Fachbereich 13
Institut für Theoretische Physik
Goethe-Universität Frankfurt am Main

Dirty Extreme Mass Ratio Inspiral

On the significance of environmental effects

Abschlussarbeit zur Erlangung des akademischen Grades
Master of Science in Physik

vorgelegt von Lukas Kerim Arda, geb. am 30.05.2000 in
Frankfurt am Main

Erstgutachterin: Prof.in Dr.in Laura Sagunski
Zweitgutachter: Dr.in Natalia Korsakova

Eingereicht am 05.09.2025

Contents

Throughout this thesis, we use geometrised units with $c = G = 1$ and the Einstein summation convention.

This thesis was written with the help of an artificial intelligence writing assistant, which was used to refine text passages[18].

Acronyms

`FEW FastEMRIWaveforms`

AGN Active Galactic Nucleus

BH Black Hole

CO Compact Object

EMRI Extreme Mass Ratio Inspiral

ESA European Space Agency

GR General Relativity

GRMHD General Relativistic Magnetohydrodynamics

GSF Gravitational Self-Force

GW Gravitational Wave

IMRI Intermeditate Mass Ratio Inspiral

ISCO Innermost Stable Circular Orbit

LISA Laser Interferometer Space Antenna

MBH Massive Black Hole

MCMC Monte Carlo Markov Chain

ODE Ordinary Differential Equation

PSD Power Spectral Density

RJ Reversible Jump

SNR Signal-to-Noise-Ratio

SSB Solar System Barycenter

TDI Time Delay Interferometry

1 Abstract

This thesis investigates the impact of environmental effects on Extreme Mass Ratio Inspirals (EMRIs), which will be observed by the future space-based gravitational wave observatory Laser Interferometer Space Antenna (LISA). Focusing on type I planetary(-like) migration in thin gaseous accretion disks and deviations from General Relativity (GR) allowing for a varying gravitational constant, we introduce an agnostic phenomenological parameterisation capable of modelling both effects. We extend existing waveforms for prograde, equatorial, and circular orbits by the environmental effect model, to create "dirty" EMRI. We perform Bayesian parameter estimation on simulated LISA data and use Monte Carlo Markov Chain (MCMC) integration to estimate evidence for Bayesian model selection. Our results demonstrate that longer observation times prior to merger enhance the detectability of environmental effects. The considered effect leaves a stronger imprint at larger separations, as quantified by the accumulated orbital phase shift. Constraints on the phenomenological model parameters enable characterisation of the accretion disk's physical properties and distinction between accretion disks and GR deviation. The Bayes factors reveal that stronger effects provide decisive support for the environmentally perturbed templates when the posterior distributions of the vacuum templates are viable. We have established dephasing as a measure of the environmental effect's strength, linking it to the threshold at which the effects become decisive in model selection.

2 Introduction

In 2015, the field of observational astrophysics underwent a paradigm shift with the first measurement of Gravitational Waves (GWs). These tiny spacetime fluctuations, corresponding to relative length contractions of $\approx 10^{-21}$, were deemed impossible to measure one century earlier, but thanks to the collaborative efforts of the scientific community, they were detected on 14th of September 2015[1]. Measuring GW enables us to perceive the universe's gravitational radiation, as if a new 'sense' has been revealed to us. Accurately measuring, understanding and interpreting GWs is one of the major challenges in modern physics.

2.1 Gravitational Waves

GR links the curvature of spacetime to mass and energy. In a simplified model, we can imagine masses curving spacetime in a similar way that bowling balls curve a bedsheet. For example, when two masses orbit each other, the spacetime geometry is periodically perturbed and GWs are emitted. These waves carry away orbital energy and angular momentum, causing the orbits to decay and the two masses to inspiral towards each other until they merge. The coalescence of two compact objects of similar mass is the class of sources that scientists first detected in 2015 [1], and that we can currently observe using the ground-based interferometers LIGO, VIRGO and KAGRA. However, this is just the beginning of GW astronomy. More GW sources are predicted, such as supernovae, non-axisymmetric rotating (neutron) stars, and stochastic signals from cosmology, which may originate from inflation or first-order phase transitions in the early universe, or from cosmic strings. Fig. 1 provides an overview of the sources of GWs and their frequency ranges, as well as of the current and future detectors.

By observing and interpreting GWs, we can refine and advance current models and explore new theories.

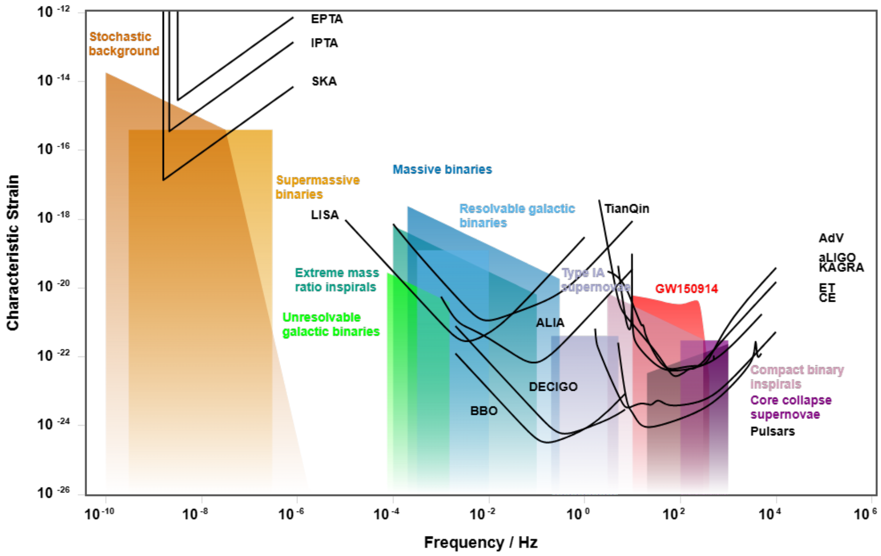

Figure 1: The ranges of GW sources and observatories. The multiple sources cover a wide frequency range, with different detectors and observatories being sensitive to different frequency intervals. (Plot created with `gwplotter.com` [44])

Theoretical Background Overview

The wave nature of GWs can be derived by a assuming a flat background metric $\eta_{\mu\nu} = \mathrm{diag}(-,+,+,+)$ with a small metric perturbation $h_{\mu\nu}$, such that the metric of the spacetime is $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$. Using this metric, we can derive a wave equation from the linearised Einstein equations

$$\Box \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}, \tag{1}$$

where the operator $\Box = \partial_\mu \partial^\mu$ is the flat space d'Alembert operator and $\bar{h}_{\mu\nu}$ is the *trace-reversed metric perturbation* [43]

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

in Lorenz gauge, where $h = \eta^{\mu\nu} h_{\mu\nu}$. In vacuum far from the source ($T_{\mu\nu} = 0$), the transverse-traceless gauge can be chosen ($h^{0\mu} = 0, h^i_i = 0, \partial^j h_{ij} = 0$, with i, j

being the spatial indices) to find plane wave solutions to $\Box\bar{h}_{\mu\nu} = 0$ for a wave propagating along the z-direction at the speed of light, which are

$$h_{ij}^{TT} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij} \cos(\omega(t-z)) \tag{3}$$

with the two GW polarization modes $h_+, h_\times$. GWs are generated by solving eq. 1 with a non-zero stress-energy tensor via a multipole expansion. Unlike in electrodynamics, the leading contribution to this expansion is the quadrupole term, which is represented with the quadrupole tensor Q_{ij}. This encodes the asymmetric, time-dependent mass distribution of the GW source. The quadrupole formula connects the quadrupole tensor and the *trace-reversed metric perturbation*, which is linked to the GW strain,

$$\bar{h}_{ij} = \frac{2}{r}\ddot{Q}_{ij}(t-r). \tag{4}$$

Only anisotropic sources, which have a non-zero and non-varying quadrupole tensor, can emit GW radiation.

Principle of GW Interferometry

GWs squeeze and stretch spacetime, effectively altering the geodesics of photons. This alteration can be measured using interferometers. In a Michelson interferometer, light from a single source is split by a beam splitter and directed into two equal-length arms. The light is then reflected at each end of the arms and recombined before reaching a detector. In the absence of GWs, the laser beams remain in phase when they recombine, suppressing the laser frequency noise common to both arms and cancelling out the phase difference. In the presence of GW, the armlengths change, and by measuring the interference pattern at the recombination, GWs can be detected. This approach has been refined since the 1970s, is currently employed in ground-based detectors such as LIGO, VIRGO and KAGRA and enabled the first direct detection of GWs in 2015 [1]. As shown in Fig. 1, these detectors cover the frequency range between 10 Hz and 10^3 Hz.

2.2 Laser Interferometer Space Antenna

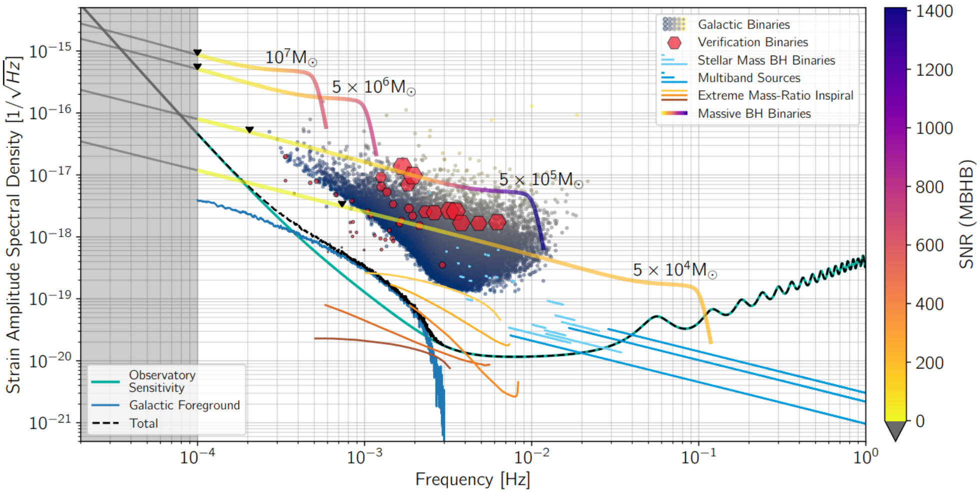

Figure 2: LISA's sensitivity and different primary source classes. LISA will be able to measure GW signals between 10^{-4} Hz and 1 Hz, covering a frequency band with a variety of sources. These include massive Black Hole (BH) binaries, EMRI, stellar-mass Black Hole (BH), Galactic binaries and verification binaries.(Reprinted from [16])

In order to observe more sources of GWs, we need to access other frequency bands with new detectors (see Fig. 1). One such band is that of the future space-based GW observatory: LISA. Measuring GWs in space enables us to avoid the limitations of terrestrial noise and significantly increase the interferometer's armlength. This will allow measurements at frequencies between 10^{-4} Hz and 1 Hz, corresponding to new classes of sources such as supermassive black hole coalescences, EMRIs, compact binary stars and cosmological signals (see Fig. 2). LISA will consist of three spacecraft on heliocentric orbits trailing Earth by 20 degrees and forming an approximately equilateral triangle with an armlength of approximately 2.5 million km (see Fig. 3). Laser beams will link the spacecraft, and each will host two freely falling test masses, one at each end of the three arms of the constellation.

Such a setup requires high-precision optical and space engineering, but the LISA Pathfinder mission in 2017 has proven its feasibility[3]. Today, LISA is

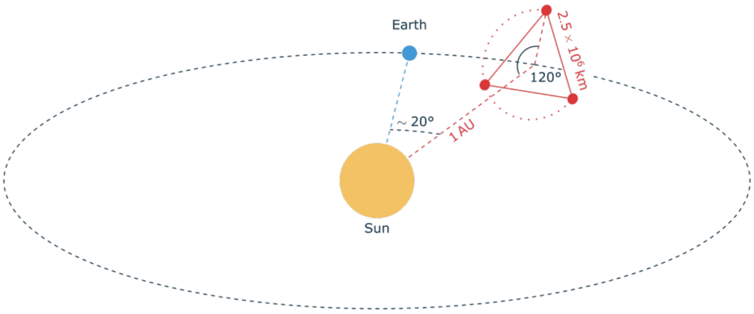

(a) Position of LISA's orbit

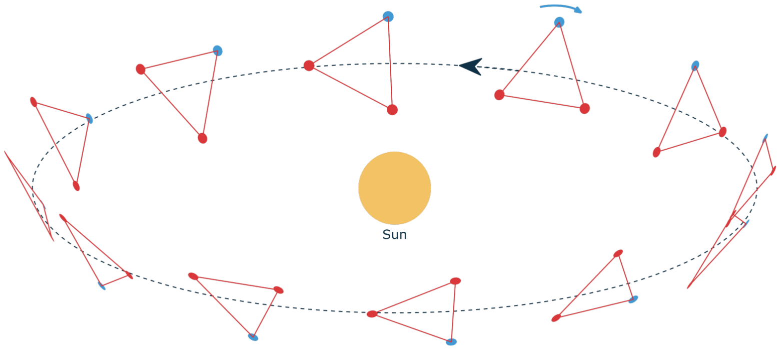

(b) Dynamics of LISA's orbit

Figure 3: The position and dynamics of LISA's orbit. Fig. 3(a) shows how the three spacecraft form an equilateral triangle, with the barycentre following Earth's orbit around the Sun. The plane of the triangle "is inclined at 60° with respect to the Ecliptic"[16]. Fig. 3(b) displays the annual clockwise rotation of the spacecraft within the constellation's plane. (Reprinted from [16])

European Space Agency (ESA)'s L3 mission, and was adopted in January 2024, with an expected launch date of 2035.

One particular challenge for LISA is the immense distances and the fact that the armlengths are unequal and vary over time (< 12m/s[16]). The distances of millions of kilometres reduce the light intensity by $\approx 10^{-9}$[16] such that reflected light would be undetectable, making a standard Michelson interferometer setup infeasible. In LISA there are three spacecraft, six optical benches, six lasers and six test masses. Each one-way link between spacecraft has its own light source, optical bench and test mass, and the interference pattern is recorded between the outgoing and received light. For such a setup where laser frequency noise doesn't cancel due to variation of arm length, laser frequency noise cancellation needs to be adapted to enable sufficient sensitivity for GW observation. The algorithm implemented for this purpose in post-processing is called Time Delay Interferometry (TDI) and has been extensively studied and probed[56]. It combines knowledge of the two phase measurements at each end of an arm in the LISA constellation, with precisely synchronised atomic clocks and time delays due to the light travel time between the spacecraft, constructing a virtual equal armlength interferometer. Therefore, it requires accurate estimates of the armlengths ($\approx$ 1m) and can then accommodate the dynamics and orbital motion of the constellation.

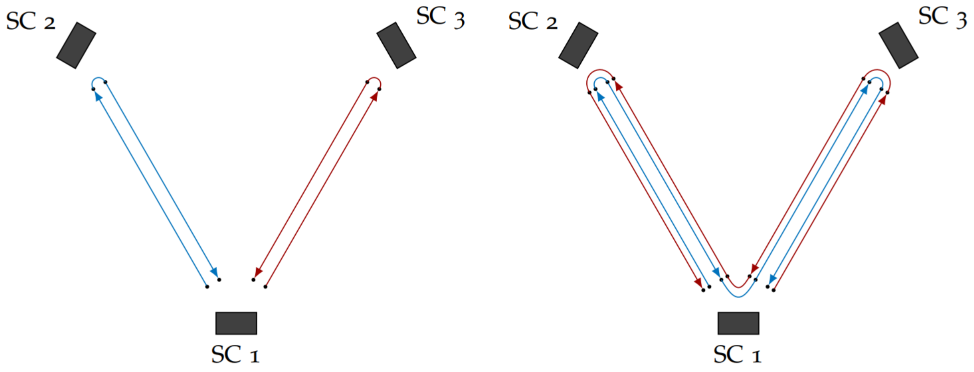

Figure 4: Schematics of TDI relaying. The figure shows the first (left) and second (right) generation approaches to TDI relaying. As the laser light is too weak to be reflected directly, it must be detected and actively re-emitted at each end of an arm. Schemes that incorporate TDI combinations with further relaying (more round-trips) can suppress noise further, with the 2nd generation of TDI displayed in the right panel being sufficient for the required sensitivity[16]. (Reprinted from [30])

In practice, each laser beam is split twice. Two of these beams leave the spacecraft and travel to the other two, where two new beams with the same phase are emitted and sent back. This process relies on sufficient phase locking between the incoming and relayed beams. The incoming beams are time-delayed by the effective round-trip travel time. They are then recombined with the local laser beam (see Fig. 4). In the classic TDI algorithm, the phase difference of each pair of adjacent round-trip laser beams forms an observable, called the TDI observables $\{X, Y, Z\}$. While these TDI observables have correlated noise properties, a set of uncorrelated TDI observables $\{A, E, T\}$ can be obtained by linear transformations [57]

$$A = \frac{1}{\sqrt{2}}(Z - X) \tag{5}$$

$$E = \frac{1}{\sqrt{6}}(X - 2Y + Z) \tag{6}$$

$$T = \frac{1}{\sqrt{3}}(X + Y + Z). \tag{7}$$

These TDI observables $\{A, E, T\}$ are the product of the L1 data analysis step in LISA "constitut[ing] the lowest-level data ready for scientific interpretation" [16]. However, the real optical bench setup and algorithm are more complex, accounting for phenomena such as orbital motion [37], the motion of the optical bench relative to the proof mass, and fluctuations between the lasers within a spacecraft.

2.3 Extreme Mass Ratio Inspirals

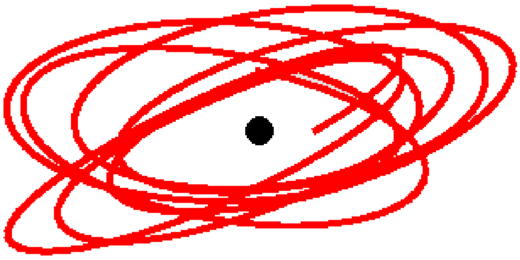

(a) EMRI orbit

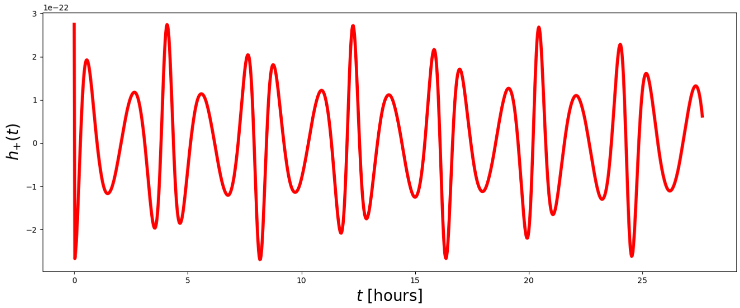

(b) EMRI GW strain

Figure 5: The complex nature of EMRI orbits and their corresponding GW signals. Fig. 5(a) shows the intriguing orbits of the Compact Object (CO) as it inspirals into the Massive Black Hole (MBH). These orbits can be retrograde, inclined and eccentric. This complex trajectory can be observed in the corresponding GW signal in 5(b). The orbital frequency, the perihelion precession frequency and the orbital plane precession frequency all evolve over time, resulting in many harmonics and a highly intricate signal structure. (Plots created with `github.com/OllieBurke/animations`)

Following the 2020 Nobel Prize observations of stellar dynamics in galactic nuclei and the imaging of black hole shadows by the Event Horizon Telescope, it is widely accepted that Massive Black Holes (MBHs) reside at the centres of most galaxies. Due to mass segregation, MBHs are surrounded by a dense cloud of stellar remnants, such as stellar-mass black holes, neutron stars and possibly white dwarfs, which we collectively call Compact Objects (COs). As

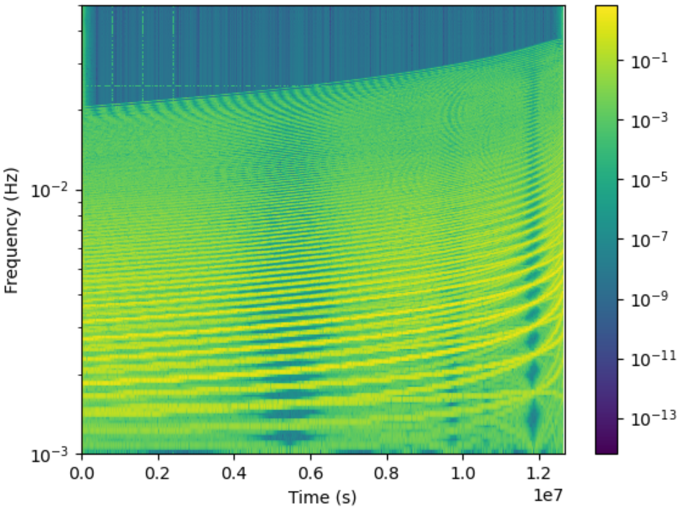

Figure 6: Time-frequency spectrogram of a typical EMRI signal. Shown here is the time-frequency decomposition of the h_+ polarisation mode. The tracks indicate that the frequencies are not constant throughout the evolution of the system. Some harmonic frequencies increase and form the typical chirp of the signal, while some decrease.

a result of their mutual interaction, a CO can be thrown towards a MBH with a small impact parameter. It loses non-negligible orbital energy in GW bursts at the periapsis, which modifies its trajectory, making it more bound. After several passages, the CO is detached from the stellar cloud (cusp) and forms a bound two-body system where the CO slowly spirals towards the MBH until the merger; an EMRI is formed.

