Gravitational Lensing of Gravitational Waves: From Detection to Cosmology and Population Inference

Eungwang Seo

Submitted in fulfilment of the requirements for the Degree of Doctor of Philosophy

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

September 2026

Call to me

And I will answer you

And tell you great and unsearchable things you do not know

— Jeremiah 33:3

Abstract

The direct detection of gravitational waves has transformed our view of the Universe, elevating them from a confirmation of General Relativity to a tool for precision astrophysics and cosmology. Unlike electromagnetic radiation, which is subject to absorption, scattering, and extinction by baryonic matter, gravitational waves propagate through the Universe largely unimpeded, carrying relatively clean information about both their sources and the intervening gravitational structures.

Gravitational lensing, another prediction of General Relativity, naturally arises along such propagation paths. Acting as a cosmic telescope, it magnifies distant signals while encoding the lens properties into the observed waveform. This phenomenon has long been exploited in electromagnetic observations to probe dark matter, the expansion history of the Universe, and the population properties of compact objects. When applied to gravitational waves, lensing provides a complementary probe of these same questions, with sensitivity to regimes that remain inaccessible to light-based observations.

In this thesis, I investigate both the detection of gravitationally lensed gravitational waves and their scientific applications. Current searches focus on two regimes: strong lensing, in which massive structures such as galaxies or clusters produce multiple images with time delays ranging from minutes to years, and microlensing, induced by smaller-scale structures such as stars, compact remnants, or dark matter subhalos, which imprint frequency-dependent interference patterns due to the wave nature of gravitational waves. Up to the third observing run of the LIGO–Virgo–KAGRA Collaboration, no compelling evidence for gravitational-wave lensing has been reported. I present results from a search for isolated point-mass microlensing in the first part of the fourth observing run, highlighting an interesting candidate event, GW231123_135430.

I then transition to a more astrophysically realistic lensing scenario. While an isolated point-mass lens provides a useful approximation for the effect of an individual microlens, microlensing is expected to occur within more complex lens environments, where compact objects can be embedded in a larger-scale gravitational potential. I therefore extend the search to scenarios involving combined strong lensing and microlensing. The complex and highly degenerate nature of such signals poses a challenge for existing model-dependent pipelines. To address this, a model-independent search framework based on residual tests is developed, enabling rapid identification of candidate events, several orders of magnitude faster than fully model-dependent analyses, while remaining sensitive to a broad range of lensing-induced deviations from unlensed waveforms.

Once detected, gravitationally lensed gravitational waves offer powerful applications in cosmology. One of the central challenges in modern cosmology is the Hubble tension, the discrepancy between measurements of the Hubble constant inferred from the early Universe and those obtained from local distance indicators. Gravitational waves provide an independent way to measure cosmic distances through their amplitude, without relying on the traditional distance ladder. In the presence of strong lensing, this capability is further enhanced: combining multiple lensed images with an accurate reconstruction of the lens mass distribution significantly improves the precision of luminosity distance estimates, offering a promising avenue to alleviate this tension.

In addition to cosmological applications, gravitational-wave lensing provides a direct means to probe the population of compact objects and the small-scale structure of matter. Because gravitational waves are insensitive to baryonic effects such as dust extinction, lensing signatures reflect the underlying gravitational field alone, enabling a mapping of the otherwise invisible components of the Universe. Such measurements constrain the abundance and mass function of intermediate-mass black holes, primordial black holes, and dark matter subhalos, thereby providing insights into both compact-object formation channels and the small-scale structure of dark matter, which are difficult to access through electromagnetic observations only.

As we look toward the next decade of gravitational-wave astronomy—with the continued operation of the LVK network and the anticipation of next-generation detectors and space-based detectors—the detection of a definitively lensed GW event is no longer a matter of if but when. This work establishes the essential methodologies required to both identify such events and fully exploit their scientific potential, paving the way toward precision cosmology and a deeper understanding of the dark Universe.

Contents

List of Tables

List of Figures

Acknowledgements

This thesis could not have been completed without faith, hope, and love. I am deeply grateful for the opportunity to have carried out the research contained within these pages. First and foremost, I offer my praise and gratitude to Almighty God, who granted me the strength to persevere and surrounded me with love and wonderful people both in Glasgow and across the world. Without His Spirit, I would not have found the hope and resilience to endure this journey.

I must express my deepest appreciation to my supervisor, Professor Martin Hendry. You have guided me with great care, shaping me into a better gravitational-wave researcher. Your unwavering support and constant encouragement enabled me to navigate and complete the demanding and often exhausting path of a PhD. I am incredibly proud to have been your student. I would also like to extend my sincere thanks to Professor Graham Woan. Whenever I felt discouraged or overwhelmed, you always provided encouragement and reassurance. I will always cherish your practical advice and your reassuring words that you would stand by me.

Beyond Glasgow, I am deeply grateful to the academic advisors. I am particularly thankful to Professor Tjonnie G. F. Li, who has provided continuous encouragement and guidance since my MPhil and has patiently supported my development as a good and confident researcher. I also thoroughly enjoyed collaborating with Professor Justin Janquart during my PhD. Although I had known you before my PhD, it was through our close collaboration that I came to truly appreciate your insight, meticulousness, and thoughtful approach to research. I have learned a great deal from working with you. I would also like to give my heartfelt thanks to Doctor Kyungmin Kim. It has been a true privilege to work together on the basis of mutual trust and to achieve fruitful results. Your comment that you genuinely enjoyed working with me remains deeply meaningful, and it has encouraged me to continue pursuing my goal of becoming a better researcher. Additionally, I wish to extend my gratitude to Professors Otto A. Hannuksela and Jose M. Ezquiaga for your academic guidance and insightful discussions. I greatly appreciate the opportunities to interact and collaborate with you during my visits to Hong Kong and Denmark, which helped open new avenues in my research and contributed significantly to my academic development.

Next, I am also fortunate to have shared my PhD journey with wonderful friends and colleagues. First, to my dear friends, Doctors Mick Wright and Ania Liu: learning about gravitational-wave lensing together was a joy—even though I certainly learned far more from you both—and your patience and encouragement kept me grounded. I must also thank the entire Institute for Gravitational Research community at the University of Glasgow.

Learning how to communicate and collaborate as a scientist among such brilliant faculty members has been a privilege, and I am grateful for the respect and kindness I was always shown. Additionally, I would also like to thank my dear friends in Room 253—Storm Colloms, Varun Deshmukh, Victoria Graham, Jessica Irwin, Ross Johnston, Natalia Koziol, August Miller, Naren Nagarajan, Thejas Seetharamu, Josh Sharkey, Leigh Smith, Federico Stachurski, and Elizabeth Todd—for making my PhD experience so warm and memorable. I truly hope our paths cross again in the future. Special thanks also go to Xiaofei Dong, Qian Hu, and Zhuotao Li, with whom I shared not only the ups and downs of PhD life, but also many conversations, meals, and moments of comfort along the way.

I would also like to thank my examiners, Professors Christopher Berry and Chris Clarkson, for asking such profound and thought-provoking questions during my viva. Although I was certainly nervous, I genuinely enjoyed engaging with your questions, and the discussion made me even more enthusiastic about the field. It reminded me how much I enjoy doing research in this area and made me feel that pursuing a PhD in this field was absolutely the right choice. Your comments on my thesis have been immensely helpful in improving it and making it a more thorough and polished piece of work. I am truly grateful for your thoughtful feedback, which will make this thesis a meaningful and lasting record of my academic journey.

To all my co-authors and collaborators with whom I have had the opportunity to publish our work, this thesis would simply not have been possible without your collective expertise and support. I look forward to continuing our work together in the years to come.

Finally, my time in Glasgow would not have been complete without Pastor Sunghoon Kim, Ms. Mi-kyung Seo, and my precious friends at the Glasgow Korean Church. The fellowship and encouragement shared within the community, grounded in God's love, became a meaningful and important part of my doctoral journey. I am deeply grateful for your support, which helped me persevere through challenges and encouraged me to live with patience and love.

Lastly, I would like to thank my mother and grandmother. Thank you for your endless love and support, which has been a constant source of strength throughout my studies abroad. I am deeply grateful for everything you have done for me, and I wish you continued health and happiness. Although we may be physically apart, you are always in my thoughts.

Declaration

I declare that, except where explicit reference is made to the contributions of others, this dissertation is the result of my own work carried out under the supervision of Professors Martin Hendry and Graham Woan at the University of Glasgow, and has not been submitted for any other degree at the University of Glasgow or any other institution. All chapters of this thesis were written with editorial input from Professor Martin Hendry.

Chapters 6 to 9 include material derived in part from papers that have been published, submitted, or are currently under review in peer-reviewed journals. Chapter 10 includes material that has neither been published nor submitted to a peer-reviewed journal. Unless explicitly stated as "Figure is from [reference]" or "Adapted from [reference]", all figures and plots in Chapters 1–10 were generated using figure-generation tools or scripts written by the author.

Abbreviations

- **aLIGO** Advanced Laser Interferometer Gravitational-Wave Observatory
- **AdV** Advanced Virgo detector
- **BAO** baryon acoustic oscillations
- **BBH** binary black hole
- **BH** black hole
- **BNS** binary neutron star
- **BWD** binary white dwarf
- **C.I.** credible interval
- **CCSN** core-collapse supernova
- **CDM** cold dark matter
- **CMB** cosmic microwave background
- **CW** continuous wave
- **DESI** Dark Energy Spectroscopic Instrument
- **DM** dark matter
- **DR** data release
- **EFE** Einstein field equation
- **EM** electromagnetic
- **EOB** effective one-body
- **FAP** false alarm probability
- **FAR** false alarm rate
- **FFT** fast Fourier transform
- **FLRW** Friedmann–Lemaître–Robertson–Walker
- **FPP** false positive probability
- **GR** general relativity
- **GRB** gamma-ray burst
- **GW** gravitational wave
- **GWTC** Gravitational-Wave Transient Catalog
- **HBA** hierarchical Bayesian analysis
- **HOM** higher-order mode
- **IMBH** intermediate-mass black hole
- **ISCO** innermost stable circular orbit
- **JPE** joint parameter estimation
- **KDE** kernel density estimate
- **ΛCDM** Lambda-cold dark matter

- **LOS** line-of-sight
- **LVK** LIGO-Virgo-KAGRA
- **MCMC** Markov Chain Monte Carlo
- **MH** Metropolis-Hastings
- **NFW** Navarro–Frenk–White
- **NIE** non-singular isothermal ellipsoid
- **NR** numerical relativity
- **NSBH** neutron star-black hole
- **PBH** primordial black hole
- **PM** point mass
- **PN** post-Newtonian
- **PNFW** pseudo-elliptical Navarro–Frenk–White
- **PSD** power spectral density
- **QNM** quasi-normal modes
- **ROC** receiver operating characteristic
- **SFR** star formation rate
- **SGWB** stochastic gravitational-wave background
- **SIE** singular isothermal ellipsoid
- **SIS** singular isothermal sphere
- **SLGW** strongly lensed gravitational wave
- **SN** supernova
- **SNR** signal-to-noise ratio
- **SPEMD** singular power-law elliptical mass distribution
- **TT** transverse-traceless

Part I

Theoretical and Methodological Framework

Chapter 1

Gravitational Waves: From Source to Detection

The exploration of the Universe relies on various physical carriers of information, often referred to as *messengers*. Electromagnetic (EM) radiation has long served as the most established observational window, spanning an extensive range from solar and galactic stellar emission to high-redshift extragalactic sources and the Cosmic Microwave Background (CMB) [1]. In addition to EM waves, neutrinos offer a unique probe into high-energy astrophysical processes [2]. While solar neutrinos have been extensively studied, providing insights into the nuclear fusion processes within the Sun [3, 4], only a small number of extragalactic neutrino sources have been identified to date [5, 6].

A transformative addition to this multimessenger landscape is gravitational waves (GWs). The existence of GWs, defined as ripples in spacetime curvature propagating at the speed of light, was fundamentally predicted by Albert Einstein [7, 8] within the framework of general relativity (GR) [9, 10]. These waves are generated by the acceleration of asymmetric mass distributions. Since gravity is intrinsically weak, as reflected in the small value of the gravitational constant G, detectable GW amplitudes are only produced by extremely dynamic, high-energy astrophysical events. Main detectable sources include catastrophic stellar collapses such as supernovae (SNe) and gamma-ray bursts (GRBs), the inspiral and merger of compact binary systems comprising neutron stars (NSs) or black holes (BHs), and the stochastic gravitational-wave background (SGWB) originating from cosmological inflation.

The experimental efforts to detect GWs began with Joseph Weber's development of resonant-mass detectors [11]. However, these early attempts did not detect any decisive GW signals, primarily because the expected event rate of detectable GW sources within the local universe was insufficient for the sensitivity of the apparatus. A pivotal turning point occurred in 1974 with the discovery of the Hulse–Taylor binary pulsar [12]. The observed orbital decay of this system matched the energy loss predicted by GR via GW emission, providing the first indirect confirmation of their existence and galvanising the development of more sensitive detection technologies.

The decades of experimental refinement culminated in September 2015, when the Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) [13] achieved the first direct observation of a GW signal, designated as GW150914 [14]. This signal, detected by both the LIGO-Livingston and LIGO-Hanford detectors, was identified as the merger of a binary black hole (BBH) system, marking a milestone in GW astronomy. Almost two years later, the two aLIGO detectors and Advanced Virgo (AdV) [15] detected GW170817, a binary neutron star (BNS) merger with EM counterparts, heralding the advent of multimessenger astronomy [16, 17]. Furthermore, [18] announced the first detections of two neutron star–black hole (NSBH) mergers, GW200105 and GW200115. Although some NSBH mergers could produce EM counterparts if the NS is tidally disrupted before being swallowed by the BH, none were confidently identified [19].

More than a decade has passed since the first detection of GWs. As of the time of writing, the LIGO-Virgo-KAGRA (LVK) Collaboration has reported a few hundred GW events [20], signifying a transition from the discovery of individual sources to an era of population-level inference.

In this chapter, we provide a comprehensive overview of the theoretical foundations of GWs and categorise the diverse astrophysical and cosmological sources that have been identified to date and can be detected in the future. Furthermore, we discuss the instrumental principles of laser interferometry and how to find GW signals in the observed data.

1.1 Gravitational Radiation in Linearised GR

Within the framework of GR, GWs can be understood as radiation in the form of propagating disturbances in the fabric of spacetime, sourced by time-varying mass–energy distributions. Although their generation occurs in regions of intense gravitational interaction, their subsequent evolution depends on the ambient spacetime through which they propagate. Specifically, while the emission process is governed by the nonlinear dynamics of the strong-field regime near the source, the subsequent propagation toward distant observers in the asymptotically flat universe, where the spacetime curvature is sufficiently weak, is well described by linear perturbations of a background metric.

1.1.1 Einstein Field Equation

GR can be concisely defined as a geometric theory of gravitation, in which the curvature of spacetime is dynamically coupled to the energy and momentum of matter and radiation via the *Einstein Field Equations* (EFE) [10].

To define the spacetime curvature, we first need the mathematical framework to describe the four-dimensional manifold of our Universe, i.e., a space with three spatial dimensions and one time dimension. Throughout this work, we adopt the metric signature $(-,+,+,+)$ and employ the Einstein summation convention, in which repeated indices—one in the upper (contravariant) and one in the lower (covariant) position—are implicitly summed over their range[1].

The metric tensor, $g_{\mu\nu}$, defines the spacetime interval and causal structure in curved spacetime. The infinitesimal displacement between two events, ds, is given by the line element:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \tag{1.1}$$

where dx^μ and dx^ν represent small displacements along the spacetime coordinates and $g_{\mu\nu}$ encodes the local geometric properties of spacetime. In the limit of flat spacetime, the metric reduces to the Minkowski metric, $\eta_{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1)$.

Within the framework of GR, spacetime is described not by Euclidean geometry but by a pseudo-Riemannian manifold, a topological space that is locally Minkowskian but globally curved [21, 22]. On this manifold, the metric $g_{\mu\nu}$ varies smoothly across the coordinates, reflecting the intrinsic curvature of spacetime. Consequently, gravity is no longer a force in the Newtonian sense, but a manifestation of spacetime geometry. To describe the dynamics of fields on such a manifold, we generalise the notion of differentiation. The covariant derivative, ∇_μ, ensures that the derivative of a tensor remains a tensor—i.e., it transforms covariantly under coordinate transformations— by accounting for the variation of the coordinate basis vectors along the manifold. Unlike the partial derivative, it accounts for the variation of the basis vectors themselves. For a vector (or (1,0) tensor) $\vec{a}$, the covariant derivative is expressed as:

$$\begin{aligned} \nabla_\mu \vec{a} &= \frac{\partial a^\lambda}{\partial x^\mu} \vec{e}_\lambda + a^\lambda \frac{\partial \vec{e}_\lambda}{\partial x^\mu} \\ &= \frac{\partial a^\lambda}{\partial x^\mu} \vec{e}_\lambda + a^\lambda \Gamma^\nu_{\mu\lambda} \vec{e}_\nu, \end{aligned} \tag{1.2}$$

where $\vec{e}_\lambda$ and $\vec{e}_\nu$ are the coordinate basis vectors of $\vec{a}$ and $\Gamma^\nu_{\mu\lambda}$ are the *Christoffel symbols of the second kind*, which encode how the coordinate basis vectors vary from point to point on the manifold. They are uniquely determined by the metric and defined as [21]:

$$\Gamma^\nu_{\mu\lambda} = \frac{1}{2} g^{\nu\sigma} \left(\partial_\mu g_{\lambda\sigma} + \partial_\lambda g_{\mu\sigma} - \partial_\sigma g_{\mu\lambda} \right), \tag{1.3}$$

where $\partial_\lambda \equiv \partial / \partial x^\lambda$.

