provides the closest comparison with the approximate bounds of the previous section. The bound on β derived by setting a Bayes factor threshold of 100 are roughly 3 times larger than the approximate bounds when b is held fixed and roughly 2 times larger

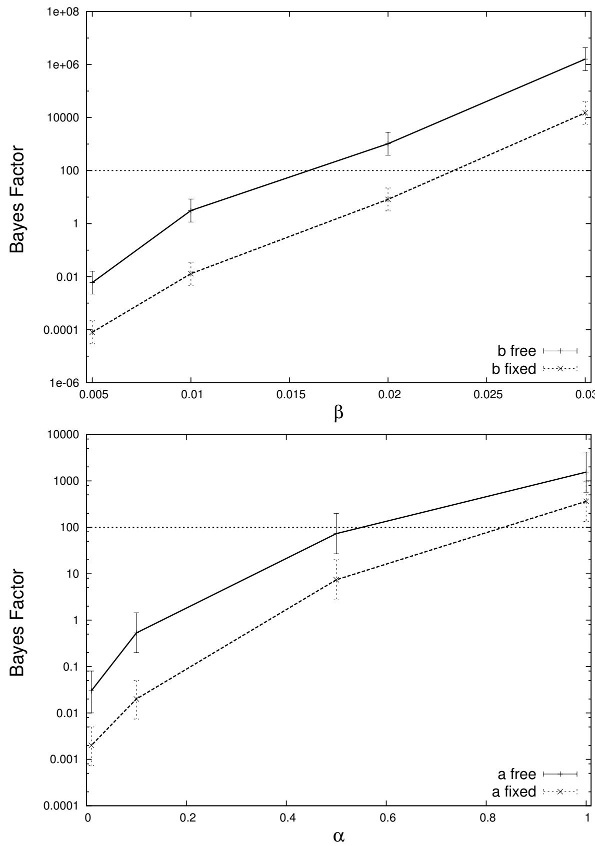

Figure 28: UPPER PANEL: Bayes factors for a SNR = 20 aLIGO ppE injection with parameters $(a, \alpha, b, \beta) = (0, 0, -3.75, \beta)$. The Bayes factors are the 'betting odds' that ppE (and not GR) is the model that accurately describes the data. As the deviation from GR gets larger, ppE becomes the preferred model.
LOWER PANEL: Bayes factors for a SNR = 20 aLIGO ppE injection with parameters $(a, \alpha, b, \beta) = (-1.5, \alpha, 0, 0)$.

when b is free to vary. The bounds on α match the approximate bounds when a is held fixed, and is slightly smaller when a is allowed to vary.

We were surprised to find that the bounds are tighter for the higher dimensional models, with (a, b) free, than for the lower dimensional models, with (a, b) fixed. To explore this more thoroughly, we performed a study where the prior on b was increased from a very small range to the full prior range. Because holding a parameter fixed is equivalent to using a delta-function prior, we expect the evidence to interpolate between the values found when b was fixed and when b was free to explore the full prior. Figure 29 confirms this expectation, and also provides an explanation for the growth in the evidence.

To understand this plot, it is helpful to look at the Laplace approximation to the evidence [14], which assumes that the region surrounding the maximum of the posterior distribution is well approximated by a multivariate Gaussian. With this assumption, the evidence is given by

$$p(d|\mathcal{H}) \approx p(d|\vec{\theta}, \mathcal{H})|_{\mathrm{MAP}} \left(\frac{\Delta V_{\mathcal{H}}}{V_{\mathcal{H}}} \right) . \tag{7.9}$$

The first term is the likelihood evaluated at the maximum of the posterior, and the second term is the ratio of the posterior volume ΔV to the prior volume V. The posterior volume can be estimated from the volume of the error ellipsoid containing 95% of the posterior probability. The ratio $\Delta V/V$ is termed the "Occam factor", and the quantity $I = \log_2(V/\Delta V)$ provides a measure of how much information has been gained about the parameters from the data. Now consider a situation where we have nested hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$, with the second hypothesis involving an additional parameter y. If the likelihood is insensitive to y then the first factor in the evidence stays the same, and since y is unconstrained, $\Delta V_y = V_y$ and the Occam factor is also unchanged. Thus, both models have the same evidence, even though one has more parameters than the other. Conversely, if the additional parameter is tightly constrained by the data, $\frac{\Delta V_y}{V_y}$ can be a very small number. In this case, the evidence for $\mathcal{H}_1$ is much reduced by the Occam factor, and the factor is referred to as an "Occam penalty."

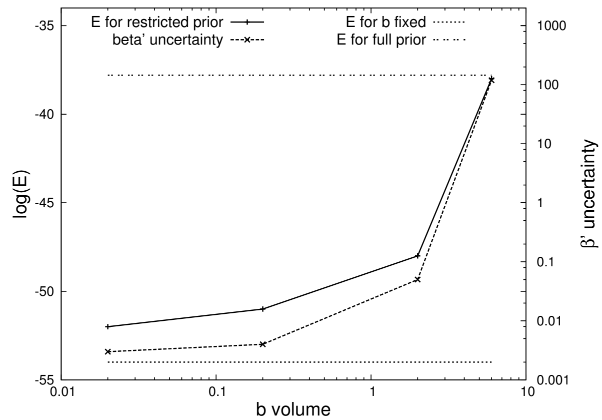

Figure 29: Here we plot the log of the evidence (E) for the ppE model characterizing a ppE injection as the prior volume on b is increased. The evidence for the ppE model increases with the prior volume on b. The growth in the evidence can be attributed to the growth in the variance of β, which lessens the severity of the 'Occam penalty' for more model parameters.

The growth in evidence for the ppE model as the prior range for b gets larger is an effect of this Occam factor, which is a ratio of the uncertainty in the recovered value of an extra parameter to the prior volume for that parameter. As the prior range on b expands, this leads to a greater variance in the recovered values for β'. Because the prior volume of β' remains unchanged, the large growth in its variance as the prior range of b is expanded leads to a large growth in the Occam factor - and thus a shrinking of the Occam penalty. As the Occam factor gets larger, so does the evidence for the ppE model. The evidence for the GR model, of course, does not depend on the priors we use for the ppE parameters, and so as the evidence for ppE grows, the Bayes factor indicates a stronger preference for ppE.

Figure 30 shows Bayes factors between the GR and ppE hypotheses for a $z = 1$ LISA source. In the upper panel, the injections were chosen with $a = 0, b = -3$ and variable β, while in the lower panel the injections were chosen with $a = 015, b = 0$ and variable α. Because LISA sources have much higher SNR, the ppE parameters are more tightly constrained, and the difference between the Bayes factors when a or b are fixed versus freely varying is less pronounced. The more rigorous bounds on α and β are both a factor of ~ 2 times weaker than those predicted by the approximate bounds, which is in line with what we found for the phase correction β in the aLIGO example. In summary, the approximate bounds provide a fair approximation to the bounds that can be derived from Bayesian model selection, and can generally be trusted to within an order of magnitude.

7.5. Parameter Biases

If we assume that Nature is described by GR, but in truth another theory is correct, this will result in the recovery of the wrong parameters for the systems we are studying. For instance, when looking at a signal that has non-zero ppE phase parameters, a search using GR templates will return the incorrect mass parameters, as illustrated in Fig. 31. Observe that as the magnitude of β is increased (thus

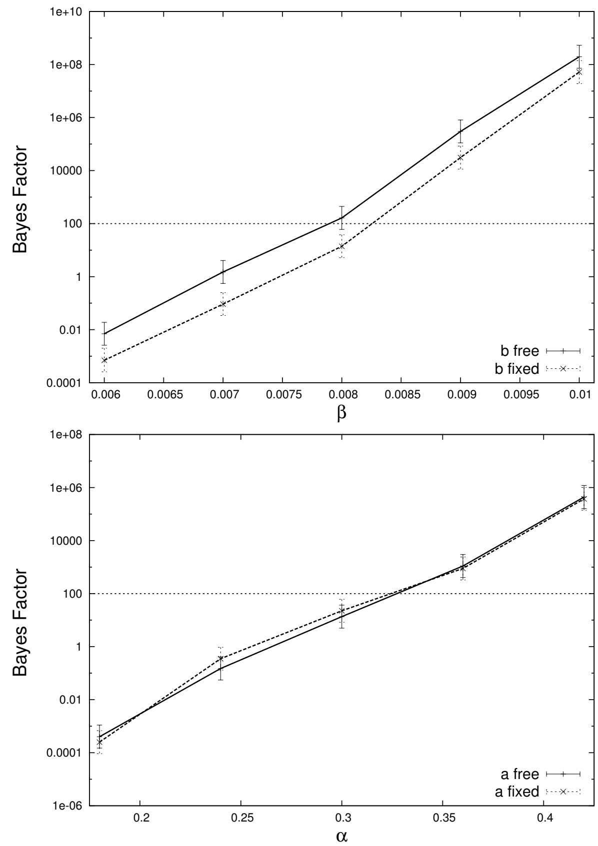

Figure 30: UPPER PANEL: Bayes factors for a $z = 1$ LISA ppE injection with parameters $(a, \alpha, b, \beta) = (0, 0, -3.0, \beta)$.
LOWER PANEL: Bayes factors for a $z = 1$ LISA ppE injection with parameters $(a, \alpha, b, \beta) = (1.5, \alpha, 0, 0)$.

increasing the Bayes factor), the error in the total mass parameter extraction grows well beyond statistical errors.

Perhaps the most interesting point to be made with this study is that the GR templates return values of the total mass that are completely outside the error range of the (correct) parameters returned by the ppE search, *even for ppE signals that are not clearly discernible from GR*. We refer to this parameter biasing as 'stealth bias', as it is not an effect that would be easy to detect, even if one were looking for it.

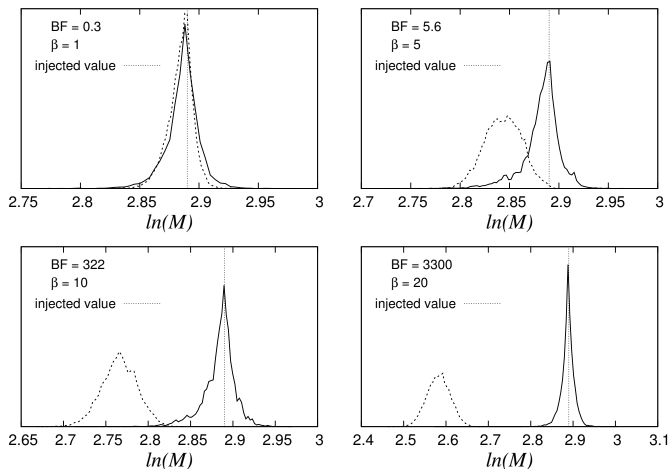

Figure 31: Histograms showing the recovered log total mass for GR (dashed) and ppE (solid) searches on ppE signals. As the source gets further from GR, the value for total mass recovered by the GR search moves away from the actual value. All signals had injected $b = -0.75$.

As an example, consider stealth bias for non-zero ppE α parameters, as illustrated in Fig. 32. As one would expect, when a GR template is used to search on a ppE signal that has non-zero ppE amplitude corrections, the parameter that is most affected is the luminosity distance. We again see the bias of the recovered parameter becoming more apparent as the signal differs more from GR[2]. For example, the recovered posterior distribution from the search using GR templates has zero weight at the correct value of luminosity distance when the Bayes factor is ~ 50. Even when the Bayes factor is of order unity, the peaks of the posterior distributions of the luminosity distance differ by approximately 10 Gpc.

[2]Here, the uncertainty in the recovered luminosity distance changes considerably between the different systems, because we held the injected luminosity distance constant instead of the injected SNR.

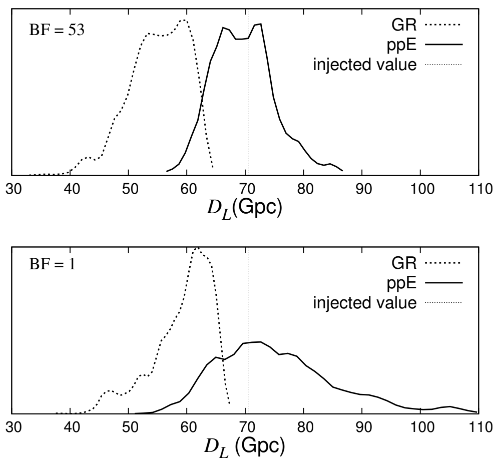

Figure 32: Histograms showing the recovered values for luminosity distance from GR and ppE searches on a LISA binary at redshift $z = 7$. Both signals have $a = 1.5$, and were injected with a luminosity distance of 70.5 Gpc. The top plot has $\alpha = 3.0$ and the bottom has $\alpha = 2.5$. As the Bayes factor favors the ppE model more strongly, the bias in the recovered luminosity distance from the GR search becomes more pronounced.