The mass ratios of EMRIs detectable by LISA range from 10^{-7} to 10^{-4}. EMRIs are key targets for LISA, whose observations will enable the study of MBHs. We will be able to measure their properties (mass and spin) with sub-percent accuracy and constrain the lower end of the MBH mass distribution ($10^4 - 10^7$ solar masses). With the expected precision of LISA's measurements, we will be able to localise the source in the sky to within a few square degrees and study the co-evolution of the MBH and its host galaxy throughout cosmic history. The CO acts as a probe of the strongly curved spacetime around the central

MBH and detecting an EMRI would therefore provide a better understanding of the strong gravity regime and allow us to test the “no-hair” theorem of GR. Gravitational wave signals of an EMRI generally comprise $\approx 10^5$ cycles, corresponding to observation times of years. However, their Signal-to-Noise-Ratio (SNR) is relatively low, and the signal's structure is highly complex, because the CO inspirals rapidly and can perform inclined, retrograde and eccentric orbits (see Fig. 5(a)). The GW signal then contains the evolution of the harmonics of three fundamental frequencies: the orbital frequency, the perihelion precession frequency and the orbital plane precession frequency (see Fig. 5(b)). This evolution is most clearly displayed in time-frequency spectrograms, such as the one shown in Fig. 6.

In the "clean" EMRI model, the CO orbits the MBH in isolation, unaffected by external forces. This is a good approximation for most EMRIs for the final years of the inspiral within the vicinity of the MBH. In this project, we aim to consider “dirty” EMRIs, where “dirty” implies the presence of a perturber influencing the orbital dynamics of the CO around the MBH that could potentially be measured.

Aims & Structure

The aim of this work is to assess the significance of environmental effects on EMRIs and to determine when models including these effects are favoured by parameter estimation. This thesis is structured as follows. First, I introduce the fundamental principles of GWs and discuss the properties and scientific potential of EMRI systems. Next, I establish the external effect influencing these systems and explain how the resulting signals are modelled. Then, I present the statistical analysis framework, its implementation, and the obtained results. Finally, I will summarise my work and provide an outlook for future research.

3 Environmental effects

GW astronomy enables us to investigate COs, such as Neutron Stars, Black Holes (BHs) and White Dwarfs, in a new way. With increasingly accurate waveform templates, we can estimate the properties of these objects with great precision, and test GR. However, current waveform templates neglect the realistic astrophysical environment and are referred to as vacuum templates. As astrophysical environments influence CO binaries on all scales, they will impact the GW signal. There are various ways in which the environment could influence these binaries, and some of these effects could be strong enough to modify the observed signal such that this effect will be measurable, and should therefore be accounted for in the waveform model.

In this chapter, we first provide a brief overview of the possible effects, before motivating and discussing the class of effects that we have employed.

Dirty EMRI

EMRI are especially susceptible to environmental effects as they perform many orbital cycles and map out large areas of the MBH vicinity[16]. The EMRI population is still relatively unconstrained, and EMRI detection rate estimates in the LISA band range from a few to thousands per year[7]. Nevertheless, one of LISA's science objectives is to "[s]tudy the properties and immediate environment of Milky Way-like MBHs using EMRIs"[16]. Such environments could include overly dense regions, such as populations of cores of massive stars or accretion disks of Active Galactic Nucleus (AGN), star clusters, dark matter halos or strong electromagnetic fields [9]. Detailed modelling of the different environmental effects would require computationally expensive *Numerical Relativity* simulations; however, some of the effects can be approximated as a correction to the vacuum GW signal's phase evolution. Tab. 1 provides an overview of the phase correction to the vacuum GW signal for some environmental effects in quasi-circular EMRI.

Correction	**$\Delta\Phi$ [rad]**
planetary migration	$< 10^4$
thin accretion disks (DF)	$\leq 10^2$
thin accretion disks (GP)	$\leq 10^{-3}$
magnetic field	10^{-4}
charge	10^{-2}
gas accretion onto the central BH	10^{-2}
thick accretion disks (DF)	10^{-9}
Dark Matter accretion onto central BH	10^{-8}
thick accretion disks (GP)	10^{-11}
Dark Matter distribution (DF)	10^{-14}
Dark Matter distribution $\rho \sim r^{-\hat{\alpha}}$ (GP)	10^{-16}
galactic Dark Matter halos	10^{-16}
cosmological constant	10^{-26}

Table 1: Table of corrections due to environmental effects. These estimates were calculated under the assumption of quasi-circular EMRI. The different effects cover a wide range of phase corrections. DF and GP stand for "dynamical friction" and "gravitational pull" respectively. (Adapted from [9] Table VI)

MBH are often found in gas-rich environments such as AGN, which could be an important formation channel for EMRIs[23][21]. Previous studies have also found, that 1 to 10% of the EMRIs that will be observed by LISA could reside in accretion disks around AGN[46]. These accretion disks could modify the orbital trajectory through hydrodynamic or gravitational effects, leading to, dynamical friction and planetary migration, for example. Another prediction of these studies is that EMRIs, which have formed in accretion disks are "likely to have circularised and aligned with (or be[en] born in) the disk"[53] [46][21].

3.1 Planetary Migration

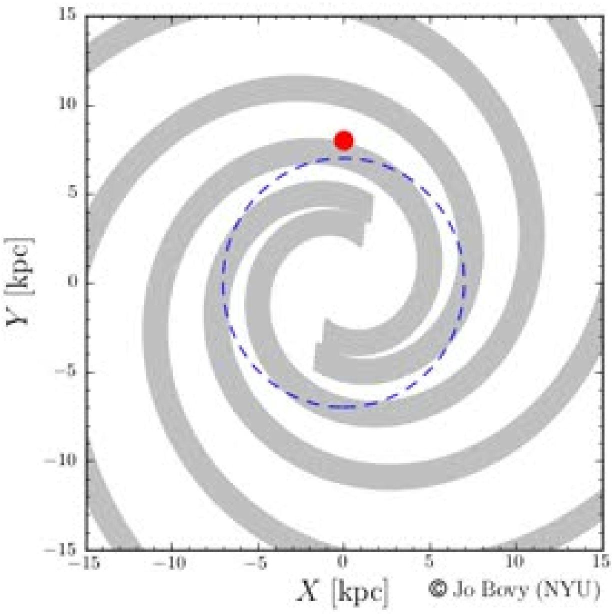

Figure 7: Schematic of spiral density waves from Lindblad resonances. The grey spiral tracks represent spiral density waves in a gaseous disk, which can be induced by an orbiting CO represented by the red dot. The dashed lines show a circular orbit, indicating that the density waves perform differential motion interior and exterior of the orbit. (Reprinted from [12])

The dominant effect identified in the study by Barausse et al.[9] is planetary(-style) migration. As a CO moves through a gaseous accretion disk, it causes density perturbations which form a *wake*, i.e. a denser region in proximity to the CO. The gravitational pull of this local overdensity on the CO can lead to dynamical friction. However, dynamical friction does not account for the dynamics and interplay of the *wake* within the disk, where it can excite Lindblad resonances that lead to spiral density waves, as shown in Fig. 7. These density waves can stretch from the inner to the outer regions of the accretion disk, resulting in differential motion in the disk's annuli, caused by the fact that the rotational velocity is a decreasing function of the radius[9]. Therefore, part of the *wake* of the CO will be inside its orbit, increasing the CO's angular momentum and causing it to migrate outwards. Another part of the *wake* of the CO will be exterior to its orbit, decreasing the CO's angular momentum and causing an inward migration. These two cases are summarised as *type I migration*, and the direction of the migration (inward or outward) depends on the balance between the two effects, requiring sophisticated calculations. However, in most cases, the torque of the outer spiral dominates, causing the CO to lose angular momentum and migrate inwards.

Another phenomenon can occur when the CO removes gas from its orbit in the disk faster than the resulting gap can be filled by viscous flow. This causes the CO to migrate inwards[40][65]. This process is called *type II migration*. However, more recent simulations show that the gas flow is highly complex, making it uncertain whether or when *type II migration* can occur [24][25][33][50]. Reviews of planetary migration can be found in [5][60]. These models have been developed to account for the dynamics of planets in protoplanetary disks[4][45]. Unfortunately, there are currently no numerical simulations of EMRI in realistic AGN accretion disks. However simulations of Intermeditate Mass Ratio Inspiral (IMRI) have proven the approximate validity of the model (within one order of magnitude)[20][19][21]. These simulations also suggest that focusing our study on inward *type I migration* in thin disks is conservative and avoids overestimating the effect. This approach allows the migration torque to be modelled using a simple power-law dependency on the orbital radius [54], which will be extremely useful. Finally, it has been found that the torque due to planetary(-style) migration is comparable in strength to the torque due to GW emission[65][40].

3.2 Accretion Disk

The physics of accretion disks is fascinating, combining the fields of electrodynamics, hydrodynamics and GR. Describing the physics of accretion disks numerically requires expensive simulations within a General Relativistic Magnetohydrodynamics (GRMHD) framework [2]. However, there are also analytical models of accretion disks, such as those developed by Shakura and Sunyaev in 1973 [52]. These models describe the accretion flow of geometrically thin, optically thick disks. They discovered that the internal torque within such disks must be proportional to the pressure and they proposed two possible descriptions. In the α description the torque is proportional to the gas pressure and the radiation pressure $t_{r\phi} = \alpha(p_{gas} + p_{rad})$, while in the β description it is only proportional to the gas pressure $t_{r\phi} = \alpha p_{gas}$, in "both cases α is the viscosity, which parametrises the complex (and uncertain) magnetohydrodynamic features of accretion disks"[53]. An agnostic model of the torque exerted on the CO of an EMRI including both types of disk, has been derived in [53], motivated by previous work [9][40][65]. This model simply adds the contributions of the disk's torque and GW emission torque

$$\dot{L} = \dot{L}_{GW} + \dot{L}_{disk}. \tag{8}$$

The term accounting for the interaction with the disk follows a simple power-law approach consisting of an amplitude A and a radial slope n_r

$$\dot{L}_{disk} = A \left(\frac{r}{10}\right)^{n_r} \dot{L}_{GW}^{(0)}, \tag{9}$$

it is scaled by the leading order circular orbit GW torque $\dot{L}_{GW}^{(0)} = -\frac{32}{5}\epsilon\, r^{-7/2}$, motivated by the estimations in [9]. In this parametrisation, the amplitude A incorporates the information on the accretion disk, specifically its type and the strength of the effect

$$A = \mathcal{C} \left(\frac{\alpha}{0.1}\right)^{n_\alpha} \left(\frac{f_{\rm Edd}}{0.1}\frac{0.1}{\epsilon}\right)^{n_{f_{\rm Edd}}} \left(\frac{M}{10^6 M_\odot}\right)^{n_M}, \tag{10}$$

where $n_\alpha, n_{f_{Edd}}, n_M$ decode the dependence on the parameters α, f_{Edd}, M for each disk model and are given in Tab. 2.

This parametrisation is only valid for geometrically thin disks and becomes invalid at the inner edge of the disk, where the gas density and the torques decrease. Our waveform models avoid this region anyway, as they only model systems up to the point at which they approach the Innermost Stable Circular Orbit (ISCO) (see Sec. 4). Conversely, this parameterisation is agnostic and can, in principle, be applied to any environmental effect where the exerted torques are proportional to the orbital separation r.

	Migration (α)	Migration (β)
$\mathcal{C}$	$7.2 \cdot 10^{-10}$	$8.1 \cdot 10^{-6}$
n_r	8	5.9
n_α	-1	-4/5
$n_{f_{Edd}}$	-3	-7/5
n_M	1	6/5

Table 2: Parameters of the disk torque model. The relevant properties of the accretion disk torque are combined into a single amplitude using the parameterisation derived in [53]. (Adapted from [53])

Detection of GR deviations

EMRIs enable us to test the regime of strong gravity and probe the theory of GR. If gravity is not fully described by GR, current waveform templates may be inadequate for analysing GW signals. This could result in biases in parameter estimation and reduced detection efficiency[17], similar to unmodelled environmental effects. To address this issue, *parametrised post-Einsteinian* waveforms have been proposed[62][17]. Some theories, including deviations from GR, modify the phase of GW compared to GR and can be parametrised using the presented torque model. For fixed values of the slope parameter n_r the amplitude A can represent the size of a deviation from GR[53]. Deriving theories that extend or deviate from GR is outside the scope of this work. However, since our aim is to investigate confusion between vacuum and perturbed GW templates in parameter estimation, we incorporate one such deviated model and its corresponding slope into our analysis.

More precisely, the class of deviations from GR we are considering allows for a time-varying gravitational constant, which would modify the waveform templates. Previous studies have found that such a deviation could be described by an additional term in a post-Newtonian expansion of order -4 [63][14][10][58]. This deviation can therefore be accounted for in our modified torque model by fixing $n_r = 4$[53].

4 Signal modelling

To prepare for the upcoming EMRI observations with LISA, we need to understand these sources and develop models for their gravitational wave signals. This section discusses how the scientific community responds to this call and the "state-of-the-art" of EMRI modelling. The presented techniques and models will enable us to prepare EMRI signal analysis and utilise them to answer research questions from astrophysics, cosmology, and fundamental physics.

In this chapter, we will outline the baseline approach, discuss the `FastEMRIWaveforms` (`FEW`) package and our implementation and finally explain how we extend the vacuum solution to introduce an environmental effect.

4.1 EMRI Gravitational Wave Strain

Realistic EMRIs can be modelled as the complex motion of the compact object in an approximately Kerr spacetime with various initial conditions as discussed in section 2.3. Modelling the EMRI's gravitational wave signal reflects this complex nature and therefore poses a significant challenge, which has been addressed by the self-force approach [8].

As discussed before 2.3, EMRIs could last 4-5 years in the detectable window of LISA. This translates into $10^4 - 10^5$ orbital cycles with periods ranging from minutes to hours. By separating these two timescales, it is possible to simplify the complexity of the EMRI's dynamics and study effects on each timescale separately. The "fast" component corresponds to the periodic motion along the orbit and is therefore called the *orbital timescale* $T_{orb} \sim M$. The "slow" component describes the gradual shift inward due to dissipative effects (radiation of GW) and is therefore called the *radiation-reaction timescale* $T_{rad} \sim M/\epsilon$, with the MBH's mass M, the CO's mass μ and the EMRI's mass ratio $\epsilon = \mu/M$.

Combining the knowledge of the behaviour on the two timescales leads to so-called *adiabatic waveforms*. This framework assumes that the orbits are quasi-stationary on the orbital timescale, such that they can always be considered bound Kerr geodesics. The effects of GW radiation are averaged on each orbit, and at adiabatic order, the secondary shifts from one geodesic to another, gradually approaching the primary. This forms the inspiral.

Most of the work in modelling GW waveforms relies on perturbative methods. EMRIs exist in the strong-field regime, where the post-Newtonian expansion used for compact binary inspirals is not applicable. Instead, it is necessary to solve the two-body problem for the perturbed Einstein equations up to linear order, with the expansion being in the mass ratio ϵ [49]. The spacetime in which the system evolves is dominated by the primary, making it essentially Kerr, with the secondary causing a small perturbation. The metric tensor can then be expanded as $g_{\mu\nu} = g_{\mu\nu}^{\text{Kerr}} + h_{\mu\nu}$. With the perturbative approach, the linearised Einstein equations can be formulated into the Teukolsky equation [55], which describes how perturbations evolve in Kerr spacetime.

The motion on the *orbital timescale* corresponds to the zeroth order in the mass ratio expansion $\mathcal{O}(\epsilon^0)$ and follows Kerr geodesics. 2nd and higher orders $\mathcal{O}(\epsilon^{2+})$ include Gravitational Self-Force (GSF) effects, which arise due to the finite size and mass of the secondary as well as the back-reaction with its gravitational field [8].

For the motion on the *radiation-reaction timescale*, the dissipation due to GW emission has to be calculated. These first-order effects in the expansion are obtained by solving the Teukolsky equation. The GW strain "at infinity" maps directly to the dissipation of the constants of motion.

This process is iterative. The Teukolsky solutions provide the radiation reaction that shifts the orbit to the next geodesic, changing the object's trajectory. This then updates the source term in the Teukolsky equation for the next step.

Each bound orbit is described by three fundamental frequencies: Ω_r (radial), Ω_θ (polar) and Ω_φ (azimuthal) [51], corresponding to the periodic motion in the respective direction. The solution to the Teukolsky equation is then written in terms of harmonics of these frequencies

$$\omega_{mkn} = n\Omega_r + k\Omega_\theta + m\Omega_\varphi \qquad , \qquad m, k, n \in \mathbb{Z}, \tag{11}$$

and the GW strain can finally be written in the *"multi-voice decomposition"*[31]

$$h(t) = h_+ - ih_\times = \frac{\mu}{d_L} \sum_{lmkn} A_{lmkn}(t)\; S_{lmkn}(\theta, \phi)\; \exp[-i\Phi_{mkn}(t)]. \tag{12}$$

Therefore in ths approach the waveform can be represented by the sum of multiple harmonics described by a complex amplitude A_{lmkn}, spin-weighted spheroidal harmonics $S_{lmkn}(\theta, \phi)$ accounting for the viewing angle on the system and the oscillatory part $\exp[-i\Phi_{mkn}]$ determined by the phase Φ_{mkn}. The factors μ and d_L correspond to the secondary's mass and the luminosity distance between the observer and the source, respectively.

To reduce the complexity of the inspiral trajectory and leverage existing waveform implementations, we assume prograde, circular and equatorial orbits. This is justified as most accretion disks lie in the equatorial plane, and we specifically want to study their influence on the EMRI 3. Further, it is predicted that EMRIs circularise throughout their evolution [16]. With these assumptions, we reduce the waveform model by two degrees of freedom (radial and polar motion, indices k,n respectively).