1. For example, for arbitrary covariant and contravariant four-vectors A_μ and B^μ, this summation is expressed as: $A_\mu B^\mu \equiv \sum_{\mu=0}^{3} A_\mu B^\mu$.

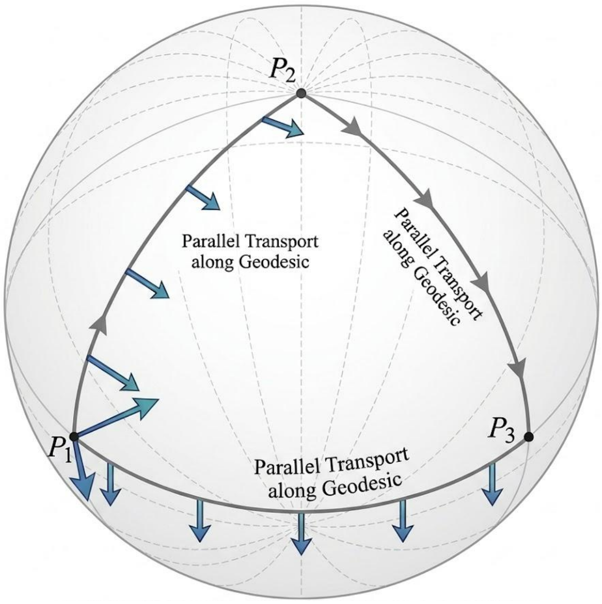

Figure 1.1: Illustration of the parallel transport of a vector along a closed path ($P_1 \rightarrow P_2 \rightarrow P_3 \rightarrow P_1$) on a curved surface.

The presence of curvature is formally indicated by the non-commutativity of covariant derivatives. As shown in Figure 1.1, a vector parallel-transported around a closed loop in a curved spacetime generally returns with a different orientation. Geometrically, this indicates that parallel transport along distinct paths between the same endpoints yields different results. This deviation is quantified by the *Riemann curvature tensor*, $R^\rho{}_{\sigma\mu\nu}$, defined as [21]:

$$R^\rho{}_{\sigma\mu\nu} = \partial_\mu \Gamma^\rho_{\sigma\nu} - \partial_\nu \Gamma^\rho_{\sigma\mu} + \Gamma^\rho_{\mu\lambda}\Gamma^\lambda_{\sigma\nu} - \Gamma^\rho_{\nu\lambda}\Gamma^\lambda_{\sigma\mu}. \tag{1.4}$$

To extract the physical components of curvature relevant to the EFE, the Riemann tensor is often contracted. The *Ricci tensor*, $R_{\mu\nu} = R^\alpha{}_{\mu\alpha\nu}$, sums over two indices, and further contraction with the metric yields the *Ricci scalar*, $R = g^{\mu\nu}R_{\mu\nu}$. While the Riemann tensor contains the full information about the spacetime curvature, the Ricci tensor and scalar provide a measure of how the local volume of a region deviates from the flat Euclidean case. However, the Ricci tensor alone cannot ensure energy–momentum conservation. According to the Bianchi identities, specifically the contracted form $\nabla_\nu(R^{\mu\nu} - \frac{1}{2}g^{\mu\nu}R) = 0$, one can construct the *Einstein tensor*, $G_{\mu\nu}$, defined as:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R. \tag{1.5}$$

The Einstein tensor is symmetric and divergence-free under covariant differentiation, which ensures local conservation of energy and momentum. It encodes the curvature of spacetime induced by matter and energy.

The distribution of matter and energy is described by the stress-energy tensor $T_{\mu\nu}$, which captures both the density and the flux of energy and momentum. For any physical matter or fluid, local conservation of energy and momentum is enforced through the vanishing covariant divergence ($\nabla_\nu T^{\mu\nu} = 0$) [21]. The consistency between the divergence-free property of $G^{\mu\nu}$ (from geometry) and $T^{\mu\nu}$ (from physics) forms the basis of the EFE:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \tag{1.6}$$

where Λ is the cosmological constant [23]. In the context of GW generation in the local universe, the effect of Λ is negligible compared to dynamical metric perturbations, and henceforth we set $\Lambda = 0$. Eq. (1.6), which are non-linear partial differential equations, dictate the reciprocal interaction between matter and geometry. In the following section, we linearise the Einstein equations and demonstrate that small metric perturbations propagate as GWs.

1.1.2 The Linearised Field Equation

In the absence of extremely compact matter sources, gravity is typically weak, and the spacetime geometry deviates only slightly from flat Minkowski spacetime. In this regime, the gravitational field can be treated as a small perturbation about a fixed background metric. This approach, known as linearised gravity, assumes that the spacetime metric differs only marginally from the Minkowski metric, $\eta_{\mu\nu}$ [24]. Accordingly, the full metric can be written as:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \tag{1.7}$$

where $h_{\mu\nu}$ represents the small metric perturbation satisfying $|h_{\mu\nu}| \ll 1$, which allows the Riemann curvature tensor in Eq. (1.4) to retain only terms up to first order in $h_{\mu\nu}$.

Using the metric in Eq. (1.7), the (0,4) Riemann curvature tensor, the Ricci tensor, and the Ricci scalar are approximated as follows:

$$R_{\alpha\mu\beta\nu} \approx \frac{1}{2}\left(\partial_\mu \partial_\beta h_{\alpha\nu} + \partial_\alpha \partial_\nu h_{\mu\beta} - \partial_\mu \partial_\nu h_{\alpha\beta} - \partial_\alpha \partial_\beta h_{\mu\nu}\right) \tag{1.8}$$

$$\begin{aligned} R_{\mu\nu} &= \eta^{\alpha\beta} R_{\alpha\mu\beta\nu} \\ &= \frac{1}{2}\left(\partial_\mu \partial_\alpha h^\alpha{}_\nu + \partial_\nu \partial_\alpha h^\alpha{}_\mu - \partial_\mu \partial_\nu h - \Box h_{\mu\nu}\right) \end{aligned} \tag{1.9}$$

$$\begin{aligned} R &= \eta^{\mu\nu} R_{\mu\nu} \\ &= \partial_\mu \partial_\nu h^{\mu\nu} - \Box h, \end{aligned} \tag{1.10}$$

where $h \equiv \eta^{\mu\nu} h_{\mu\nu}$ is the trace of the perturbation and $\Box \equiv \partial_\alpha \partial^\alpha = \nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}$ is the d'Alembertian operator. From these linearised curvature quantities, the Einstein tensor can be constructed to first order in $h_{\mu\nu}$ using Eq. (1.5). To simplify its structure, it is convenient to introduce the *trace-reversed metric perturbation* [24]:

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} h. \tag{1.11}$$

Note that $\bar{h} \equiv \eta^{\mu\nu}\bar{h}_{\mu\nu} = -h$. By substituting $\bar{h}_{\mu\nu}$ into the linearised forms of the Ricci tensor (Eq. (1.9) and scalar (Eq. (1.10)), the first-order Einstein tensor is written as [24]:

$$G_{\mu\nu} = -\frac{1}{2}\left(\Box \bar{h}_{\mu\nu} + \eta_{\mu\nu}\partial_\rho\partial_\sigma \bar{h}^{\rho\sigma} - \partial_\nu\partial_\rho \bar{h}^{\rho}{}_{\mu} - \partial_\mu\partial_\rho \bar{h}^{\rho}{}_{\nu}\right) + \mathcal{O}(h^2). \tag{1.12}$$

In the linearised regime, a coordinate transformation $x'^\mu = x^\mu + \xi^\mu$, with $\partial_\nu \xi^\mu \sim \mathcal{O}(h)$, acts as a gauge transformation of the metric perturbation: $h_{\mu\nu} \to h_{\mu\nu} - \partial_\mu \xi_\nu - \partial_\nu \xi_\mu$. To decouple the differential equations in Eq. (1.12), we impose the *harmonic gauge condition* (also known as the Lorenz gauge), $\partial^\nu \bar{h}_{\mu\nu} = 0$, which eliminates the second, third, and fourth terms on the right-hand side. In the presence of a source $T_{\mu\nu}$, the EFE in Eq. (1.6) reduce to uncoupled wave equations:

$$\Box \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}. \tag{1.13}$$

These equations, known as the *linearised field equations*, show that in the weak-field limit, the metric of spacetime behaves like a relativistic field that propagates at the speed of light, as the d'Alembertian operator directly relates the spatial and temporal second derivatives through c. In this regime, $T_{\mu\nu}$ acts as the source of gravitational radiation, analogous to how charge density acts as a source for EM waves.

1.1.3 Generation of GWs

To understand how astrophysical sources generate GWs, we need to solve the linearised field equations in the presence of a non-zero stress-energy tensor $T_{\mu\nu}$ (Eq. (1.13)) [24]. The inhomogeneous wave equation can be solved using the retarded Green's function, which accounts for the finite propagation speed of gravity, ensuring that the perturbation at a spacetime point $(t, \vec{x})$ is determined by the state of the source at an earlier retarded time, $t_r = t - |\vec{x} - \vec{y}|/c$. Here, $\vec{x}$ and $\vec{y}$ denote the spatial coordinates of the field point and the source point, respectively. The solution is given by [24]:

$$\bar{h}_{\mu\nu}(t, \vec{x}) = \frac{4G}{c^4}\int \frac{T_{\mu\nu}(t - \frac{|\vec{x}-\vec{y}|}{c}, \vec{y})}{|\vec{x} - \vec{y}|} d^3\vec{y}, \tag{1.14}$$

For most astrophysical sources, the internal velocities are non-relativistic ($v \ll c$), the observer is located in the far-field region ($|\vec{x}-\vec{y}| \gg |\vec{y}|$), and the source size is much smaller than the wavelength of the emitted GWs, so that it may be treated as effectively point-like. Under these assumptions, we can perform a multipole expansion of the retarded solution in Eq. (1.14).

While the conservation of energy and momentum forbids monopole and dipole gravitational radiation, the leading-order contribution arises from the higher-order moments of the stress-energy tensor. In the non-relativistic limit, the energy density component, $T^{00} \sim \rho c^2$, is the dominant term, while the momentum density $T^{0i} \sim \rho c v$ and the stress tensor $T^{ij} \sim \rho v^2$ are suppressed by factors of (v/c) and $(v/c)^2$, respectively. By applying the identity derived from the local conservation law $\partial_\mu T^{\mu\nu} = 0$, the volume integral of the stress components T^{ij} can be related to the second time derivative of the mass distribution as follows [24]:

$$\begin{aligned} \int T^{ij} d^3\vec{y} &= \frac{1}{2c^2}\frac{d^2}{dt^2}\int T^{00} y_i y_j d^3\vec{y} \\ &= \frac{1}{2}\frac{d^2}{dt^2} M_{ij}, \end{aligned} \tag{1.15}$$

where M_{ij} is called the *mass quadrupole moment.* In the slow-motion and far-field limits, the spatial components of $\bar{h}_{\mu\nu}$ are dominated by the second time derivative of this quadrupole moment. Substituting Eq. (1.15) into Eq. (1.14), the quadrupole radiation is written as:

$$\bar{h}_{ij}(t,\vec{x}) \approx \frac{2G}{c^4 r}\frac{d^2}{dt^2} M_{ij}(t_r). \tag{1.16}$$

The number of independent components of $\bar{h}_{\mu\nu}$ in Eq. (1.13) is initially 16. Considering the symmetry of the tensor, along with the Lorenz gauge and the vacuum condition ($\Box \bar{h}_{\mu\nu} = 0$), this number reduces to 6. Nevertheless, $\bar{h}_{\mu\nu}$ still contains unphysical degrees of freedom due to coordinate invariance. This residual gauge freedom can be further reduced by adopting the *transverse-traceless (TT) gauge*, in which the perturbation is purely spatial ($h_{0\mu} = 0$), traceless ($h^i_i = 0$), and transverse ($\partial^j h_{ij} = 0$). The TT gauge effectively eliminates all non-dynamical components, leaving the gravitational radiation with only two independent polarisation states, h_+ (plus) and $h_\times$ (cross). These polarisations encode the intrinsic geometry of the radiation and serve as the fundamental observables for dynamic astrophysical sources. For example, the amplitude and phase evolution of these two polarisations directly reflect the orbital dynamics of a compact binary, as will be discussed in Section 1.2.1.

The physical TT components of the far-field gravitational radiation can be extracted by projecting onto the plane orthogonal to the propagation direction. The projection is performed using the Lambda tensor [24, 25], constructed from the transverse projector, $P_{ij} = \delta_{ij} - n_i n_j$, where n_i is the unit vector in the direction of wave propagation:

$$\Lambda_{ij,kl} \equiv P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl} \,. \tag{1.17}$$

Applying the Lambda tensor to Eq. (1.16), and noting that the traceless condition ensures $\bar{h}^{\rm TT}_{ij} = h^{\rm TT}_{ij}$, the spatial components of the metric perturbation in the TT gauge can be written as:

$$h^{\rm TT}_{ij}\left(t, x^i\right) = \Lambda_{ij,kl} \frac{2G}{c^4 r} \frac{d^2}{dt^2} Q_{kl}(t_r). \tag{1.18}$$

Here, $Q_{kl} \equiv M_{kl} - \frac{1}{3}\delta_{kl} M_{mm}$ is called the *reduced quadrupole moment.*

1.1.4 GW Propagation

In vacuum, far from the source ($T_{\mu\nu} \approx 0$), Eq. (1.13) reduces to the homogeneous wave equation describing the propagation of the perturbation:

$$\Box \bar{h}_{\mu\nu} = 0. \tag{1.19}$$

We consider monochromatic plane-wave solutions travelling in the z-direction with an angular frequency ω. In the TT gauge, the general solution is given by [24]:

$$h^{\rm TT}_{\mu\nu}(t, z) = A^{\rm TT}_{\mu\nu} \cos\left[\omega\left(t - \frac{z}{c}\right)\right], \tag{1.20}$$

where $A_{\mu\nu}$ is a constant 4×4 symmetric amplitude tensor containing 2 independent components, given by:

$$A^{TT}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{1.21}$$

The components A_{xx} and A_{xy} are conventionally denoted h_+ and $h_\times$, respectively, corresponding to the two polarisations of a GW waveform introduced above.

Consequently, the metric perturbation can be expressed as:

$$h^{\rm TT}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos[\omega\left(t - \frac{z}{c}\right)]. \tag{1.22}$$

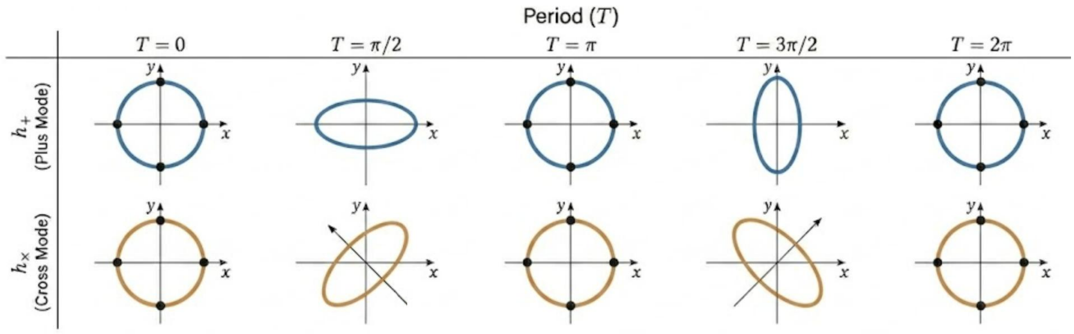

Figure 1.2: Illustration of the two GW polarisations ($h_+, h_\times$), showing the deformation of a ring of freely falling test particles over one wave period (T).

Eq. (1.22) shows that, despite the initial complexity of the 16-component metric tensor, the physical content of gravitational radiation in GR is fully captured by only two transverse, quadrupolar degrees of freedom. As the wave propagates along the z-direction, the two polarisation modes remain transverse and induce tidal deformations in the perpendicular plane. Figure 1.2 illustrates this effect for a ring of freely falling test particles, which is periodically distorted into an ellipse as the wave passes. The h_+ and $h_\times$ components correspond to two independent quadrupolar deformation patterns rotated by $\pi/4$ with respect to each other.

1.2 GW Source

While the previously established theoretical framework indicates that any asymmetric acceleration of mass generates GWs, this section focuses on signals that can be observed. Since gravity is inherently weak, spacetime responds only slightly even to large amounts of energy ($c^4/8\pi G$), detection is only feasible when the GWs originate from systems characterised by large masses and relativistic acceleration. We therefore describe the astrophysical and cosmological sources most likely to generate observable signals, classifying them according to their emission mechanisms and waveform characteristics.