7.6. More Complicated Waveforms

The simplest ppE waveform family presented in the introduction and in Chapter 5 is not sufficiently complex to represent a realistic alternative gravity theory. This is because modified gravity theories will differ from GR by an infinite series of terms in both the amplitude and the phase. We expect that an alternative theory of gravity will give rise to waveforms where the amplitude and phase depend on one or more fundamental coupling constants multiplied by functions of the system parameters. Thus, if one wishes to use a ppE-type template to inject non-GR signals, one must consider more complex ppE models, such as Eq. (46) in [163], namely Eq. (**??**) with the replacements [163]

$$\alpha u^a \to \sum_{i=0}^{N} \alpha_i u^{a_i}\,, \quad \beta u^b \to \sum_{i=0}^{N} \beta_i u^{b_i}\,, \tag{7.10}$$

where the α, β's depend on a universal coupling constant κ, and functions of the system parameters $\vec{\lambda}$:

$$\begin{aligned} \alpha_i(\kappa, \vec{\lambda}) &= \kappa \sum \phi_i(\vec{\lambda}) \\ \beta_i(\kappa, \vec{\lambda}) &= \kappa \sum \theta_i(\vec{\lambda})\,. \end{aligned} \tag{7.11}$$

The functions $\phi_i(\vec{\lambda}), \theta_i(\vec{\lambda})$ can be computed for specific theories, but their general form is unknown. So while κ takes a single value for a particular theory, the (α_i, β_i) constants will vary from detection to detection depending on the masses, spins and other parameters that describe the system. In some theories there will be more than one additional coupling constant κ, but here we will assume that one sector of the modified theory dominates and consider only a single series of correction terms. With a large number of high SNR detections, it may be possible to infer the functional form of $(\phi_i(\vec{\lambda}), \theta_i(\vec{\lambda}))$. However, since our immediate concern is in deciding if the data is consistent with the prediction of GR, we will argue that it is best to use a much simpler waveform for the initial tests.

The ppE exponents (a_i, b_i) are real numbers that give the effective PN order at which the non-GR modification enters the signal, while the ppE amplitude parameters (α_i, β_i) are real numbers that indicate the strength of the modification, in turn controlled by the overall coupling strength κ. In principle, we could extend the sum in Eq. (7.10) to infinity, but in practice, realistic detectors are sensitive only to a finite number of terms in the phase and amplitude. The injected signals then consist of a GR waveform with its amplitude and phase modified by a series of ppE corrections.

Several simplifications can be made to the general waveform presented above. First, for quasi-circular inspiral signals, Chatziioannou, *et al* [34] have argued that analyticity demands that the exponents (a_i, b_i) take on integer values with possible logarithmic corrections (just as the PN expansion in GR comes in integer powers of u and products of integer powers of u with $\log u$, where recall here that u is related to the orbital velocity). Second, ground-based advanced detectors will be of limited sensitivity, rarely being sensitive to more than the first three terms in the PN expansion, and usually being much more sensitive to the phase evolution than they are to the amplitude evolution. Thus, we choose to simplify the analyses by truncating the sum at three terms and setting $\alpha_i = 0$. The injections are then given by Eq. (**??**) but with the replacement

$$\beta u^b \to \sum_{i=0}^{2} \beta_i u^{b+i} = \beta_b u^b + \beta_{b+1} u^{b+1} + \beta_{b+2} u^{b+2} \,, \tag{7.12}$$

and $\alpha = 0$, where in the last equality the Einstein summation convention is not assumed. Written in this way, β_b is always proportional to u^b for any b. Previous investigations have been restricted to signals with only one ppE correction injected, which reduces to Eq. (7.12) when one retains only the first term in the sum. As argued above, this is far from realistic for a modified gravity injection and we will show that the higher-order terms can have a significant effect on the analysis.

Ultimately the claim that a detection is in agreement (or conflict) with GR comes down to model selection. Does GR describe the data best or does another model do a better job? In Bayesian statistics [40, 38, 51], model selection is performed via the calculation of the Bayes factor, which is simply the "betting odds" of one model against another. For instance, if the Bayes factor between GR and a non-GR model is 100, and you originally gave both possibilities equal odds, then there is a 100:1 odds ratio that GR better describes the data than the other model. In this case, you would be well-advised to put your money on GR. There is no prescription for deciding what Bayes factor is required before we should consider one model "right" and another "wrong". However, in the case of a well-tested theory like GR being brought into question by, for instance, a GW signal, it is likely that the scientific community would require a detection that gives us a fairly high Bayes factor in favor of the non-GR model to overcome the prior belief in GR being the correct theory. In order to determine whether more ppE terms in an injection affect the detectability of a deviation from GR, we need to see how these different types of injections affect the Bayes factor.

In the following sections, Bayes factors are calculated using the Savage-Dicke density ratio [54, 40] and/or Reversible Jump Markov Chain Monte Carlo (RJMCMC) [114, 40]. Both of these methods are discussed in Chapter CHAPTER.

All tests in this section use GWs emitted by a NS-NS binary with $\approx 1.4M_{\odot}$ component masses in the inspiral phase with SNR ~ 12. We model all waveforms with a quadrupolar, adiabatically quasi-circular waveform, with a 3.5PN-accurate phasing, but neglecting PN amplitude correction and spin effects, and truncating all evolution at the Schwarzschild test-particle innermost stable circular orbit. The waveforms are then described by nine source-parameters: the chirp and the reduced mass; the time and phase of coalescence; two sky-position angles; the inclination angle and the GW polarization angle; and the luminosity distance (see [38] for a similar waveform prescription). In addition to these we have the ppE phase parameters of Eq. (7.12). We consider a three detector network of second-generation detectors, such as aLIGO at Hanford, aLIGO at Livingston, and aVirgo, with identical broadband-configuration spectral densities, assuming the noise to be Gaussian and stationary. Table 8 shows the system parameters for all systems studied in this section (masses are listed in solar masses, and luminosity distances are in megaParsecs).

In what follows, we examine two factors that influence the outcome - the signs of the different phase corrections, and their relative magnitude. We begin by exploring the effects of injecting phase corrections with the same or differing signs. In particular, we study the effect that this relative sign has on the detectability of non-GR behavior. We then explore the difference between non-GR phase corrections that either shrink in magnitude at higher PN order, stay at approximately the same magnitude, or grow in magnitude at higher PN order.

Signal	$m_1(M_\odot)$	$m_2(M_\odot)$	$\log(D_L)$(Mpc)	t_c	β_{-3}	β_{-2}	β_{-1}
One ppE Term	1.62	1.73	3.96	5.58	0.01	0.0	0
Alternating Sign	1.62	1.73	3.96	5.58	0.01	-0.04	0
Same Sign	1.62	1.73	3.96	5.58	0.01	0.04	0
Sub-Critical	1.62	1.73	3.96	5.58	0.01	0.005	0
Critical	1.62	1.73	3.96	5.58	0.01	0.08	0
Super-Critical	1.62	1.73	3.96	5.58	0.01	0.25	0
GR Source	1.45	1.43	3.41	8.76	0	0	0

Table 8: Source parameters for sources used in Fig. 33 (top), Fig. 34 (middle) and Figs. 25, 26, and 35 (bottom). All sources had $(\alpha, \cos\delta, \phi_L, \cos\theta_L) = (1.0, 0.77, 1.76, -0.43)$.

We begin by examining how the relative sign of the phase corrections affects the detectability of departures from GR. To do this, we consider three non-GR injections:

- **Case i.** A ppE waveform with a single ppE phase term ($b = -3$), with magnitude controlled by β_{-3}.
- **Case ii.** A ppE waveform with two ppE phase terms ($b = -3$ and $b = -2$), with β_{-3} and β_{-2} of the same sign.
- **Case iii.** A ppE waveform with two ppE phase terms ($b = -3$ and $b = -2$), with β_{-3} and β_{-2} of different sign.

We choose these values of b because β_b is already well-constrained by binary pulsar observations for $b < -5$, as demonstrated in [159, 38]. The $b = -3$ terms correspond to non-GR corrections at the first post-Newtonian (1PN) level, and the $b = -2$ terms correspond to a 1.5PN correction. Case (i) is the type of injection that has been explored in previous work. Cases (ii) and (iii) include higher-order phase corrections, but differ in their relative sign.

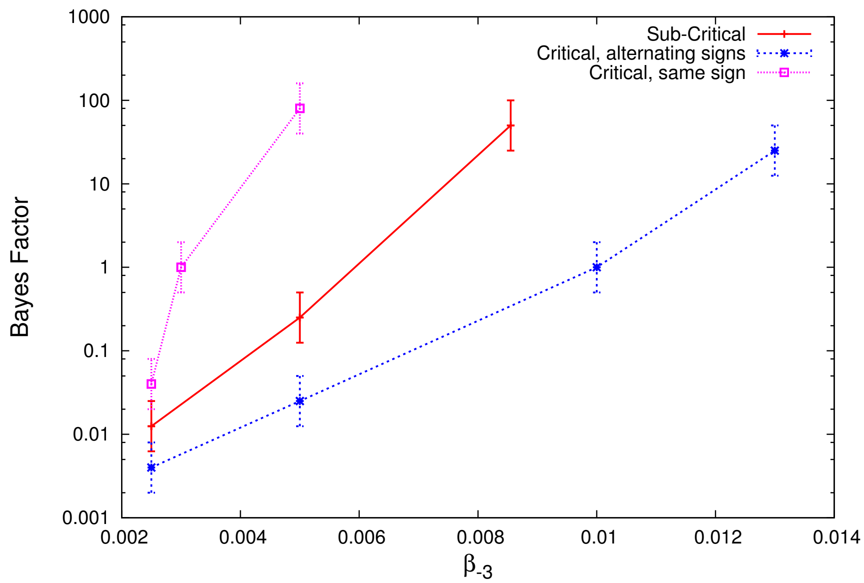

Figure 33: Bayes factors between a GR model and a one-parameter ppE model for three different ppE signal injections. The dotted (magenta) line corresponds to an injection with the two positive ppE terms $\beta_{-3} > 0$ and $\beta_{-2} > 0$ (case ii), the solid (red) line corresponds to the single, positive ppE term $\beta_{-3} > 0$ (case i), and the dashed (blue) line corresponds to the two ppE terms of alternating sign $\beta_{-3} > 0$ and $\beta_{-2} < 0$ (case iii). System parameters for the systems studied here are listed in Table 8. As expected, the signal with ppE terms of alternating sign is harder to distinguish from GR, as evidenced by its Bayes factor growing the slowest with the magnitude of β_{-3}.

Figure 33 shows the Bayes factors between GR and a one-parameter ppE template family with $b = -3$ and ppE parameter β_{-3} for the three injections discussed above. The error bars in this figure are estimated by calculating the Bayes factors using multiple MCMC runs with different random seeds. The spread in the calculated values are reflected in the error bars. Observe that when the injection contains ppE corrections of the same sign (dotted, magenta curve), these add up to make the signal more discernible from GR. In this case, the Bayes factor becomes larger than 10, i.e. crosses our threshold for detectability, for the smallest value of β_{-3}. Therefore, if (β_b, β_{b+1}) share the same sign, we can detect deviations from GR with lower strengths than if there were only one phase correction. On the other hand, observe how when the non-GR signal contains alternating sign GR modifications (dashed blue line), these have the effect of partially canceling the non-GR effect out. In this case, the Bayes factor crosses 10 for a much larger value of β_{-3}. Therefore, if the corrections have alternating signs, e.g. if (β_b, β_{b+1}) have different signs, then our ability to detect departures from GR is reduced. The sign of the ppE amplitude exponent also affects the PDFs of the recovered β_i parameters, as we will see below.

The relative magnitudes of the terms also affects the analysis. Concentrating on the multi-term ppE models of Eq. (7.12), we define three cases, depending on the relative magnitude of these exponents in the series expansion:

- **Sub-Critical Case**: Injections where the ppE terms get smaller as the PN order increases, i.e. $\beta_b > \beta_{b+1} > \beta_{b+2}$.
- **Critical Case**: Injections where the ppE terms remain of about the same size as the PN order increases, i.e. $\beta_b \sim \beta_{b+1} \sim \beta_{b+2}$.
- **Super-Critical Case**: Injections where the ppE terms get bigger as the PN order increases, i.e. $\beta_b < \beta_{b+1} < \beta_{b+2}$.

Obviously, there are an infinite number of ways to choose how large the β_i constants are relative to each other, but the classification defined above provides a useful summary. More concretely, we here define *sub-critical*cases as those where the ppE terms injected have $\beta_{n+1} < (u_{\rm max})^{-b_n}$, where $u_{\rm max} = \pi \mathcal{M} f_{\rm max}$. Similarly, *critical* cases are defined such that $\beta_{n+1} \approx (u_{\rm max})^{-b_n}$, while *super-critical* cases have $\beta_{n+1} > (u_{\rm max})^{-b_n}$.