A challenge in computing these waveforms is the spin-weighted spheroidal harmonics $S_{lm}(\theta, \phi)$, which emerge from the way perturbations are decomposed in the Teukolsky equation. For Schwarzschild spacetime (i.e. non-spinning primaries), the spheroidal harmonics reduce to spherical harmonics $S_{lm} = Y_{lm}$. However, determining these factors in Kerr spacetime complicates the waveform computation and can be avoided by introducing a new set of amplitudes. The amplitudes A_{lm} and spheroidal harmonics S_{lm} become new amplitudes C_{lm} and are expressed in spherical harmonics Y_{lm}. The exact derivation can be found in [64]. The GW characteristic from eq. 12 finally becomes

$$h(t) = \frac{\mu}{d_L} \sum_{lm} C_{lm}(t)\; Y_{lm}(\theta)\; \exp[-i\Phi_m(t)]. \tag{13}$$

Unlike in Schwarzschild spacetime, where the orbits are fully described by energy E and angular momentum L alone, in Kerr spacetime a third constant arises due to a hidden symmetry related to the Killing tensor, the Carter constant Q [13]. It carries information about the inclination of the orbit and the motion in the latitudinal/polar direction. The parametrisation of an orbit in constants of motion $\{E, L, Q\}$ and orbital parameters $\{p, e, \iota\}$ (semi-latus rec-

Parameter	**Definition**	**Valid range in `FEW`**
M	Mass of primary	$[10^4, 5 \cdot 10^7]$
μ	Mass of secondary	$[1, 10^4]$
a	Dimensionless spin of primary	$[0, 0.99]$
p_0	Initial semilatus-rectum	$[7 \cdot (6 + 2 \cdot e_0) - 41.9, 16 + 2 \cdot e_0]$
e_0	Initial eccentricity	$[0, 0.75]$
$x_{I,0}$	Cosine of initial inclination angle	$[-1, 1]$
d_L	Luminosity distance	-
q_S	Sky location polar angle	$[0, \pi]$
ϕ_S	Sky location azimuthal angle	$[0, 2\pi]$
q_K	Primary spin polar angle	$[0, \pi]$
ϕ_K	Primary spin azimuthal angle	$[0, 2\pi]$
Φ_{ϕ_0}	Initial azimuthal phase	$[0, 2\pi]$
Φ_{θ_0}	Initial polar phase	$[0, 2\pi]$
Φ_{r_0}	Initial radial phase	$[0, 2\pi]$

Table 3: EMRI parameters and their valid ranges in `few`

tum, eccentricity and inclination angle of orbital plane respectively) is equivalent.

A generic EMRI signal in its source frame can be modelled by 9 parameters, which we call the *intrinsic* parameters. In the Solar System Barycenter (SSB) frame, distance d_L, sky position (q_S, ϕ_S) and spin orientation (q_K, ϕ_K) of the MBH are added (the *extrinsic* parameters), amounting to a total of 14 parameters, as shown in Tab. 3.

4.2 `FastEMRIWaveforms`

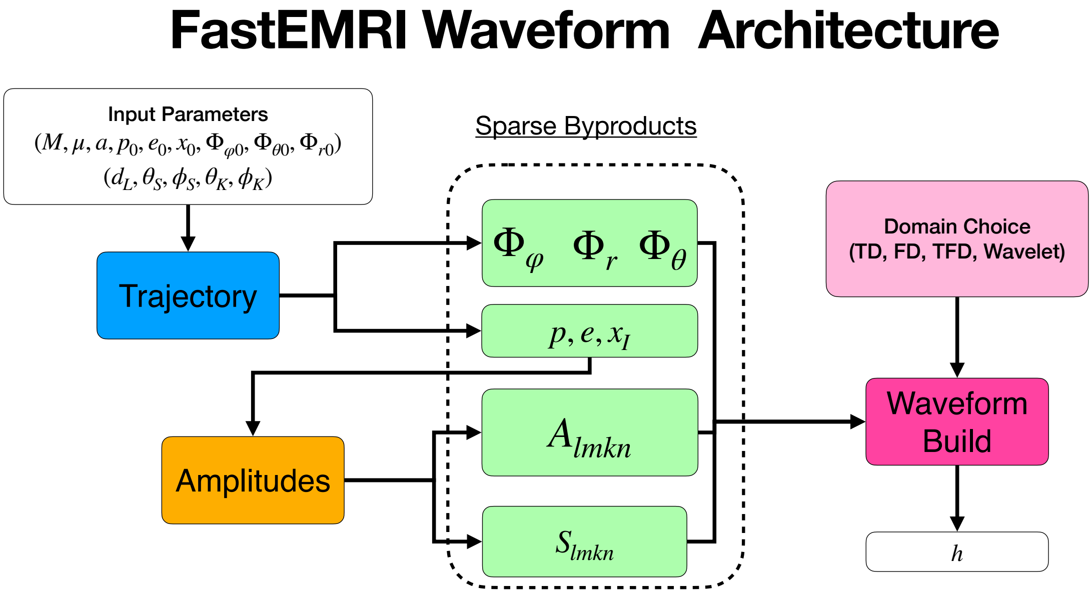

Figure 8: General architecture of `FEW`. A trajectory is calculated with the input parameters, which is then used to interpolate the amplitudes of the "multi-voice decomposition". Finally, amplitudes and orbital phases are summed to form the waveform approximation. (Reprinted from [38])

The `FEW` package [38] [15] is a potent tool within the `BlackHolePerturbationToolKit` [11]. It provides efficient EMRI waveform generation and has implemented different models for several spacetimes and specific orbits to high precision. Additionally, it is GPU compatible and accelerates the computation time for waveform generation, enabling us to create waveforms of 4-year inspirals within $\sim$ 10 - 100 ms. This software is crucial to this work, and in the following section, we will sketch the fundamental functionality of the package and the waveform implementation we are utilising.

Fig. 8 displays the general scheme of a waveform computation. First, a set of intrinsic parameters is passed to a chosen trajectory function (spacetime and orbital assumption). This returns the timeseries of the phases $\{\Phi_\varphi, \Phi_r, \Phi_\theta\}$ and the orbital parameters $\{p, e, x_I\}$. The latter are used to interpolate the Teukolsky amplitudes A_{lmkn} and S_{lmkn} (or in our model C_{lm}) from a precomputed grid. It is obvious that this introduces errors and that the interpolation method as well as the grid of precomputed data have to be carefully chosen, which was studied in [31]. Finally, all the factors of the "multi-voice decomposition" are multiplied and summed up, forming the EMRI waveform. The output can be either in the source or detector frame.

4.3 Implementation of Non-Vacuum Trajectories

The trajectory function is at the heart of the waveform computation as it produces the phase evolution factor and is used to interpolate the amplitude factors in the *multi-voice decomposition*. It is precisely what we are modifying to include environmental effects.
In `FEW` trajectories are computed by solving a set of coupled Ordinary Differential Equations (ODEs)

$$\dot{p} = \epsilon\, f_p(a, p, e, x_I) \tag{14}$$

$$\dot{e} = \epsilon\, f_e(a, p, e, x_I) \tag{15}$$

$$\dot{x}_I = \epsilon\, f_{x_I}(a, p, e, x_I) \tag{16}$$

$$\dot{\Phi}_{\varphi,\theta,r} = \Omega_{\varphi,\theta,r}(a, p, e, x_I)/M, \tag{17}$$

where the dot corresponds to the total time derivative $\frac{d}{dt}$ and the pre-factor to the mass-ratio $\epsilon = \frac{\mu}{M}$. The f_p, f_e, f_{x_I} are the orbital elements' fluxes due to the emission of GW. Like the Teukolsky amplitudes, they are interpolated from a precomputed grid in the equivalent constants of motion parameterisation $\{E, L, Q\}$. The derivation scheme is illustrated in [27]. For the case of circular equatorial orbits, the ODE system simplifies significantly as the evolution of eccentricity and inclination becomes negligible. The semi-latus rectum p becomes equivalent to the radial separation of the two objects r. This allows to express the evolution of p in terms of one parameter. With the chain rule, it is possible to express the ODEs as

$$\dot{p} = \epsilon \frac{dL}{dt}\frac{dp}{dL} = \epsilon \frac{\dot{L}}{L'} \tag{18}$$

$$\dot{\Phi}_\varphi = \Omega_\varphi/M. \tag{19}$$

Here we denote the derivative with respect to p with a prime. The analytical expression for the derivative L' is obtained from the `KerrGeodesics` Mathematica package [59] within `BlackHolePerturbationToolKit` [11]. It simplifies to

$$\frac{dp}{dL} = \frac{dL}{dp}^{-1} = \left(\frac{-3a^3 + a^2(8-3p)\sqrt{p} + (-6+p)p^{5/2} + 3ap(-2+3p)}{2 \cdot (2a + (-3+p)\sqrt{p})^{3/2} \cdot p^{7/4}} \right)^{-1} \tag{20}$$

in the equatorial circular case. The flux of the angular momentum $\dot{L}$ can be transformed into a flux of orbital energy $\dot{E}$ and follows the simple relation $\dot{L} = \dot{E}/\Omega_\varphi$, with $\Omega_\varphi = (r^{3/2} + a)^{-1}$ for Kerr circular orbits. The flux of orbital energy $\dot{E}$ is, as mentioned before, interpolated from a precomputed grid.

In section 3 we discussed that environmental effects can be modelled by an agnostic power-law approach

$$\dot{L} = \dot{L}_{GW} + \dot{L}_{env} \tag{21}$$

$$\dot{L}_{env} = A \left(\frac{r}{10}\right)^{n_r} \dot{L}^{(0)}_{GW}, \tag{22}$$

where $\dot{L}^{(0)}_{GW} = -\frac{32}{5}\epsilon\, r^{-7/2}$ is the leading order circular orbit GW angular momentum flux. Now, modelling the environmental effect as an additional flux of angular momentum L is simple and means just adding a term to the ODE [53]. The modified ODEs including environmental effects are given by

$$\dot{p} = \epsilon \frac{\dot{L}_{GW} + \dot{L}_{env}}{L'} \tag{23}$$

$$\dot{\Phi}_\varphi = \Omega_\varphi / M. \tag{24}$$

We have now achieved a nested implementation of the environmental effects model on the trajectory computation. The trajectory is the first step towards the GW strain. The introduced environmental effect perturbs this trajectory and its imprint will propagate into the waveform.

LISA Observables

LISA will not directly measure the gravitational strain. Through the preprocessing of raw data through the LISA analysis flow and the application of the TDI algorithm, we will obtain the $\{A, E, T\}$ observables (see Sec. 2.2 and [16]). We model this via the `fastlisaresponse` package [37].

5 Bayesian Inference

This thesis aims to conduct a statistical analysis of EMRI GW signals. Using Bayesian model selection techniques, we will assess whether it is possible to distinguish between the observed data of an EMRI that has been affected by an environmental effect and one that has not. We employ Bayesian parameter estimation of competing models through stochastic sampling, followed by Bayesian model selection.

This chapter aims to explain and document the entire stochastic sampling process for this project. The following sections will first present the theoretical framework, followed by the numerical implementation.

5.1 Statistical Inference and Bayes' Theorem

Statistics can be divided into two primary branches: descriptive and inferential statistics. Descriptive statistics are a means of describing and summarising data without making predictions or generalisations. Conversely, inferential statistics operates under the assumption that the data is a sample of a larger population. This sample is utilised to make predictions and generalisations about the population, incorporating uncertainty when drawing conclusions. In each of these branches there are two general approaches: frequentist and Bayesian. In the frequentist approach, probabilities are defined as the occurrences of events divided by the repetitions of the experiment. This approach is based on the convergence law, namely that the experiment is repeated infinitely many times in order to determine the 'true' value, and does not incorporate prior knowledge (e.g. physical or non-physical regions in the parameter space). Furthermore, conditional probabilities of data in the context of a hypothesis or model can be considered. However, it is impossible to discuss the probabilities of a model or hypothesis being correct, or of parameters taking certain values, given some data, as models or hypotheses are not the outcomes of repeatable experiments. The rejection or acceptance of models and hypotheses is only possible under the assumption of confidence intervals or significance.

In contrast, the Bayesian approach offers a higher degree of flexibility in these terms. The analytical expression of Bayes' theorem follows from fundamental properties of conditional probability. The probability of the event A occurring, given the event B, is expressed as follows:

$$P(A|B) = \frac{P(A,B)}{P(B)}, \tag{25}$$

where A, B denotes both events occurring. Similarly

$$P(B|A) = \frac{P(B,A)}{P(A)}, \tag{26}$$

where $P(A,B) = P(B,A)$. Solving both expressions for the probability of the intersection, $P(A,B)$, and equating them, we arrive at

$$P(A|B)P(B) = P(B|A)P(A) \tag{27}$$

and finally at Bayes' theorem

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{28}$$

Another common way to write the Bayes' theorem emphasises the posterior distribution of parameters θ for a given model $\mathcal{M}$ describing the data d is

$$p\left(\vec{\theta}|d,\mathcal{M}\right) = \frac{\mathcal{L}(d|\vec{\theta},\mathcal{M})\pi(\vec{\theta}|\mathcal{M})}{z(d|\mathcal{M})}, \tag{29}$$

with $\mathcal{L}(d|\vec{\theta},\mathcal{M})$ describing the likelihood of observing the data for a given set of parameters in the given model, $\pi(\vec{\theta}|\mathcal{M})$ the prior probability of these parameters in the given model and the evidence $z(d|\mathcal{M})$ which is the probability of observing the data for the given model and i.e. the normalisation factor for the product in the nominator[48].

Here, we can also understand why Bayes' theorem proves to be useful. We have a hypothesis of some parameter values describing a GW signal, and start with the prior probability of these parameter values. We then observe some data

and update our posterior probability of the parameter values given the data, accordingly, by incorporating the likelihood of the data given the parameter values. The posterior becomes more accurate and can be updated as more data becomes available.
When analysing GW data, the Bayesian approach is preferred over the frequentist approach, as there is only one measurement and we can not repeat the observation multiple times. It can handle high-dimensional GW models, and it incorporates information from theory and simulations. Further, it provides full posteriors, which correspond to the estimates of the parameters, where the uncertainty correspond to the uncertainty of the measurement. The Bayesian approach allows testing and comparison of different models and excels at dealing with noise in the data, as it can include stochastic noise models, while the frequentist approach relies on simplifying assumptions about the noise.

5.2 Stochastic Sampling

The posterior distribution provides the estimates and uncertainties on the model parameters, but can not be solved analytically as it is complex and high-dimensional. Therefore, stochastic sampling is employed, which is the process of selecting members of a statistical population to form a subset that represents its characteristics. Combined with Bayesian inference, it approximates the posterior distribution. We use MCMC methods to sample from the product distribution of the likelihood and prior, which is proportional to the posterior distribution

$$p\left(\vec{\theta}|d,\mathcal{M}\right) \propto \mathcal{L}(d|\vec{\theta},\mathcal{M})\pi(\vec{\theta}|\mathcal{M}). \tag{30}$$

Monte Carlo methods are an umbrella term for computational algorithms that use random sampling to estimate numerical quantities. In our case, these quantities correspond to the parameter value estimates for a given model. A Markov chain is a sequence of points in parameter space, where the probability of each point depends solely on the previous point. These chains are designed such that, once they have converged, they represent the posterior distribution. To construct such a chain, we employ Metropolis–Hastings algorithms. The scheme of this fairly simple method is visualised in Fig. 9.

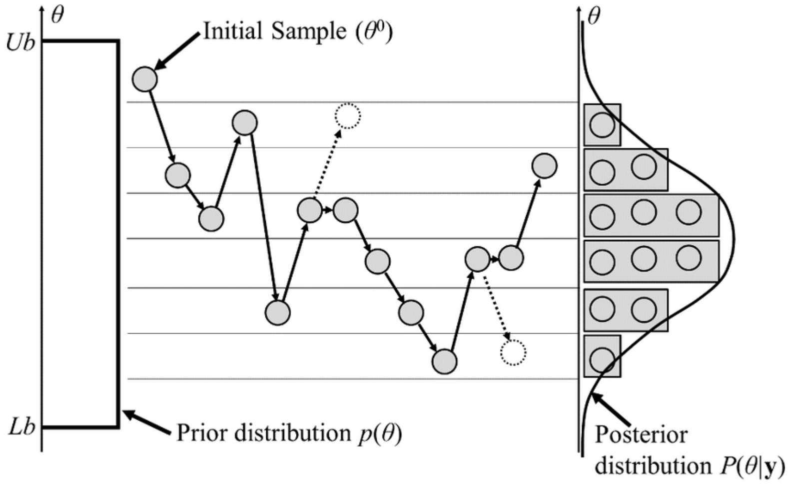

Figure 9: The MCMC Metropolis-Hastings algorithm scheme of approximating the posterior probability distribution. Starting on the left with a uniform prior distribution, the boundaries on the explorable parameter space are set. Then a random value is drawn as a starting point, and at each iteration, a new value is proposed. Based on the ratio of the likelihood values of the current and proposed state, the proposal is accepted or rejected. Recording the accepted states and binning them finally produces a histogram which corresponds to the parameter estimate. (Reprinted from [42])

Based on the current state in the parameter space, a new state, called the 'proposal', is drawn from a distribution q. Then, the acceptance probability is calculated as the ratio of the posterior probabilities of the current and proposed states [41]

$$\alpha = \min\left(1, \frac{p\left(\vec{\theta}_p|d,\mathcal{M}\right)}{p\left(\vec{\theta}_t|d,\mathcal{M}\right)}\frac{q(\vec{\theta}_p)}{q(\vec{\theta}_t)}\right) \tag{31}$$

$$= \min\left(1, \frac{\mathcal{L}(d|\vec{\theta}_p,\mathcal{M})}{\mathcal{L}(d|\vec{\theta}_t,\mathcal{M})}\frac{\pi(\vec{\theta}_p|\mathcal{M})}{\pi(\vec{\theta}_t|\mathcal{M})}\frac{q(\theta_p)}{q(\theta_t)}\right) \tag{32}$$

where the index t denotes the state at the current iteration and the index p the proposed state. Bayes' theorem is employed to rewrite the term, cancelling out the evidence in the denominator. In the case of uniform priors, the probabilities are constant across the parameter space and would also cancel out. While it is more probable for a chain to move to higher likelihood states, it can still move to lower states as the probabilities are not constant. This process continues iteratively, exploring the likelihood surface weighted by the prior. Fig. 10 depicts this process for a multi-dimensional parameter space.

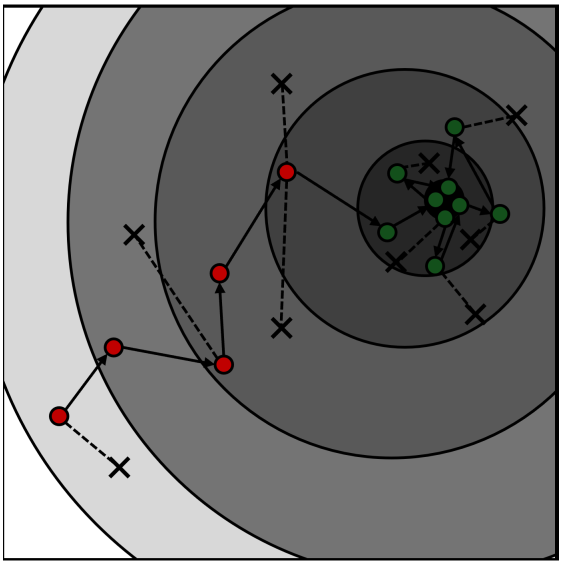

Figure 10: Simplified scheme of a chain exploring a 2D parameter space. The shaded regions represent a likelihood gradient, and the surrounding box the limitations of uniform priors. We observe a general tendency to move towards higher likelihood states, with rejected proposals represented by the crosses. After a "burn-in" phase (red states), the chain explores the high likelihood region (green states). Those states are used to produce histograms as in Fig. 9.

`eryn`

The employed sampler to perform the parameter estimation is `eryn` [34][36][26]. This advanced, extended version of the popular `emcee` sampler implements parallel tempering, multiple model types, and Reversible Jump (RJ)-MCMC techniques. The framework is designed to allow users to implement their own setups, include new proposals and combine moves.
To initialise the sampler, you need to define the data, the likelihood function, the prior, the starting points for the chains and the proposal. The likelihood function establishes the link between sampling and the posterior probability in Bayes' theorem. The prior constrains the explorable parameter space, and its width and distribution type depend on the task at hand. In cases where the general population is unknown or poorly constrained, as is often the case in GW astronomy and particularly for EMRIs, uniform distributions are typically chosen, though arbitrary distributions can also be used. The chains are commonly placed at random starting points within the prior volume; for parameter estimation, they should be placed close to the true value. The proposal determines how the chains explore the parameter space; a variety of options are available in `eryn`. Tempering or RJ can also be set up optionally.