1.2.1 Compact Binary Coalescences

Among GW sources, compact binary coalescences (CBCs) are currently the most frequently and robustly detected. A CBC system consists of two compact objects, such as BHs and NSs, orbiting one another and decaying via GWs. CBCs are categorised depending on their components. Namely, BBH, NSBH, and BNS [2].

2. Additionally, there are binary white dwarfs (BWDs), but their orbital frequencies are too low to be detected by ground-based interferometers. Despite that, BWDs will be the most ubiquitous continuous sources in the millihertz band for future space-based detectors [26, 27, 28].

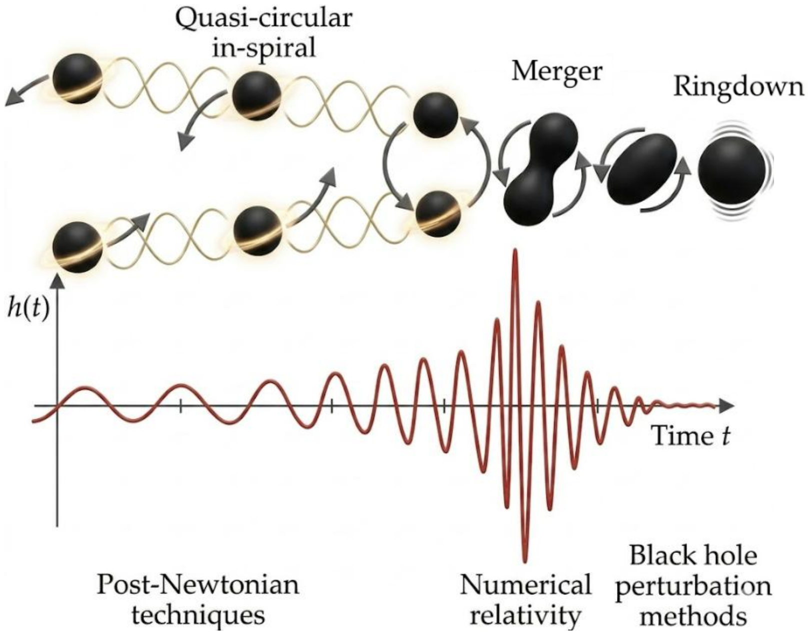

Figure 1.3: In-spiral, merger, and ringdown phases of a compact binary coalescence. Each phase is described by a different theoretical framework: post-Newtonian theory for the in-spiral, numerical relativity for the merger, and black hole perturbation theory for the ringdown. Adapted from [29].

The orbital evolution of a CBC is deterministic and can be divided into three distinct phases (see Figure 1.3) [29]. The first phase is the *in-spiral* phase, characterised by a gradual orbital decay driven by the emission of GW. To leading order, this phase can be analytically described using the quadrupole approximation for a quasi-circular binary [30]. For a system with component masses m_1 and m_2 located at a luminosity distance d_L, the emitted GW strains in the time domain are expressed as:

$$\begin{aligned} h_+(t) &= \frac{4}{d_L}\left(\frac{G\mathcal{M}_c}{c^2}\right)^{5/3}\left(\frac{\pi f}{c}\right)^{2/3}\frac{1+\cos^2\iota}{2}\cos(2\pi f t+\phi_0), \\ h_\times(t) &= \frac{4}{d_L}\left(\frac{G\mathcal{M}_c}{c^2}\right)^{5/3}\left(\frac{\pi f}{c}\right)^{2/3}\cos\iota\sin(2\pi f t+\phi_0), \end{aligned} \tag{1.23}$$

where f is the GW frequency, which is twice the orbital frequency, $f = 2f_{\text{orb}}$, ι is the inclination angle between the binary's orbital angular momentum and the observer's line of sight, ϕ_0 is the initial orbital phase, and $\mathcal{M}_c$ is the chirp mass defined as:

$$\mathcal{M}_c = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}}. \tag{1.24}$$

By equating the energy radiated by GWs to the loss of orbital binding energy, the frequency evolution with time scales as $\dot{f} \propto \mathcal{M}_c^{5/3} f^{11/3}$, implying that lower-mass systems, such as BNS, evolve much more slowly. Consequently, they remain within the sensitive frequency band of ground-based detectors for a significantly longer duration, reaching much higher frequencies before coalescence compared to more massive BBHs.

The analytical description in Eq. (1.23) remains valid until the binary separation reaches the innermost stable circular orbit (ISCO). To accurately track the phase accumulation near the ISCO, however, higher-order post-Newtonian (PN) expansions of GR are required, as the leading-order quadrupole approximation becomes insufficient in the strong-field regime [31]. For a non-spinning binary with total mass $M = m_1 + m_2$, the ISCO is located at $R_{\text{ISCO}} = 6GM/c^2$, corresponding to a maximum in-spiral frequency of:

$$f_{\text{ISCO}} \simeq \frac{c^3}{6^{3/2}\pi GM}. \tag{1.25}$$

The ISCO marks the definitive end of the stable inspiral. Beyond this point, the system enters the merger phase, where the strain amplitude reaches its maximum. This phase cannot be described analytically and requires full numerical relativity (NR) simulations for accurate modelling [29].

Finally, the newly formed remnant black hole enters the ring-down phase. During this stage, the distorted BH settles into a stationary state by emitting GWs, described by *quasi-normal modes* (QNMs) [32, 33]. They are called quasi-normal because the BH is a dissipative system. Unlike the standard normal modes of a closed system, QNMs are exponentially damped over time as energy and angular momentum radiate away. Measuring the specific frequencies and decay times of these modes is known as BH spectroscopy, enabling precision tests of the no-hair theorem in the strong-field regime [34].

1.2.2 Other sources

Other GW sources beyond CBCs are briefly introduced for pedagogical purposes, although no direct detections have yet been confirmed.

Firstly, in GW astronomy, continuous waves (CWs) are persistent, quasi-monochromatic signals with nearly constant frequency and amplitude over observational timescales. The primary astrophysical candidates include rapidly spinning, isolated NSs with non-axisymmetric deformations (often referred to as *mountains*), which may be sustained by strong magnetic stresses or elastic strains in the solid crust [24], as well as millisecond pulsars in the Milky Way center [35] and annihilating ultralight boson clouds around rotating BHs [36, 37, 38].

The main challenge in detecting CWs lies in their extremely weak strain amplitudes. Since the signal is buried well below the detector noise, detection does not rely on transient high-SNR events but instead on coherent integration over months to years. This approach requires precise matched filtering that accounts for Doppler modulations due to the Earth's rotation and orbital motion, making CW searches computationally demanding.

Secondly, GW bursts are short-duration transient signals, typically lasting from milliseconds to seconds, for which accurate waveform models are not available. The canonical sources are core-collapse supernovae (CCSNe) [39, 40]. In these systems, GW emission arises from complex, asymmetric fluid motions, including neutrino-driven convection, standing accretion shock instabilities, and oscillation modes of the proto-neutron star. While bulk motion produces low-frequency emission ($f \lesssim 200$ Hz), higher-frequency components can arise from gravity-mode (g-mode) oscillations. Other burst candidates include magnetar flares [41, 42] and cosmic string cusps [43][3].

Since reliable templates for the bursts are unavailable, matched-filtering techniques are not applicable. Instead, burst searches rely on unmodelled or weakly modelled methods. A standard approach employs coherent network algorithms, such as coherent WaveBurst (cWB) [44], which identify excess power consistent across multiple detectors while suppressing instrumental noise artifacts.

Lastly, the SGWB is an incoherent superposition of unresolved signals originating from all directions. It is often compared to the CMB, but contains both astrophysical and cosmological contributions [45].

The astrophysical SGWB arises from the cumulative emission of distant, unresolved sources such as CBCs, CCSNe, and other transients throughout cosmic history [46, 47, 48]. A major development occurred in 2023, when several pulsar timing array collaborations reported strong evidence for a nanohertz SGWB, broadly consistent with expectations for a population of inspiraling supermassive BBH [49, 50, 51, 52].

The cosmological SGWB originates from the early Universe and carries information from epochs preceding the formation of the CMB [45, 53]. Proposed sources include amplified quantum tensor fluctuations during inflation [54, 55], first-order phase transitions [56, 57], and cosmic strings [58, 59]. Detecting this component remains a major goal of modern cosmology, as it would provide a direct observational probe of fundamental physics at extremely high energies.

3. Short GRBs are primarily associated with BNS or NSBH mergers rather than unmodelled burst sources.

1.3 Detection Principles

The fundamental principle of GW detection rests on measuring the geodesic deviation between free-falling test masses [21]. To study the physical effects of a passing GW, we consider two particles situated in the TT frame. Particle 1 is placed at the origin $(t, 0, 0, 0)$ and Particle 2 is located at $(t, \varepsilon, 0, 0)$ with $\varepsilon \ll 1$, and both particles initially at rest at $t = 0$. When a GW propagates along the z-direction, the proper distance Δl between the two particles is given by the line element:

$$\begin{aligned}
\Delta l = \int \sqrt{|ds^2|} &= \left| \int g_{\mu\nu} dx^\mu dx^\nu \right|^{1/2} \\
&= \int_0^\varepsilon |g_{xx}|^{1/2}\, dx \\
&= \int_0^\varepsilon \left|\eta_{xx} + h_{xx}^{\mathrm{TT}}\right|^{1/2} dx \\
&\simeq \left[1 + \frac{1}{2} h_{xx}^{\mathrm{TT}}(x=0)\right] \varepsilon.
\end{aligned} \tag{1.26}$$

The Δl is not constant but varies in time because h_{xx}^{TT} itself is time-dependent. As discussed in the previous section, the metric perturbation propagates with distinct plus $(+)$ and cross $(\times)$ polarisation states. Consequently, the proper distance between the test particles does not change arbitrarily, but oscillates harmonically, exhibiting a characteristic quadrupolar modulation.

To analyse this effect locally, we adopt the *proper detector frame*, a locally inertial coordinate system centered on the detector in which the metric reduces to the flat Minkowski form at the origin, and its first derivatives vanish [24, 60]. In this frame, as the GW propagates across the particles at c, the components of the metric tensor are perturbed. For ground-based detectors, the instrument size is much smaller than the typical GW wavelength. Hence, the spatial gradient of the metric perturbation across the detector is negligible, allowing the local strain to be approximated as a spatially uniform and purely time-dependent strain, $h(t)$ [24, 61]. In this section, we review the principles of GW detection with interferometers, along with the methods used to extract and process the signals.

1.3.1 Laser Interferometry

To translate the strain $h(t)$ to physical length changes of the detector arms, the geometric response of the interferometer to an incoming GW must be specified. The measured strain corresponds to the projection of the GW metric perturbation onto the detector tensor D^{ij} [62]:

$$h(t) = D^{ij} h_{ij}(t), \tag{1.27}$$

where $h_{ij}(t)$ denotes the metric perturbation of the GW and the detector tensor is constructed from the unit vectors $\hat{u}$ and $\hat{v}$ pointing along the two interferometer arms as follows:

$$D_{ij} = \frac{1}{2}\left(u_i u_j - v_i v_j\right). \tag{1.28}$$

To explicitly calculate the projection in Eq. (1.27), h_{ij} is rotated from its propagation frame into the Earth-based detector frame, which depends on the source's sky location, typically defined by its right ascension (α) and declination (δ), and the polarisation angle (ψ) of the GW [61]. As shown in Eq. (1.22), the GW signal consists of two independent polarisation states, h_+ and $h_\times$, and the angular coefficients resulting from the geometric projection are defined as the *antenna pattern functions*, F_+ and $F_\times$, which quantify the directional sensitivity of the detector to each polarisation state.

When the two interferometer arms are orthogonal ($\hat{u} \cdot \hat{v} = 0$), as is the case for the aLIGO, AdV, and KAGRA [63, 64], the resulting antenna patterns are expressed as:

$$\begin{aligned} F_+(\alpha,\delta,\psi) &= \frac{1}{2}\left(1+\cos^2\alpha\right)\cos 2\delta\cos 2\psi - \cos\alpha\sin 2\delta\sin 2\psi\ , \\ F_\times(\alpha,\delta,\psi) &= \frac{1}{2}\left(1+\cos^2\alpha\right)\cos 2\delta\sin 2\psi + \cos\alpha\sin 2\delta\cos 2\psi. \end{aligned} \tag{1.29}$$

Since the antenna patterns are continuous functions of the sky coordinates, ground-based interferometers are effectively un-targeted, all-sky monitors. However, their sensitivity is highly anisotropic, peaking at zenith and nadir and dropping to zero at specific locations along the detector plane. This limitation can be mitigated by combining multiple detectors into a network. The differing orientations and geographical locations of the detectors lead to complementary antenna patterns, thereby reducing directional blind spots and improving the overall sky coverage and detection sensitivity.

For convenience, the antenna response is often expressed in terms of an *amplitude pattern*, defined as $R(\alpha,\delta) \equiv \sqrt{< F_+^2 >_\psi + < F_\times^2 >_\psi}$, which characterises the direction-dependent sensitivity of the detector independent of polarisation [65]. Figure 1.4 shows the corresponding amplitude patterns for a single detector and for the current ground-based detector network, including the operating LIGO and Virgo interferometers, as well as the forthcoming KAGRA detector.

Ultimately, the strain $h(t)$ measured by a single detector is the linear combination of the two independent GW polarisations, weighted by the corresponding antenna pattern functions:

$$h(t) = F_+(\alpha,\delta,\psi)h_+(t) + F_\times(\alpha,\delta,\psi)h_\times(t). \tag{1.30}$$

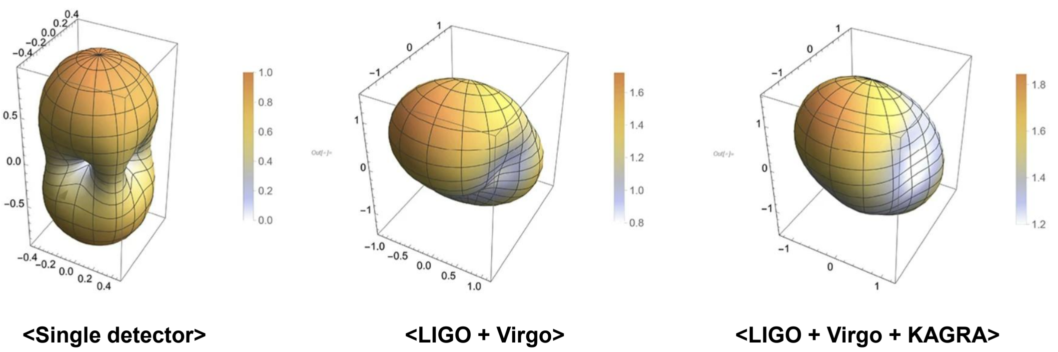

Figure 1.4: Illustration of the amplitude patterns, for a single detector, the LIGO–Virgo network, and the LIGO–Virgo–KAGRA network. The colour bar indicates the direction-dependent sensitivity, i.e., $R(\alpha, \delta)$. Figure is from [65].

For a detector with an arm length L, the metric perturbation induced by the GW produces a fractional change in the proper length, $\delta L(t)$, which can be approximated as:

$$\delta L(t) \approx \frac{1}{2} h(t) L. \tag{1.31}$$

Typical astrophysical sources, such as CBCs, produce a dimensionless strain amplitude on the order of $h \sim 10^{-21}$ when reaching Earth [24]. For a kilometer-scale detector ($L \sim 10^3$ m), this strain corresponds to an absolute length variation of approximately 10^{-18} m, a scale significantly smaller than the diameter of a proton. Measuring length changes at this sub-nuclear scale requires differential optical phase measurements with laser interferometers [66].

To measure the tiny changes in arm length induced by a passing GW, current ground-based detectors employ a power-recycled Michelson interferometer configuration [67]. In this setup, a coherent laser beam is divided by a beam splitter and directed along two orthogonal arms of length L. Figure 1.5 shows the simplified schematic of the current LIGO detectors.

GW's quadrupolar nature differentially stretches one arm while compressing the other by $\delta L(t)$. When the laser beam travels down an arm and reflects back, it covers a total round-trip distance of $2L$. Consequently, the length variation introduces a total path difference of $2\delta L$. Multiplying this path difference by the wavenumber ($k = 2\pi/\lambda_{\text{laser}}$) yields the relative phase shift $\Delta\phi$ between the recombining beams as follows:

$$\Delta\phi(t) = \frac{4\pi}{\lambda_{\text{laser}}} \delta L(t), \tag{1.32}$$

where λ_{laser} is the wavelength of the laser. The phase shift modulates the interference pattern observed at the photodetector, converting the spacetime strain into a measurable optical power fluctuation [66].