An alternative and roughly equivalent way to define these three different cases is by the number of *useful cycles* of phase [47] that accumulate during the signal for each correction to the phase. The number of useful cycles is defined via

$$N_{useful} \equiv \left(\int_{F_{\rm min}}^{F_{\rm max}} \frac{df}{f} \frac{a^2(f)}{S_n(f)} \frac{d\phi}{2\pi df} \right) \left(\int_{F_{\rm min}}^{F_{\rm max}} \frac{df}{f} \frac{a^2(f)}{f S_n(f)} \right)^{-1} \tag{7.13}$$

where $|\tilde{h}(f)|^2 = N(f) a^2(f)/f^2$ is the squared modulus of the frequency domain GW signal, and $N(f) = (f/2\pi)(d\phi/df)$. This quantity tells us about the phase accumulated from each PN (or ppE) term during the course of the signal, weighted by the sensitivity of the detector to different parts of frequency space. Tables of the number of useful cycles of phase for each system analyzed in this section are included herein. "Sub-Critical" signals are those for which the number of useful cycles due to the non-GR phase corrections decreases at higher order. "Critical" signals have roughly the same number of useful cycles at each order. "Super-critical" signals have larger numbers of useful cycles from the non-GR phase at higher orders.

Figure 34 shows the PDFs of the recovered β_{-3} parameter for a one ppE parameter template family, with injections given by sub-critical, critical and super-critical versions of cases (ii) and (iii). These PDFs are computed using a MCMC approach.

Signal	ϕ_{-3}	ϕ_{-2}	ϕ_{-1}
Convergant	0.312	0.012	0
Critical	0.312	0.194	0
Super-Critical	0.312	0.607	0

Table 9: Number of useful cycles from the different injected ppE terms - Fig 33 and Fig 34.

The top panel of this figure shows the PDFs for β_{-3} given an super-critical injection, the middle panel given a critical injection, and the bottom panel given a sub-critical injection. The left and right panels correspond to injections with the same (left) or alternating (right) signs. When there is as much or more weight at $\beta_{-3} = 0$ in the PDF's as there was in the prior probability density, this indicates that GR is the preferred model. In our case, the prior probability for β_{-3} is flat between -5.0 and 5.0, and so the prior probability density at all points, including $\beta_{-3} = 0$, is 0.1. When the posterior density at $\beta_{-3} = 0$ is less than 0.1, an alternative model is preferred.

Figure 34 reveals several interesting facts. First, note that in the sub- and super-critical injection cases, the sign of the βs is irrelevant: in both cases most of the weight is outside $\beta_{-3} = 0$. Second, we can see that in the sub-critical case, the second ppE term ($b = -2$) is very sub-dominant to the first term, and so its sign has little impact on the results. Third, observe that in the critical injection case, when the βs have alternating signs, the modified gravity effects partially cancel out, yielding a β_{-3} PDF with non-negligible weight at the GR value. It is clear from these studies that neglecting higher-order phase corrections can seriously bias our assessment of

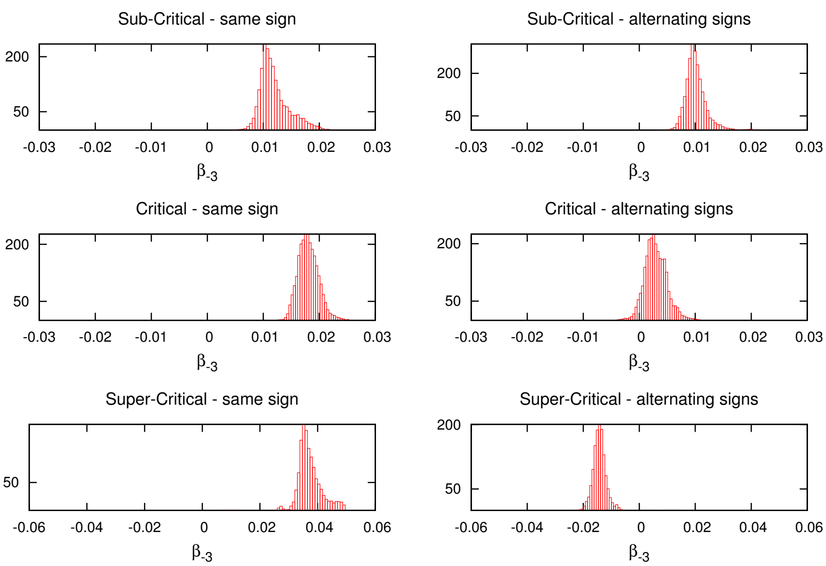

Figure 34: The PDF's for β_{-3} in a one-parameter ppE template recovered from MCMC searches on injections containing two ppE parameters ($b = -3$ and $b = -2$). In all injections, $\beta_{-3} = 0.01$, but the value of β_{-2} varies between cases. The plots on the left are for injections containing two ppE parameters of the same sign, and on the right of opposite signs. The more weight in the PDF at $\beta = 0$, the lower the Bayes factor in favor of a non-GR signal. In the critical case, we find that alternating signs in the phase corrections can cause a non-GR signal to be indistinguishable from a GR one. In the sub- and super-critical cases, this does not occur. System parameters for this figure are the same as in Figure 33, also listed in Table 8, and the useful cycles of phase are in listed in Table **??**.

our ability to test GR with GW signals. For the "Critical" case, our ability to detect departures from GR is enhanced if the terms have the same sign, and diminished if the signs alternate.

7.7. Optimal Model Selection

We have seen that it is important to consider multi-term ppE signal injections when assessing the bounds we will be able to place on alternative gravity theories. The question still remains, however, as to what type of templates we should use to recover

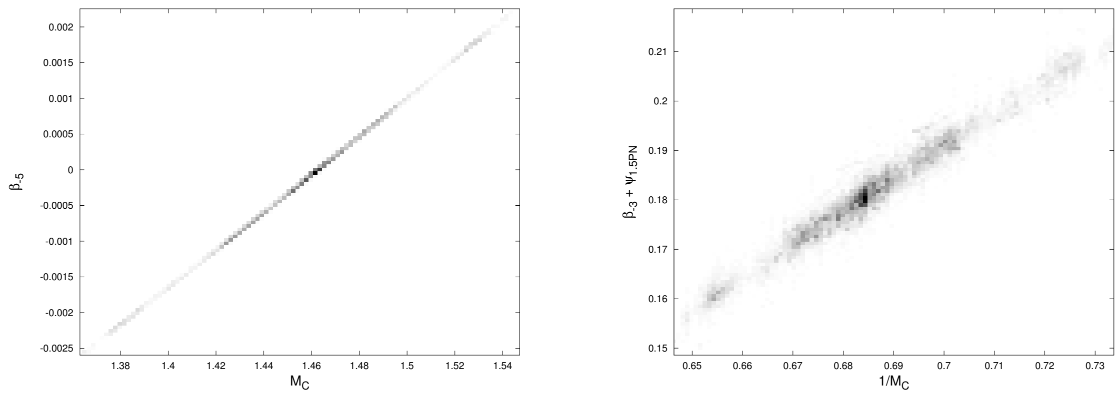

Figure 35: Correlation between the β_{-5} ppE parameter and the chirp mass (left panel) and the β_{-3} parameter and the inverse chirp mass (right panel) for an injection including two PN phase terms as well as two ppE phase corrections. The parameters are restricted only by their prior ranges.

such signals. In this section we address this question by first showing that adding too many parameters to the templates is counter-productive. Then we determine the optimal ppE template family for detecting departures from GR described by the more realistic multi-term ppE signal injection model.

7.7.1. Overfitting

One may consider using a ppE template with many ppE phase and amplitude terms in the sums of Eq. (7.10). For example, one could include as many ppE phase terms as there are in the GR PN series, but this is far from ideal. The reason is clear: if we include the same number of free ppE parameters in our phase model as we have phase terms that are functions of system parameters, then there is no way to constrain any of them. In other words, the ppE phase terms will have a 100%

correlation with the standard GR system parameters that form the coefficients of the GR PN phase.

As a simple example, consider the possibility of detecting a non-GR signal that includes ppE corrections at $b = -5$ (a so-called *Newtonian* ppE correction) and $b = -3$ (a 1PN ppE correction). We will truncate our injection at 1PN order for this example, which implies that the GW phase contains two standard PN terms that are functions of the system parameters, and two free ppE terms. Figure 35 shows that there is a 100% correlation between these PN and ppE parameters.

These types of correlations are commonly encountered in GW data analysis, but they may not be widely appreciated by theoretical model builders. We can understand this correlation analytically as follows. Let us write the simplified ppE template Fourier phase $\Psi_{\rm ppE}(f)$ as follows

$$\begin{aligned}\Psi_{\rm ppE}(f) = & \left[\frac{3}{128}\left(\pi\mathcal{M}\right)^{-5/3} + \beta_{-5}\left(\pi\mathcal{M}\right)^{-5/3}\right] f^{-5/3} \\ & + \left[\frac{3}{128\eta^{2/5}\pi\mathcal{M}}\left(\frac{3715}{756} + \frac{55}{9}\eta\right) + \frac{\beta_{-3}}{\pi\mathcal{M}}\right] f^{-1}\,. \end{aligned} \tag{7.14}$$

where we have expanded out the definition of u. Clearly, we can rescale β_{-5} by a constant and β_{-3} by a function of η, and then also adjust $\mathcal{M}$ and η, to recover the same value of the Fourier phase. This shows a direct correlation between these parameters. Figure 35 demonstrates how such a correlation manifests itself in the posterior distributions.

This argument can be extended to whatever PN order we choose. If we include the same number of ppE terms as PN terms in our model, then we will not be able to place bounds on *any* parameter, let alone use the results as a test of GR. It is also true, however, that ppE models that include more ppE terms will be able to achieve a better overall fit of whatever signal we happen to detect, just as any model with extra parameters can typically fit data better than a simpler model. In the next section we explore the tradeoff between these two effects. We also attempt to determine what types of signals are best to analyze using more complex ppE models, and what types are better served with a simple ppE model.

7.7.2. Inclusion of Spin

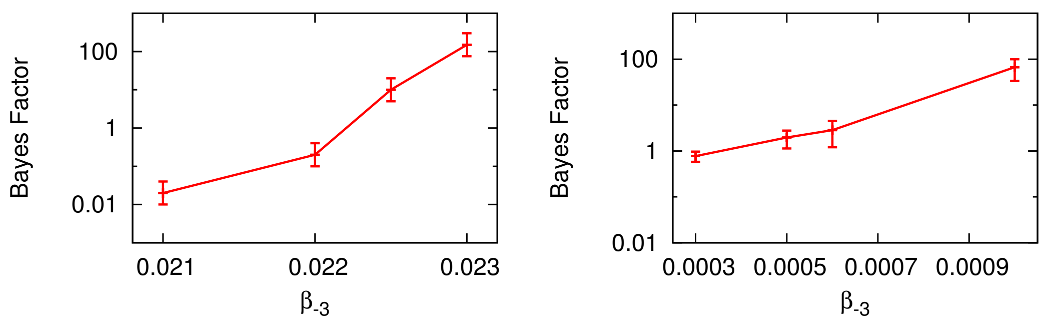

Figure 36: Bayes factors of a $b = -3$ ppE model versus GR. The injected signals in both cases were non-spinning, critical ppE injections, with the value of β_{-3} plotted on the x axis. The left panel shows Bayes factors for templates that include aligned spin parameters, and the right panel is for templates with no spin parameters. The degeneracy between the 1.5PN spin term and the β_{-3} ppE amplitude parameter significantly weakens the bounds.

There are many potential effects, both astrophysical and purely gravitational, that will make it more difficult to test GR. For instance, the presence of accretion disks [160, 83], the presence of a third companion [157], the unknown effects of the neutron star equations of state [154, 153], etc. To illustrate how these types of effects can hinder our ability to test the nature of gravity, in Figure 36 we have plotted the Bayes factors between a $b = -3$ ppE model and GR applied to critical ppE injections. In the top panel, we included the standard PN terms for aligned spins in the phase for both the GR and ppE waveform models, which introduces two new parameters. The correlation between these spin parameters and the β_{-3} ppE parameter causes the detection threshold for β_{-3} to be larger by a factor of ~ 20 compared to the case in which the spins are held fixed to zero.

The inclusion of spin effects when testing GR has been explored before in the context of particular theories of gravity. Using systems with aligned spins degrades the bounds due to correlation between the spin parameters and the alternative theory parameters [19]. Including spin precession effects [152] restores the bounds to levels closer to what is found for systems without spin [147], as recently explained in [165]. Using waveforms that include additional structure such as higher harmonics of the the orbital frequency can also improve the bounds on alternative theory parameters [13]. Thus, the situation shown in Fig. 36 should be considered a worst-case scenario. Throughout the rest of this chapter, we hold spins fixed to zero. This means that our actual bounds on the ppE strength parameters are probably a little optimistic, but it does not change the conclusions we draw about model selection.

7.7.3. Parsimonious Fitting

Source	$m_1(M_\odot)$	$m_2(M_\odot)$	$\log(D_L)$(Mpc)	t_c	β_{-3}	β_{-2}	β_{-1}
Sub-Critical	1.42	1.73	3.83	3.5	0.003	0.003	0.003
Critical	1.52	1.33	3.9	3.5	0.0006	0.018	0.54
Super-Critical	2.04	1.34	3.86	3.5	0.0007	0.07	7.0

Table 10: Source parameters for Figures 37 and 38. The β_b values listed are for a particular case - the ratio between different β_b values was kept constant for each injected signal. The ratio for sub-critical was $\times 1.0$, critical was $\times 30$, and super-critical was $\times 100$. All sources had $(\alpha, \cos\delta, \phi_L, \cos\theta_L) = (1.42, 0.87, 2.5, 0.43)$.