Likelihood Function

In Metropolis-Hastings algorithms, proposals are accepted or rejected based on the posterior probability of the proposed and current states. As posterior probabilities are inaccessible, the criterion is rewritten in terms of the product of the likelihood and the prior, using Bayes' theorem (see eq. 32). The likelihood describes the probability of the data d given a set of parameters θ in a gravitational wave strain $h(\theta)$. In the simplified case, the data d is assumed to be a sum of gravitational wave strain h projected onto the sensitive access of the detector and processed through time delay interferometry and noise n, such that $d = h + n$. Therefore, the probability of observing the data d given the strain h is equal to the probability of the noise being $n = d - h$ ($p(d|h) = p(n = d - h)$). Thus, the likelihood is computed based on the noise residuals $n = d - h$. The noise is typically assumed to be Gaussian and stationary. Under this approximation, the noise in the frequency domain is uncorrelated for each frequency bin, and the variance is proportional to $(1/2)S_n(f)$ [43], where $S_n(f)$ is the detector's Power Spectral Density (PSD) [6]. The corresponding Gaussian probability distribution then becomes

$$p(n) \propto \exp\left(-\frac{1}{2}\int_{-\infty}^{\infty}\frac{|\tilde{n}(f)|^2}{(1/2)S_n(f)}\mathrm{d}f\right) \tag{33}$$

We can generalise by defining the *(noise) weighted inner product*

$$\begin{aligned}(a(t)|b(t)) &= \Re\int_{-\infty}^{\infty}\frac{\tilde{a}^*(f)\tilde{b}(f)}{(1/2)S_n(f)}\mathrm{d}f \qquad (34)\\ &= 4\,\Re\int_{0}^{\infty}\frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}\mathrm{d}f, \qquad (35)\end{aligned}$$

where $a(t)$ and $b(t)$ are real functions such that $\tilde{a}(-f) = \tilde{a}^*(f)$ [43]. As real GW data is discrete, this integral is approximated as a sum over frequency bins

$$(a|b) \approx 4\,\Re\sum_k\frac{\tilde{a}^*(f_k)\tilde{b}(f_k)}{S_n(f_k)}\Delta f. \tag{36}$$

Using the definition in eq. 35 and the probability distribution function from eq. 33 we can write the likelihood function. `eryn` and most other MCMC samplers deal with the logarithm of the likelihood, which is then expressed as

$$\ln\mathcal{L} \propto -\frac{1}{2}(d-h|d-h) \propto -\frac{1}{2}(n|n). \tag{37}$$

This procedure has the advantage of directly incorporating the noise model and comparing the powers of the signal and noise curves in all frequency bins. It accounts for the fact that frequencies contribute differently to signal detection due to varying noise and sensitivity levels. This makes the likelihood sensitive to how well the signal model $h(\theta)$ explains the data relative to the expected noise. For LISA, the data is additionally modulated by the detector's response function, as discussed in Sec. 2.2.

Proposal

The choice of proposal is crucial for convergence speed and parameter space exploration. Simple proposal techniques include Gaussian moves, while more sophisticated ones include ensemble stretch moves. Still, there are many other moves.

Gaussian moves are a classic Metropolis–Hastings move, whereby each walker explores the likelihood surface independently. The proposal state is created by adding a small perturbation to the current state. This perturbation is drawn from a Gaussian normal distribution centred on the current state with a predefined variance. One can choose to perturb all dimensions of the state simultaneously, sequentially, or randomly select a dimension to be perturbed. The choice of variance effectively controls the convergence time and exploration of the parameter space.

A simple example is a unimodal distribution with a strictly monotonous likelihood gradient. If the variance chosen for a dimension is too small, the chain will converge slowly to the maximum likelihood point and explore the region around this maximum very narrowly. This would lead to a high acceptance rate. Conversely, if the variance is too large, proposals may overshoot the maximum likelihood point and be more likely to be rejected because they do not indicate a higher likelihood than the current state. This would lead to a low acceptance rate. The dependence on the initial choice of variances for each dimension can be avoided by adapting the variance dynamically based on the spread of previous states. While this proposal is simple, robust and easy to implement, it is not suitable for complex, high-dimensional or multimodal distributions.

More sophisticated proposal methods use information from other chains to suggest new states based on the collective behaviour of the chains as a whole. Affine invariant methods, for instance, demonstrate consistent performance when the parameter space undergoes transformations such as rotation, scaling and stretching. This property is particularly useful for complex parameter spaces, such as those of EMRIs, where parameters are highly correlated and have varying scales. These methods are implemented in `eryn`, and the red-blue stretch move is a default choice [28].

The stretch move proposes a new state by moving along the line connecting it to another chain chosen at random and then scaling it by a random factor Z with

$$Z \propto g(z) \propto \begin{cases} \frac{1}{\sqrt{z}} & \text{if } z \in \left[\frac{1}{a}, a\right], \\ 0 & \text{otherwise,} \end{cases} \tag{38}$$

where a is the stretch scale parameter. It is typically set to $a = 2$ with a strict lower bound of $a > 1$. This parameter allows for adjusting the acceptance rate: a smaller a would lead to a proposal closer to the current state, and vice versa. The stretch move effectively "kicks" the chains towards each other. This proposal is accepted or rejected based on the Metropolis-Hastings criterion 32. The stretch move scheme is illustrated in Fig. 11.

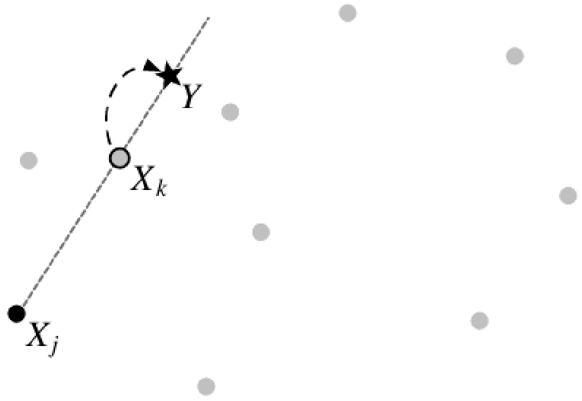

Figure 11: Schematic of the stretch move. A new state Y is proposed for state X_j along the line connecting it to state X_k. The stretch parameter a is always larger than 1, such that the proposed state is outside the space between X_j and X_k. By employing this proposal, the chains share information between each other, and the ensemble moves as a group. (Reprinted from [28])

In `eryn`, this proposal is extended to a red-blue move. In red-blue moves, the chains are split into groups, each of which is updated based on the states of the chains in the other group. The advantages of this setup are that autocorrelation between the samples is reduced and the exploration is enhanced, and it

leverages parallelisation since the two groups can be updated simultaneously. Further details of this framework can be found in [26].
Regardless of the proposal method, the acceptance rate should be tuned to be around $\approx 40 - 60\%$. This rate is crucial for convergence and exploration; however, adapting it is sometimes insufficient for efficient sampling of the parameter space[34].

Parallel Tempered MCMC

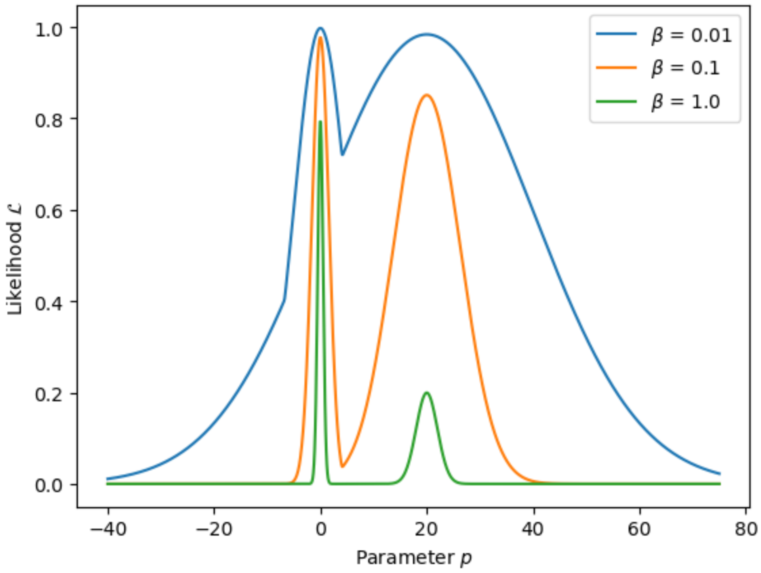

Figure 12: Schematic of tempered likelihood surface. Higher temperatures (lower β) smoothen out the likelihood surface. The low likelihood gap between the two peaks becomes less prominent, allowing chains to be more mobile as the difference in likelihood decreases. Proper normalisation is unnecessary since these are density functions, and only the ratio of likelihoods for a given temperature affects the acceptance rate (see eq. 32).

Complex likelihood surfaces may have multiple maxima, which should be identified by exploring the parameter space. However, if the initial states are close to the true value (as is the case for parameter estimation), the chains may fail to explore the likelihood surface; they won't identify secondary maxima or other degeneracies. Chains can become trapped at local likelihood maxima,

which is one of the key challenges in parameter estimation. This challenge is even more significant on multimodal likelihood surfaces, as the surface will not be fully explored and convergence can suffer.
Parallel-tempered MCMC methods address this issue by "smoothing out" the likelihood surface, enabling the chains to traverse it more easily. This involves transforming the posterior target distribution to the *power posterior target distribution*, essentially exponentiating the likelihood with the inverse temperature. Then we have

$$p_\beta\left(\vec{\theta}|d,\mathcal{M}\right)=\frac{\mathcal{L}(d|\vec{\theta},\mathcal{M})^\beta\pi(\vec{\theta}|\mathcal{M})}{z_\beta(d|\mathcal{M})}, \tag{39}$$

with the inverse temperature $\beta = 1/T$. For $T = 1$ this is the target posterior density, and in the limit $T \to \infty$ it is the prior density. All temperatures in between gradually flatten ('smooth out') the target likelihood landscape, until it becomes completely flat at an infinite temperature (see Fig. 12).
This allows the walkers to be mobile and increases the acceptance rate.

Temperature Swap

In `eryn`'s parallel tempering, ensembles of walkers are created in parallel at different temperatures. Hot chains can explore the parameter space more broadly and then propagate the information down to cold chains. This allows unknown regions of high likelihood to be discovered by hot chains and explored by cold chains. Information about these regions is passed on by proposing swaps on the temperature ladder. Technically speaking, we sample the posterior at different temperatures as shown in Eq. 39. The temperature swap acceptance probability between chain i and j is the ratio of those posteriors [34]

$$\alpha_{T,i\to j}=\min\left(1,\frac{p_T\left(\vec{\theta}_i|d,\mathcal{M}\right)}{p_T\left(\vec{\theta}_j|d,\mathcal{M}\right)}\right) \tag{40}$$

$$=\min\left(1,\left(\frac{\mathcal{L}(d|\vec{\theta}_i,\mathcal{M})}{\mathcal{L}(d|\vec{\theta}_j,\mathcal{M})}\right)^{\beta_i-\beta_j}\right) \tag{41}$$

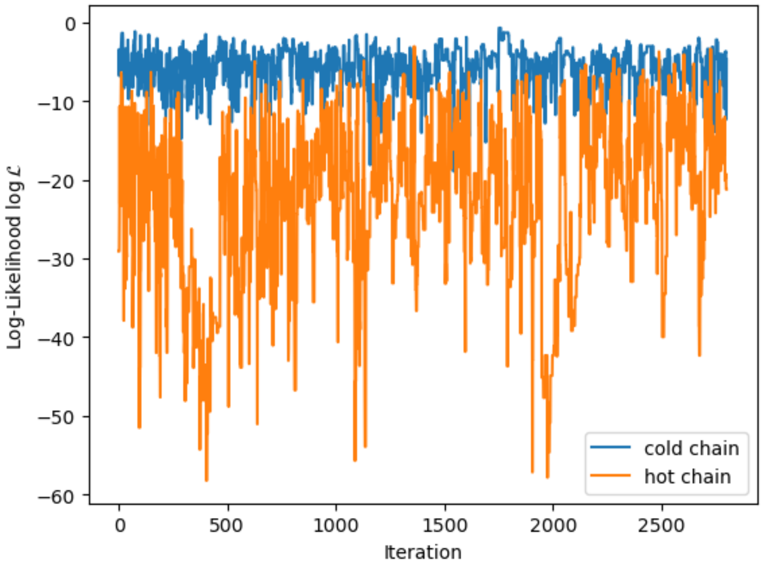

Figure 13: Hot and cold chains have different mobility on the likelihood surface. The cold chain stays close to the likelihood maximum, while the hot chain explores regions further away. The "flattened" likelihood surface of a hot chain increases the acceptance probability of states distant from the "true" value.

where $\beta_i = 1/T_i$ is the inverse temperature and $p(d|\theta_i)$ is the likelihood. Swapping is done iteratively, starting with the hottest two chains. This ensures that information on better likelihood values is "passed down" to the cold chains.
Let's consider two cases to illustrate the principle. When a cold chain proposes to swap with a hot chain (i.e. cold to hot), the exponent is positive $\beta_i - \beta_j > 0$. This makes swaps more likely if the likelihood ratio $p(d|\theta_i)/p(d|\theta_j)$ is favourable; for example, if the cold chain has a better likelihood, the hot chain will inherit this state. When a hotter chain proposes a swap with a colder chain (i.e. hot to cold), the exponent is negative $\beta_i - \beta_j < 0$, which suppresses swaps. A swap is only accepted if the likelihood ratio significantly favours the hotter chain, e.g. if the hot chain's likelihood is significantly better than the cold chain's. This ensures that high-likelihood regions propagate to colder chains while protecting the cold chain's refined state.

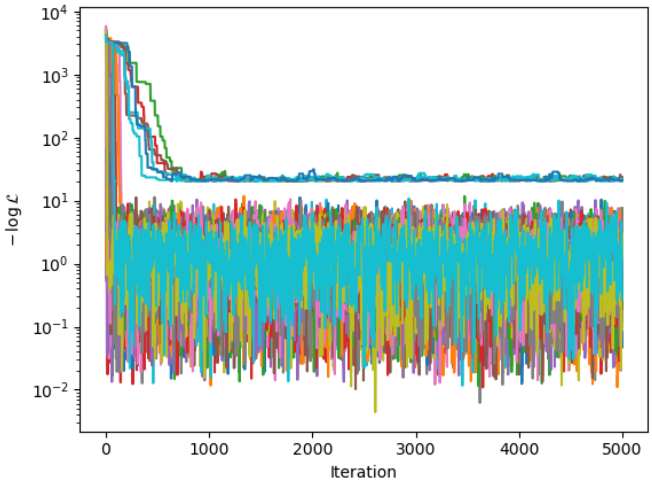

(a) Likelihood of cold Markov chains

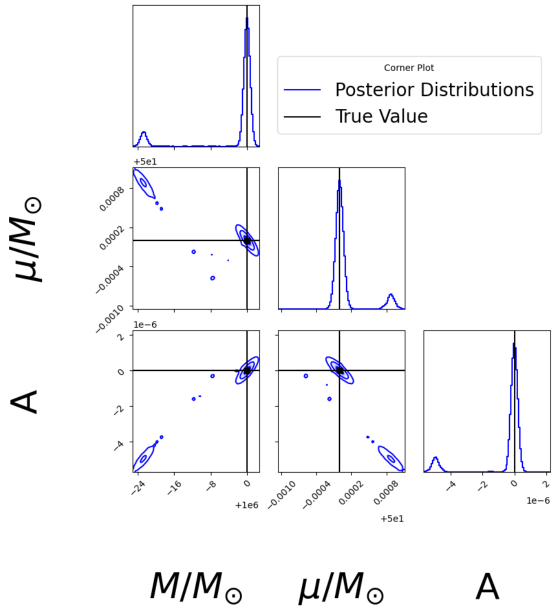

(b) Corner plot of posterior distribution

Figure 14: Example of the parameter estimation with `eryn` of three parameters for a typical EMRI. Fig. 14(a) shows how the cold chains take a few iterations to move towards the likelihood maxima, before exploring the surrounding parameter space. We can see that through the parallel tempering, the chains discover a secondary maximum with a lower likelihood. Fig. 14(b) displays the corresponding corner plot. One maximum around the true value is found, as well as one secondary maximum. As no iterations were discarded during the "burn-in" phase, points appear in the corner plot between the two maxima.

Adapting Temperatures

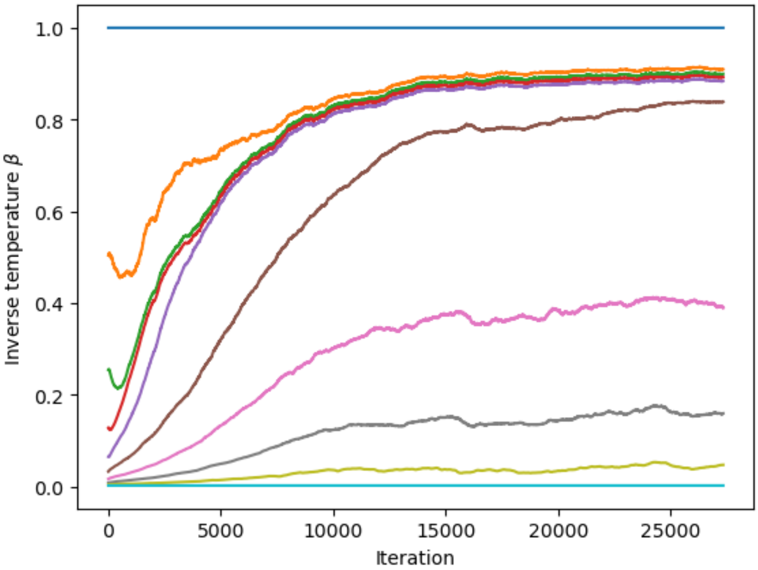

Figure 15: Adaptation of the temperature scale. The (inverse) temperatures are tuned until the temperature swap acceptance rate between neighbouring chains reaches an equilibrium. The highest and the lowest temperatures remain fixed.

In order to achieve information flow between the chains, the temperature ladder must be set up correctly. It is initialised with inverse temperatures β ranging from 0 to 1, and the interval is divided into log-equidistant values depending on the number of temperatures. In `eryn`, the temperatures can then be adjusted so that the swap acceptance rate is equal between all pairs of neighbouring temperature chains.

The adaptation rate can be customised to allow tuning based on the desired behaviour. The details of this implementation can be found in [34]. Following an initial burn-in phase, adaptation of temperatures should cease in a stable configuration. The final part of the analysis is then carried out. This provides accurate posterior sampling and allows for evidence calculation via thermodynamic integration or the stepping-stone algorithm, which require constant temperature.

5.3 Evidence Calculation

Starting from an environmentally perturbed EMRI, we aim to assess the quality of the parameter estimation of two competing models: one which includes an environmental effect and one which doesn't (i.e. a vacuum case). Finally, we want to perform Bayesian model selection to quantify the support of one model over the other.

This process requires an approximation of the evidence (marginalised likelihood). It is the integral of the product of the likelihood and the prior probability over the prior parameter space

$$z(d|\mathcal{M}) = \int d\vec{\theta}\, \mathcal{L}(d|\vec{\theta}, \mathcal{M}) \pi(\vec{\theta}|\mathcal{M}). \tag{42}$$

Often, there is no analytical closed-form expression for this integral, so numerical methods must be employed. There are a variety of options to obtain this quantity.

MCMC Integration

The integrand in eq. 42 is the product of the likelihood and the prior distribution, which is precisely what we are exploring using MCMC sampling. Once the chains have converged, we can approximate the integral through MCMC integration, essentially summing the product of the likelihood and the prior over the samples, and then normalising by the number of samples. This approach is particularly straightforward in our setup, since the likelihood values are stored and the priors are independent and uniform, thus forming a constant factor.

$$z(d|\mathcal{M}) \approx \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}_i(d|\vec{\theta}_i, \mathcal{M}) \pi_i(\vec{\theta}_i|\mathcal{M}) \tag{43}$$

Unfortunately, this approach has its disadvantages. We rely on converged chains that have thoroughly explored the prior parameter space. If we miss part of this space — for example, if the chains are not converged or the exploration is insufficient — the calculated evidence will be biased, which is true for all other methods.