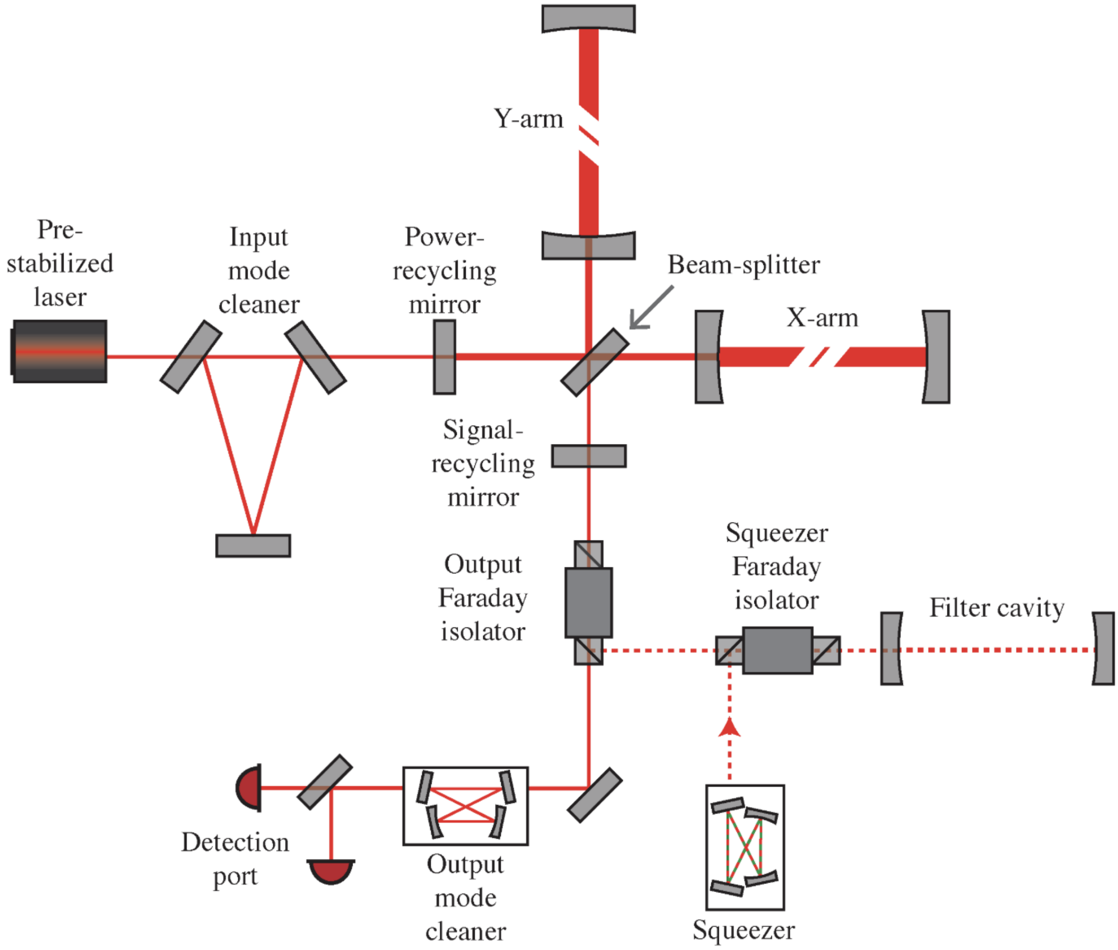

Figure 1.5: Simplified optical layout of the LIGO detectors during the fourth observing run, showing the squeezed-light injection, main interferometer, and output signal path. The main interferometer is a dual-recycled Michelson interferometer with Fabry–Pérot arm cavities. Figure is from [68].

However, the strain $h(t)$ is not a static deformation but oscillates with a period $T_{\mathrm{GW}} = 1/f_{\mathrm{GW}}$. If a photon remains in the interferometer arm for longer than half of this period ($T_{\mathrm{GW}}/2$), the GW phase changes from stretching to compressing, so the accumulated optical phase shift begins to cancel. To maximise the measurable signal, the photon's round-trip travel time ($\tau = 2L/c$) must perfectly match a half-period of the GW:

$$\frac{2L}{c} = \frac{T_{\mathrm{GW}}}{2} \quad \Rightarrow \quad L = \frac{c}{4f_{\mathrm{GW}}} = \frac{\lambda_{\mathrm{GW}}}{4}. \tag{1.33}$$

During the late inspiral and merger, the GW frequencies of stellar-mass BBHs and BNSs sweep through the 10–10^4 Hz band [24]. To optimise the detector for a typical 100 Hz signal within this band, the ideal arm length would be $L = (3\times10^5\ \mathrm{km/s})/(4\times100\ \mathrm{Hz}) = 750\ \mathrm{km}$. Constructing a vacuum baseline of this scale is geographically infeasible.

To overcome the physical length problem, modern interferometers incorporate Fabry-Perot cavities within each arm [69]. By placing highly reflective mirrors at both ends of the arms, the laser beam is trapped and traverses the baseline multiple times before recombining at the beam splitter. As the cavity stores photons, the laser power is enhanced, effectively increasing the optical path length from the physical length L to an effective length L_{eff}:

$$L_{\text{eff}} \approx \frac{2\mathcal{F}}{\pi} L, \tag{1.34}$$

where $\mathcal{F}$ denotes the finesse of the cavity, a dimensionless parameter determined by the mirrors' reflectivity. By effectively extending the arm length, the accumulated optical phase shift is amplified, thereby enhancing the detector signal-to-noise ratio (SNR).

However, as the effective length increases, the photon storage time in the cavity increases as well. As noted above, when the storage time becomes comparable to the GW period, the positive and negative strains begin to cancel the accumulated optical phase shift, which sets a characteristic frequency scale, the *cavity pole frequency*. For GW signals below this pole frequency, the phase shift is coherently amplified. For signals above the pole frequency, the detector's response begins to attenuate. Therefore, by carefully designing the finesse $\mathcal{F}$, the cavity pole frequency can be tuned so that the detector's peak sensitivity accurately aligns with the target 10–10^4 Hz band, achieving optimal observational capability without requiring impractically long arm lengths [66, 70].

1.3.2 Signal Processing

In practice, any physical measurement is fundamentally limited by instrumental and environmental noise. In laser interferometry, sensitivity at low frequencies (< 20 Hz) is primarily limited by seismic noise, in the mid-frequency range (50–300 Hz) by thermal noise from the mirrors and suspensions, and at high frequencies (> 500 Hz) by quantum shot noise of the laser beam [71]. The cumulative noise from these sources produces a time series vector $n(t)$. To estimate the noise background, we assume the noise is stationary and follows a Gaussian distribution [72]. Under this assumption, the noise exhibits no correlation between different frequency bins. Consequently, the ensemble average of its Fourier components is represented by a diagonal matrix [24]:

$$\langle \tilde{n}^*(f)\tilde{n}(f')\rangle = \frac{1}{2} S_n(f)\delta(f - f'), \tag{1.35}$$

where $\tilde{\square}$ denotes the Fourier transform, $\square^*$ indicates the complex conjugate, and δ is the Dirac delta function. The term $S_n(f)$ is a real, symmetric function with dimensions of Hz^{-1}.

In practical experiments, data is continuously recorded over a finite observation time T, which sets the frequency resolution of the Fourier transform to $\Delta f = 1/T$. Evaluating the Eq (1.35) at $f = f'$ gives the variance:

$$\begin{aligned} \left\langle |\tilde{n}(f)|^2 \right\rangle &= \frac{1}{2} S_n(f) T \\ &= \frac{1}{2} \frac{S_n(f)}{\Delta f}. \end{aligned} \tag{1.36}$$

Since we assume the noise is stationary, the ensemble average of the noise variance in the time domain remains constant:

$$\begin{aligned} \left\langle n^2(t) \right\rangle &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} df\, df' \left\langle \tilde{n}^*(f) \tilde{n}(f') \right\rangle e^{2\pi i (f-f')t} \\ &= \frac{1}{2} \int_{-\infty}^{\infty} S_n(f) df \\ &= \int_0^{\infty} S_n(f) df, \end{aligned} \tag{1.37}$$

where the function $S_n(f)$ defines the one-sided *power spectral density* (PSD) [24]. Throughout this thesis, we utilise the one-sided PSD rather than its two-sided counterpart, as it is strictly a function of physical frequencies. The square root of the PSD, $\sqrt{S_n(f)}$, is the amplitude spectral density (ASD), which characterises the detector's sensitivity and carries units of $\mathrm{Hz}^{-1/2}$. Figure 1.6 shows representative sensitivity curves of the past and future LIGO detectors.

The data stream $d(t)$ recorded by the interferometer consists of the GW strain $h(t)$ and the detector noise $n(t)$:

$$d(t) = h(t) + n(t). \tag{1.38}$$

In reality, $n(t)$ is strongly *coloured*, with an amplitude that varies by orders of magnitude across frequency. Before signal analysis, $d(t)$ needs to be conditioned to mitigate instrumental artifacts and to account for the frequency-dependent noise [14, 72, 73]. Narrow-band notch filters are first applied to remove known instrumental spectral lines, after which $d(t)$ is whitened to flatten the noise spectrum. Whitening is performed in the frequency domain by dividing the Fourier-transformed data $\tilde{d}(f)$ by the detector's ASD, $\sqrt{S_n(f)}$:

$$\tilde{d}_w(f) = \frac{\tilde{d}(f)}{\sqrt{S_n(f)}}. \tag{1.39}$$

This normalisation suppresses frequency bands with large noise and amplifies those where the detector is most sensitive. The resulting noise $n_w(t)$ is approximately white, with a nearly flat power spectrum and unit variance.

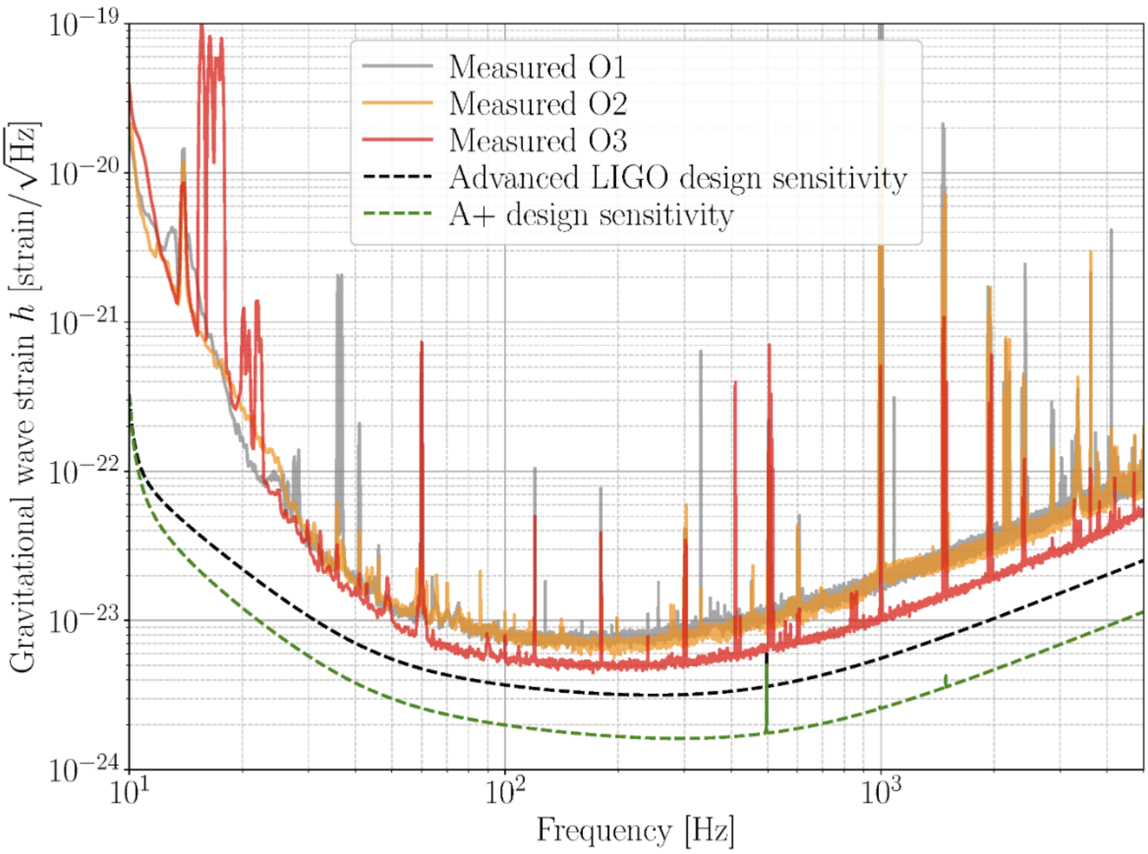

Figure 1.6: Sensitivity curves (noise ASD) of the LIGO Hanford detector during the first to third observing runs (O1-O3), together with the design sensitivity of Advanced LIGO and the projected A+ upgrade. Figure is from [68].

Following the data conditioning, the next step is to extract $h(t)$ from the processed data. As previously established, the detector noise is modelled as stationary and Gaussian. Under this assumption, the likelihood of observing a data stream $d(t)$ given a waveform template $h_T(\boldsymbol{\theta})$ parameterised by $\boldsymbol{\theta}$ can be written in logarithmic form [74, 75], as follows:

$$\log\left(\mathcal{L}_{\text{gw}}(d \mid \boldsymbol{\theta})\right) \propto -\frac{1}{2}\langle d - h_T(\boldsymbol{\theta}) \mid d - h_T(\boldsymbol{\theta})\rangle, \tag{1.40}$$

where $\langle \ldots \rangle$ denotes the noise-weighted inner product [76]:

$$\langle a \mid b \rangle = 2\int_0^\infty \tilde{a}_w^*(f)\tilde{b}_w(f) + \tilde{a}_w(f)\tilde{b}_w^*(f)df. \tag{1.41}$$

Eq. (1.40) can be interpreted as the log-likelihood of the signal hypothesis, $\mathcal{H}_{\text{gw}}$. The null hypothesis, $\mathcal{H}_0$, corresponds to the case in which the data consists entirely of instrumental noise ($d = n$). Under this hypothesis, the log-likelihood simplifies to:

$$\log \mathcal{L}_0 = -\frac{1}{2}\langle d \mid d\rangle. \tag{1.42}$$

Hence, the log-likelihood ratio between the signal hypothesis and the null hypothesis is given by:

$$\log \Lambda \equiv \log \mathcal{L}_{\mathrm{gw}}(d \mid \boldsymbol{\theta}) - \log \mathcal{L}_0 = \langle d \mid h_T(\boldsymbol{\theta})\rangle - \frac{1}{2}\langle h_T(\boldsymbol{\theta}) \mid h_T(\boldsymbol{\theta})\rangle. \tag{1.43}$$

This ratio effectively cancels the purely data-dependent noise term.

Two SNR measures are useful for characterising signal detection. The optimal SNR, ρ_{opt}, represents the maximum achievable SNR for a perfectly matched signal in stationary Gaussian noise, that is its intrinsic strength. The matched-filter SNR, ρ_{mf}, quantifies how strongly the data correlate with the expected template and therefore provides a measure of the presence of a signal in the data [24, 77]. They are defined as:

$$\begin{aligned} \rho_{\mathrm{opt}} &\equiv \sqrt{\langle h_T(\boldsymbol{\theta}) \mid h_T(\boldsymbol{\theta})\rangle}, \\ \rho_{\mathrm{mf}} &\equiv \frac{\langle d \mid h_T(\boldsymbol{\theta})\rangle}{\sqrt{\langle h_T(\boldsymbol{\theta}) \mid h_T(\boldsymbol{\theta})\rangle}}. \end{aligned} \tag{1.44}$$

Substituting these definitions into Eq. (1.43), the log-likelihood ratio is reformulated in terms of SNRs:

$$\log \Lambda = \rho_{\mathrm{opt}}\rho_{\mathrm{mf}} - \frac{1}{2}\rho_{\mathrm{opt}}^2. \tag{1.45}$$

For a given template, where ρ_{opt} is fixed, Eq. (1.45) shows that maximising the log-likelihood ratio is mathematically equivalent to maximising ρ_{mf}. Therefore, matched filtering is defined as the process of cross-correlating the continuous data stream $d(t)$ against a discrete bank of waveform templates $h(\boldsymbol{\theta})$ to find the exact parameter set that maximises ρ_{mf}.

In practice, matched filtering is computationally demanding, and several techniques are employed to mitigate this challenge. For example, the `GstLAL` pipeline [78] filters the data primarily in the time domain. To handle the large number of templates, `GstLAL` applies Singular Value Decomposition (SVD) [79, 80] to compress the template bank into a smaller set of orthogonal basis vectors, substantially reducing the computational cost and enabling the generation of low-latency triggers required for rapid multi-messenger alerts. Alternatively, the `PyCBC` pipeline [81] performs matched filtering in the frequency domain using Fast Fourier Transforms (FFTs), enabling the matched-filter inner products in Eq. (1.43) to be computed efficiently over large template banks.

After the matched-filtering procedure identifies candidate triggers, additional signal-consistency tests are applied before a GW signal can be confidently established as an astrophysical event. This is necessary because transient instrumental artifacts can also produce high-SNR triggers that mimic signals from CBCs. For example, the `PyCBC` pipeline employs a χ^2 consistency test to evaluate whether the frequency evolution of a trigger is consistent with that of a true CBC waveform. The resulting χ^2 value is then used to reweight the matched-filter SNR [82, 83], thereby down-ranking high-SNR triggers caused by instrumental glitches.