We now study what types of ppE templates are best suited for detecting a GR deviation. In particular, we examine whether using one-term or two-term ppE templates works better. For this analysis, we inject ppE signals containing three phase terms, and attempt to recover them using one- and two-parameter ppE templates. We calculate Bayes factors between the ppE models against the GR model to see which model is best suited to detecting departures from GR. Because of our strong prior belief in the validity of GR, a Bayes factor significantly greater than unity would be necessary to convince us that a new theory of gravity is needed.

Let us then consider three different ppE injections, starting at 1 PN order ($b = -3, -2, -1$), a sub-critical, a critical and a super-critical one, each for a NS-NS inspiral, with parameters listed in Table 10 . We explore these simulated signals with a MCMC algorithm, using a one- and a two-term ppE model. The one-term ppE models are allowed to choose between phase exponents $b = -3$ and $b = -2$, while the two-term models are allowed to choose between the pairs $(-3, -2)$ and $(-2, -1)$

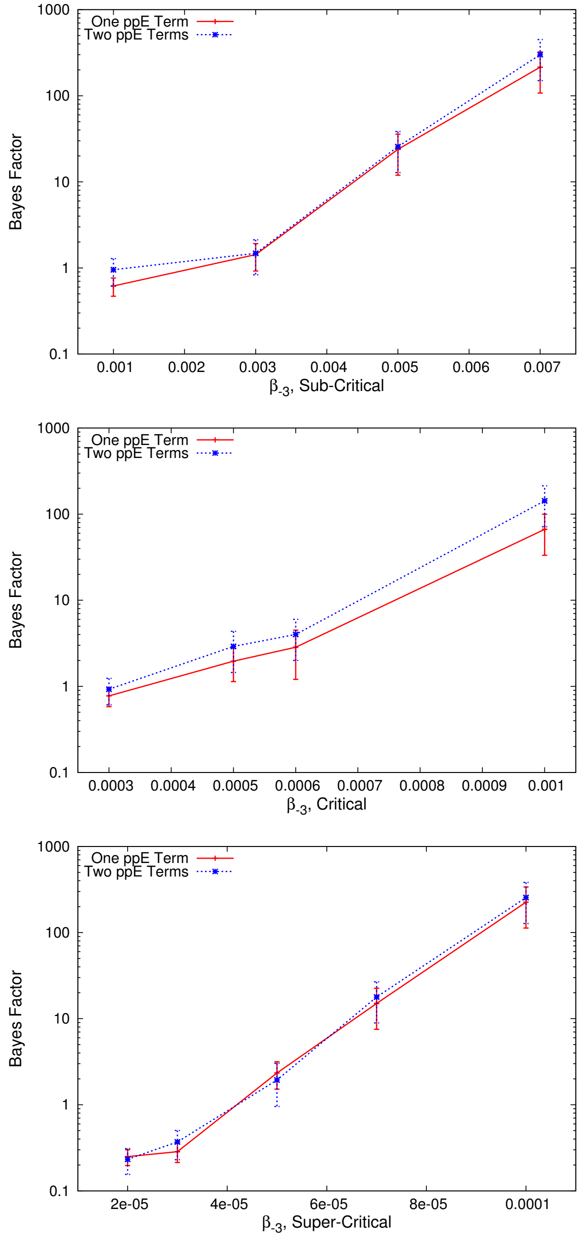

Figure 37: Bayes factors for one-term (solid red) and two-term (dashed blue) ppE templates for a sub-critical (top-left), critical (top-right) and super-critical (bottom) ppE injection as a function of the injected value of β_{-3}. System parameters are listed in Table 10, and useful cycles of phase in Table 11. In the sub- and super-critical cases, both models perform equally well at detecting a deviation from GR. In the critical case, the two-term model slightly out-performs the one-term model.

- *i.e.* the two terms must differ by a single power of u, and models with exponents $(-3, -1)$ are not allowed.

The Bayes factors between the one-term ppE model and GR (red solid curve) and between the two-term ppE model and GR (blue dashed curve) are shown in Fig. 37 as a function of the injected value of β_{-3} for a sub-critical (top-left panel), critical (top-right panel) and super-critical (bottom panel) injection. These Bayes Factors are again calculated using the Savage-Dicke density ratio. Calculating the posterior density at a $\beta_i = 0$ from a Markov chain involves counting the number of points in the chain that fall within the histogram bin containing $\beta_i = 0$, and so the error bars reflect the counting error involved in this process, as well as the spread in BF values calculated from multiple MCMC runs on the same signal but with different random seeds. Observe that the only injections for which two-term ppE templates consistently outperform one-term ppE templates are the critical ones. Even in this case, however, the preference is not large; the curves track each other very well in all cases. Therefore, our results indicate that the one-term ppE templates are sufficient for searching for deviation from GR in GW data.

Once a deviation from GR has been definitively detected, the next step is to learn as much about the signal as possible, in order to give theorists as much guidance as possible in their attempts to build an alternative theory of gravity. The information we could hope to extract from the type of analysis we have described in this chapter is the structure of the series of phase corrections - do they enter at a certain PN power and then fade away? Or do they enter at that power and grow more important at higher orders in the expansion? Figure 38 plots the posterior distribution of the five models under consideration derived using a RJMCMC [40] analysis. In RJMCMC, moves are proposed between models of different dimensionality according to the Metropolis-Hastings ratio:

$$\alpha = \min \left\{ 1, \frac{p(\vec{\lambda})_Y p(s|\vec{\lambda}_Y) q(\vec{u}_Y)}{p(\vec{\lambda})_X p(s|\vec{\lambda}_X) q(\vec{u}_X)} |\mathbb{J}| \right\} \tag{7.15}$$

Here, model X and model Y differ by some number of parameters, $q(\vec{u})$ is the distribution for random numbers chosen to generate the extra parameters, and $|\mathbb{J}|$ is the Jacobian of the two sets of parameters, which compensates for the difference in dimensionality. When using this Hastings ratio as an acceptance probability, we can allow our chains to explore the full space of allowed ppE models, both one- and two-term families, and use these to generate PDF's for the models themselves. The ratio of the heights of the PDF for model X and model Y is equal to the Bayes Factor between X and Y.

Signal	ϕ_{-3}	ϕ_{-2}	ϕ_{-1}
Convergant	0.109	0.008	0.0005
Critical	0.024	0.051	0.085
Super-Critical	0.024	0.181	1.047

Table 11: Number of useful cycles from the different injected ppE terms - Fig. 37

To generate Figure 38, we have run an RJMCMC search on three different types of signals - one sub-critical, one critical, and one super-critical - and plot the number of iterations that the chains spent in each of the five different models. These five models include two ppE models with only one phase correction, ($b = -3$ and $b = -2$), two ppE models with two phase corrections, ($b = -3 + b = -2$ and $b = -2 + b = -1$), and GR. We find that, although there are some slight differences between the different models, in all cases we cannot draw meaningful distinctions between the different ppE models. The strongest Bayes Factor between two models is in the sub-critical case, where the Bayes Factor between the $b = -3$ only model and the $b = -2$ only model is ≈ 5. While this does show some preference for the first model, it is not a strong preference, and so we would not want to use this result to draw conclusions about the underlying theory of gravity. In summary - even though these signals are clearly differentiable from GR (all have Bayes Factors of ≈ 100 in favor of the ppE models), the four different ppE models perform almost as well in fitting the signal. This means that if we hope to gain more information about the underlying nature of an alternative gravity theory, we would need higher SNR signals and/or multiple detections. On a more hopeful note, it means that our ability to detect a deviation from GR is not strongly dependent on which particular ppE template we choose to use in our analysis.

7.8. Detecting GWs from Scalar-Tensor Gravity

The ppE waveforms cover all known inspiral waveforms from specific alternative theories of gravity that are analytic in the frequency evolution of the GWs [38]. There are, however, known theories for which the frequency evolution of GWs is *not* analytic

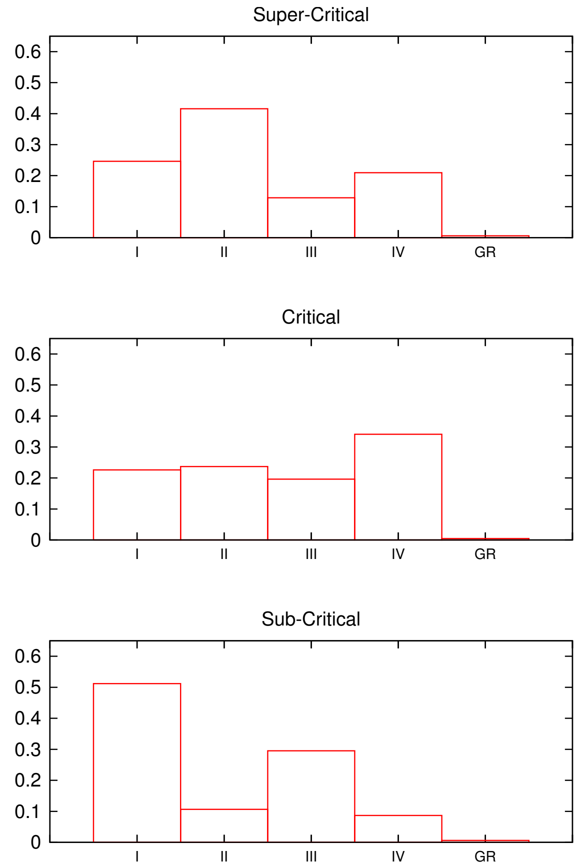

Figure 38: Posterior distributions for the four different ppE models, generated by RJMCMC. The top two panels show the distribution for a sub-critical injection, the middle two for a critical injection, and the bottom two for an super-critical injection. All systems are NS-NS binaries with Bayes Factors of 100 favoring ppE over GR. System parameters are in Table 10. Model I has $b = -3$, model II has $b = -2$, model III has $b = -3$ *and* $b = -2$, and model IV has $b = -2$ *and* $b = -1$. The y axis shows the percentage of iterations that the chain spent in each model, and the Bayes Factors between two models are simply the ratios of the percentages. Because the Bayes Factors are not large enough, these results indicate that we would not be able to make confident statements about the type of non-GR signal we had observed with this type of analysis.

– in particular, ST theories of gravity in which spontaneous scalarization can occur, and theories of gravity that contain a massive scalar field [7, 32, 16, 121].

The theories that include spontaneous scalarization can be defined with a generic ST action of the form [16, 121]

$$S = \int d^4x \frac{\sqrt{-g}}{2\kappa}\left[\phi R - \frac{\omega(\phi)}{\phi}\partial_\mu\phi\partial^\mu\phi\right] + S_M. \tag{7.16}$$

where $\kappa = 8\pi G$, R is the Ricci scalar, g is the determinant of the metric, ϕ is the gravitational scalar, and S_M is the matter action.

In this type of theory, neutron stars that are not initially scalarized can acquire a scalar charge when the system reaches a high enough binding energy – i.e. once the orbital frequency is high enough. This spontaneous change leads to the "turning on" of dipole GW radiation once the merging stars get close enough together. This radiation, in the ppE scheme, goes as $b = -7$, a low PN order effect – lower order, in fact, than even the Newtonian term. This means that binary pulsar measurements have placed very tight restrictions on the possible strength of this dipole radiation at low frequencies. However, we do not yet have any measurements of binary systems at high frequencies. It is therefore not impossible that signals of this type, that have no dipole radiation at low frequencies, but significant dipole radiation at high frequencies, could be detected by GW experiments.

Another type of ST theory that produces this type of radiation is that in which the scalar field, ϕ, has a mass [7, 32]. The gravitational radiation due to this type of theory has been calculated in [21]. The phase is altered from the GR expression, and is equal to

$$\begin{aligned}\psi(f) = \psi_{\mathrm{GR}}(f) + \xi\Gamma^2\nu&\left[\frac{5}{462}u^{-11} - \frac{\nu}{1632}\eta^{12/5}u^{-17}\right] \\ &\times \Theta(2\pi f - m_s) \\ &+ \xi\mathcal{S}^2\left[\frac{25\nu}{1248}\eta^{8/5}u^{-13} - \frac{5}{84}\eta^{2/5}u^{-7}\right] \\ &\times \Theta(\pi f - m_s).\end{aligned} \tag{7.17}$$

In this expression, m_s is the mass of the scalar field, $\eta = (m_1m_2)^{3/5}/(m_1+m_2)^{1/5}$ is the symmetric mass ratio and the other quantities are given by

$$\begin{aligned}\xi &= \frac{1}{2+\omega_{\mathrm{BD}}}, \\ \Gamma &= 1 - 2\frac{s_1m_2 + s_2m_1}{M}, \\ \mathcal{S} &= s_2 - s_1, \\ \nu &= 5.60 \times 10^{-21}\left(\frac{m_s}{10^{-20}eV}\frac{M}{M_\odot}\right),\end{aligned} \tag{7.18}$$

where s_1 and m_1 are the sensitivity and mass of body 1, M is the total mass of the system , and ω_{BD} has been constrained by the Cassini spacecraft such that $\omega_{\mathrm{BD}} \geqslant 40000$.