Power Posterior Methods

Since we used parallel tempering to improve the sampling process, we obtained the *power posterior target distribution* from eq. 39 at different temperatures. These can be used to obtain more accurate approximations of the evidence using either the thermodynamic integration algorithm [41] or the stepping-stone algorithm [61]. However, there is a trade-off, as the accuracy of the algorithms increases with the number of temperatures. In our implementation, an equal number of chains is required for each temperature; therefore, each additional temperature multiplies the number of chains. This is particularly costly for high-dimensional parameter spaces, as is the case here. Both algorithms also require constant temperatures. In the employed framework, however, the temperatures are adaptive. To use the thermodynamic integration or stepping-stone algorithm, this adaptation must be stopped, which should only occur once the chains have approximately converged. As the convergence depends on the parameters of each set, stopping the adaptation process would require fine-tuning.

Normalising Flows

Normalising flows can also be used to estimate evidence. This machine learning approach effectively estimates the probability density of a distribution. A normalising flow maps an unknown distribution onto a normal distribution via multiple layers of bijective (invertible) transformations[47]. In theory, therefore, a normalising flow could approximate any distribution using normal mapping, provided that a sufficient number of layers are employed and the training samples accurately represent the unknown distribution.

Therefore, if the flow approximates the distributions well, it is a powerful tool, enabling sampling from the normalised posterior. We can then obtain the evidence by dividing the unnormalised distribution (likelihood multiplied by the prior) by the normalised distribution obtained from the normalising flow. This yields the desired normalisation constant– the evidence [48]

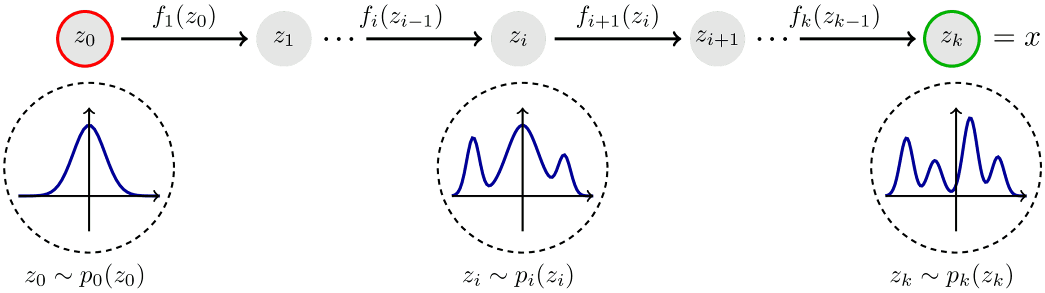

Figure 16: Concept of a normalising flow. The simple distribution is mapped to a complex one by applying a series of bijective/invertible transformations. (Reprinted from [32])

$$z(d|\mathcal{M}) \approx \frac{\mathcal{L}(d|\vec{\theta}, \mathcal{M})\pi(\vec{\theta}|\mathcal{M})}{p_{NF}\left(\vec{\theta}|d, \mathcal{M}\right)}. \tag{44}$$

where $p_{NF}\left(\vec{\theta}|d, \mathcal{M}\right)$ represents the learned/approximated posterior probability of a sample set of parameters $\vec{\theta}$. If the normalising flow were an exact mapping, calculating the evidence with just one sample would suffice. In practice, however, one averages evidence calculations for a set of samples to estimate the evidence.

$$z(d|\mathcal{M}) \approx \frac{1}{N}\sum_{i}^{N} \frac{\mathcal{L}_i(d|\vec{\theta}_i, \mathcal{M})\pi_i(\vec{\theta}_i|\mathcal{M})}{p_{NF,i}\left(\vec{\theta}_i|d, \mathcal{M}\right)}. \tag{45}$$

These samples should be generated using the normalising flow. The normalising flow is trained on the states of the MCMC chains. Therefore, to obtain a reliable approximation of the unknown distribution, it is necessary to have converged chains and accurate training data. Although this approach yields more accurate approximations of the evidence, it requires manual adjustment of the training process and significant computing resources given the scale of our analysis.

5.4 Bayesian Model Selection

Model selection is a well-studied and fundamental problem in statistics. In frequentist statistics, it is typically evaluated using p-values and significance tests. However, these methods fail to incorporate additional information (priors) and can only accept or reject a hypothesis. For this project, we aim to use a more precise, continuous method that quantifies the support for competing models $\mathcal{M}_i$ through their evidence and the ratio between them, – the Bayes factor

$$BF = \frac{z(d|\mathcal{M}_1)}{z(d|\mathcal{M}_2)}, \tag{46}$$

with $z(d|\mathcal{M}_i)$ being the evidence of posterior target distribution of model i. At a high level, a Bayes factor of $BF > 1$ means that the data are more likely to be explained by $\mathcal{M}_1$ than by $\mathcal{M}_2$. The Bayes factor naturally penalises complex models with more parameters, because their probability is spread across a larger parameter volume, thus decreasing the evidence. Only when the model offers a significantly better description of the data, the Bayes factor favours the complex model with the additional parameters. For a more precise interpretation, thresholds have been defined for Bayes factor values. A commonly cited table is given in Table 4.

Table 4: Bayes factor interpretation scale (Adapted from [35])

$\log_{10} BF$	BF	Strength of Evidence
0 to 1/2	1 to 3.2	Not worth more than a bare mention
1/2 to 1	3.2 to 10	Substantial
1 to 2	10 to 100	Strong
> 2	> 100	Decisive

Savage-Dickey Density Ratio

For nested models, the Bayes factor can be approximated using the Savage–Dickey density ratio. Models are nested when the parameters of one model are a subset of the parameters of another, or, put simply, when one model incorporates additional parameters, making it a "special case" of the other. The Savage–Dickey density ratio can be applied when fixing the additional parameters in the complex model to specific values yields the simple model. Then the Bayes factor is given by the ratio of the posterior and prior densities of the complex model, evaluated at the specific value that produces the simple model.

To illustrate this, consider two models, M_0 and M_1. Both models include a set of parameters θ and are distinguished by θ_0, which is allowed to vary in M_1 and fixed to be $\theta_0 = 0$ in M_0. The ratio was derived in [22] and is given by

$$BF = \frac{p(\theta_0 = 0, \theta_{1,2,..} | d, M_1)}{p(\theta_0 = 0, \theta_{1,2,..} | M_1)}, \tag{47}$$

where the model M_1's posterior and prior distributions are evaluated at the slice of the parameter space of M_1 where $\theta_0 = 0$. To apply the method, the prior of M_1 has to be continuous at $\theta_0 = 0$ (such that the density ratio is well defined) and the posterior of M_1 needs to be accurate, which is challenging with poor MCMC convergence. The advantage of this method over the standard approach of computing the Bayes factor is that it only requires the posterior and prior of one model (M_1) at the fixed point of interest, which is set to zero.

RJ-MCMC

In a RJ-MCMC set-up, the Bayes factor can be approximated as the ratio of iterations within each model. The trans-dimensional nature of RJ-MCMC makes it especially applicable to nested models [29]. We mention this because this class of MCMC algorithms is particularly important for LISA as it allows us to address the problem of an unknown number of signals in the data. However, there are a few prerequisites to meet to use this fairly simple approximation. For further reading, there is an extensive discussion of this method in [34].

6 Results

This thesis project aimed to determine when environmental effects are significant and when perturbed and unperturbed EMRIs are distinguishable. Two types of effect were considered: migration in thin accretion disks and modifications to GR due to a varying gravitational constant, which have both been studied in previous work [53]. We extend those studies by implementing more accurate trajectories and waveforms developed in [39] and extending the statistical analysis. We investigated the ability to detect and characterise the perturbing effect. Furthermore, we quantified model distinguishability by computing Bayes factors and established a link to the dephasing. This chapter presents these results and discusses their implications.

6.1 Numerical Setup

To investigate the significance of the environmental effect, we first define 21 EMRI systems with their corresponding parameter sets. We randomly draw parameter sets from the detectable and physically plausible parameter regions, which are further constrained by the available waveform model[39], given in Tab. 5.

Parameter	**Definition**	**Random draw interval limits**
M	Mass of primary	$[10^4, 5 \cdot 10^7]$
μ	Mass of secondary	$[1, 10^4]$
a	Dimensionless spin of primary	$[0.1, 0.99]$
p_0	Initial semilatus-rectum	$9, 35$
q_S	Sky location polar angle	$[0, 0.9 \cdot \pi]$
ϕ_S	Sky location azimuthal angle	$[0, 0.9 \cdot 2\pi]$
q_K	Primary spin polar angle	$[0, 0.9 \cdot \pi]$
ϕ_K	Primary spin azimuthal angle	$[0, 0.9 \cdot 2\pi]$
Φ_{ϕ_0}	Initial azimuthal phase	$[0, 0.9 \cdot 2\pi]$
A	Amplitude of perturbing effect	$[10^{-6}, 10^{-3}]$
n_r	Slope parameter of perturbing effect	$\{4, 5.9, 8\}$

Table 5: Limits of the random draw intervals for the parameter sets.

We confirm that their mass ratio is within the EMRI range $\epsilon \in [10^{-4}, 10^{-7}]$ and select the 21 systems so that they cover dephasings in the range $\Delta\Phi \in [10^{-1}, 10^{5}]$. Finally, we adjust the luminosity distance d_L of each system, such that the corresponding SNR is 30 for an observation time of 4 years before the merger. These parameter sets can be found in Tab. 7 in the appendix. We then inject these parameters into the waveform generator and project the gravitational wave strain h onto the LISA constellation to obtain the A, E, T observables, which define the data. To simulate the estimation for 1, 2, 3, and 4 years of observation data before the merger, we adjust the initial separation p_0 to ensure the data has the correct length (correct observation time). This slightly decreases the SNR, but we confirmed that it always exceeds 20, which is the minimum detection limit for LISA [16]. We then proceed to perform Bayesian parameter estimation on this data using both the vacuum template and the perturbed template, which takes the environmental effect into account. Our MCMC setup uses a red-blue stretch move proposal with stretch parameter $a = 1.4$. This value produces the desired acceptance rates between 0.4 and 0.6. We use 25 chains to ensure that there are twice as many chains as there are dimensions in the estimation, and to enable the red-blue proposal. We initialise an adaptive temperature ladder with three temperatures, where the lowest and highest correspond to $\beta = 1$ and $\beta = 0$. Finally, we run the algorithm for 3500 iterations to ensure that our chains are much longer than a hundred times the autocorrelation time. The priors are defined as in Tab. 6 with the δ parameter being set to $\delta \in \{0.01, 0.1, 0.5\}$ depending on the posterior uncertainty. We initialise the chains at starting points close to the "true" values.

Parameter	**Prior distribution**
$\log M$	$U[(1-\delta)\cdot\log\bar{M},(1+\delta)\cdot\log\bar{M}]$
$\log\mu$	$U[(1-\delta)\cdot\log\bar{\mu},(1+\delta)\cdot\log\bar{\mu}]$
a	$U[(1-\delta)\cdot\bar{a},(1+\delta)\cdot\bar{a}]$
p_0	$U[(1-\delta)\cdot\bar{p}_0,(1+\delta)\cdot\bar{p}_0]$
d_L	$U[0.,10.]$
$\cos\theta_S$	$U[-0.999999,0.999999]$
ϕ_S	$U[0,2\pi]$
$\cos\theta_K$	$U[-0.999999,0.999999]$
ϕ_K	$U[0,2\pi]$
Φ_{ϕ_0}	$U[0,2\pi]$
A	$U[10^{-3},10^{3}]$
n_r	$U[-10,20]$

Table 6: Priors for the parameter estimation. Barred values correspond to the injection parameters, and U denotes the uniform distribution. The δ parameter controls the width of the uniform prior distribution. It is set to a value $\delta \in \{0.01, 0.1, 0.5\}$ to ensure that the priors are broad enough to contain the full posterior, but small enough to reduce computational time.

This process returns the MCMC chains, which form the unnormalised posterior distributions. These distributions are then used to compute the evidence and ultimately the Bayes factors. Fig. 17 summarises the multiple frameworks discussed in the previous sections.

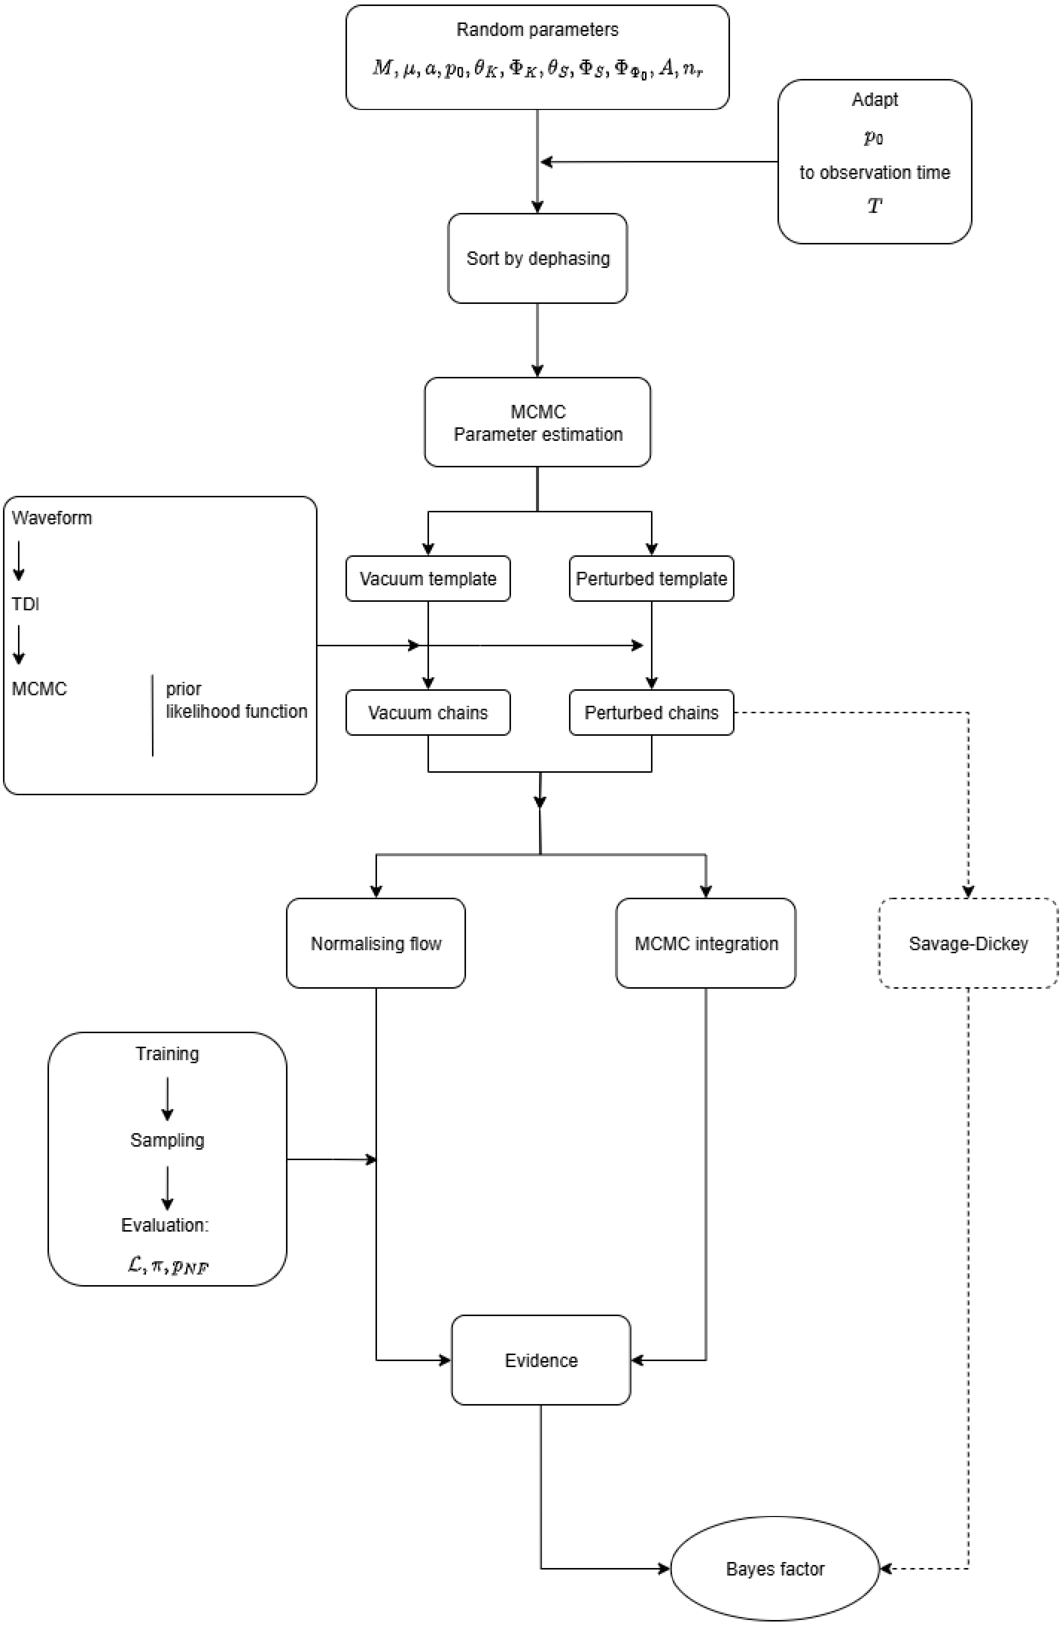

Figure 17: Flowchart of Bayes factor computation for model selection.

Dominant Regime of Environmental Effect

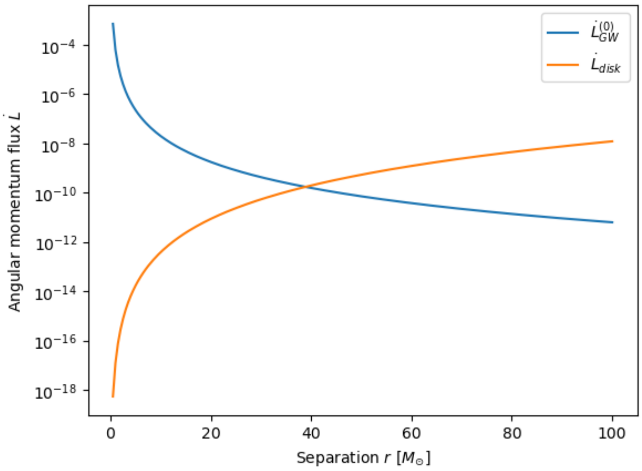

Figure 18: Example comparison of angular momentum fluxes. In the power law model, the dissipation due to the migration effect dominates at large separations of the two objects, i.e. in the early phase of the inspiral. The provided example uses the typical values from [53][39] with mass ratio $\epsilon = 10^{-6}$, amplitude $A = 1.92 \cdot 10^{-5}$ and power law slope $n_r = 8$.

The flux of angular momentum continuously decreases the orbital separation, driving the inspiral. For our model there are two phenomena, which cause such an angular momentum dissipation. The emission of GW and the torque on the secondary due to migration effect in the accretion disk. The latter is modelled as given in eq. 22. Because the considered parameterisations have positive power-law exponents n_r and amplitudes A, the dissipation is proportional to the separation between the objects. Further, the effect is scaled by the leading-order circular orbit GW angular momentum dissipation. Therefore, the environmental effect dominates the dissipation at large separations or early times and becomes negligible in the last part of the inspiral at small separations. This behaviour is illustrated in Fig. 18. This radial dependency allows predictions about the significance of the environmental effect on the inspiral. If it is observed close to the merger (i.e. at small separations), the environmental effect is negligible and may not be detectable. Therefore, in this study, we assume that the signal begins at least one year prior to the merger.