The statistical significance of the remaining candidates is then evaluated by comparing them against the expected behaviour of detector noise. This significance is typically quantified using detection metrics such as the false alarm rate (FAR), which denotes how often detector noise is expected to produce triggers with ranking statistics (e.g., SNR) greater than that of the candidate [72, 84]. To estimate these quantities, an empirical background distribution of the ranking statistic is constructed using time-shift analyses, in which the data streams from different detectors are artificially offset by time intervals larger than the inter-detector GW travel time. Any coincident triggers identified in these shifted analyses therefore arise purely from accidental noise coincidences rather than true astrophysical signals. Repeating this procedure many times produces a large background sample from which the false alarm rate can be estimated as:

$$\mathrm{FAR}(\rho^*) = \frac{N_{\mathrm{bg}}(\rho > \rho^*)}{T_{\mathrm{bg}}}, \tag{1.46}$$

where ρ^* is the threshold ranking statistic, N_{bg} is the number of background triggers with ranking statistic exceeding ρ^*, and T_{bg} is the effective background livetime accumulated from all time shifts.

Another commonly used measure of statistical significance is p_{astro}, which quantifies the probability that a candidate is of genuine astrophysical origin rather than a noise-induced artifact [85]. Unlike the FAR, which characterises the rate of noise triggers exceeding a given ranking statistic, p_{astro} estimates the relative probability that a candidate belongs to the astrophysical-signal population instead of the terrestrial-noise population. Formally, for a candidate with a ranking statistic x, p_{astro} is is modelled using a Poisson mixture framework [86, 87]:

$$p_{\mathrm{astro}}(\rho) = \frac{\Lambda_h P_h(\rho)}{\Lambda_h P_h(\rho) + \Lambda_n P_n(\rho)}, \tag{1.47}$$

where $P_h(\rho)$ and $P_n(\rho)$ denote the probability density functions of the ranking statistic for astrophysical signals and terrestrial noise triggers, respectively. The parameters Λ_h and Λ_n represent the expected occurrence rates of signal and noise events in the analysed dataset.

In practice, the noise distribution $P_{\mathrm{n}}(\rho)$ is empirically estimated using background analyses such as the time-shift procedures employed in FAR estimation, whereas the signal distribution $P_h(\rho)$ is derived from simulated astrophysical populations together with detector selection effects. Consequently, p_{astro} naturally depends not only on the candidate ranking statistic itself, but also on the relative abundance of expected astrophysical mergers and noise transients.

Candidates with sufficiently low FAR values and high p_{astro} are identified as confident GW detections and are subsequently reported in official GW transient catalogues. In the LVK transient catalogues, events are typically classified as confident detections if they satisfy criteria such as FAR < 1 yr^{-1} and $p_{\mathrm{astro}} > 0.5$. For particularly important discoveries, such as the first observation of a new class of GW sources (e.g., GW150914), more stringent significance requirements are imposed, often corresponding to an effective significance exceeding the 5σ level [14].

Taken together, the sequence of matched filtering, signal-consistency tests, and statistical significance estimation enables modern GW searches to robustly distinguish true astrophysical signals from detector noise. Once a confident detection has been established, the next step is to infer the source's physical properties using Bayesian inference techniques, which will be introduced in Chapter 4.

Chapter 2

Gravitational Lensing of Astrophysical Signals

Another important prediction of GR is *gravitational lensing* [88, 89]. Since mass determines the curvature of spacetime, any signal travelling through it, such as EM waves or GWs, follows a curved path. This phenomenon is termed *lensing* because the gravitational field of a massive foreground object bends the trajectories of travelling signals, much like a conventional glass lens refracts light in classical optics (see Figure 2.1). As a result, the source appears deflected from its true position, and its signal can be magnified (appearing brighter or louder) or demagnified (appearing dimmer or quieter) to the observer.

Figure 2.1: Representative gravitational lensing of background galaxies. Figure credit: ESA/Webb, NASA & CSA, H. Atek, and M. Zamani.

The concept of gravitational lensing was famously proven in 1919 when Arthur Eddington, Frank Dyson, and Charles Davidson observed the deflection of starlight passing near the Sun during a total solar eclipse [90]. Based on GR, Albert Einstein predicted that a ray passing a point mass M at an impact parameter b would be deflected by an angle $\alpha = 4GM/c^2b$. The angle is exactly twice the value expected from Newtonian mechanics, and its observational confirmation solidified lensing as a key physical phenomenon.

EM lensing is now a well-established observational probe across a wide range of physical scales. On small object scales, microlensing events enable the detection of exoplanets and compact objects that emit little light [91, 92, 93]. On galactic scales, strong lensing of distant quasars [94, 95] and supernovae [96, 97, 98] produces multiple repeated images and observable time delays. Measuring these time delays provides a direct method to constrain the Hubble constant, independent of the local cosmic distance ladder. Furthermore, weak lensing is widely used to map the large-scale distribution of dark matter (DM) across galaxy clusters [99, 100]. At the largest cosmological scales, the lensing of the CMB by intervening large-scale structures helps constrain the sum of neutrino masses, the amplitude of matter density fluctuations, and the Hubble constant [101].

GW lensing is emerging as a practical observational frontier. With the rapidly increasing number of GW detections, the probability of observing lensed GW events is steadily growing. Because GWs pass through intervening matter without being absorbed or scattered, their lensing provides an exceptionally clean probe for determining the mass distribution of the lens. Looking ahead, the joint detection of a lensed GW and its lensed EM counterpart, known as *multi-messenger lensing*, is a major scientific objective [102]. Since both signals travel the same lensed path, comparing their arrival times can directly test the fundamental speed of gravity. Additionally, combining the luminosity distance extracted from the GW signal with the redshift and lensing time delays from the EM counterpart would yield a highly precise, independent measurement of the Universe's expansion rate.

In this chapter, we review the theoretical framework of gravitational lensing. We begin with the general formalism, describing how signals are modified by lensing and how lensing probabilities are quantified. We then introduce standard lens models for characterising mass distributions. Finally, we present the formalism for EM and GW lensing, highlighting the fundamental differences in the measurement and analysis of these two classes of signals.

2.1 General Lensing Theory

Within the framework of GR, a weak gravitational field can be treated as an optical medium characterised by an effective index of refraction. For a weak, stationary field, the spacetime metric can be approximated in terms of the Newtonian gravitational potential Φ, with $\Phi \ll c^2$. The corresponding metric then takes the form as follows [103]:

$$ds^2 = -\left(1 + \frac{2\Phi}{c^2}\right) c^2 dt^2 + \left(1 - \frac{2\Phi}{c^2}\right) d\mathbf{r}^2, \tag{2.1}$$

where $\mathbf{r}$ denotes the three-dimensional spatial coordinates. A photon travels along a null geodesic ($ds^2 = 0$), and its effective coordinate speed, $v = |d\mathbf{r}/dt|$, is modified by Φ. Expanding to the first order in Φ/c^2, the speed becomes $v \approx c(1 - 2\Phi/c^2)$, which defines an effective refraction index for the spacetime [103]:

$$n = \frac{c}{v} \approx 1 - \frac{2\Phi}{c^2}. \tag{2.2}$$

The physical extent of the lens along the line of sight is negligibly small compared to the distances between the observer, the lens, and the source. The lens can therefore be treated under the *thin-lens approximation*, in which the continuous bending of the light ray is approximated as an instantaneous deflection on a two-dimensional lens plane [103]. Figure 2.2 illustrates the geometry of gravitational lensing for an astrophysical source (here, a BBH), with the observer located at the origin of the coordinate system. The true angular position of the source is denoted by $\boldsymbol{\beta}$, while the observed angular position of the image on the lens plane is $\boldsymbol{\theta}$. Using the angular-diameter distances, the corresponding transverse physical coordinates on the source and lens planes are $\boldsymbol{\eta} = D_S\boldsymbol{\beta}$ and $\boldsymbol{\xi} = D_L\boldsymbol{\theta}$, respectively.

The total travel time of a signal along a path P from the source to the observer, with the refractive index n in Eq. (2.2), is given by:

$$t_{\text{tot}} = \int_P \frac{n}{c} dl = \frac{1}{c} \int_P dl - \frac{2}{c^3} \int_P \Phi dl, \tag{2.3}$$

where dl denotes the infinitesimal element of Euclidean path length along P.

In the absence of a lens ($\Phi = 0$), the spacetime reduces to Minkowski space and the refractive index becomes $n = 1$. The total travel time of a signal along the unperturbed straight path P_0 is therefore:

$$t_0 = \frac{1}{c} \int_{P_0} dl. \tag{2.4}$$

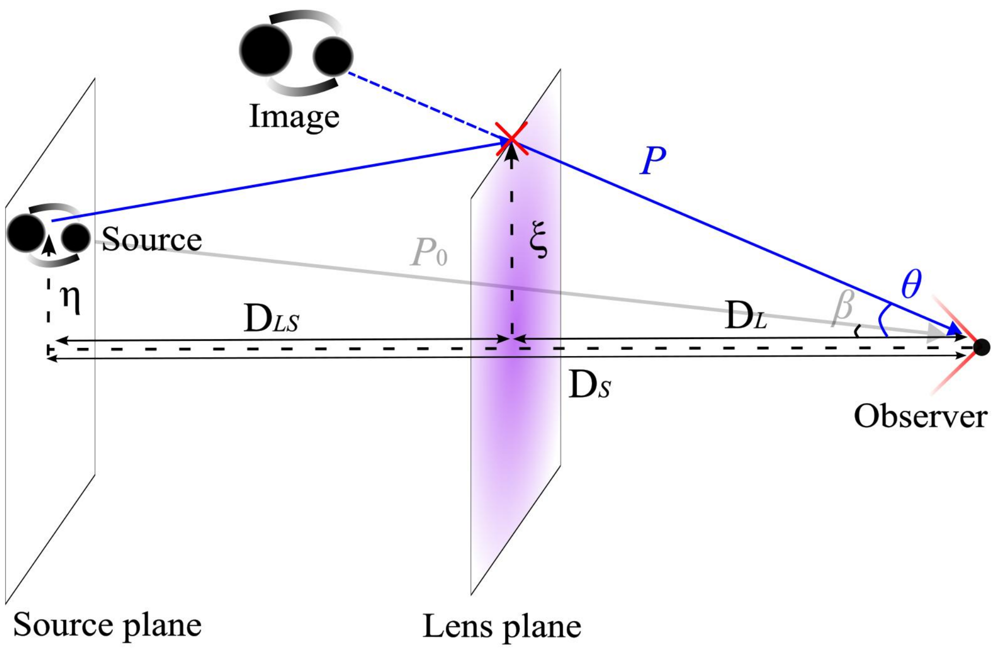

Figure 2.2: Schematic of the gravitational lensing system. ξ is the distance between lens and image positions, the so-called impact parameter, and η is the source position from the line of sight. D_L and D_S are the angular diameter distances from the observer to the lens plane and the source, respectively, and their difference is D_{LS}. Adapted from [104].

The time delay relative to this unperturbed propagation is then given by:

$$\Delta t = t_{\text{tot}} - t_0 = \underbrace{\frac{1}{c}\left(\int_P dl - \int_{P_0} dl\right)}_{\Delta t_{\text{geom}}} \underbrace{-\frac{2}{c^3}\int_P \Phi dl}_{\Delta t_{\text{grav}}} . \tag{2.5}$$

The first term in Eq. (2.5) corresponds to the geometric time delay, Δt_{geom}. Since the lens deflects the signal, it travels a longer physical distance along the curved path P compared to the unperturbed straight line P_0. The physical length of P_0 from the observer to the source is given by:

$$\begin{aligned} l_0 &= \sqrt{D_S^2 + |\boldsymbol{\eta}|^2} \\ &\approx D_S + \frac{|\boldsymbol{\eta}|^2}{2D_S}. \end{aligned} \tag{2.6}$$

Note that the angular separations are much smaller than the distances between the observer, lens, and source, and thus the small-angle approximation can be applied.

On the other hand, the length of P is given by:

$$\begin{aligned} l &= \sqrt{D_L^2 + |\boldsymbol{\xi}|^2} + \sqrt{D_{LS}^2 + |\boldsymbol{\eta} - \boldsymbol{\xi}|^2} \\ &\approx D_L + \frac{|\boldsymbol{\xi}|^2}{2D_L} + D_{LS} + \frac{|\boldsymbol{\eta} - \boldsymbol{\xi}|^2}{2D_{LS}}. \end{aligned} \tag{2.7}$$

Hence, the excess physical path length introduced by the deflection can be written as:

$$\begin{aligned}\Delta l &\equiv l - l_0 \\ &= \frac{|\boldsymbol{\xi}|^2}{2D_L} + \frac{|\boldsymbol{\eta} - \boldsymbol{\xi}|^2}{2D_{LS}} - \frac{|\boldsymbol{\eta}|^2}{2D_S} \\ &= \frac{D_L}{2}|\boldsymbol{\theta}|^2 + \frac{|D_S\boldsymbol{\beta} - D_L\boldsymbol{\theta}|^2}{2D_{LS}} - \frac{D_S}{2}|\boldsymbol{\beta}|^2 \\ &= \frac{D_L D_S}{2D_{LS}}\left(|\boldsymbol{\theta}|^2 + |\boldsymbol{\beta}|^2 - 2\boldsymbol{\theta}\cdot\boldsymbol{\beta}\right) \\ &= \frac{D_L D_S}{2D_{LS}}|\boldsymbol{\theta} - \boldsymbol{\beta}|^2.\end{aligned} \tag{2.8}$$

Since the deflection occurs at the lens plane at redshift z_L, the time delay is multiplied by a factor of $(1 + z_L)$ to account for cosmic expansion. Δt_{geom} is finally given by:

$$\begin{aligned}\Delta t_{\text{geom}} &= \frac{(1 + z_L)\Delta l}{c} \\ &= \frac{1 + z_L}{c}\frac{D_L D_S}{D_{LS}}\frac{1}{2}|\boldsymbol{\theta} - \boldsymbol{\beta}|^2.\end{aligned} \tag{2.9}$$

The second term in Eq. (2.5) represents the gravitational time delay (t_{grav}), also known as the Shapiro delay [105], which arises from the propagation of the signal through the gravitational potential well of the lens, resulting in an additional coordinate time relative to propagation through flat spacetime. In realistic lensing configurations, the deflection angle is very small, so the curved trajectory P differs only slightly from the unperturbed straight path. As a result, the line element along the true path dl can be approximated by the line-of-sight coordinate dz, which is also known as the *Born approximation* [103, 106]. The gravitational time delay can then be written as:

$$t_{\text{grav}} \approx -\frac{2}{c^3}\int_{-\infty}^{\infty}\Phi(D_L\boldsymbol{\theta}, z)dz, \tag{2.10}$$

where the three-dimensional Newtonian potential is explicitly expressed as $\Phi = \Phi(D_L\boldsymbol{\theta}, z)$. To mathematically relate Eq. (2.10) to Eq. (2.9), we define a dimensionless, projected lensing potential, which is given by:

$$\psi(\boldsymbol{\theta}) \equiv \frac{2D_{LS}}{c^2 D_L D_S}\int_{-\infty}^{\infty}\Phi(D_L\boldsymbol{\theta}, z)dz. \tag{2.11}$$

Using Eq. (2.11) and accounting for cosmic expansion in the same manner as for the geometric time delay, Eq. (2.10) can be rewritten as:

$$t_{\text{grav}} = -\frac{1 + z_L}{c}\frac{D_L D_S}{D_{LS}}\psi(\boldsymbol{\theta}). \tag{2.12}$$

By summing Eq. (2.9) and Eq. (2.12), the total time delay for any given source position $\boldsymbol{\beta}$ and image position $\boldsymbol{\theta}$ is as follows:

$$\Delta t(\boldsymbol{\theta}, \boldsymbol{\beta}) = \frac{1+z_L}{c} \frac{D_L D_S}{D_{LS}} \left[\frac{1}{2} |\boldsymbol{\theta} - \boldsymbol{\beta}|^2 - \psi(\boldsymbol{\theta}) \right]. \tag{2.13}$$

In geometrical optics, the propagation of radiation is governed by Fermat's principle, which states that a signal follows paths for which the total arrival time (emission time plus travel time) is stationary, i.e., at a local minimum, maximum, or saddle point. Mathematically, this implies that a first-order perturbation of the true path produces no change in the total travel time ($\delta t = 0$), assuming the emission time is the same in all cases. Under the thin-lens approximation, finding the stationary paths is equivalent to finding the positions on the lens plane, where the gradient of the arrival time function vanishes ($\nabla_{\boldsymbol{\theta}} t_{\text{tot}} = 0$), which provides the fundamental condition for image formation [103, 107].