The non-GR portions of the signal are not present until a frequency $f = m_s/2\pi$. This means that, in order for the non-GR signal to be detectable in the aLIGO band, i.e. $f \sim 100$ Hz, the mass of the scalar field must be approximately $m_s \sim 10^{-13}$ eV. This leads to a value of $\nu \sim 10^{-14}$, which implies that the phase terms that are multiplied by ν are highly suppressed in this frequency range, leaving only dipole deviations from GR.

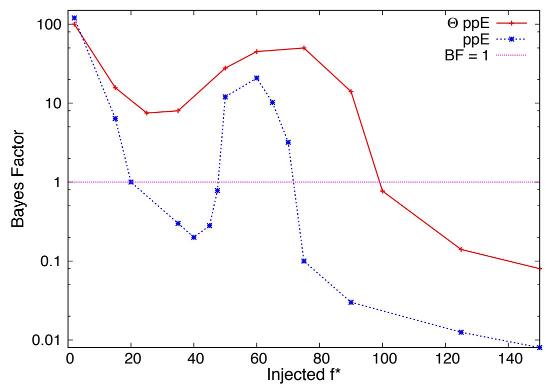

Figure 39: BF between GR and ppE models for injections of the form in Eq. (7.19), for varying values of f^*. The dashed (blue) line is for a standard and simplest, inspiral ppE template, and the solid (red) line is one that has been modified with a Heaviside function. Both types of templates show the same general behavior, and both are successful at detecting deviations from GR for certain ranges of f^*.

Thus both of these types of ST GW signals can be approximated, in ppE notation, as

$$\tilde{h}(f) = \tilde{h}_{\rm GR} e^{i\Theta(f-f^*)\beta u^{-7}}. \tag{7.19}$$

Here, $\Theta(f - f^*)$ is the Heaviside function, and f^* is the frequency at which the dipole radiation 'turns on.' Clearly, as f^* goes to infinity, the signal in Eq. (**??**) becomes a GR signal, and when f^* is only a few Hz, the signal is vastly different from GR.

Our first task, then, is to determine at what value of f^* the standard ppE templates, with $b = -7$, will detect the non-GR character of the signal. To find this value, we inject signals of the form in Eq. (7.19) with $\beta = 10^{-6}$, and varying values of f^*, ranging from 5 Hz to 150 Hz. These signals are SNR 12 inspirals from neutron star binary systems with masses of $1.4 - 2.0$ solar masses, using the neutron-star-binary optimized aLIGO noise curve. We then use Markov Chain-Monte Carlo (MCMC) techniques to recover the signals and calculate the Bayes factor (BF) between GR and non-GR models. The BF between two models, A and B, is the ‘betting odds’ that model A provides a better description of the data than model B. If the BF of A vs. B is greater than one, the data shows a preference for model A. In this paper, we compute BFs using the Savage-Dicke density ratio [135], in which the prior weight at $\beta = 0$ is compared to the posterior weight. A decrease in probability density at this point corresponds to a preference for a non-GR model. The BFs between GR and non-GR are plotted in Fig. 39. When the BF is above 1, the model selection process favors a non-GR signal.

The overall behavior in Fig. 39 is as expected. For low f^*, the signal is clearly non-GR, and when f^* is very high, GR is favored. There is an unexpected region in the middle, however, in which the BF grows with f^*. We can further investigate this region by looking at the posterior distributions for the ppE strength parameter, β, plotted in Fig. 40. At first, there is only one peak in the posterior, and it is centered at the correct, injected value of β. As f^* becomes larger, a secondary maximum appears, centered at an incorrect value of β. For some values of f^*, these two maxima fit the data equally well, which leads to significant posterior weight at $\beta = 0$, leading to a BF that favors GR. As f^* grows, the secondary maximum becomes a better fit to the data, until eventually the GR value of $\beta = 0$ wins out. This can be understood by noting the relationship between the BFs and the aLIGO noise curve we have used. The BF is largest when f^* is in the region of highest sensitivity for the detector.

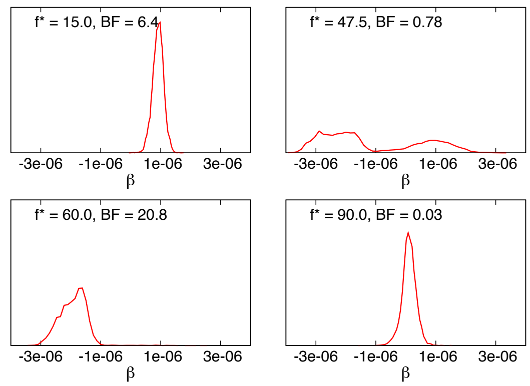

Figure 40: Posterior distributions for β, recovered using standard ppE templates. The injected signal was of the form in Eq. (7.19), with $\beta = $ 1e-06. If there is little weight in the posterior at $\beta = 0$, the signal is detectable as non-GR. In the top left panel, f_{min} is low, and β is recovered at the correct value. In the bottom right panel, f_{min} is very high, and the GR model is clearly favored. In the bottom right panel, the signal is clearly non-GR, but the recovered value for β is incorrect. Finally, in the top right panel, two peaks in the posterior are clearly visible – one mode near the correct value of β, and one at the incorrect, negative value. Because the chain swaps between the two peaks, there is significant weight at $\beta = 0$, and this signal is not detectable as non-GR.

This behavior can be understood further by examining the correlations between β and the other source parameters, for example in Fig. 41. This plot shows the two-dimensional posterior distribution for β and the chirp mass, $\mathcal{M}$. In this example, there is clear correlation visible between the two parameters, and two peaks in the posterior are clearly visible. While these two peaks both represent good fits to the data, there is also significant weight between them, which means significant weight at $\beta = 0$, and thus a BF that favors GR.

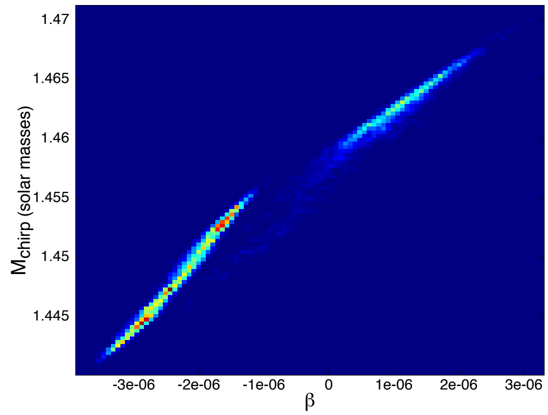

Figure 41: The correlation between β and $\mathcal{M}$, generated from a signal of the form in Eq. (7.17), with $f^* = 47.5$. The two separate maxima in the likelihood are clearly visible, as well as the strong correlation between these two parameters.

In addition to testing the standard ppE templates, we also use templates of the form in Eq. (7.19), in which we allow f^* to be a parameter that is determined by the data. We use these templates to recover the same signals injected in the previous study, and again use MCMC techniques to calculate the BF between GR and non-GR models. The results are also plotted in Fig. 39. These enhanced templates show the same qualitative behavior as the standard ppE templates, although their overall performance is better. This is to be expected, as this template family can perfectly match the injected signals.

An additional point of interest for the enhanced templates is the precision with which we are able to recover the injected parameter, f^*. The uncertainty in the recovered f^* is plotted as a function of the injected f^* in Fig. 42. The precision with which we can measure this parameter depends on the BF, as expected. Even so, for a large range of values, we are able to recover this parameter quite accurately. If a

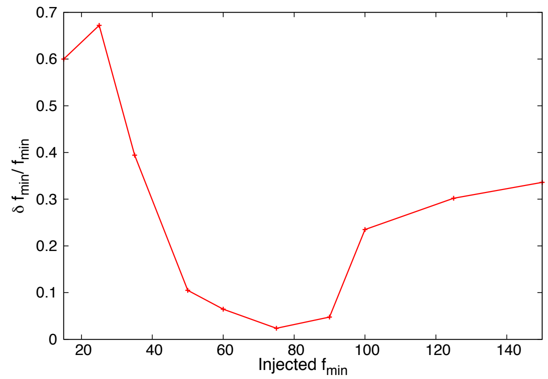

Figure 42: Fractional uncertainty in the recovered value of f^*, for different injected values of f^*. The uncertainty is inversely proportional to the BF in favor of the non-GR model – i.e. when the BF indicates a clear departure from GR, the f^* parameter is recovered with high accuracy.

signal of this type were detected in the data, information about f^* would be useful for theorists attempting to learn about the underlying gravitational theory.

7.9. Inspiral-Merger-Ringdown Signals

In this section, we analyze some of the issues that can arise from using inspiral-only templates in tests of GR, as well as some of the science we can do by including merger and ringdown in our analysis. We explore the problem of testing GR using full GW injections - i.e., injected waveforms that include merger and ringdown in addition to inspiral. In the first subsection, we investigate what can happen when we use inspiral-only templates to extract full signals. In the second subsection, we consider a family of ppE templates that includes non-GR parameterizations of the merger and ringdown stages.

7.9.1. Extracting with Inspiral-only Templates

Typically, GWs from binary systems are talked about as if they have three discrete parts – inspiral, merger, and ringdown (IMR). The inspiral is the part of the waveform that is generated while the two bodies are still widely separated, and thus slowly spiraling towards each other due to the emission of GWs. The merger is the most difficult part of the signal to model analytically, and is the portion in which the two bodies are very near each other and moving very quickly – eventually becoming one object. Finally, the ringdown stage is produced after the two bodies have merged, as the resulting object relaxes to its final state.

The inspiral portion of GW signals has been calculated in several alternative gravity theories [163, 18, 46, 19, 147, 143, 151, 150] (for a recent review, see [165]). These calculations were then used as motivation for the inspiral waveforms in the ppE family. The merger phase, unfortunately, has not been calculated in any non-GR theories – GWs outside of GR may even lack ringdowns altogether, and almost certainly would have different relationships between the system parameters and the quasi-normal modes [20, 55, 70]. Even within GR, the merger stage must be calculated numerically, and connected phenomenologically to the analytic solutions for inspiral and ringdown [117, 4]. This means that we lack theoretical motivation for non-GR merger and ringdown templates, and thus when testing GR we usually choose to use only the inspiral portion of the signal where we have an analytic expression for GR waveforms, and well-motivated templates for alternative gravity.

We often discuss the three stages of GWs as if they are clearly separable, but the transition from inspiral to merger and merger to ringdown is a somewhat arbitrary distinction. One common choice is to take the transition from inspiral to merger to be the frequency of a test-particle at the innermost stable circular orbit (ISCO) of a Schwarzschild black hole, $f_{\rm ISCO} = 1/(6^{3/2}\pi M)$. For full waveforms that include the merger and ringdown, the transition from inspiral to merger is smooth, and can begin to have effects earlier or later than this, depending on the system. When using inspiral-only templates, as is commonly done in GW data analysis [1, 3, 2], the question of when (i.e. at what frequency) to terminate the waveforms is not a trivial one.

In order to use the full, three-stage signal model for detection and characterization of GWs, we need an efficient template family that can capture the full signal. In this paper, we use PhenomC waveforms in our analyses of full IMR signals [117, 4]. In these waveforms, the inspiral stage is modeled within the PN approximation. The merger stage has been studied numerically [77, 72, 15, 30, 108], and is approximated analytically. The final stage, ringdown, is modeled from black hole perturbation theory analytically. These three pieces of the waveform are stitched together with matching procedures and calibrated against full numerical results to produce full, IMR waveforms.

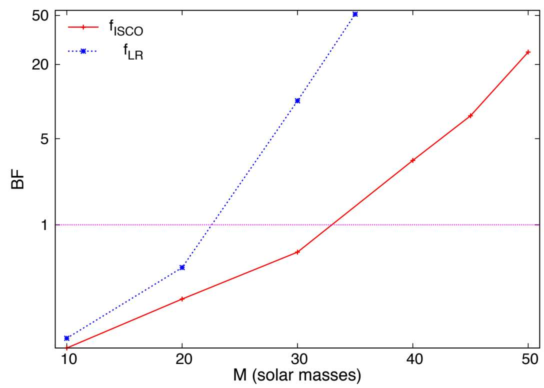

Figure 43: BFs between GR and ppE templates. The injected signals were GR, PhenomC waveforms, and they were recovered using inspiral-only ppE waveforms. The dashed (blue) line shows the BFs calculated by using the frequency at the light ring as as the cutoff frequency for the waveforms. The solid (red) line shows BFs calculated by using $f_{\rm ISCO}$ as the cutoff frequency. A BF larger than 1 indicates a preference for the non-GR model.

Our goal is to determine where inspiral ends in a self-consistent way – one that does not lead to significant biases in the recovered parameters. When using GR inspiral templates to fit a full GR IMR signal, the biases arise primarily in the recovered value of the total mass. This is because the cutoff frequency of the template is determined by the total mass, and so the inspiral waveform "stretches" to fit some of the merger power by changing this parameter. Although this type of bias is clearly undesirable, the true danger arises when using ppE inspiral signals to fit the full waveform. In this case, not only the total mass, but the ppE strength parameter, β, change to attempt to fit some of the merger/ringdown power. This can lead to a recovered value of β that is not consistent with GR, and thus the claim of a detection of a deviation from GR in the GW signal.