6.2 Dephasing

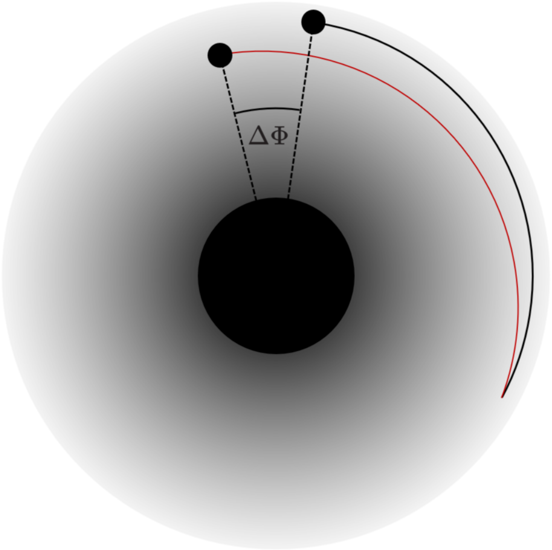

Figure 19: Schematic representation of dephasing. The same system evolves differently in the presence of an environment. The red line depicts the perturbed trajectory and the black one the unperturbed case. The dephasing is the difference in the azimuthal phase.

The first measure to quantify this perturbation is through the so-called dephasing. It describes the azimuthal phase difference between a system which evolves with an environmental effect and one without (see scheme in Fig. 19). Dephasing accumulates during the inspiral, and the value considered in this work is the cumulative dephasing up to the end of the inspiral (i.e. when the secondary reaches the primary's ISCO and "plunges"). Due to its cumulative nature, dephasing depends not only on the amplitude A and slope n_r of the perturbative effect but also on the observation time T (see Fig. 20), or equivalently on the initial separation p_0. Conversely, the environmental effect decreases during the inspiral (see Fig. 18), such that the dephasing only varies comparatively slowly for most of the inspiral.

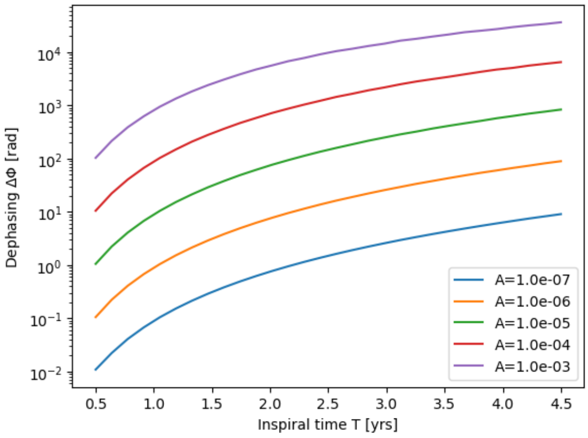

(a) Variable amplitude parameter A, fixed n_r

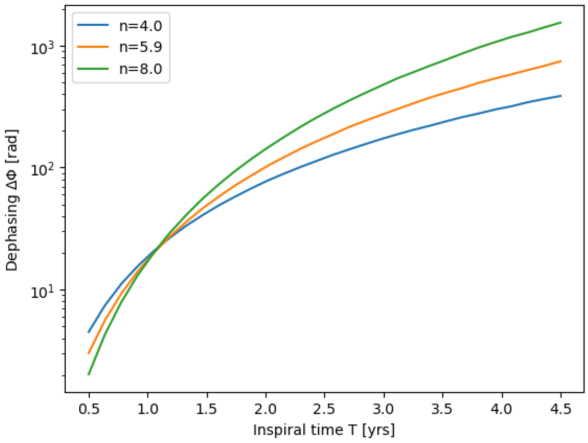

(b) Variable slope parameter n_r, fixed A

Figure 20: Dephasing accumulates over observation time. While the amplitude controls the magnitude of the dephasing (see Fig. 20(a)), the slope parameter controls how much dephasing is accumulated in the early and late phases of the inspiral(see Fig. 20(b)).

Dephasing is a quick and accessible measure to compare the sensitivity of systems to the environmental effect. We use it to identify candidate systems for the model selection process. The proposed parameter sets and the corresponding dephasings after four years of signal before plunge are given in Tab. 7. Still, the dephasing is not directly observed and is an insufficient measure to distinguish environmentally perturbed EMRIs from vacuum EMRIs.

6.3 Parameter Estimation

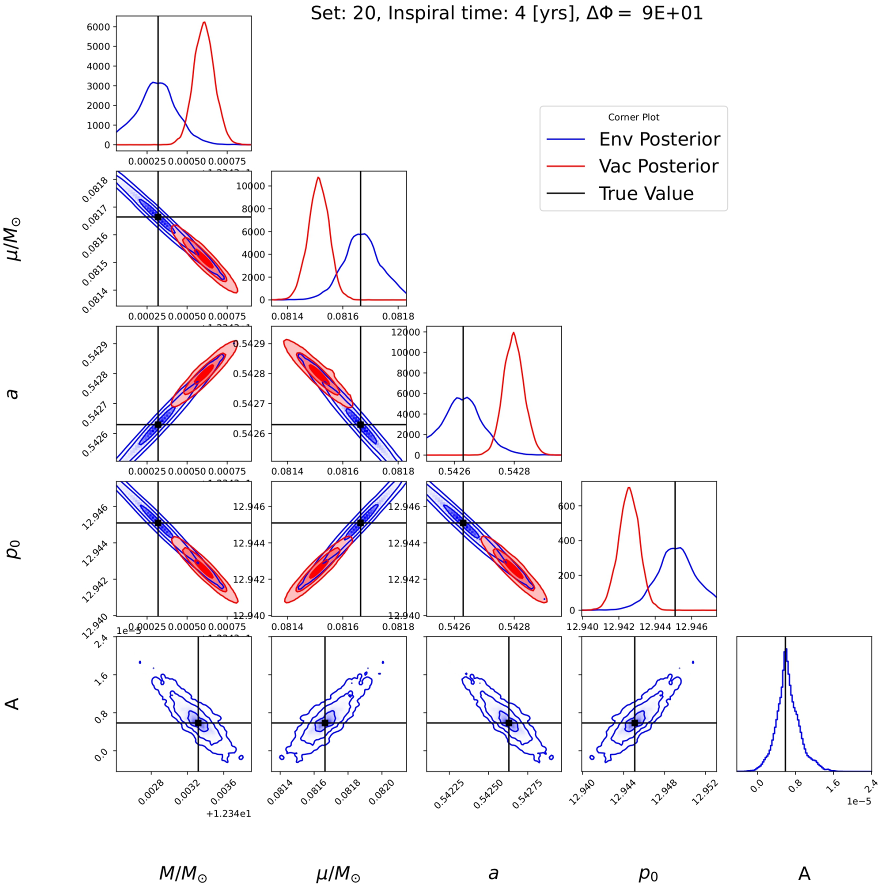

Figure 21: Bias in the parameter estimation due to an environmental effect. We present the results of the parameter estimation for set 20 with four years of observation time. The two templates converge and produce an offset in the posterior distributions. The vacuum template produces a biased estimate when attempting to account for the increased inspiral rate.

The MCMC sampling yields the approximate posterior distributions of the model parameters. Because the data is simulated, we have access to the true values and use the sampling results to understand the parameter's variances, degeneracies and the overall behaviour of the parameter space. We run the selected parameter for 1, 2, 3 and 4 years before the plunge and examine the posterior marginal distributions.

We can evaluate the impact of environmental effects on parameter estimation by comparing the posterior distributions of the vacuum and environmental templates. We reproduce the behaviour mentioned in [53]: "In particular, unaccounted (inward) migration leads to an overestimation of the mass and spin of the primary, as it increases the rate of inspiral." Additionally, we find that the vacuum template converges on a lower initial separation and a lower secondary mass. An example of this behaviour is evident in the parameter estimation result of set 20 for four years of observation time (see Fig. 21). These results suggest that neglecting the environmental effect can introduce biases in the parameter estimation for certain EMRI systems.

Furthermore, we found that, for the systems with high dephasings, the vacuum template often fails to estimate the parameters, i.e. produce a consistent posterior distribution. The posterior density distributions of these systems are uniform, as can be seen in Fig. 22. This will be very important later on when computing the Bayes factors.

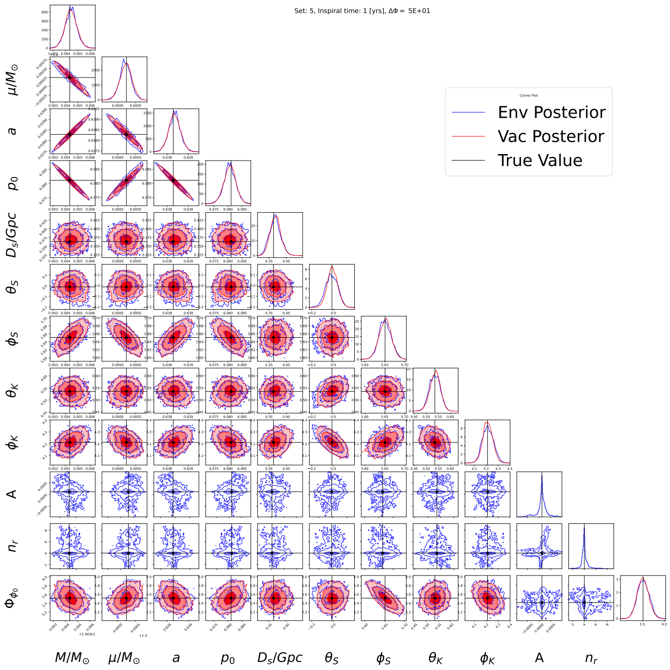

(a) Converged

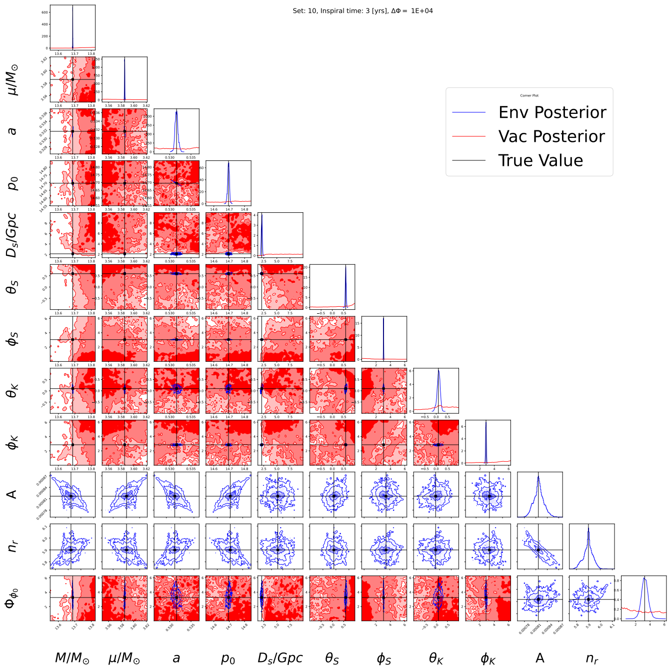

(b) Not converged

Figure 22: Corner plots of converged and non-converged parameter estimations. Fig. 22(a) displays the posterior distribution of parameter set 5 (dephasing $\Delta\Phi \approx 50$) with 1 year of data recovered using the environmental and vacuum templates. All parameters exhibit a distinct maximum, and the contours lie within the prior range. Fig. 22(b) shows the posterior distribution of parameter set 10 (dephasing $\Delta\Phi \approx 10^4$) with 3 years of data. The vacuum template fails to identify the posterior distribution. In contrast, the environmental effect template converges on the true values.

Detection and Constraints on the Amplitude of the Effect

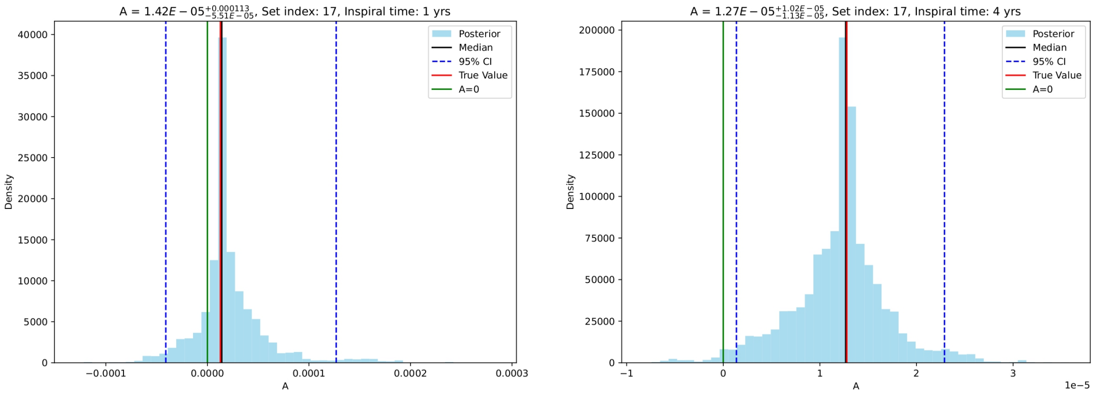

(a) 1 year of data

(b) 4 years of data

Figure 23: Histograms of amplitude parameter A estimation. We present example histograms that demonstrate how longer observation times allow for more precise constraints on the amplitude of the perturbing effect. The 95% credible intervals for parameter set 17 also illustrate the necessary observation time required to rule out an amplitude of $A = 0$, which corresponds to no effect. For this set with a dephasing of ≈ 100 [rad], around four years of signal data prior to the merger are required to rule out the no effect hypothesis and claim detection.

Using the results of the parameter estimation, we can examine the constraints on the amplitude A. This amplitude includes the properties of the central engine, the AGN, as well as properties of the disk (see Sec. 3). Constraining A constrains these properties. As explained in [53], an additional electromagnetic counterpart observation can be leveraged to break the degeneracy and constrain the AGN and disk parameters individually.

To detect the effect, we examine the posterior distributions of the amplitude parameter estimations. We can claim detection when the hypothesis of $A = 0$ is excluded from the 95% credible interval of the distribution, following the analysis in [53]. This implies that the estimation is inconsistent with an amplitude of $A = 0$. This approach is illustrated in Fig. 23 for one system.

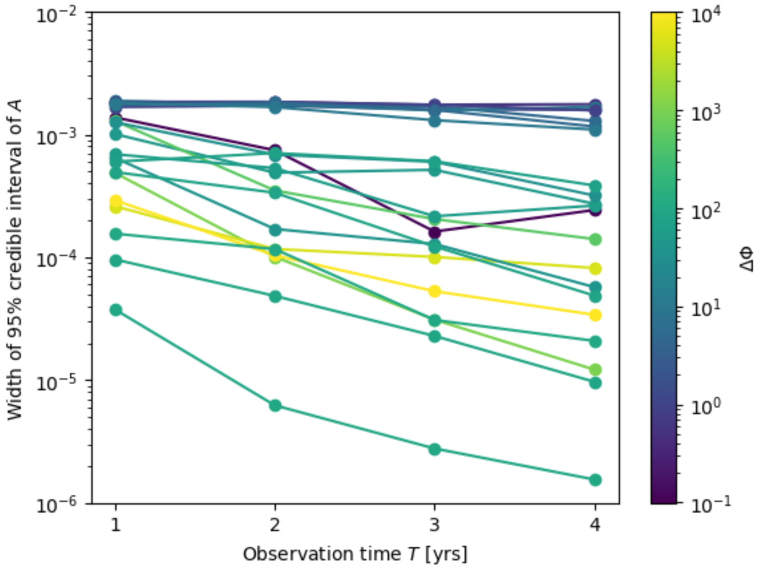

Figure 24: The width of the 95% credible interval on A as a function of the observation time. The longer a system is observed before merging, the more accurately the amplitude of the perturbative effect can be constrained. This enables us to constrain the properties of the accretion disk further.

We observe that the credible interval shrinks and the amplitude becomes better constrained with longer observation times. This enables us to impose tighter constraints on the properties of the disk (see Fig. 24 for comparison of all systems). This is because longer observation times correspond to larger separations, which is exactly where the environmental effect is strongest. Therefore if there is data available from this region, it allows us to constrain the amplitude effectively.

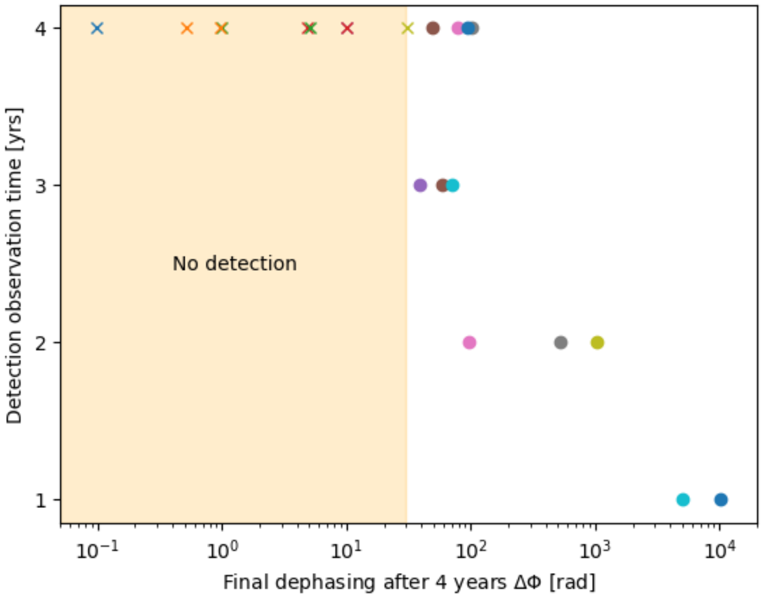

Figure 25: The required observation time for the detection of an environmental effect. We observe a correlation between strong effects with large dephasing and a short required observation time for detection. Below the threshold of ≈ 30[rad] (indicated by the orange region), it is not possible to claim detection, even with 4 years of observational data. These systems are marked with a "$\times$".

We investigate how much observation time is required to detect an environmental effect. We found a correlation between dephasing and the observation time required for detection. This correlation is depicted in Fig. 25. Furthermore, we found that, for small dephasing $\Delta\Phi \lesssim 30$, the hypothesis $A = 0$ lies within the 95% credible interval, even with 4 years of observation time, and detection cannot be claimed. This suggests that weak effects are hard to detect. It provides an approximate threshold for the amount of dephasing needed for detection within the mission lifetime of LISA.

Identifying the Perturbative Effect

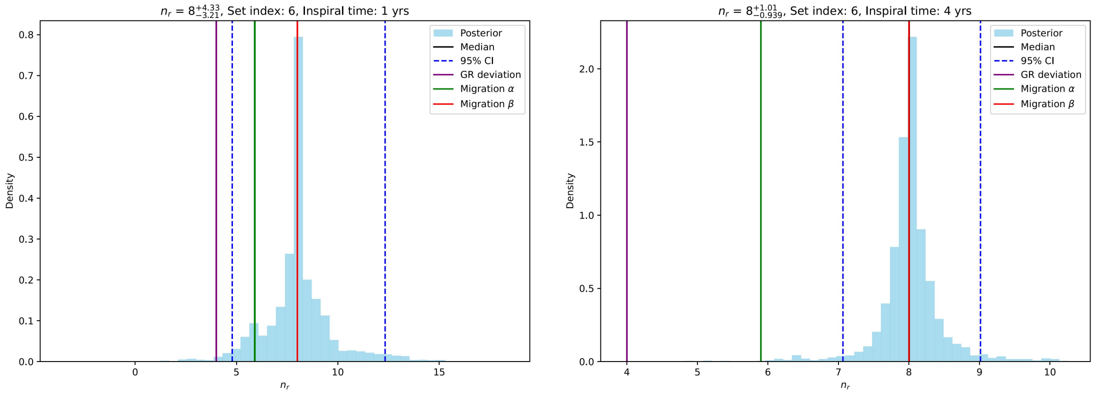

(a) 1 year of data

(b) 4 years of data

Figure 26: Histograms of the estimation of the model parameter n_r. We present an example of how the model identification increases with observation time. The histograms and 95% credible regions of the parameter set 6 are shown for 1 and 4 years of observed data before the merger. As observation time increases and more data becomes available, the value of n_r becomes increasingly constrained. This means that with 1 year of data, it is not possible to identify the type of perturbing effect, whereas with 4 years, it is.