2.1.1 Lens Equation

The terms enclosed within the brackets in Eq. (2.13) define the Fermat potential, $T(\boldsymbol{\theta}, \boldsymbol{\beta})$ [103]:

$$T(\boldsymbol{\theta}, \boldsymbol{\beta}) = \frac{1}{2} |\boldsymbol{\theta} - \boldsymbol{\beta}|^2 - \psi(\boldsymbol{\theta}). \tag{2.14}$$

As mentioned above, Fermat's principle applied to this potential, $\nabla_{\boldsymbol{\theta}} T = 0$, provides the condition for image formation. Defining the deflection angle as the gradient of the projected lensing potential, $\boldsymbol{\alpha}(\boldsymbol{\theta}) \equiv \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta})$, this condition reduces to the fundamental *lens equation*:

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta}). \tag{2.15}$$

It is convenient to express the lens equation in dimensionless form, typically by normalising with the Einstein radius, defined as:

$$\theta_E = \sqrt{\frac{4GM}{c^2} \frac{D_{LS}}{D_L D_S}}, \tag{2.16}$$

where M is the lens mass. Introducing dimensionless coordinates, $\boldsymbol{y} = \boldsymbol{\beta}/\theta_E$ and $\boldsymbol{x} = \boldsymbol{\theta}/\theta_E$, Eq. (2.15) reduces to:

$$\boldsymbol{y} = \boldsymbol{x} - \boldsymbol{\alpha}(\boldsymbol{x}), \tag{2.17}$$

where $\boldsymbol{\alpha}(\boldsymbol{x}) \equiv \frac{D_{\rm L} D_{\rm LS}}{\theta_E D_{\rm S}} \boldsymbol{\alpha}(\boldsymbol{\theta})$ is called the *scaled deflection angle*. The lens equation can be solved analytically only in simple cases. In most situations, it must be solved numerically.

2.1.2 Multiple Images and Magnification

Since the deflection angle is generally non-linear, Eq. (2.17) can yield multiple solutions for $\boldsymbol{\theta}$ given a single $\boldsymbol{\beta}$, leading to the formation of multiple lensed images. For a point-like source, each image corresponds to a single position on the lens plane. In contrast, for an extended source, such as a background galaxy, the finite spatial extent implies that different regions of the source experience different deflection angles, $\boldsymbol{\alpha}(\boldsymbol{x})$. This differential deflection distorts the morphological shape of the observed images.

Since the apparent size of the source is typically much smaller than the scale over which the deflection angle changes, the first-order approximation of the lens equation is highly accurate. The resulting linear mapping between infinitesimal angular elements on the source plane and their corresponding areas on the lens plane is described by the Jacobian matrix of the lens equation, which is given by [103, 107]:

$$\begin{aligned}\mathcal{A}_{ij} &= \frac{\partial y_i}{\partial x_j} \\ &= \delta_{ij} - \frac{\partial^2 \psi}{\partial x_i \partial x_j},\end{aligned} \tag{2.18}$$

where δ_{ij} is the Kronecker delta and the second term is the Hessian matrix of $\psi(\boldsymbol{x})$. The trace of the Hessian corresponds to the two-dimensional Laplacian of the projected lensing potential, $\nabla^2\psi = \psi_{11} + \psi_{22}$, which is proportional to the dimensionless convergence, $\kappa \equiv \Sigma/\Sigma_{\text{crit}}$, where Σ is the lens surface mass density [106]. Σ_{crit} is the critical surface mass density that depends on the observer-lens-source geometry [103]:

$$\Sigma_{\text{crit}} = \frac{c^2}{4\pi G}\frac{D_s}{D_l D_{ls}}. \tag{2.19}$$

According to the two-dimensional Poisson equation, the relationship between κ and ψ is given by:

$$\begin{aligned}\kappa &= \frac{1}{2}\nabla^2\psi \\ &= \frac{1}{2}(\psi_{11} + \psi_{22}).\end{aligned} \tag{2.20}$$

The trace-free part of the Hessian matrix characterises the anisotropic distortion of images, and is described by the gravitational shear γ, which has two independent components:

$$\gamma_1 = \frac{1}{2}(\psi_{11} - \psi_{22}), \quad \gamma_2 = \psi_{12} = \psi_{21}. \tag{2.21}$$

By using the linear relations in Eqs. (2.20) and (2.21), the Jacobian matrix in Eq. (2.18) can be rewritten as:

$$\begin{aligned}\mathcal{A} &= \begin{pmatrix} 1-\psi_{11} & -\psi_{12} \\ \psi_{21} & 1-\psi_{22} \end{pmatrix} \\ &= \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix}.\end{aligned} \tag{2.22}$$

Convergence describes the isotropic magnification of the image, acting as a uniform scaling of the source image in all spatial directions. Consequently, an intrinsically circular source image with radius r remains circular under the action of convergence, although its size is rescaled. In contrast, the shear components γ_1 and γ_2 induce anisotropic distortions of the image. Shear stretches the image along one principal axis while compressing it along the orthogonal direction, thereby transforming an intrinsically circular source into an ellipse [103, 106].

According to Liouville's theorem, gravitational lensing conserves the source's surface brightness [103]. Since photons are neither created nor destroyed during the deflection process, any observed magnification in the total flux is a purely geometric effect. Consequently, the local magnification factor μ at any specific position on the lens plane is determined entirely by the ratio of the observed solid angle to the intrinsic solid angle of the source. This area ratio corresponds to the inverse of the Jacobian determinant:

$$\begin{aligned}\mu &= \frac{1}{\det \mathcal{A}} \\ &= \frac{1}{(1-\kappa)^2 - (\gamma_1^2 + \gamma_2^2)}.\end{aligned} \tag{2.23}$$

The magnification μ can take both positive and negative values, with its algebraic sign encoding the parity of the image. A positive μ indicates that the image preserves the orientation of the source, even if its overall shape is severely stretched by shear. Conversely, a negative μ corresponds to an image with reversed parity, appearing as a mirror reflection of the intrinsic source. Note that $\mu = \mu(\boldsymbol{x})$ varies across the lens plane as a function of the image-plane coordinate $\boldsymbol{x}$. For an extended source with a finite size, different regions of the galaxy undergo different degrees of magnification. Therefore, the total magnification of an extended source must be calculated by integrating this local magnification factor over the entire intrinsic surface brightness distribution of the source.

The points on the lens plane where $\det \mathcal{A} = 0$, correspond to locations at which the magnification formally diverges ($\mu \to \infty$). The locus of such points forms continuous closed curves on the lens plane, known as *critical curves* [107]. Mapping these curves onto the source plane through Eq. (2.17) produces corresponding boundaries called *caustics*. Sources located close to a caustic therefore appear highly magnified, and extended sources in this regime often produce strongly elongated structures. When a source crosses a caustic, the number of images changes by two; that is, a new pair of images is either created or annihilated.

The condition det $\mathcal{A}$ can be understood more clearly by expressing the magnification in terms of the eigenvalues of the Jacobian matrix. From Eq. (2.23), the magnification can be rewritten as:

$$\begin{aligned}\mu &= \frac{1}{(1-\kappa-\gamma)(1-\kappa+\gamma)} \\ &= \frac{1}{\lambda_t \lambda_r},\end{aligned} \tag{2.24}$$

where $\gamma = \sqrt{\gamma_1^2 + \gamma_2^2}$ denotes the total shear magnitude. The quantities $\lambda_t = 1 - \kappa - \gamma$ and $\lambda_r = 1 - \kappa + \gamma$ represent the tangential and radial eigenvalues of the Jacobian matrix, respectively. The corresponding magnifications are $\mu_t = \lambda_t^{-1}$ and $\mu_r = \lambda_r^{-1}$. Critical curves arise where either $\lambda_t = 0$ (tangential critical curves) or $\lambda_r = 0$ (radial critical curves) on the lens plane [103]. Mapping these curves onto the source plane yields the corresponding caustics. When an extended source lies on or near a tangential caustic, the produced images are extremely magnified and stretched in the azimuthal direction around the lens. In contrast, a source approaching a radial caustic produces images that are elongated along the radial direction toward the lens center. Figure 2.3 shows how the configuration of lensed images depends on the source position relative to the caustics.

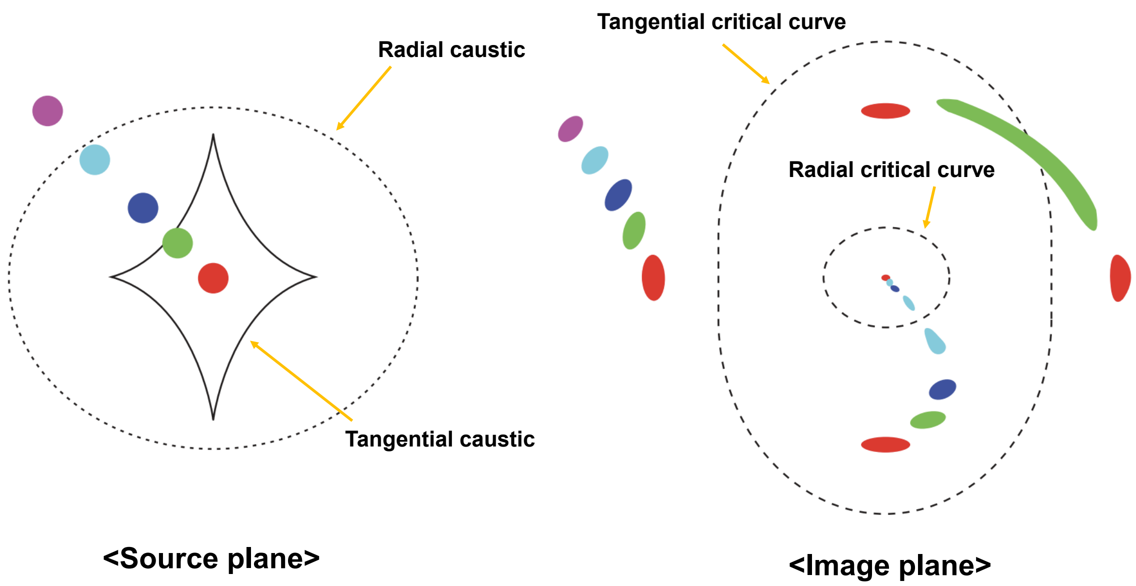

Figure 2.3: Example of critical curves and caustics in the image and source planes, respectively. When a source crosses a caustic, multiple images are produced; near the caustic, the images are highly magnified and distorted. Adapted from [108].

In the presence of a central singularity in the lens mass distribution, the radial critical curve can collapse to the singular origin ($\boldsymbol{x} = 0$). Mapping this central singularity onto the source plane produces a distinct boundary known as a *pseudo-caustic*, or a *cut* [103, 109]. Unlike a standard caustic, crossing a cut changes the total number of images by one rather than two.

2.1.3 Time Delay Surface and Image Type

Rewriting the total arrival time delay into a dimensionless form gives:

$$\begin{aligned}\Delta t(\boldsymbol{x}) &= \frac{1+z_L}{c}\frac{D_L D_S}{D_{LS}}\left[\frac{1}{2}|\boldsymbol{x}-\boldsymbol{y}|^2 - \psi(\boldsymbol{x})\right] \\ &= D_{\Delta t} T(\boldsymbol{x}, \boldsymbol{y}),\end{aligned} \tag{2.25}$$

where $D_{\Delta t} \equiv \frac{1+z_L}{c}\frac{D_L D_S}{D_{LS}}$ is the *time-delay distance*, a cosmological distance scale determined by the geometry of the universe. For a lensing system producing multiple images, one of the physical observables is the relative time delay $\Delta t_{ij} = \Delta t(\boldsymbol{x}_i) - \Delta t(\boldsymbol{x}_j)$ between image i and image j. If a transient source, such as a GW signal, is strongly lensed and produces multiple images, the measured time delays can be used to infer $D_{\Delta t}$, thereby constraining cosmological parameters. This approach is known as *time-delay cosmography* [110, 111, 112].

The dimensionless Fermat potential $T(\boldsymbol{x}, \boldsymbol{y})$, which represents the time delay surface on the lens plane, can also be used to classify the local topology of each lensed image. To do so, we evaluate the Hessian matrix of the Fermat potential, $H_{ij} = \partial^2 T/\partial x_i \partial x_j$. Note that the Hessian of the time delay surface is mathematically identical to $\mathcal{A}$ in Eq. (2.18):

$$H_{ij} = \delta_{ij} - \frac{\partial^2 \psi}{\partial x_i \partial x_j} = \mathcal{A}_{ij}. \tag{2.26}$$

The signs of the eigenvalues of the Hessian matrix in Eq. (2.26) determine the local curvature of the arrival-time surface and classify the lensed images into three types [103, 107]. A *Type I image* forms at a local minimum of the arrival-time surface, where both eigenvalues are positive. In this case, det $\mathcal{A} > 0$ and tr $\mathcal{A} > 0$, yielding an image with positive parity that corresponds to the first arriving signal. A *Type II image* forms at a saddle point, where the eigenvalues have opposite signs. Consequently, det $\mathcal{A} < 0$, producing an image with negative parity that arrives after the Type I image. A *Type III image* forms at a local maximum of the arrival-time surface, where both eigenvalues are negative. Although det $\mathcal{A} > 0$ and the parity remains positive, the negative trace (trA<0) reflects that the local curvature is concave, corresponding to a longer effective light-travel path. Hence, the Type III image arrives last. The classification and properties of these image types are summarised in Table 2.1.

The image types also impose a constraint on the global properties of the lens system, formulated as the *odd number theorem* [113]. For any bounded, non-singular, and transparent mass distribution, the Euler characteristic of the continuous time-delay surface requires that the number of minima (n_I), saddle points (n_{II}), and maxima (n_{III}) satisfy the index theorem relation:

$$n_I + n_{III} - n_{II} = 1. \tag{2.27}$$

As a result, the total number of images, $N = n_I + n_{II} + n_{III}$, must be odd. In observational astrophysics, however, an even number of images (e.g., doubles or quads) is frequently detected. Type III images typically form within the dense central region of the lens, where κ is extremely high and the local magnification approaches zero (see Eq. (2.23)), Such highly demagnified images fall below the detection threshold, leaving only an even number of visible images.

2.1.4 Lensing Probability

Gravitational lensing is a rare phenomenon, reflecting the low average mass density of the Universe relative to the critical density. Nevertheless, the sparsity of effective lenses makes the observed frequency of lensing events a sensitive probe of the cosmic mass distribution and the underlying cosmology. Comparing the theoretically expected probability of lensing with observational detection rates, therefore, provides a fundamental test of cosmological models. In this context, we review how to calculate lensing probability.

We define the lensing probability as the fraction of the sky covered by the effective lensing cross-sections of all intervening mass distributions [114]. For a background source located at a given redshift z_S, the lensing probability is formally expressed as the optical depth, $\tau(z_S)$. Conceptually, $\tau(z_S)$ is the ratio of the total solid angle subtended by all lenses to the solid angle of the whole sky. Mathematically, it is calculated by integrating the product of the comoving number density of lenses, $n(z_L)$, and their respective lensing cross-sections, $\sigma_{\text{lens}}(z_L)$, over the cosmological volume. By integrating over the differential comoving volume dV_c and normalising by the solid angle of the whole sky (4π steradians), the optical depth is derived as [114]:

$$\tau(z_S) = \frac{1}{4\pi} \int_0^{z_S} n(z_L)\sigma_{\text{lens}}(z_L)\frac{dV_c}{dz_L}dz_L. \tag{2.28}$$

Table 2.1: Classification of lensed images based on the local topology of the time-delay surface. The physical properties of each image type are strictly determined by the eigenvalues, determinant (det $\mathcal{A}$), and trace (tr$\mathcal{A}$) of the Hessian matrix.

Image Type	Stationary Point	Eigenvalues	det $\mathcal{A}$	tr$\mathcal{A}$	Parity	Arrival Sequence
Type I	Minimum	Both positive	> 0	> 0	Positive	First
Type II	Saddle point	Opposite signs	< 0	–	Negative	Intermediate
Type III	Maximum	Both negative	> 0	< 0	Positive	Last

In a more realistic case, where the lens population spans a range of masses, $n(z_L)$ is replaced by the differential mass function $dn(z_L)/dM_L$. The optical depth then involves an additional integration over the lens mass M_L as follows:

$$\tau(z_S) = \frac{1}{4\pi}\int_0^{z_S}\int \sigma_{\rm lens}(M_L, z_L)\frac{dn(z_L)}{dM_L}\frac{dV_c}{dz_L}dM_L dz_L. \tag{2.29}$$

Since $\tau \ll 1$ for typical astronomical lenses, the optical depth is directly interpreted as the lensing probability:

$$\begin{aligned} P(\text{lens}|z_S) &= 1 - \exp\tau(z_S) \\ &\approx \tau(z_S). \end{aligned} \tag{2.30}$$

Equation (2.30) defines the lensing probability for a single source located at a specific redshift z_S. Background sources are distributed across a range of redshifts, and the distribution is characterised by a normalised probability density function, $p_{\rm src}(z_S)$. To calculate the expected lensing probability for an entire population of sources observed up to a maximum redshift $z_{\rm max}$, the individual optical depth for a specific redshift must be marginalised over this source distribution. The population-averaged optical depth is expressed as [114]:

$$\bar{\tau} = \int_0^{z_{\rm max}} \tau(z_{S,i})p_{\rm src}(z_{S,i})dz_{S,i}, \tag{2.31}$$

where $\tau(z_{S,i})$ and $p_{\rm src}(z_{S,i})$ are optical depth for a source at $z_{S,i}$ and distribution of sources at $z_{S,i}$. Consequently, if an observational survey contains a total of $N_{\rm unlens}$ background sources within a given cosmic volume, the expected total number of lensed sources, $N_{\rm lens}$, is given by:

$$N_{\rm lens} = N_{\rm unlens}\bar{\tau}. \tag{2.32}$$

2.2 Lens Model

To relate the lensing effects introduced in the previous section to the physical parameters of the corresponding lens, a lens mass profile (hereafter referred to as the *lens model*) must be specified, which determines the projected lensing potential in Eq. (2.11). We introduce several lens models used throughout this thesis. We begin with a basic axially symmetric mass profile, which serves as an idealised model for proof-of-concept studies. Under axial symmetry, the two-dimensional lens equation reduces to a one-dimensional form, which significantly simplifies the analysis of lensing effects. We then consider more realistic non–axially symmetric mass profiles that account for ellipticity or additional features arising from the surrounding environment of the lens.