Figures 43 and 44 show some of the consequences that can arise when using inspiral-only templates to analyze a signal that includes merger and ringdown. To generate both figures, we injected PhenomC waveforms, with total mass ($M = m_1 + m_2$) beginning at $M = 10M_\odot$, up to $M = 50M_\odot$. We then recovered these signals using inspiral-only GR templates, as well as inspiral-only ppE templates with $b = -2$, which corresponds to a 1.5 PN correction to the GW phase. We used two different determinations of the cutoff frequency for the inspiral waveforms. In one set of runs, we used $f_{\rm IM} = f_{\rm LR}$, where $f_{\rm LR}$ is the frequency of a test-particle at the light ring of a Schwarzschild black hole. For the other set, we used $f_{\rm IM} = f_{\rm ISCO}$, the ISCO frequency. The choice of cut-off frequency had a significant effect on our results, signaling a departure from GR for sufficiently massive systems, even when the injection had none. In all cases, we again use the neutron-star-binary optimized aLIGO noise curve.

In Fig. 43, we plot the BF of ppE vs. GR templates recovered from these signals. When the BF is larger than 1, the non-GR model is preferred. Even though all injections are GR signals, the BF shows a preference for non-GR models for $M > 30M_\odot$ when the light ring frequency is used to terminate the waveforms, and for $M > 40M_\odot$ when the transition frequency for the waveforms is set to $f_{\rm ISCO}$. If we are not careful to use inspiral-only templates only for low-mass systems, then, we could mistakenly claim the detection of a deviation from GR.

This growth in BF in favor of non-GR models is accompanied by a growing bias in the recovered value for M. This is illustrated in Fig. 44, where we plot $\delta M/M$, where

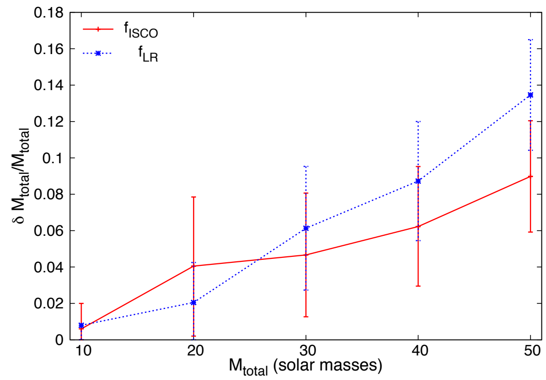

Figure 44: The bias in total mass, M, recovered when using an inspiral-only, GR signal to fit an IMR, GR signal. The dashed (blue) line was calculated using the light ring to determine cutoff frequency, and the solid (red) line used the ISCO. The error bars show the 1-σ limits for the recovered values. For high-mass systems, the bias nears 10%. For lower-mass systems, the recovered mass is very close to the injected value.

δM is the difference between the recovered best-fit value for M and the injected value. For the systems we analyzed, the discrepancy between the recovered and injected masses was never larger than 8%, even though, as expected, the templates using a cutoff frequency different from the injected transition frequency performed worse than the others.

It is not surprising that analyses using inspiral-only templates are only dependable for low-mass systems. As indicated above, the transition from the inspiral portion of the waveform to the merger portion depends upon the total mass of the system in question. This means that for a low-mass system, most of the SNR of the signal is contained in the inspiral, whereas for a high-mass system, a large fraction is in the merger and ringdown. Table 12 lists the inspiral/merger transition frequency and the fraction of SNR contained in the inspiral for systems of varying total masses.

Total Mass ($M_\odot$)	$f_{\rm ISCO}$ (Hz)	% SNR before ISCO
10	879	100
20	220	93
30	147	82
35	126	76
40	110	72
45	98	66
50	88	61

Table 12: ISCO frequency and percentage of total SNR before the ISCO for systems of different total mass.

The fact that using inspiral-only templates to fit IMR signals will lead to parameter biases has been understood for some time [91, 90, 27, 43, 106, 31], and to this point the method for avoiding these biases has been to use this type of signal only when analyzing low-mass systems. Here we present a simple, two-stage technique that allows us to use some higher mass systems in inspiral-only analyses:

- Stage I: run a standard inspiral-only template analysis on the full signal, using the self-consistently determined frequency $f_{\rm ISCO}$ as a cut-off.
- Stage II: using the (biased) value of M recovered in Stage I, low-pass filter the data to remove everything above a frequency corresponding to $r = 10M$, where r is the separation distance of the binary, and re-run the analysis.

The results of this type of analysis are shown in Fig. 45. Here we show the posterior distribution for M and β for both Stage I and Stage II, for the case where

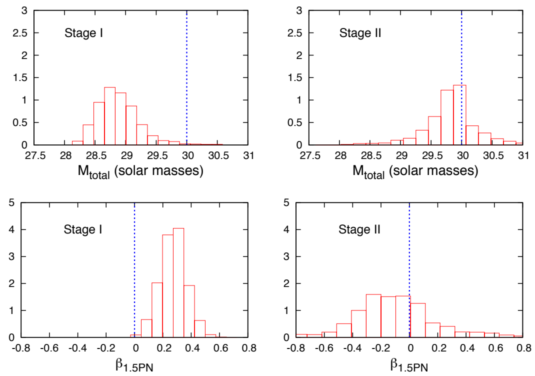

Figure 45: The posterior distributions for M and β for a 30 $M_{\odot}$ system, from Stage I in the left panels and Stage II in the right panels. The bias in recovery of M is removed in Stage II, as is the model preference for ppE over GR. The injected value is shown by the vertical line in each panel.

$M = 30M_{\odot}$. By using the conservative cutoff frequency associated with $r = 10M$ in our analysis, we are able to remove the bias in the recovered total mass, as well as avoiding making false claims that the signals are not consistent with GR. One more feature to note is that the distributions recovered in Stage II are broader than those in Stage I – this is because a part of the SNR of the signal is discarded when we make the cutoff frequency of our analysis so low. Because of this effect, this type of procedure is only useful for high total mass systems that also have a high SNR.

The BF calculations are further illustrated in Fig. 46, where we show the BF between GR and ppE for GR injections of varying M, first for Stage I using $f_{\rm ISCO}$ as the cutoff frequency, then for Stage II, using $r = 10M$. The more conservative analysis never results in the erroneous favoring of a non-GR model.

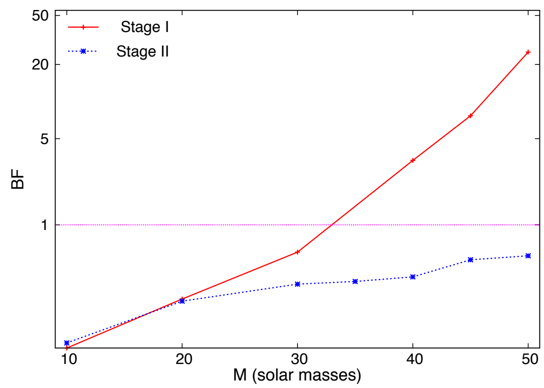

Figure 46: BFs between GR and ppE templates, from Stage I (solid/red), using $f_{\rm ISCO}$ as the cutoff frequency, and from Stage II (dashed/blue), using $r = 10M$ to calculate the cutoff frequency. All signals were GR signals. In Stage II, the model selection process always favors GR.

7.9.2. Extracting with IMR Templates

Because it is not possible to analyze systems of all masses with inspiral-only templates, the next problem we investigate is in using the full IMR waveforms to test GR. As we have mentioned, there are currently no concrete examples of merger/ringdown waveforms in non-GR gravity theories. This unfortunately means that we do not have strong theoretical motivations for what non-GR templates should look like for these parts of the signal. We do know, however, that by adding some flexibility to GR templates, via introducing parameters to the merger and ringdown stages, we will be able to fit a wider class of signals than with GR templates alone. We cannot at present know if the flexibility is enough to fit all possible non-GR signals, but if the extra parameters are recovered at their GR values, we can at the very least say that the data is consistent with GR.

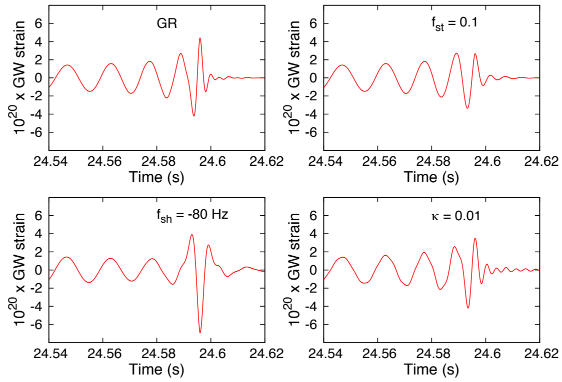

Figure 47: Time-domain waveforms generated using the parameterization in Eq. (7.20), for an SNR 30 signal with total mass $M = 50M_{\odot}$. Top left: GR waveform. Top right: $f_{\rm stretch} = 0.1$: the merger portion of the waveform is compressed, but the frequency at which merger begins and the structure of the ringdown are unaffected. Bottom left: $f_{\rm shift} = -80$ Hz: the beginning of merger is shifted to a lower frequency by 80 Hz, but the duration of merger and the ringdown structure are unaffected. Bottom right: $\kappa = 0.01$: merger is unaffected, but ringdown is changed such that the decay is much slower than in GR.

This was exactly the philosophy followed in Ref. [163] when proposing the ppE template family. That paper, in fact, proposed a variety of families, including an IMR one. Restricting attention to a simplified version of such ppE IMR family [163], we will consider the following templates:

$$\tilde{h}(f) = h(f) = \begin{cases} h_{\rm GR}(\exp i\beta u^b), & \text{if } f < f_{IM} \\ A_M f^{-2/3}\exp i\delta, & \text{if } f_{IM} \leqslant f < f_{MR} \\ \frac{A_R}{1+4\pi^2\tau^2(f-f_{MR})^2} & \text{if } f_{MR} \leqslant f \end{cases}, \tag{7.20}$$

The inspiral portion is a standard PN inspiral with the inclusion of a single ppE phase term; we here restrict attention to the $b = -2$ case. The merger comes from an analytic fit to numerical data, and includes two matching parameters, A_M and δ, to ensure continuity at the transition point. The ringdown is a single quasi-normal mode, with matching parameter A_R to ensure continuity, real frequency f_{MR} and decay time τ. In GR, the decay time can be modeled via [163]

$$\frac{1}{\tau^{\rm GR}} = \frac{0.51\eta^2 + 0.077\eta + 0.022}{\pi M}, \tag{7.21}$$

while the transition frequencies between inspiral/merger and merger/ringdown can be modeled as [163]

$$\begin{aligned} f_{IM}^{\rm GR} &= \frac{0.29\eta^2 + 0.045\eta + 0.096}{\pi M} \\ f_{MR}^{\rm GR} &= \frac{0.054\eta^2 + 0.09\eta + 0.19}{\pi M}. \end{aligned} \tag{7.22}$$

To model deviations away from the GR expectation, we include four parameters that encode non-GR effects. These are β, the usual ppE phase parameter from the inspiral portion, and the following three new non-GR parameters:

- $f_{\rm shift}$ shifts the beginning of merger:

  $f_{IM} = f_{IM}^{\rm GR} + f_{\rm shift}$

- $f_{\rm stretch}$ stretches the merger:

  $f_{MR} = (f_{MR}^{\rm GR} - f_{IM}^{\rm GR}) f_{\rm stretch} + f_{IM}$

- κ adjusts the value of τ from its GR value:

  $\tau^2 = \tau_{\rm GR}^2/\kappa$

The effects of these new parameters on the time-domain waveforms are illustrated in Fig. 47. GR is recovered when $(\beta, f_{\rm shift}, f_{\rm stretch}, \kappa) = (0, 0, 1, 1)$.

In analogy with our previous work [38], we assess how well these templates could be used to test GR by determining the range of values for each parameter that are consistent with a GR signal. We do this by injecting a GR signal of the form of

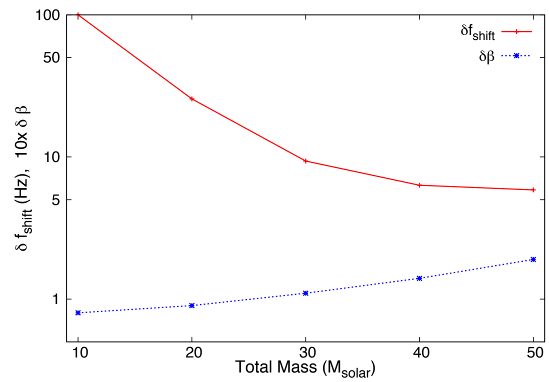

Figure 48: Uncertainty in the recovered value of $f_{\rm shift}$ (solid/red) and $10\times$ the uncertainty in the recovered value of β (dashed/blue), for different injected values of M. The uncertainty in $f_{\rm shift}$ decreases as the total mass increases and the merger-ringdown portion of the waveform becomes more important. The uncertainty in β increases due to correlations between the two parameters. All injected waveforms were GR signals.