The parameter estimation further yields the credible intervals for the slope parameter n_r. As n_r is physically constrained to the values ($n_r \in \{4, 5.9, 8\}$), the estimation results can be used to determine the underlying process. By comparing the credible regions, we can establish whether the processes could be confused. We claim identification of the perturbative process when two of the three possible values of n_r are excluded from the 95% credible region. As with detection, we can examine how much observation time is necessary to resolve the confusion. We found that identification improves with observation time, with an example given in Fig. 26.

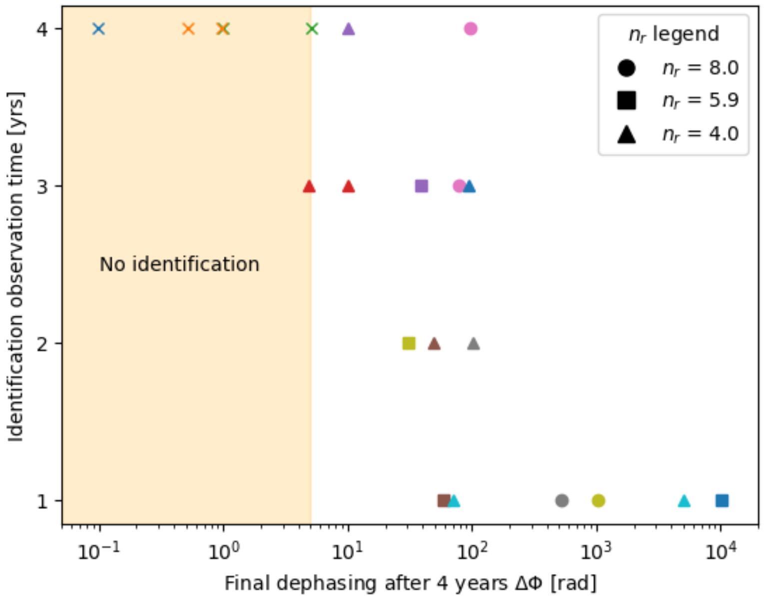

Figure 27: The required observation time before merger for identification as a function of the dephasing. We observe a loose correlation between large dephasing and a short observation time required for identification. Below the threshold of ≈ 5 [rad] (indicated by the orange region), identification is impossible, even with 4 years of observational data. These systems are marked with a "$\times$". Furthermore, we observe that the identification is independent of the process type (i.e. the slope parameter n_r).

Furthermore, we found that a larger dephasing (i.e. a stronger effect) is only weakly correlated with the required observation time for identification. This behaviour is independent of the value of the slope parameter n_r, as shown in Fig. 27. We also found that it is impossible to identify the process below a dephasing of ≈ 5[rad].

6.4 Bayes Factors

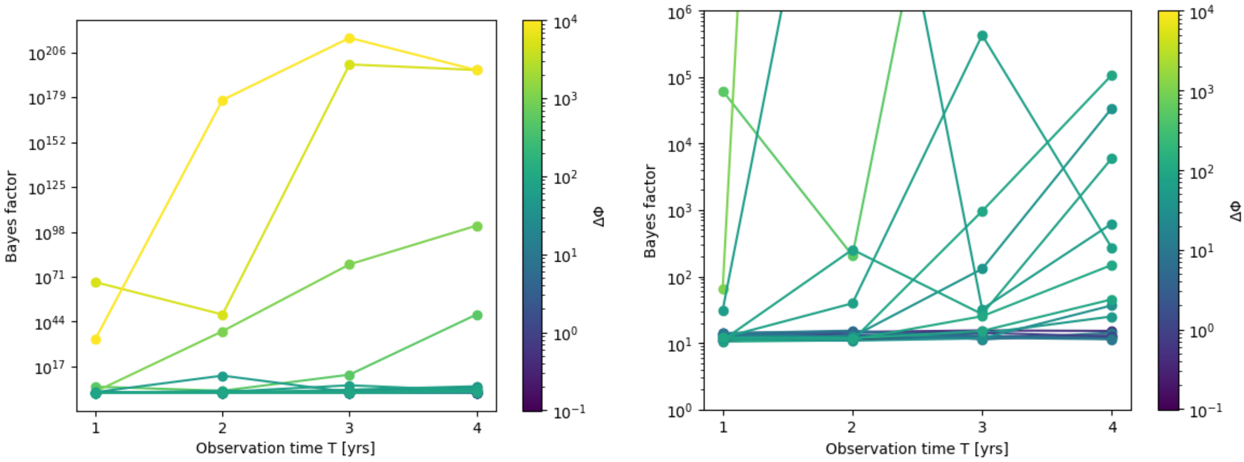

(a) All Bayes factors

(b) Zoomed-in on reasonable Bayes factors range

Figure 28: Bayes factors for different observation times. Fig. 28(a) shows how the systems with large dephasing (indicated by the colorbar) can produce unreasonably high Bayes factors. This is due to a failure in the parameter estimation using the vacuum template. Zooming in on the region with reasonable Bayes factors in Fig. 28(b), we find that some systems, with comparably low dephasing, demonstrate an approximately constant Bayes factor of ≈ 15 for all observation times. Other systems with intermediate dephasing show significantly increasing Bayes factors as the observation time increases. However, for some systems, this increase is not monotonous, suggesting mistuned parameter estimation.

We quantified the difference in support for the models with and without an environmental effect using the Bayes factor, which is the ratio of the evidence. We obtained the evidence for each parameter estimation via MCMC integration. We then examined the trends in the Bayes factors for the different observation times. We found that, for some systems, the Bayes factor remains approximately constant regardless of the observation time. For other systems, the Bayes factor steadily increases with observation time. These relations between the Bayes factor and observation time are presented in Fig. 28. We found that systems with approximately constant Bayes factors have a low dephasing and are therefore only mildly influenced by the environmental effect. In contrast, systems with increasing Bayes factors exhibit greater dephasing and are thus more strongly affected by the environmental effect.

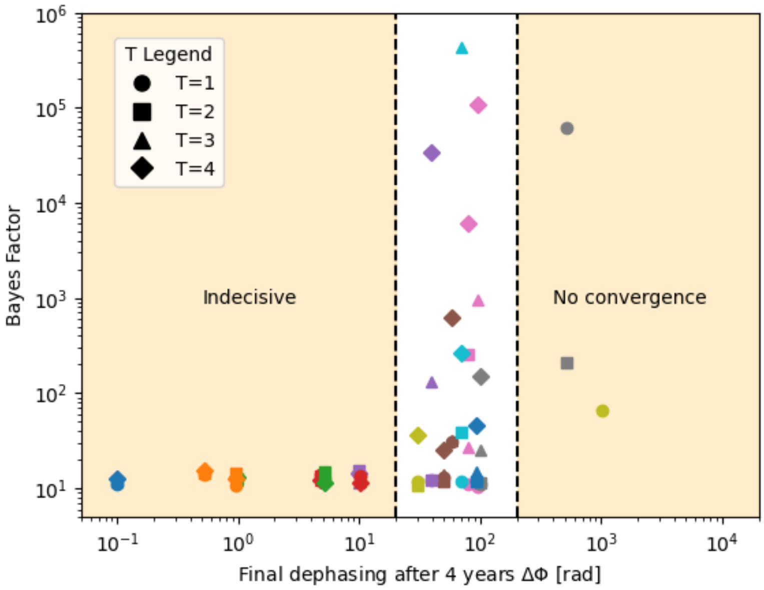

Figure 29: Bayes factors for different dephasing and observation times. We identify three regimes of low, moderate and high dephasing. In the low dephasing regime, the Bayes factors are strongly in favour of the model with an environmental effect, and remain roughly constant across different observation times. In the moderate regime, sufficient observation time is required for the Bayes factors to become decisively in favour of the model with an environmental effect. In the high dephasing regime, the Bayes factors decisively favour the model with an environmental effect. However, in this regime, the vacuum model fails to estimate the parameters, and the exact value of the Bayes factor is not statistically meaningful.

This led to the discovery that our considered systems, which exhibit effects ranging from weak to strong effects (as quantified by dephasing), can be divided into three regimes. The first regime is characterised by low dephasing between 0.1 [rad] and 20 [rad], corresponding to weak effects. Here, the Bayes factor is ≈ 15 for all observation times, thus providing strong evidence for the model that considers the environmental effect, but not being decisive (see Tab. 4). In the intermediate dephasing regime, between 20 [rad] and 200 [rad], the Bayes factors vary from ≈ 15 to ≈ 6000 depending on the observation time. With sufficiently long observation times, we can conclude that the evidence decisively favours the model with an environmental effect. In the third regime, dephasing exceeds 500 [rad] and the Bayes factors reach values of up to 10^{46}. These values are unreasonable and arise because the model that considers the vacuum case, fails to find a maximum likelihood point within the prior range. Consequently, the evidence for these parameter estimations is extremely low. In this regime, the model with an environmental effect outperforms the vacuum model, but the statistical analysis is not meaningful. These results are shown in Fig. 29.

7 Summary

In this thesis, I laid out the perspectives of observing GW originating from EMRI with the future space-based GW observatory LISA, explaining the fundamentals of the observatory and the observables that will form the data. I also discussed the scientific potential of EMRI observations, which will enable us to measure the properties and population of MBHs as well as their evolution throughout cosmic history. This will allow us to study the strong regime of gravity and test GR. I also explained how environmental effects can impact the dynamics of EMRIs and modify the emitted GW signal. In this thesis, I considered the expected dominant environmental effect: planetary(-like) migration in gaseous accretion disks. I outlined the fundamental mechanism behind this effect and introduced an agnostic parameterisation. I then discussed the theoretical basis of EMRI waveform modelling and the current state of numerical implementation. I demonstrated how the agnostic environmental effect parameterisation can be used to extend vacuum waveform models and thus model "dirty" EMRIs. Previous research has shown that unaccounted environmental effects can lead to biases in parameter estimation. The objective of this thesis project was therefore to assess the significance of environmental effects. To achieve this, I explained the principles of Bayesian inference and parameter estimation, and performed Bayesian model selection.

As accurate generic waveform models are not yet available, I worked under the simplifying assumptions of prograde, equatorial and circular orbits, as motivated by previous studies[46][21]. I simulated noise-free LISA data of "dirty" EMRIs performing type I migration in a thin accretion disk, as well as EMRIs that are described by a deviation from GR, namely a varying gravitational constant. I then performed parameter estimation using vacuum and perturbed templates. Using posterior distributions, I estimated the evidence for each template via MCMC integration and performed Bayesian model selection, i.e. I computed Bayes factors. I repeated this analysis for different strengths of effect and different EMRI parameters in the parameter space, which will be detectable by LISA.

The key results are as follows:

A longer observation time before the merger corresponds to a larger initial separation. Since planetary migration dominates at large separations, it leaves a stronger imprint on the signal. The accumulated shift in orbital phase is an easily accessible initial measure of the strength of this imprint, and therefore of the environmental effect.
Tighter constraints on the model parameters can be obtained by observing the system for a longer period before the merger. Constraints on the amplitude A also constrain the properties of the accretion disk (see Fig. 23). I found that weak effects can be undetectable, indicating that the credible interval of the amplitude A is consistent with no effect ($A = 0$). Constraints on the slope n_r allow distinction between α disks, β disks and the GR deviation. Once again, I found that weak effects do not permit such identification, implying that the credible intervals are consistent with multiple of the three possible values for the slope parameter. Both detection and identification require less observation time with increasingly strong effects. This is because a strong effect has a sufficiently strong impact on the signal, even in the final part of the inspiral, where the torque from GW emission dominates.

This behaviour is also reflected in the results of the Bayesian model selection. I found that Bayes factors increased with observation time, indicating strong support for the environmentally perturbed template. I identified a region where the Bayes factors decisively favoured the environmentally perturbed template. The effect must be strong enough to result in less accurate modelling of the waveform with the vacuum template, but not so strong that the vacuum template fails to identify the posterior distribution. In the latter case, the Bayes factors would be meaningless, as no evidence of the vacuum template could be calculated.

In this thesis project, I established dephasing as an initial measure of the strength of an environmental effect. I tested the point at which environmental effects become significant enough to make Bayesian model selection decisive, and established a link to dephasing.

8 Outlook

In general, there are two ways in which the project could be adapted to create more realistic scenarios. One approach would be to use realistic data or consider more generic signals, and identify areas where the imprints of environmental effects and deviations from GR can be confused.

More realistic data would be a superposition of many signals with underlying noise from multiple sources[16], and the parameters of the GW sources would be unknown. The first step would be to detect an EMRI in this data. Even a single source and simplified noise assumptions could pose a significant problem. Once an EMRI has been detected in such data, an initial estimate of its parameters can be made. However, this estimate would be biased by environmental or post-Einsteinian effects, so the parameter estimation process would need to distinguish between them. Thus Bayesian model selection would need to be performed again.

Alternatively, we could consider more generic signals by allowing for generic trajectories of the CO and considering other types of environmental effects, such as interaction with the stellar cloud around the MBH, dynamical friction effects, or others listed in Tab. 1. EMRIs could perform inclined, eccentric and retrograde orbits; therefore, incorporating generic trajectories is necessary. Such complex motion translates into complex GW signals, which are very challenging to model efficiently. As the EMRI population is not well constrained all types of scenarios need to be considered.

From a technical perspective, the statistical analysis performed in this study could be refined further. Bayesian model selection relies on estimates of the evidence for the posterior distributions of different templates. We calculated these quantities approximately through MCMC integration, which is easily accessible and enabled us to investigate multiple systems in a reasonable amount of time. However, this approach results in decreased accuracy. Using normalising flow techniques would allow us to achieve greater accuracy and estimate errors in the evidence and the Bayes factors. Another approach would be to adapt the sampling so that the Savage-Dickey density ratio could be applied to estimate the Bayes factor, thus bypassing the challenge of evidence calculation entirely. Lastly, the RJ-MCMC setup, which has gained popularity in GW data analysis, could be employed to address the issue of multiple competing and possibly nested models, as well as providing easily accessible estimates of the Bayes factor.

References

[1] B. P. Abbott et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". In: *Physical Review Letters* 116.6 (Feb. 2016). ISSN: 1079-7114. DOI: `10.1103/physrevlett.116.061102`. URL: `http://dx.doi.org/10.1103/PhysRevLett.116.061102`.

[2] Marek A. Abramowicz and P. Chris Fragile. "Foundations of Black Hole Accretion Disk Theory". In: *Living Reviews in Relativity* 16.1 (Jan. 2013). ISSN: 1433-8351. DOI: `10.12942/lrr-2013-1`. URL: `http://dx.doi.org/10.12942/lrr-2013-1`.

[3] M. Armano et al. "In-depth analysis of LISA Pathfinder performance results: Time evolution, noise projection, physical models, and implications for LISA". In: *Physical Review D* 110.4 (Aug. 2024). ISSN: 2470-0029. DOI: `10.1103/physrevd.110.042004`. URL: `http://dx.doi.org/10.1103/PhysRevD.110.042004`.

[4] Philip J. Armitage. "Dynamics of Protoplanetary Disks". In: *Annual Review of Astronomy and Astrophysics* 49.1 (Sept. 2011), pp. 195–236. ISSN: 1545-4282. DOI: `10.1146/annurev-astro-081710-102521`. URL: `http://dx.doi.org/10.1146/annurev-astro-081710-102521`.

[5] Philip J. Armitage. *Lecture notes on the formation and early evolution of planetary systems*. 2017. arXiv: `astro-ph/0701485 [astro-ph]`. URL: `https://arxiv.org/abs/astro-ph/0701485`.

[6] Stanislav Babak, Martin Hewitson, and Antoine Petiteau. *LISA Sensitivity and SNR Calculations*. 2021. arXiv: `2108.01167 [astro-ph.IM]`. URL: `https://arxiv.org/abs/2108.01167`.

[7] Stanislav Babak et al. "Science with the space-based interferometer LISA. V. Extreme mass-ratio inspirals". In: *Physical Review D* 95.10 (May 2017). ISSN: 2470-0029. DOI: `10.1103/physrevd.95.103012`. URL: `http://dx.doi.org/10.1103/PhysRevD.95.103012`.

[8] Leor Barack and Adam Pound. "Self-force and radiation reaction in general relativity". In: *Reports on Progress in Physics* 82.1 (Nov. 2018), p. 016904.

ISSN: 1361-6633. DOI: 10.1088/1361-6633/aae552. URL: http://dx.doi.org/10.1088/1361-6633/aae552.

[9] Enrico Barausse, Vitor Cardoso, and Paolo Pani. "Can environmental effects spoil precision gravitational-wave astrophysics?" In: *Physical Review D* 89.10 (May 2014). ISSN: 1550-2368. DOI: 10.1103/physrevd.89.104059. URL: http://dx.doi.org/10.1103/PhysRevD.89.104059.

[10] Riccardo Barbieri et al. "Constraining the evolution of Newton's constant with slow inspirals observed from spaceborne gravitational-wave detectors". In: *Physical Review D* 107.6 (Mar. 2023). ISSN: 2470-0029. DOI: 10.1103/physrevd.107.064073. URL: http://dx.doi.org/10.1103/PhysRevD.107.064073.

[11] *Black Hole Perturbation Toolkit*. (bhptoolkit.org).

[12] Jo Bovy. *What is an orbital resonance?* https://www.astro.utoronto.ca/~bovy/resonance.html. Accessed: 2025-09-02. 2025.

[13] Brandon Carter. "Global Structure of the Kerr Family of Gravitational Fields". In: *Phys. Rev.* 174 (5 Oct. 1968), pp. 1559–1571. DOI: 10.1103/PhysRev.174.1559. URL: https://link.aps.org/doi/10.1103/PhysRev.174.1559.

[14] Katie Chamberlain and Nicolás Yunes. "Theoretical physics implications of gravitational wave observation with future detectors". In: *Physical Review D* 96.8 (Oct. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.96.084039. URL: http://dx.doi.org/10.1103/PhysRevD.96.084039.

[15] Alvin J. K. Chua et al. "Rapid Generation of Fully Relativistic Extreme-Mass-Ratio-Inspiral Waveform Templates for LISA Data Analysis". In: *Physical Review Letters* 126.5 (Feb. 2021). ISSN: 1079-7114. DOI: 10.1103/physrevlett.126.051102. URL: http://dx.doi.org/10.1103/PhysRevLett.126.051102.

[16] Monica Colpi et al. *LISA Definition Study Report*. 2024. arXiv: 2402.07571 [astro-ph.CO]. URL: https://arxiv.org/abs/2402.07571.

[17] Neil Cornish et al. "Gravitational wave tests of general relativity with the parameterized post-Einsteinian framework". In: *Physical Review D* 84.6 (Sept. 2011). ISSN: 1550-2368. DOI: `10.1103/physrevd.84.062003`. URL: `http://dx.doi.org/10.1103/PhysRevD.84.062003`.

[18] *DeepL Write*. An AI-powered writing assistant for text refinement and language improvement. DeepL SE. 2025. URL: `https://www.deepl.com/write` (visited on 09/04/2025).

[19] A Derdzinski et al. "Evolution of gas disc–embedded intermediate mass ratio inspirals in theLISAband". In: *Monthly Notices of the Royal Astronomical Society* 501.3 (Dec. 2020), pp. 3540–3557. ISSN: 1365-2966. DOI: `10.1093/mnras/staa3976`. URL: `http://dx.doi.org/10.1093/mnras/staa3976`.