In this section, we restrict our discussion to single-plane lensing. Although multiple-plane lensing can occur, its probability is expected to be low. Such configurations can nevertheless be important for breaking certain degeneracies in cosmological inference.

2.2.1 Point Mass

The point-mass (PM) model is one of the simplest mass profiles for compact objects whose physical size is negligible compared to the corresponding Einstein radius, such as stars, stellar remnants, and dark matter clumps. For an isolated point mass, the convergence reduces to a two-dimensional Dirac delta function at the origin. From Eq. (2.20), the lensing potential is given by:

$$\psi(\boldsymbol{x}) = \ln|\boldsymbol{x}|, \tag{2.33}$$

where $\boldsymbol{x}$ are the image positions normalised by θ_E in Eq. (2.16). The dimensionless lens equation for the PM model is thus given by:

$$\begin{aligned} y &= x - \nabla\psi(x) \\ &= x - \nabla(\ln x) \\ &= x - \frac{1}{x}. \end{aligned} \tag{2.34}$$

The number of solutions of Eq. (2.34) is always two, which means that two lensed images are always created by a PM lens regardless of the value of y.

The individual magnification factors of the two images can be analytically calculated using Eqs. (2.23), and are given by [103]:

$$\mu_{\pm} = \frac{1}{2} \pm \frac{y^2+2}{2y\sqrt{y^2+4}}. \tag{2.35}$$

The positive root corresponds to a Type I image, while the negative root corresponds to a Type II image. In practice, the intrinsic flux of the background source is unknown. Instead, the relative magnification factor (or flux ratio) between the secondary and primary images, $\mu_{\rm rel} \equiv |\mu_-|/|\mu_+|$, is a directly measurable quantity and is given by:

$$\mu_{\rm rel} = \frac{2\left(y^2+2\right)}{y\sqrt{y^2+4}+y^2+2} - 1. \tag{2.36}$$

Furthermore, the relative time delay Δt between the two images can be calculated using Eq. (2.25), which can be directly measured for transient lensed sources. For the PM model, this is given by [103]:

$$\Delta t = \frac{4GM_L}{c^3}(1+z_L)\left[\frac{y\sqrt{y^2+4}}{2} + \ln\left(\frac{\sqrt{y^2+4}+y}{\sqrt{y^2+4}-y}\right)\right]. \tag{2.37}$$

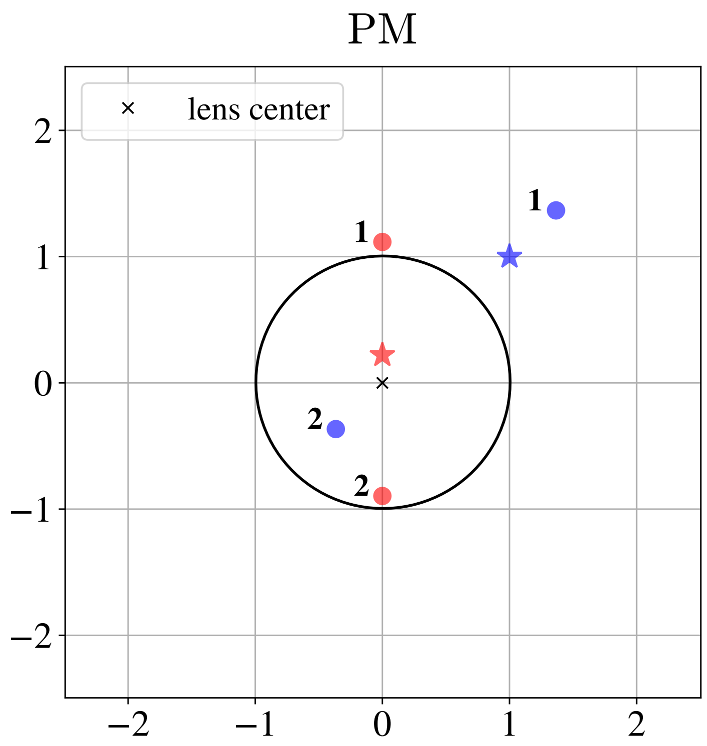

Figure 2.4: A critical line (solid black line) of a PM lens in units of normalised image and source positions. Stars and dots represent the sources and images, respectively, with each set of images and its corresponding source distinguished by colour. The numbers next to the images indicate their order of arrival.

Figure 2.4 shows the tangential critical curve of a PM lens, along with the source positions and the corresponding image positions. Since the PM model contains a singularity at the lens center, the radial critical line degenerates into the origin. Accordingly, the caustic of a PM lens is a single point at $\boldsymbol{y} = (0, 0)$. When a source is perfectly aligned with this point caustic ($y \to 0$), the two images merge to form a symmetric Einstein ring with formally infinite magnification. Conversely, in the weak-field limit where the source is significantly misaligned ($y \gg 1$), the solutions approach $\mu_+ \to 1$ and $\mu_- \to 0$. The primary image then closely resembles the unlensed source, while the secondary image becomes negligibly faint, making macroscopic lensing effects unobservable.

2.2.2 Singular Isothermal Sphere

While the PM model is suitable for compact objects, the mass distributions of galaxies and clusters are extended and dominated by DM halos. The singular isothermal sphere (SIS) is the simplest analytical profile used to model these galactic-scale lenses. It is physically motivated by the approximately flat rotation curves observed in spiral galaxies and by modelling the lens as an isotropic, isothermal collisionless system (stars or DM) with a constant one-dimensional velocity dispersion, σ_v [115].

The three-dimensional mass density of the SIS model is given by [103]:

$$\rho(r) = \frac{{\sigma_v}^2}{2\pi G r^2}. \tag{2.38}$$

The corresponding two-dimensional surface mass density is therefore:

$$\Sigma(\xi) = \frac{\sigma_v^2}{2G\xi}, \tag{2.39}$$

where ξ denotes the physical distance from the lens center on the lens plane. The Einstein radius for the SIS model is defined as [103]:

$$\theta_{E,\mathrm{SIS}} = 4\pi \left(\frac{\sigma_v}{c}\right)^2 \frac{D_{LS}}{D_S}, \tag{2.40}$$

which can be calculated using stellar dynamics, while the Einstein radius in Eq. (2.16) is specified directly by the lens mass. The convergence then simplifies to $\kappa = \Sigma/\Sigma_{\mathrm{cr}} = \theta_E/(2\theta)$. Using Eq. (2.20), the dimensionless lensing potential becomes linear in the radial coordinate:

$$\psi(x) = |x|, \tag{2.41}$$

where $x = \theta/\theta_{E,\mathrm{SIS}}$. Correspondingly, the deflection angle produced by an SIS lens is independent of y, and its magnitude is equal to $\theta_{E,\mathrm{SIS}}$.

Like the PM model, the dimensionless lens equation for the SIS model admits an analytical form:

$$\begin{aligned} y &= x - \nabla\psi(x) \\ &= x - \nabla x \\ &= x - 1. \end{aligned} \tag{2.42}$$

As can be seen, Eq. (2.42) has two solutions if $|y| < 1$, whereas it has a unique solution if $|y| \geq 1$ (see Figure 2.5). When $|y| < 1$, the magnifications for the two images are given by:

$$\mu_+ = 1 + \frac{1}{y}, \quad \mu_- = 1 - \frac{1}{y}. \tag{2.43}$$

As in the PM model, the primary image has positive parity, while the secondary image has negative parity. The relative magnification factor and the time delay between the two images are then given by:

$$\mu_{\mathrm{rel}} = \frac{1-y}{1+y}, \quad \Delta t = \frac{4GM_L(1+z_L)}{c^3}\,[2y]\,. \tag{2.44}$$

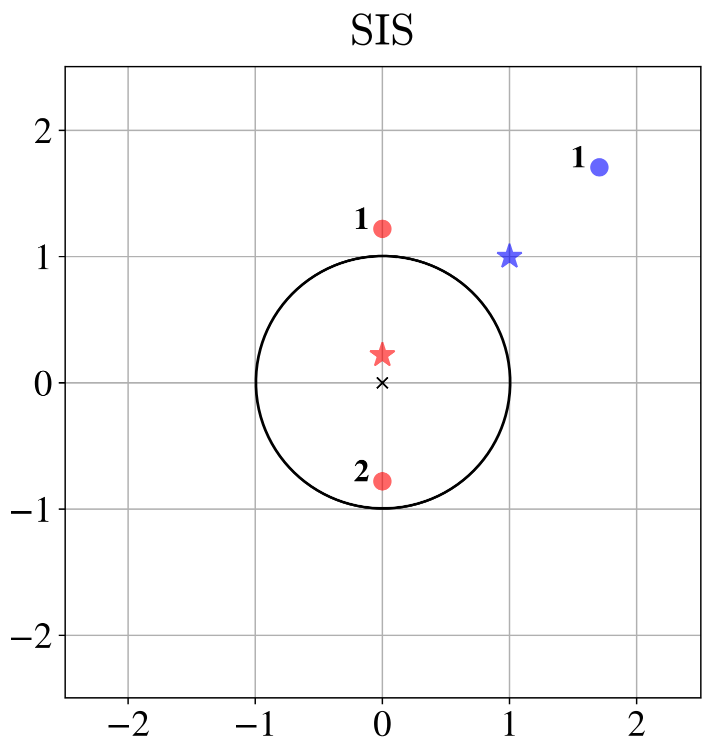

Figure 2.5: Same configuration as in Figure 2.4, but for an SIS lens.

When the source crosses the boundary $y = 1$ and moves outside the Einstein ring, the position of the secondary image approaches the coordinate origin ($x \to 0$). At this boundary, μ_- decreases to zero, and the circle defined by $y = 1$ on the source plane acts as the cut. Once the source lies at $y > 1$, the secondary image disappears as it approaches the central singularity.

2.2.3 Navarro-Frenk-White Profile

While the SIS model provides a useful approximation for galactic halos, high-resolution cosmological N-body simulations of cold dark matter (CDM) demonstrate that DM halos do not follow an isothermal density profile [116]. The universally adopted density distribution for virialised DM halos is the Navarro-Frenk-White (NFW) profile [117].

The three-dimensional mass density of the NFW profile is characterised by a shallow inner core that transitions into a steeper outer envelope, which is defined as [117]:

$$\rho(r) = \frac{\rho_s}{\frac{r}{r_s}\left(1 + \frac{r}{r_s}\right)^2} , \tag{2.45}$$

where r_s is the scale radius and ρ_s is the characteristic density of the dark halo. Since the total mass of the NFW halo is formally divergent, the halo is conventionally parameterised by a specific mass M_{200}[1] and the concentration parameter $c = r_{200}/r_s$, where r_{200} is the radius within which the mean density is 200 times the critical density of the universe. Hence, ρ_s in Eq. (2.45) is expressed by [117]:

$$\rho_s = \frac{M_{200}}{4\pi r_{200}^3} \frac{c_{200}^3}{[\ln(1 + c_{200}) - c_{200} / (1 + c_{200})]}. \tag{2.46}$$

Adopting r_s as the normalising constant, the dimensionless convergence $\kappa(x)$ can be derived analytically. Introducing a dimensionless characteristic convergence $\kappa_s = \rho_s r_s / \Sigma_{\text{crit}}$, the convergence is expressed as a piecewise function [118]:

$$\kappa(x) = \begin{cases} \frac{2\kappa_s}{x^2-1}\left[1 - \frac{1}{\sqrt{1-x^2}}\text{arccosh}\left(\frac{1}{x}\right)\right] & (x < 1) \\ \frac{2\kappa_s}{3} & (x = 1) \\ \frac{2\kappa_s}{x^2-1}\left[1 - \frac{1}{\sqrt{x^2-1}}\arccos\left(\frac{1}{x}\right)\right] & (x > 1). \end{cases} \tag{2.47}$$

It is useful to calculate the mean surface mass density enclosed within the radius x, defined as $\bar{\kappa}(x) = \frac{2}{x^2}\int_0^x x'\kappa(x')dx'$ to derive the deflection angle $\alpha(x)$ and the shear $\gamma(x)$ for axially-symmetric lens models. For the NFW profile,

$$\bar{\kappa}(x) = \begin{cases} \frac{4\kappa_s}{x^2}\left[\frac{1}{\sqrt{1-x^2}}\text{arccosh}\left(\frac{1}{x}\right) + \ln\left(\frac{x}{2}\right)\right] & (x < 1) \\ 4\kappa_s\left[1 + \ln\left(\frac{1}{2}\right)\right] & (x = 1) \\ \frac{4\kappa_s}{x^2}\left[\frac{1}{\sqrt{x^2-1}}\arccos\left(\frac{1}{x}\right) + \ln\left(\frac{x}{2}\right)\right] & (x > 1). \end{cases} \tag{2.48}$$

Then, the magnitude of the shear is fundamentally related to the convergence by $\gamma(x) = \bar{\kappa}(x) - \kappa(x)$, and the dimensionless deflection angle is given by $\alpha(x) = x\bar{\kappa}(x)$.

Unlike the PM and SIS models, the NFW profile does not admit analytical solutions to the lens equation, requiring numerical calculations to determine image positions, individual magnifications, and relative time delays. Regarding the image formation, the NFW profile produces distinct critical curves on the lens plane and corresponding caustics on the source plane (see Figure 2.6). Since the NFW profile has no central singularity, the central deflection angle vanishes. The softness of the structure allows the NFW profile to form a radial critical curve at a finite distance from the center, which maps to a corresponding radial caustic on the source plane located at a specific boundary, $y = y_{\text{cr}}$.

1. Technically, M_{200} defines the mass enclosed within r_{200}. While the original NFW formalism treats M_{200} as equivalent to the virial mass M_{vir}, the exact virial overdensity factor, $\Delta_{\text{vir}} \neq 200$, inherently depends on the adopted cosmological model and redshift.

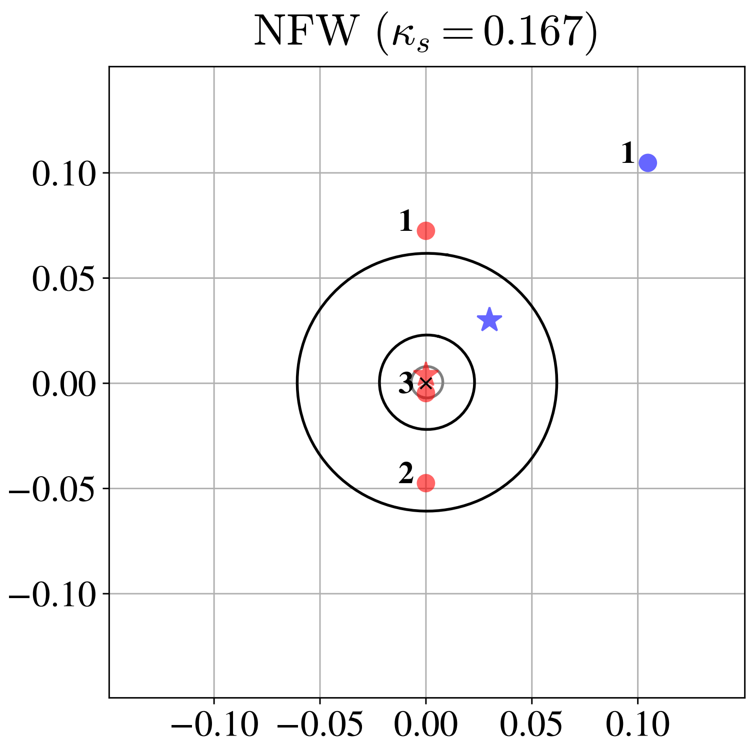

Figure 2.6: Same configuration as in Figure 2.4, but for an NFW lens ($\kappa_s = 0.167$). Unlike the PM and SIS models, which lack radial caustics, the NFW profile exhibits both tangential and radial caustics, and therefore two corresponding critical curves. The radial caustic is depicted by the solid gray line.

The total number of images is determined by the source position relative to the radial caustic. When the source is located inside the caustic ($|y| < y_{cr}$), the lens equation yields three distinct solutions. Due to the soft central profile of the NFW halo, the innermost image is preserved, giving rise to a highly demagnified Type III image near the lens center. Conversely, when the source is located outside the radial caustic ($|y| > y_{cr}$), the two inner images merge and disappear, leaving only a single image.

2.2.4 Singular Isothermal Ellipsoid

To construct a more realistic model of galactic-scale lenses, the assumption of axial symmetry must be relaxed. A straightforward extension of the SIS is the singular isothermal ellipsoid (SIE), which incorporates ellipticity into the projected two-dimensional mass distribution [119].