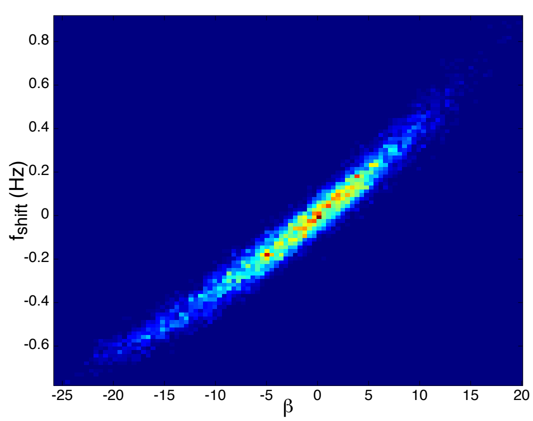

Figure 49: Correlation between the ppE phase parameter β, and the parameter $f_{\rm shift}$, that controls the start of the merger phase. This correlation leads to an increase in the uncertainty in our recovery of β for systems of increasing M.

Eq. (7.20) with $(\beta, f_{\rm shift}, f_{\rm stretch}, \kappa) = (0, 0, 1, 1)$, in this case with SNR ~ 25, and running an MCMC analysis with templates also of the form of Eq. (7.20) but with free $(\beta, f_{\rm shift}, f_{\rm stretch}, \kappa)$ to produce posterior distributions for these parameters. The results we obtain are shown in Figs. 48 and 50.

In Fig. 48, we plot the uncertainty in the recovered values for f_{shift} and β, as a function of M. If we had detected a signal of the form injected, then our results indicate we would have been able to exclude the region f_{shift} and β space above the curves shown in Fig. 48. As expected, with increasing total mass the transition between inspiral and merger has a larger effect on the signal, and so f_{shift} is better constrained. The uncertainty in β, on the other hand, grows with increasing M. This can be understood by examining the correlation between these two parameters. As seen in Fig. 49, which shows the two-dimensional posterior distribution for β and f_{shift} for a system of total mass 50 $M_{\odot}$, the correlation is high.

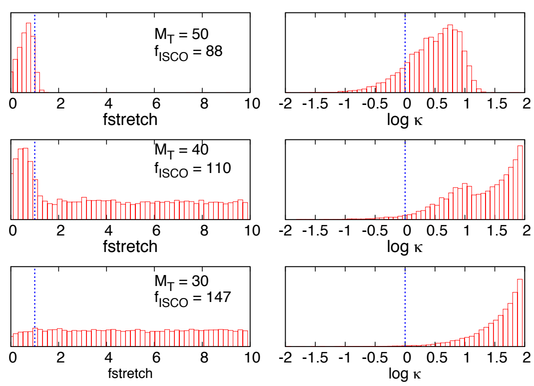

Figure 50: Posterior distributions for the parameters κ and f_{stretch} for various values of M. As the total mass increases, the parameters go from being completely unconstrained to well-measured by the data. All injected signals were GR signals. The vertical line in each panel indicates the injected, GR value for that parameter.

For the other two parameters, κ and f_{stretch}, the posteriors themselves are plotted in Fig. 50. We do so because, as can be easily seen in the figure, the posterior distributions for these parameters are highly non-Gaussian for low-mass systems, and so

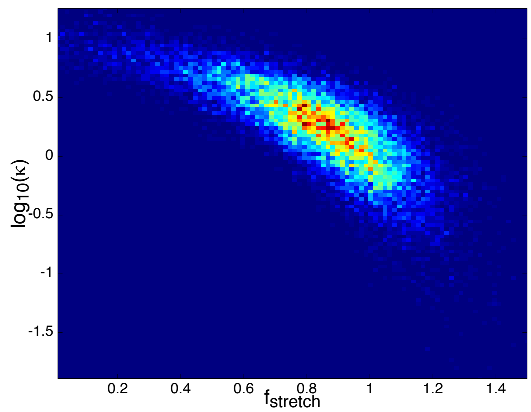

Figure 51: Correlation between the parameter τ, which affects the ringdown phase, and the parameter f_{stretch}, that controls the length of the merger phase.

an estimate of the uncertainty is somewhat meaningless. The same general pattern is still apparent, however. For low-mass binaries, κ and f_{stretch} are essentially unconstrained within their prior ranges. As the mass of the system grows, the precision with which we could measure these two parameters (or analogously exclude non-GR deviations) increases. The peak of the distribution in each parameter, however, is not centered precisely on the GR value. This is because of correlation between the two parameters, which is shown in Fig. 51 via the two-dimensional posterior distribution of κ and f_{stretch} at $M = 50M_{\odot}$.

The systems for which we could use this parameterization to test GR are those for which the new parameters can be constrained. From Figs. 49 and 50, we can conclude that:

(i) Non-GR deviations to the merger and ringdown can be detected for total masses at or above $50M_{\odot}$,

(ii) Non-GR deviations to the inspiral become less detectable for larger total mass binaries.

(iii) All non-GR parameters that characterize deviation to the merger and ringdown ($f_{\rm shift}$, $f_{\rm stretch}$, and κ) can be constrained with IMR ppE templates.

These conclusions, of course, depend on the assumptions made in our analysis, chief among which are SNR $\sim$ 25, and neglecting spins and eccentricity. Including the latter, or studying signals with lower SNR will likely weaken the degree to which we can detect non-GR deviations.

8. SUMMARY

The first detection of gravitational waves will be a triumph for the scientific community, and yet another vindication of the predictions of general relativity. Beyond this simple confirmation of Einsteins predictions, the most exciting science that this detection will lead to depends on the extraction of information from the GW signals. In particular, parameter estimation using sophisticated data analysis techniques will allow us to measure the masses, spins, and locations of black holes and other astronomical objects that we have not been able to learn about using electromagnetic signals alone. We will then be able to use these measured values to aid in our understanding of myriad problems in astrophysics, from the formation history of galaxies to the very nature of gravity itself.

The research in this thesis far has focused on the latter problem - how can we use GWs to test whether GR is the correct theory of gravity that describes our universe? We have focused on template-based search techniques in which we know the form of the signal we are looking for, and used MCMC techniques to perform Bayesian model selection based on these templates. In standard GW analyses, we assume that the signals that will be present in our data are described by GR, and so use waveform templates to extract these signals that are derived within GR. In order to test GR with GWs, we must have a template family that allows for deviations from GR in the data. Given these templates, we then have two models to describe the GWs we will detect, which can broadly be described as GR and not GR.

The template family that we have used as the not GR model in was developed under the parameterized post-Einsteinian (ppE) formalism. These templates introduce parameterized deviations from GR in both the amplitude and the phase of GWs. These non-GR parameters can then be measured using GW data. If they are found to have a value that is not consistent with GR, we can claim the detection of a deviation from GR.

The ppE templates are not the first parameterization of non-GR gravity that have been developed. In order to take advantage of Solar System experiments to test GR, Nordtvedt and Will developed the parameterized post-Newtonian (ppN) system in the 1970s. Similarly, in order to use measurements of binary pulsar systems, scientists in the 90s and early 2000s developed the parameterized post-Keplerian (ppK) system. Although they are optimized for different situations, the ppE, ppN, and ppK formalisms are all ultimately testing the same underlying theory of gravity. In Chapter 5, we developed this connection.

The rest of the original work in this thesis has focused on using the ppE templates to test GR, from a data analysis perspective. This work is summarized in Chapter 7. The first step was to determine what types of deviations from GR would be detectable using the advanced LIGO (aLIGO) detector network. We accomplished this by injecting GR signals and characterizing them using ppE templates and MCMC search methods. The recovered ranges of the non-GR parameters then indicate which values of those parameters are consistent with GR. This range also tells us how strongly these parameters could be constrained by GW measurements using ppE templates. We found, not surprisingly, that for low-order deviations from GR, i.e. those that are significant even in systems where the velocity is low, the constraints are much stronger than for high-order deviations. Significantly, however, we found that high-order deviations from GR will be better constrained by GW detections than by any current experimental bounds [38].

This last point is significant for other reasons, as well. Because current experiment only constrains deviations from GR in the low-velocity limit, it is fully possible that there are large deviations in the high-velocity regime. In the context of GWs, this means that there could be large differences between the GR prediction and the actual signal at high post-Newtonian (PN) order in the phase. In [116], we have shown that there are, in fact, signals that are perfectly consistent with all current experimental bounds that would be so poorly matched by GR templates that they could be missed entirely by a standard data analysis pipeline that assumes GR is correct. This can be ameliorated, however, by using a simple ppE template instead.

The simplest form of the ppE templates includes just one correction to the phase of GWs, characterized by a strength parameter and a PN order. Just like the PN expansion in GR, however, we know that any alternative theory of gravity will lead to an infinite number of phase corrections. This raises the question of what type of template is ideal for detecting non-GR signals - is it better to use a template that includes multiple corrections to the phase? In [115], we answered this question using

Bayesian model selection by injecting a non-GR signal that included multiple changes to the GR phase, and recovering it using ppE templates with either one or two phase corrections. We found that there was no strong preference between the different ppE models, and so a simple, one-parameter template family should be sufficient in using to test GR.

In addition to being able to detect deviations from GR that are truly contained in our GW data, it is also necessary to be confident that we are not claiming a detection of a deviation from GR when none are present. Of particular concern is the fact that, lacking knowledge of merger and ringdown in non-GR theories, we often use inspiral-only templates when testing GR. Inspiral-only is a rather ambiguous term, however, as it is not at all clear where the distinction between inspiral and merger should be made. Furthermore, even if we are using inspiral-only templates in our searches, the data will still contain the merger and ringdown signals from the system we have detected. In [116], we examined the effect of using inspiral-only templates to characterize signals that contain merger and ringdown. We found that it is indeed possible to falsely measure a deviation from GR when using this technique. We also found, however, that there is a simple technique that can be applied when characterizing inspiral-merger-ringdown signals that eliminates this danger.

In general, this work has shown that there are extremely bright prospects for testing GR with GW observations. In the next few years the first GWs will be detected, and we can begin the work of either vindicating Einstein's theory - or proving that it is incomplete.

REFERENCES CITED

[1] Abadie, J. et al. 2012, Phys.Rev., D85, 082002

[2] Abbott, B. et al. 2009, Phys.Rev., D79, 122001

[3] —. 2009, Phys.Rev., D80, 047101

[4] Ajith, P., Hannam, M., Husa, S., Chen, Y., Bruegmann, B., et al. 2011, Phys.Rev.Lett., 106, 241101

[5] Alexander, S., Finn, L. S., & Yunes, N. 2008, Phys. Rev, D78, 066005

[6] Alexander, S. & Yunes, N. 2009, Phys.Rept., 480, 1

[7] Alsing, J., Berti, E., Will, C. M., & Zaglauer, H. 2012, Phys.Rev., D85, 064041

[8] Amaro-Seoane, P., Aoudia, S., Babak, S., Binetruy, P., Berti, E., et al. 2012, Class.Quant.Grav., 29, 124016

[9] Anderson, W., Brady, P., Chin, D., Creighton, J., Riles, K., & Whelan, J. 2002, LIGO Techinal Report, LIGO-T010110-00-Z

[10] Arun, K. G., Iyer, B. R., Qusailah, M. S. S., & Sathyaprakash, B. S. 2006, Phys. Rev., D74, 024006

[11] —. 2006, Class. Quant. Grav., 23, l37

[12] Arun, K. G. & Will, C. M. 2009, Class. Quant. Grav., 26, 155002

[13] Arun, K. G. & Will, C. M. 2009, Classical and Quantum Gravity, 26, 155002

[14] Azevedo-filho, A. 1994, in In de Mantaras (Morgan Kaufmann), 28–36

[15] Baker, J. G., Centrella, J., Choi, D.-I., Koppitz, M., & van Meter, J. 2006, Phys. Rev. Lett., 96, 111102

[16] Barausse, E., Palenzuela, C., Ponce, M., & Lehner, L. 2013, Phys.Rev., D87, 081506

[17] Bekenstein, J. D. 2004, Phys. Rev., D70, 083509

[18] Berti, E., Buonanno, A., & Will, C. M. 2005, Phys. Rev., D71, 084025

[19] —. 2005, Class. Quant. Grav., 22, S943

[20] Berti, E., Cardoso, V., & Starinets, A. O. 2009, Class. Quant. Grav., 26, 163001

[21] Berti, E., Gualtieri, L., Horbatsch, M., & Alsing, J. 2012, Phys.Rev., D85, 122005

[22] Bertotti, B., Iess, L., & Tortora, P. 2003, Nature, 425, 374

[23] Blanchet, L. 2006, Living Reviews in Relativity, 9

[24] —. 2013

[25] Blandford, R. & Teukolsky, S. 1976, Astrophysical Journal, 205

[26] Bojowald, M. & Hossain, G. M. 2008, Phys.Rev., D77, 023508

[27] Bose, S., Gopakumar, A., & Tessmer, M. 2008

[28] Brans, C. & Dicke, R. H. 1961, Phys. Rev., 124, 925

[29] Burgay, M. & Others. 2003, Nature., 426, 531

[30] Campanelli, M., Lousto, C. O., Marronetti, P., & Zlochower, Y. 2006, Phys. Rev. Lett., 96, 111101