[20] A M Derdzinski et al. "Probing gas disc physics with LISA: simulations of an intermediate mass ratio inspiral in an accretion disc". In: *Monthly Notices of the Royal Astronomical Society* 486.2 (Apr. 2019), pp. 2754–2765. ISSN: 0035-8711. DOI: `10.1093/mnras/stz1026`. eprint: `https://academic.oup.com/mnras/article-pdf/486/2/2754/30256792/stz1026.pdf`. URL: `https://doi.org/10.1093/mnras/stz1026`.

[21] Andrea Derdzinski and Lucio Mayer. "In situextreme mass ratio inspirals via subparsec formation and migration of stars in thin, gravitationally unstable AGN discs". In: *Monthly Notices of the Royal Astronomical Society* 521.3 (Mar. 2023), pp. 4522–4543. ISSN: 1365-2966. DOI: `10.1093/mnras/stad749`. URL: `http://dx.doi.org/10.1093/mnras/stad749`.

[22] James M. Dickey and B. P. Lientz. "The Weighted Likelihood Ratio, Sharp Hypotheses about Chances, the Order of a Markov Chain". In: *The Annals of Mathematical Statistics* 41.1 (1970), pp. 214–226. DOI: `10.1214/aoms/1177697203`. URL: `https://doi.org/10.1214/aoms/1177697203`.

[23] Alexander J Dittmann and M Coleman Miller. "Star formation in accretion discs and SMBH growth". In: *Monthly Notices of the Royal Astronomical Society* 493.3 (Feb. 2020), pp. 3732–3743. ISSN: 1365-2966. DOI: `10.1093/mnras/staa463`. URL: `http://dx.doi.org/10.1093/mnras/staa463`.

[24] Paul C. Duffell et al. "THE MIGRATION OF GAP-OPENING PLANETS IS NOT LOCKED TO VISCOUS DISK EVOLUTION". In: *The Astrophysical Journal Letters* 792.1 (Aug. 2014), p. L10. DOI: `10.1088/2041-8205/792/1/L10`. URL: `https://dx.doi.org/10.1088/2041-8205/792/1/L10`.

[25] Christoph Dürmann and Wilhelm Kley. "The accretion of migrating giant planets". In: *Astronomy amp; Astrophysics* 598 (Feb. 2017), A80. ISSN: 1432-0746. DOI: `10.1051/0004-6361/201629074`. URL: `http://dx.doi.org/10.1051/0004-6361/201629074`.

[26] Daniel Foreman-Mackey et al. "emcee: The MCMC Hammer". In: *Publications of the Astronomical Society of the Pacific* 125.925 (Mar. 2013), p. 306. DOI: `10.1086/670067`. arXiv: `1202.3665 [astro-ph.IM]`.

[27] Kostas Glampedakis, Scott A. Hughes, and Daniel Kennefick. "Approximating the inspiral of test bodies into Kerr black holes". In: *Physical Review D* 66.6 (Sept. 2002). ISSN: 1089-4918. DOI: `10.1103/physrevd.66.064005`. URL: `http://dx.doi.org/10.1103/PhysRevD.66.064005`.

[28] Jonathan Goodman and Jonathan Weare. "Ensemble samplers with affine invariance". In: *Communications in Applied Mathematics and Computational Science* 5.1 (2010), pp. 65–80. DOI: `10.2140/camcos.2010.5.65`. URL: `https://msp.org/camcos/2010/5-1/p04.xhtml`.

[29] PETER J. GREEN. "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination". In: *Biometrika* 82.4 (Dec. 1995), pp. 711–732. ISSN: 0006-3444. DOI: `10.1093/biomet/82.4.711`. eprint: `https://academic.oup.com/biomet/article-pdf/82/4/711/699533/82-4-711.pdf`. URL: `https://doi.org/10.1093/biomet/82.4.711`.

[30] Olaf Hartwig. *Instrumental modelling and noise reduction algorithms for the Laser Interferometer Space Antenna*. 2021.

[31] Scott A. Hughes et al. "Adiabatic waveforms for extreme mass-ratio inspirals via multivoice decomposition in time and frequency". In: *Physical Review D* 103.10 (May 2021). ISSN: 2470-0029. DOI: `10.1103/physrevd.103.104014`. URL: `http://dx.doi.org/10.1103/PhysRevD.103.104014`.

[32] Janosh. *Normalizing Flow Diagram*. https://github.com/janosh/diagrams/blob/main/assets/normalizing-flow/normalizing-flow-hd.png. Accessed: 2025-07-04. 2025.

[33] Kazuhiro D. Kanagawa, Hidekazu Tanaka, and Ewa Szuszkiewicz. "Radial Migration of Gap-opening Planets in Protoplanetary Disks. I. The Case of a Single Planet". In: *The Astrophysical Journal* 861.2 (July 2018), p. 140. ISSN: 1538-4357. DOI: 10.3847/1538-4357/aac8d9. URL: http://dx.doi.org/10.3847/1538-4357/aac8d9.

[34] Nikolaos Karnesis et al. "Eryn: a multipurpose sampler for Bayesian inference". In: *Monthly Notices of the Royal Astronomical Society* 526.4 (Sept. 2023), pp. 4814–4830. ISSN: 1365-2966. DOI: 10.1093/mnras/stad2939. URL: http://dx.doi.org/10.1093/mnras/stad2939.

[35] Robert E. Kass and Adrian E. Raftery and. "Bayes Factors". In: *Journal of the American Statistical Association* 90.430 (1995), pp. 773–795. DOI: 10.1080/01621459.1995.10476572. eprint: https://www.tandfonline.com/doi/pdf/10.1080/01621459.1995.10476572. URL: https://www.tandfonline.com/doi/abs/10.1080/01621459.1995.10476572.

[36] Michael Katz, Nikolaos Karnesis, and Natalia Korsakova. *mikekatz04/Eryn: first full release*. Version v1.0.0. Mar. 2023. DOI: 10.5281/zenodo.7705496. URL: https://doi.org/10.5281/zenodo.7705496.

[37] Michael L. Katz et al. "Assessing the data-analysis impact of LISA orbit approximations using a GPU-accelerated response model". In: *Physical Review D* 106.10 (Nov. 2022). ISSN: 2470-0029. DOI: 10.1103/physrevd.106.103001. URL: http://dx.doi.org/10.1103/PhysRevD.106.103001.

[38] Michael L. Katz et al. "Fast extreme-mass-ratio-inspiral waveforms: New tools for millihertz gravitational-wave data analysis". In: *Physical Review D* 104.6 (Sept. 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.104.064047. URL: http://dx.doi.org/10.1103/PhysRevD.104.064047.

[39] Hassan Khalvati et al. "Impact of relativistic waveforms in LISA's science objectives with extreme-mass-ratio inspirals". In: *Physical Review D* 111.8 (Apr. 2025). ISSN: 2470-0029. DOI: 10.1103/physrevd.111.082010. URL: http://dx.doi.org/10.1103/PhysRevD.111.082010.

[40] Bence Kocsis, Nicolás Yunes, and Abraham Loeb. "Observable signatures of extreme mass-ratio inspiral black hole binaries embedded in thin accretion disks". In: *Physical Review D* 84.2 (July 2011). ISSN: 1550-2368. DOI: `10.1103/physrevd.84.024032`. URL: `http://dx.doi.org/10.1103/PhysRevD.84.024032`.

[41] Nicolas Lartillot and Hervé Philippe. "Computing Bayes Factors Using Thermodynamic Integration". In: *Systematic Biology* 55.2 (Apr. 2006), pp. 195–207. ISSN: 1063-5157. DOI: `10.1080/10635150500433722`. eprint: `https://academic.oup.com/sysbio/article-pdf/55/2/195/26557316/10635150500433722.pdf`. URL: `https://doi.org/10.1080/10635150500433722`.

[42] Jaewook Lee, Woosuk Sung, and Joo-Ho Choi. "Metamodel for Efficient Estimation of Capacity-Fade Uncertainty in Li-Ion Batteries for Electric Vehicles". In: *Energies* 8.6 (2015), pp. 5538–5554. ISSN: 1996-1073. DOI: `10.3390/en8065538`. URL: `https://www.mdpi.com/1996-1073/8/6/5538`.

[43] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford University Press, 2007. ISBN: 978-0-19-171766-6, 978-0-19-852074-0. DOI: `10.1093/acprof:oso/9780198570745.001.0001`.

[44] C J Moore, R H Cole, and C P L Berry. "Gravitational-wave sensitivity curves". In: *Classical and Quantum Gravity* 32.1 (Dec. 2014), p. 015014. ISSN: 1361-6382. DOI: `10.1088/0264-9381/32/1/015014`. URL: `http://dx.doi.org/10.1088/0264-9381/32/1/015014`.

[45] Sijme-Jan Paardekooper et al. *Planet-Disk Interactions*. 2022. arXiv: `2203.09595 [astro-ph.EP]`. URL: `https://arxiv.org/abs/2203.09595`.

[46] Zhen Pan and Huan Yang. "Formation rate of extreme mass ratio inspirals in active galactic nuclei". In: *Physical Review D* 103.10 (May 2021). ISSN: 2470-0029. DOI: `10.1103/physrevd.103.103018`. URL: `http://dx.doi.org/10.1103/PhysRevD.103.103018`.

[47] George Papamakarios et al. "Normalizing Flows for Probabilistic Modeling and Inference". In: *Journal of Machine Learning Research* 22.57 (2021), pp. 1–64. URL: `http://jmlr.org/papers/v22/19-1028.html`.

[48] Alicja Polanska et al. *Learned harmonic mean estimation of the Bayesian evidence with normalizing flows*. 2025. arXiv: `2405.05969 [astro-ph.IM]`. URL: `https://arxiv.org/abs/2405.05969`.

[49] Adam Pound and Barry Wardell. "Black Hole Perturbation Theory and Gravitational Self-Force". In: *Handbook of Gravitational Wave Astronomy*. Springer Singapore, 2021, pp. 1–119. ISBN: 9789811547027. DOI: `10.1007/978-981-15-4702-7_38-1`. URL: `http://dx.doi.org/10.1007/978-981-15-4702-7_38-1`.

[50] Chiara E Scardoni et al. "Type II migration strikes back – an old paradigm for planet migration in discs". In: *Monthly Notices of the Royal Astronomical Society* 492.1 (Dec. 2019), pp. 1318–1328. ISSN: 1365-2966. DOI: `10.1093/mnras/stz3534`. URL: `http://dx.doi.org/10.1093/mnras/stz3534`.

[51] W Schmidt. "Celestial mechanics in Kerr spacetime". In: *Classical and Quantum Gravity* 19.10 (Apr. 2002), pp. 2743–2764. ISSN: 0264-9381. DOI: `10.1088/0264-9381/19/10/314`. URL: `http://dx.doi.org/10.1088/0264-9381/19/10/314`.

[52] N. I. Shakura and R. A. Sunyaev. "Black holes in binary systems. Observational appearance." In: 24 (Jan. 1973), pp. 337–355.

[53] Lorenzo Speri et al. "Probing Accretion Physics with Gravitational Waves". In: *Physical Review X* 13.2 (June 2023). ISSN: 2160-3308. DOI: `10.1103/physrevx.13.021035`. URL: `http://dx.doi.org/10.1103/PhysRevX.13.021035`.

[54] Hidekazu Tanaka, Taku Takeuchi, and William R. Ward. "Three-Dimensional Interaction between a Planet and an Isothermal Gaseous Disk. I. Corotation and Lindblad Torques and Planet Migration". In: *The Astrophysical Journal* 565.2 (Feb. 2002), p. 1257. DOI: `10.1086/324713`. URL: `https://dx.doi.org/10.1086/324713`.

[55] Saul A. Teukolsky. "Perturbations of a Rotating Black Hole. I. Fundamental Equations for Gravitational, Electromagnetic, and Neutrino-Field Perturbations". In: *Astrophysical Journal* 185 (Oct. 1973), pp. 635–648. DOI: `10.1086/152444`.

[56] Massimo Tinto and Sanjeev V. Dhurandhar. “Time-delay interferometry”. In: *Living Reviews in Relativity* 24.1 (Dec. 2020), p. 1. ISSN: 1433-8351. DOI: `10.1007/s41114-020-00029-6`. URL: `https://doi.org/10.1007/s41114-020-00029-6`.

[57] Michele Vallisneri. “Synthetic LISA: Simulating time delay interferometry in a model LISA”. In: *Physical Review D* 71.2 (Jan. 2005). ISSN: 1550-2368. DOI: `10.1103/physrevd.71.022001`. URL: `http://dx.doi.org/10.1103/PhysRevD.71.022001`.

[58] Ziming Wang et al. “Simultaneous bounds on the gravitational dipole radiation and varying gravitational constant from compact binary inspirals”. In: *Physics Letters B* 834 (Nov. 2022), p. 137416. ISSN: 0370-2693. DOI: `10.1016/j.physletb.2022.137416`. URL: `http://dx.doi.org/10.1016/j.physletb.2022.137416`.

[59] Niels Warburton et al. *KerrGeodesics*. Version 0.9.0. 2023. DOI: `10.5281/zenodo.8108265`. URL: `https://doi.org/10.5281/zenodo.8108265`.

[60] William R. Ward. “Protoplanet Migration by Nebula Tides”. In: *Icarus* 126.2 (1997), pp. 261–281. ISSN: 0019-1035. DOI: `https://doi.org/10.1006/icar.1996.5647`. URL: `https://www.sciencedirect.com/science/article/pii/S001910359695647X`.

[61] Wangang Xie et al. “Improving Marginal Likelihood Estimation for Bayesian Phylogenetic Model Selection”. In: *Systematic Biology* 60.2 (Dec. 2010), pp. 150–160. ISSN: 1063-5157. DOI: `10.1093/sysbio/syq085`. eprint: `https://academic.oup.com/sysbio/article-pdf/60/2/150/24552358/syq085.pdf`. URL: `https://doi.org/10.1093/sysbio/syq085`.

[62] Nicolás Yunes and Frans Pretorius. “Fundamental theoretical bias in gravitational wave astrophysics and the parametrized post-Einsteinian framework”. In: *Physical Review D* 80.12 (Dec. 2009). ISSN: 1550-2368. DOI: `10.1103/physrevd.80.122003`. URL: `http://dx.doi.org/10.1103/PhysRevD.80.122003`.

[63] Nicolás Yunes, Frans Pretorius, and David Spergel. “Constraining the evolutionary history of Newton’s constant with gravitational wave observations”. In: *Physical Review D* 81.6 (Mar. 2010). ISSN: 1550-2368. DOI:

10.1103/physrevd.81.064018. URL: `http://dx.doi.org/10.1103/PhysRevD.81.064018`.

[64] Nicolás Yunes et al. "Extreme mass-ratio inspirals in the effective-one-body approach: Quasicircular, equatorial orbits around a spinning black hole". In: *Physical Review D* 83.4 (Feb. 2011). ISSN: 1550-2368. DOI: 10.1103/physrevd.83.044044. URL: `http://dx.doi.org/10.1103/PhysRevD.83.044044`.

[65] Nicolás Yunes et al. "Imprint of Accretion Disk-Induced Migration on Gravitational Waves from Extreme Mass Ratio Inspirals". In: *Physical Review Letters* 107.17 (Oct. 2011). ISSN: 1079-7114. DOI: 10.1103/physrevlett.107.171103. URL: `http://dx.doi.org/10.1103/PhysRevLett.107.171103`.

9 Acknowledgement

I would like to thank my supervisors at the Astroparticle and Cosmology Laboratory APC in Paris, Natalia Korsakova and Quentin Baghi, for their support and patience. I would further like to thank the whole Gravitation group at APC for accommodating me and becoming my friends along the way. I also like to thank Laura Sagunski, who enabled my stay abroad by helping me obtain a scholarship and who kindly agreed to review this thesis. I also want to thank Lorenzo Speri and Hassan Khalvati for the useful discussions.
Finally, I want to thank my family and friends for their unconditional support and the love they showed me.

10 Eigenständigkeitserklärung

Hiermit erkläre ich, dass ich die Arbeit selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel verfasst habe. Alle Stellen der Arbeit, die wörtlich oder sinngemäß aus Veröffentlichungen oder aus anderen fremden Texten entnommen wurden, sind von mir als solche kenntlich gemacht worden. Ferner erkläre ich, dass die Arbeit nicht - auch nicht auszugsweise - für eine andere Prüfung verwendet wurde.

Frankfurt am Main, den 05.09.2025

A Parameter Sets

Set	$\Delta\Phi$	$M/M_\odot$	$\mu/M_\odot$	a	p_0	D_s/Gpc	θ_S	ϕ_S	θ_K	ϕ_K	Φ_{ϕ_0}	**A**	n_r
0	0.098933	22418563	1357.302	0.761757	7.690761	0.168044	0.548768	3.652331	2.15918	2.543616	0.202715	7.34E-06	5.9
1	0.523677	24164740	21.29367	0.0311	6.413126	0.00075	1.130588	0.347047	0.768215	2.329809	0.313648	4.42E-05	4
2	0.979508	18759732	6.86151	0.179946	5.801451	0.000963	2.176784	1.593124	0.315807	4.534196	0.777893	1.64E-06	4
3	4.815271	7314824	26.32854	0.741846	5.438068	0.214022	1.372386	4.355944	1.522358	0.939964	1.031934	0.000206	4
4	10.00362	12831792	103.9956	0.206238	7.337099	0.019296	0.458888	3.617416	1.097212	3.93202	4.985544	0.00042	4
5	49.11035	1244624	4.056184	0.638284	7.960268	0.363133	1.583717	5.6507	1.005856	4.211529	5.515511	0.000118	4
6	95.59698	319140.3	3.573816	0.882943	14.16507	0.363117	1.197431	4.341244	0.47146	0.606024	1.08961	1.05E-06	8
7	517.6456	3549434	201.3383	0.796188	11.73546	1.032049	2.37392	2.289245	2.677491	3.473567	5.57356	0.000393	8
8	1024.816	357592.5	2.002429	0.036858	13.07368	0.362066	1.969308	4.311553	2.74117	3.459257	2.382526	9.32E-05	8
9	4988.031	1544475	83.64545	0.952836	14.05438	1.612572	0.445894	5.418243	0.690651	2.375298	0.525891	0.000724	4
10	10107.11	879695.1	36.07532	0.53161	15.74777	2.111487	0.939184	3.042049	1.472884	2.85118	3.256173	0.000817	5.9
11	0.957627	17341148	5.967198	0.064204	6.18116	0.000621	2.681882	4.113389	0.030759	3.841514	2.79715	2.03E-05	4
12	5.166436	7585480	36.1331	0.232795	7.256523	0.052618	1.794968	5.562349	1.498194	1.634896	0.694756	0.000262	5.9
13	10.14056	6368842	8.61162	0.288568	6.291576	0.028238	0.699124	1.178115	1.325155	5.654642	4.48611	0.000319	4
14	39.16409	4365714	403.5097	0.032295	13.92274	0.166814	0.766112	4.914105	2.26859	0.722508	0.704489	4.32E-05	5.9
15	58.56316	4867876	215.4171	0.822736	10.14368	0.607181	1.140157	4.038511	1.548025	2.5511	3.270218	0.000148	5.9
16	78.86281	2344804	98.99892	0.2268	13.10716	0.76419	2.082908	5.632099	0.481758	2.000576	0.89094	3.69E-05	8
17	100.8696	727545.3	17.48519	0.529976	14.49971	1.277957	0.471162	2.515964	1.818283	4.940436	2.078411	1.28E-05	4
18	30.51608	4238863	107.5062	0.791365	9.211959	0.909786	2.777658	2.971076	2.176275	3.902703	3.973612	0.00013	5.9
19	70.03691	5203144	266.4053	0.789177	10.41477	0.906208	0.344877	0.529888	2.644924	1.687361	4.768129	0.000156	4
20	93.22843	229422.1	1.085091	0.542629	12.94509	0.185128	0.442751	2.967745	1.766549	1.123101	0.083594	5.86E-06	4

Table 7: The injected parameter sets. The initial separation value, p_0, is adjusted for each estimation run to account for observation times of $T \in \{1, 2, 3, 4\}$ years.