[31] Canitrot, P. 2001, Phys. Rev. D, 63, 082005

[32] Cardoso, V., Chakrabarti, S., Pani, P., Berti, E., & Gualtieri, L. 2011, Phys.Rev.Lett., 107, 241101

[33] Chatziioannou, K., Klein, A., Yunes, N., & Cornish, N. 2013

[34] Chatziioannou, K., Yunes, N., & Cornish, N. 2012, Phys.Rev., D86, 022004

[35] Choudhury, S., Joshi, G. C., Mahajan, S., & McKellar, B. H. 2004, Astropart.Phys., 21, 559

[36] Collaboration, L. S. 2013, Schematic of LIGO detector

[37] Copeland, E. J., Sami, M., & Tsujikawa, S. 2006, Int.J.Mod.Phys., D15, 1753

[38] Cornish, N., Sampson, L., Yunes, N., & Pretorius, F. 2011, Phys.Rev., D84, 062003

[39] Cornish, N. J. 2010

[40] Cornish, N. J. & Littenberg, T. B. 2007, Physical Review D (Particles, Fields, Gravitation, and Cosmology), 76, 083006

[41] Cornish, N. J. & Rubbo, L. J. 2003, Phys. Rev., D67, 022001

[42] Cutler, C. 1998, Phys. Rev., D57, 7089

[43] Cutler, C. & Vallisneri, M. 2007, Phys. Rev., D76, 104018

[44] Damour, T. & Deruelle, N. 1985, Annales de l'institut Henri Poincare (A) Physique theoretique, 43, 107

[45] Damour, T. & Esposito-Farese, G. 1992, Class. Quantum Grav., 9, 2093

[46] —. 1998, Phys. Rev., D58, 042001

[47] Damour, T., Iyer, B. R., & Sathyaprakash, B. S. 2000, Phys. Rev. D, 62, 084036

[48] Damour, T. & Taylor, J. H. 1991, NASA STI/Recon Technical Report N, 92, 19818

[49] —. 1992, Phys. Rev. D, 45, 1840

[50] Deffayet, C., Dvali, G., Gabadadze, G., & Vainshtein, A. I. 2002, Phys.Rev., D65, 044026

[51] Del Pozzo, W., Veitch, J., & Vecchio, A. 2011, ArXiv e-prints

[52] Detweiler, S. 1979, ApJ, 234, 1100

[53] Dicke, R. 1960, ApJ, 28, 344

[54] Dickey, J. M. 1971, Ann. Math. Statist., 42, 204

[55] Dreyer, O. et al. 2004, Class. Quant. Grav., 21, 787

[56] Dvali, G., Gabadadze, G., & Porrati, M. 2000, Phys.Lett., B485, 208

[57] Einstein, A. 1915, Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin, 844

[58] Epstein, R. 1977, ApJ, 216, 92

[59] Faraoni, V. & Gunzig, E. 1999, Int.J.Theor.Phys., 38, 217

[60] Faraoni, V., Gunzig, E., & Nardone, P. 1999, Fund.Cosmic Phys., 20, 121

[61] Felice, A. D. & Tsujikawa, S. 2010, Living Reviews in Relativity, 13

[62] Ferdman, R., van Haasteren, R., Bassa, C., Burgay, M., Cognard, I., et al. 2010, Class.Quant.Grav., 27, 084014

[63] Freire, P. C. & Wex, N. 2010

[64] Freise, A. & Strain, K. A. 2010, Living Reviews in Relativity, 13

[65] Gair, J. R. & Porter, E. K. 2012

[66] Garay, L. J. & Garcia-Bellido, J. 1993, Nucl.Phys., B400, 416

[67] Girolami, M. & Calderhead, B. 2011, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 73, 123

[68] Godone, A., Novero, C., & Tavella, P. 1995, Phys.Rev., D51, 319

[69] Goldhaber, A. S. & Nieto, M. M. 1974, Phys. Rev. D, 9, 1119

[70] Gossan, S., Veitch, J., & Sathyaprakash, B. 2012, Phys.Rev., D85, 124056

[71] GREEN, P. J. 1995, Biometrika, 82, 711

[72] Hannam, M. 2009, Class.Quant.Grav., 26, 114001

[73] Haugan, M. P. 1985, ApJ, 296, 1

[74] Hayama, K. & Nishizawa, A. 2013, Phys.Rev., D87, 062003

[75] Heckel, B. 2013, Plus and Cross Polarizations

[76] Hellings, R. W. 1978, Phys. Rev. D, 17, 3158

[77] Hinder, I. 2010, Class.Quant.Grav., 27, 114004

[78] Iorio, L. 2007, Astrophys.Space Sci., 312, 331

[79] Jacobson, T. 2008

[80] Jeffreys, H. 1951, Zeitschrift Naturforschung Teil A, 6, 471

[81] Keppel, D. & Ajith, P. 2010, Phys. Rev., D82, 122001

[82] Key, J. S. & Cornish, N. J. 2010

[83] Kocsis, B., Yunes, N., & Loeb, A. 2011, Phys.Rev., D84, 024032

[84] Kogan, I. I., Mouslopoulos, S., & Papazoglou, A. 2001, Phys.Lett., B503, 173

[85] Kramer, M., Stairs, I. H., Manchester, R. N., McLaughlin, M. A., Lyne, A. G., Ferdman, R. D., Burgay, M., Lorimer, D. R., Possenti, A., D'Amico, N., Sarkissian, J. M., Hobbs, G. B., Reynolds, J. E., Freire, P. C. C., & Camilo, F. 2006, Science, 314, 97

[86] Kramer, M. et al. 2006, Science, 314, 97

[87] Lanahan-Tremblay, N. & Faraoni, V. 2007, Class.Quant.Grav., 24, 5667

[88] Lee, K. J., Jenet, F. A., & Price, R. H. 2008, ApJ, 685, 1304

[89] Li, T., Del Pozzo, W., Vitale, S., Van Den Broeck, C., Agathos, M., et al. 2012, Phys.Rev., D85, 082003

[90] Lindblom, L. 2009

[91] Lindblom, L., Baker, J. G., & Owen, B. J. 2010, Phys.Rev., D82, 084020

[92] Littenberg, T. B. 2009, PhD thesis, Montana State University, Bozeman

[93] Littenberg, T. B. & Cornish, N. J. 2009, Phys. Rev., D80, 063007

[94] Maggiore, M. 2008, Gravitational waves: theory and experiments (Oxford: Oxford Univ. Press)

[95] Manchester, R. 2013

[96] Manchester, R., Hobbs, G., Bailes, M., Coles, W., van Straten, W., et al. 2012

[97] McLaughlin, M. 2013

[98] Merkowitz, S. M. 2010, Living Reviews in Relativity, 13

[99] Milgrom, M. 1983, ApJ, 270, 371

[100] Milgrom, M. 1983, ApJ, 270, 365

[101] Mishra, C. K., Arun, K. G., Iyer, B. R., & Sathyaprakash, B. S. 2010

[102] Nissanke, S., Holz, D. E., Hughes, S. A., Dalal, N., & Sievers, J. L. 2010, Astrophys.J., 725, 496

[103] Nordtvedt, K. 1968, Phys. Rev., 169, 1017

[104] Nordtvedt, K. J. & Will, C. M. 1972, ApJ, 177, 775

[105] Owen, B. J. 1996, Phys. Rev. D, 53, 6749

[106] Pan, Y., Buonanno, A., Boyle, M., Buchman, L. T., Kidder, L. E., et al. 2011, Phys.Rev., D84, 124052

[107] Pitkin, M., Reid, S., Rowan, S., & Hough, J. 2011, Living Reviews in Relativity, 14

[108] Pretorius, F. 2005, Phys. Rev. Lett., 95, 121101

[109] Prince, T. A., Tinto, M., Larson, S. L., & Armstrong, J. W. 2002, Phys. Rev., D66, 122002

[110] Rosen, N. 1974, Annals of Physics, 84, 455

[111] Rubbo, L. J., Cornish, N. J., & Poujade, O. 2004, Phys. Rev., D69, 082003

[112] Ruggiero, M. L. & Tartaglia, A. 2005, Phys.Rev., D72, 084030

[113] Salgado, M., Rio, D. M.-d., Alcubierre, M., & Nunez, D. 2008, Phys.Rev., D77, 104010

[114] Sambridge, M., Gallagher, K., Jackson, A., & Rickwood, P. 2006, Geophysical Journal International, 16, 528

[115] Sampson, L., Cornish, N., & Yunes, N. 2013, Phys.Rev., D87, 102001

[116] —. 2013, Manuscript Submitted for Publication

[117] Santamaria, L., Ohme, F., Ajith, P., Bruegmann, B., Dorband, N., et al. 2010, Phys.Rev., D82, 064016

[118] Sathyaprakash, B. & Schutz, B. F. 2009, Living Reviews in Relativity, 12

[119] Scharre, P. D. & Will, C. M. 2002, Phys. Rev., D65, 042002

[120] Shapiro, I. 1964, Phys. Rev. Lett., 13, 789

[121] Shibata, M., Taniguchi, K., Okawa, H., & Buonanno, A. 2013

[122] Soldner, J. 1804, Berliner Astronomisches Jahrbuch, 161

[123] Sopuerta, C. F. 2010, GW Notes, Vol. 4, p. 3-47, 4, 3

[124] Sopuerta, C. F. & Yunes, N. 2009, Physical Review D (Particles, Fields, Gravitation, and Cosmology), 80, 064006

[125] Stairs, I. H. 2003, Living Reviews in Relativity, 6

[126] Stavridis, A. & Will, C. M. 2009, Phys. Rev., D80, 044002

[127] Talmadge, C., Berthias, J. P., Hellings, R. W., & Standish, E. M. 1988, Phys. Rev. Lett., 61, 1159

[128] Taylor, J. H. & Weisberg, J. M. 1982, , 253, 908

[129] Ter Braak, C. J. 2006, Statistics and Computing, 16, 239

[130] Ter Braak, C. J. & Vrugt, J. A. 2008, Statistics and Computing, 18, 435

[131] Tinto, M. & da Silva Alves, M. E. 2010, Phys. Rev., D82, 122003

[132] V., L. D. 1972, Bayesian statistics: a review (Philadelphia, USA: SIAM)

[133] Vainshtein, A. 1972, Physics Letters B, 39, 393

[134] Vallisneri, M. 2011, Phys.Rev.Lett., 107, 191104

[135] Verdinelli, I. & Wasserman, L. 1995, Journal of the American Statistical Association, 90, pp. 614

[136] Wagner, T., Schlamminger, S., Gundlach, J., & Adelberger, E. 2012, Class.Quant.Grav., 29, 184002

[137] Weinberg, M. D. 2009

[138] Weisberg, J. M., Nice, D. J., & Taylor, J. H. 2010, , 722, 1030

[139] Wex, N. & Kramer, M. 2007, Mon.Not.Roy.Astron.Soc.

[140] Will, C. M. 1971, ApJ, 163, 611

[141] Will, C. M. 1993, Theory and experiment in gravitational physics (Cambridge, UK: Cambridge University Press)

[142] —. 1994, Phys. Rev., D50, 6058

[143] —. 1998, Phys. Rev., D57, 2061

[144] —. 2006, Living Reviews in Relativity, 9

[145] Will, C. M. & Nordtvedt, K. J. 1972, ApJ, 177, 757

[146] Will, C. M. & Poisson, E. 2011, book in progress

[147] Will, C. M. & Yunes, N. 2004, Class. Quant. Grav., 21, 4367

[148] Williams, J., Turyshev, S., & Boggs, D. 2005, Phys.Rev., Lett93

[149] Wyithe, J. S. B. & Loeb, A. 2003, Astrophys.J., 590, 691

[150] Yagi, K., Stein, L. C., Yunes, N., & Tanaka, T. 2012, Phys.Rev., D85, 064022

[151] —. 2013, Phys.Rev., D87, 084058

[152] Yagi, K. & Tanaka, T. 2009

[153] Yagi, K. & Yunes, N. 2013

[154] —. 2013

[155] Yoder, C., Williams, J., Dickey, J., B.E., S., R.J., E., & B.D., T. 1983, Nature, 303

[156] Yunes, N., Arun, K. G., Berti, E., & Will, C. M. 2009

[157] Yunes, N., Coleman Miller, M., & Thornburg, J. 2011, Phys.Rev., D83, 044030

[158] Yunes, N. & Hughes, S. A. 2010, Phys.Rev., D82, 082002

[159] Yunes, N. & Hughes, S. A. 2010, Phys. Rev. D, 82, 082002

[160] Yunes, N., Kocsis, B., Loeb, A., & Haiman, Z. 2011, Phys.Rev.Lett., 107, 171103

[161] Yunes, N., O'Shaughnessy, R., Owen, B. J., & Alexander, S. 2010, Phys. Rev., D82, 064017

[162] Yunes, N. & Pretorius, F. 2009, Physical Review D (Particles, Fields, Gravitation, and Cosmology), 79, 084043

[163] —. 2009, Phys.Rev., D80, 122003

[164] Yunes, N., Pretorius, F., & Spergel, D. 2009

[165] Yunes, N. & Siemens, X. 2013

[166] Yunes, N. & Stein, L. C. 2011, Phys.Rev., D83, 104002