For an SIE lens, the major axis lies along the x-axis, and the dimensionless convergence is given by [119]:

$$\kappa(x_1, x_2) = \frac{1}{2\sqrt{x_1^2 + x_2^2/q^2}}, \tag{2.49}$$

where q denotes the axis-ratio between the major and minor axes of the lens. Note that we adopt the same normalisation constant as $\theta_{E,\mathrm{SIS}}$ for the dimensionless coordinates. The corresponding surface mass density, Σ, retains the singular r^{-1} radial profile of the isothermal sphere while exhibiting elliptical isodensity contours. The components of the deflection angle, $\alpha(x_1)$ and $\alpha(x_2)$, can be obtained analytically for the SIE model [119]. For an axis ratio $q < 1$, they are expressed in closed form as:

$$\alpha(x_1), \alpha(x_2) = \frac{q}{\sqrt{1-q^2}} \left(\text{arcsin}\left(\frac{\sqrt{1-q^2}x_1}{\sqrt{x_1^2+x_2^2}} \right), \ \text{arsinh}\left(\frac{\sqrt{1-q^2}}{q} \frac{x_2}{\sqrt{x_1^2+x_2^2}} \right) \right). \tag{2.50}$$

Similar to the SIS model, the central singularity of the SIE lens prevents the formation of a Type III image. However, the introduction of ellipticity fundamentally modifies the topological structure of the lens mapping. In the SIS model, there exists a single radial cut and a point-like tangential caustic at the lens center. By contrast, in the SIE model, the point-like caustic expands into an extended, diamond-shaped tangential caustic, commonly referred to as an astroid caustic, accompanied by an outer, elliptical radial cut. The source position relative to these caustics can make double or quadruple images (see Figure 2.7). Specifically, if the source lies within the radial cut but outside the central astroid caustic, the lens equation yields two solutions. When the source is located within both the interior of the astroid and the radial cut, the lens equation admits four solutions, producing four distinct, highly magnified images, typically arranged in a cross or fold configuration around the lens center.

For moderate ellipticities, the astroid caustic is entirely enclosed within the radial cut. However, for highly elliptical mass distributions, specifically $q \lesssim 0.3942$, the vertices of the astroid caustic along the major axis extend beyond the boundary of the radial cut [119]. These exposed regions are referred to as naked cusps. When the source lies within a naked cusp, the lens equation yields exactly three solutions [2].

2.2.5 Singular Power-law Elliptical Mass Distribution

The SIE model can be generalised by allowing for deviations from an isothermal density profile. Introducing a variable power-law slope for the mass distribution leads to the singular power-law elliptical mass distribution (SPEMD) model [120]. Using $r_{E,\mathrm{SIS}}$, as the normalisation constant, the dimensionless convergence is defined as [120]:

$$\kappa(x_1, x_2) = \frac{2-\gamma}{2} \left(\frac{q}{qx_1^2 + x_2^2} \right)^{\frac{\gamma}{2}}, \tag{2.51}$$

where γ denotes the power-law slope of the two-dimensional surface mass density.

2. Since the naked cusp lies outside the radial caustic, any light ray that would otherwise produce a counter-image on the opposite side of the lens is intercepted by the central singularity. As a result, the fourth image does not form.

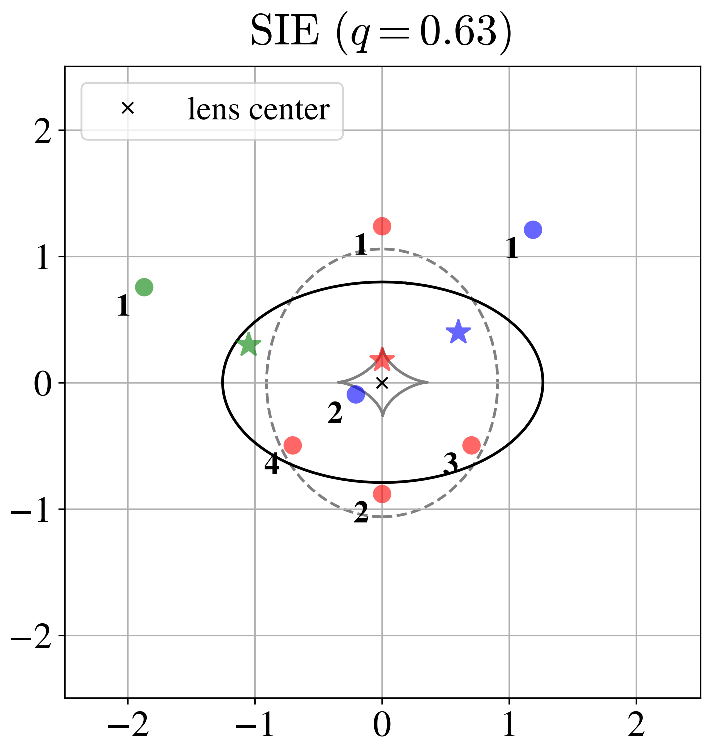

Figure 2.7: Same configuration as in Figure 2.6, but for an SIE lens ($q = 0.63$). Introducing ellipticity breaks the spherical symmetry of the SIS model, leading to an elliptical radial pseudo-caustic (cut). Note that no true radial caustic is formed due to the central singularity of this model. The cut is depicted by the dashed gray line.

Similar to the NFW profile, computing the lensing effects of the SPEMD requires computationally expensive numerical integration. However, [121] derived a semi-analytic expression for the deflection angle by employing a complex-variable formalism. By defining the deflection angle in the complex plane as $\alpha = \alpha_1 + i\alpha_2$ [122, 123], where the real and imaginary parts correspond to the Cartesian components of the deflection angle at the image position (x_1, x_2), the deflection angle is given by [121]:

$$\begin{aligned} \alpha_{\text{SPEMD}}(x_1, x_2) = \frac{2\sqrt{q}}{1+q} \left(\frac{\sqrt{q}}{\sqrt{q^2 x_1^2 + x_2^2}} \right)^{\gamma - 1} \\ \times \, \mathrm{e}^{\mathrm{i}\varphi} {}_2F_1 \left(1, \frac{\gamma}{2}; 2 - \frac{\gamma}{2}; -\frac{1-q}{1+q} \mathrm{e}^{\mathrm{i}2\varphi} \right) , \end{aligned} \tag{2.52}$$

where $\varphi = \arctan(x_2/qx_1)$, describing the orientation of the ellipse relative to the observer's reference axis, and ${}_2F_1$ denotes the Gaussian hypergeometric function. By using Eq. (2.52), the lens equation can be solved on the complex plane, from which the corresponding magnification factors and relative time delays can be derived.

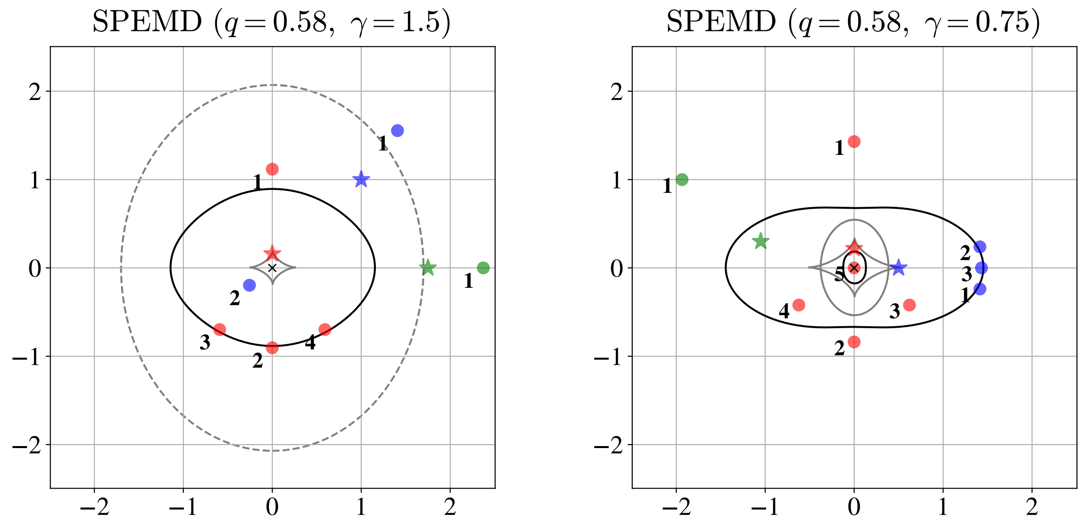

Figure 2.8: Same configuration as in Figure 2.7, but for two SPEMD lenses with the same axis ratio ($q = 0.58$) and different power-law slopes: $\gamma = 1.5$ (steep) and $\gamma = 0.75$ (shallow).

The SPEMD model exhibits a variety of image-formation characteristics. For steep profiles ($\gamma \geq 1$), including the SIE limit ($\gamma = 1$), the lens structure is similar to that of the SIE model. The radial cut encloses an inner astroid tangential caustic, and consequently, no Type III image forms. In this regime, the SPEMD lens produces either double or quadruple images depending on the source position relative to the caustics. Note that the inner astroid caustic never extends beyond the radial cut, unlike the SIE case, and therefore triple-image configurations do not occur.

For shallow profiles ($\gamma < 1$), the convergence diverges at the lens center, while the enclosed mass scales such that the deflection angle smoothly approaches zero. The radial cut is therefore replaced by a well-defined radial caustic, and the lens produces either triple or quintuple images depending on the source's alignment with the caustics. Two representative examples of image configurations and caustic structures for the SPEMD model are shown in Figure 2.8.

2.2.6 Pseudo-elliptical Navarro-Frenk-White Profile

Analogous to the introduction of ellipticity in the SIS model to obtain the SIE model, the pseudo-elliptical NFW (PNFW) profile is defined using the dimensionless elliptical coordinate system, $x_\epsilon = \sqrt{x_{1\epsilon}^2 + x_{2\epsilon}^2} = \sqrt{q x_1^2 + \frac{x_2^2}{q}}$ [124]. We adopt the PNFW profile instead of the exact elliptical NFW profile (e.g., [125]) because the latter requires computationally expensive numerical integration [126], whereas the PNFW model introduces ellipticity directly into the lensing potential rather than the surface mass density. This approach provides fully analytic expressions for the deflection angle and magnification factors, significantly accelerating the computation of lensing effects.

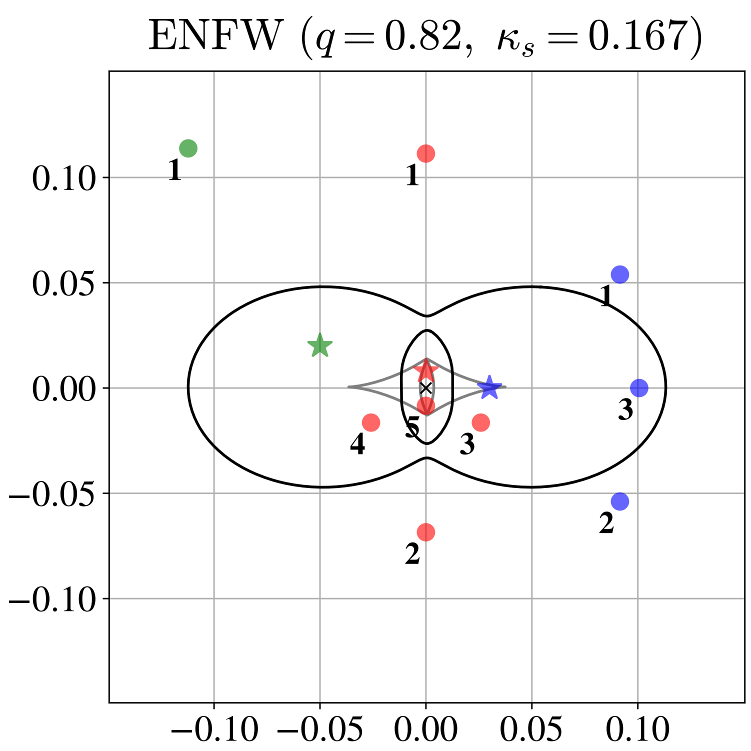

Figure 2.9: Same configuration as in Figure 2.7, but for a PNFW lens ($q = 0.82$, $\kappa_s = 0.167$).

However, this analytical efficiency comes with a physical limitation. Since the ellipticity is imposed on the potential, the corresponding true surface mass density obtained via the Poisson equation becomes strongly distorted for lenses with a low axis ratio. For axis ratios of $q \lesssim 0.6$, the projected mass distribution degenerates into an unphysical dumbbell shape and may exhibit regions of negative mass density along the minor axis. Therefore, the PNFW profile is valid only for moderate ellipticities.

The deflection angle of the PNFW profile, normalised by r_s of the corresponding spherical NFW profile, can be written as [124]:

$$\begin{aligned}\alpha_{\text{PNFW}}(x_1, x_2) = {}& \alpha_{\text{NFW}}\left(x_\epsilon\right) \sqrt{q} \cos\phi_\epsilon \\ & + \alpha_{\text{NFW}}\left(x_\epsilon\right) \frac{\sin\phi_\epsilon}{\sqrt{q}},\end{aligned} \tag{2.53}$$

where $\alpha_{\text{NFW}}(x_\epsilon) = x_\epsilon \bar{\kappa}(x_\epsilon)$, with $\bar{\kappa}$ given in Eq. (2.48), and $\phi_\epsilon = \arctan\left(x_{2\epsilon}/x_{1\epsilon}\right)$ is the elliptical angle, analogous to φ in Eq. (2.52). The PNFW profile exhibits image-formation characteristics analogous to the shallow SPEMD model ($\gamma < 1$), that is producing either triple or quintuple images depending on the source position relative to the tangential and radial caustics. Figure 2.9 shows an example of how a source splits into multiple images by an PNFW lens.

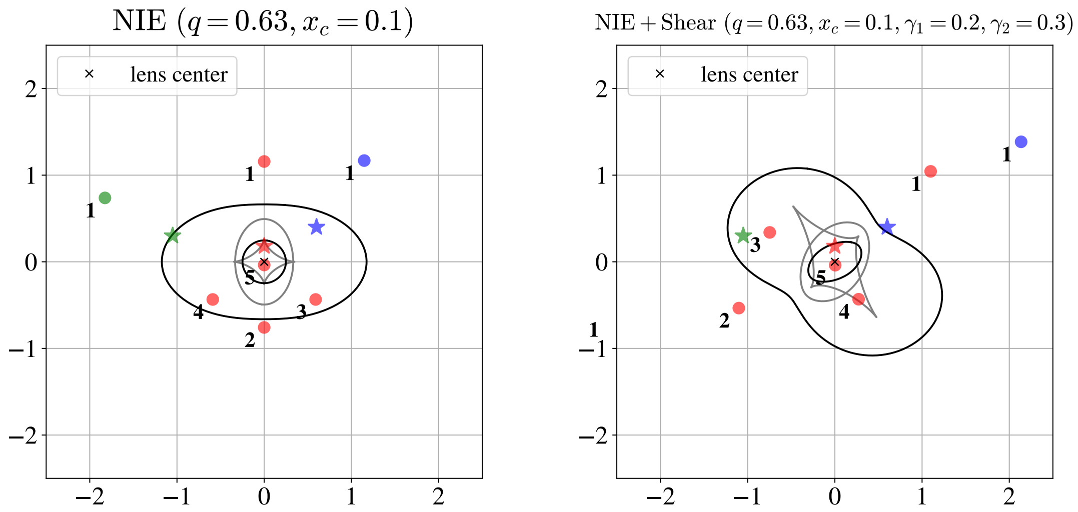

Figure 2.10: *Left panel*: Same source positions and lens parameters as the SIE lens shown in Figure 2.7, with a core radius of $x_c = 0.1$ introduced, resulting in an NIE lens. *Right panel*: The same NIE lens as in the left panel, with an additional external shear field $(\gamma_1, \gamma_2) = (0.2, 0.3)$.

2.2.7 Beyond the Fiducial Lens Models

While low-mass lenses are well approximated by the PM model, analyses of galaxy-scale lensing systems often employ simplified density profiles, such as the SIE model, which exhibits a central singularity. However, realistic galactic mass distributions structurally deviate from these idealised models. The introduction of additional structural components modifies the mapping defined by the lens equation, thereby altering the resulting image configurations, magnification factors, and relative time delays. Two extensions are commonly introduced to model realistic galaxies: (1) a finite core radius and (2) external shear.

First, as expected from baryonic physics and DM halo dynamics, the central mass density of galaxies does not physically diverge to a singularity [127]. Introducing a dimensionless finite core radius ($x_c \equiv \theta_c/\theta_E$) removes the central singularity, softening the central density profile. For the SIE model, this modification yields the *non-singular isothermal ellipsoid* (NIE) model [119], where the dimensionless convergence is redefined as:

$$\kappa(x_1, x_2) = \frac{1}{2\sqrt{x_c^2 + x_1^2 + x_2^2/q^2}}. \tag{2.54}$$

Similar to the NFW profile, the NIE model always produces an odd number of images and can generate a highly demagnified central Type III image. The corresponding deflection angles are derived by incorporating x_c into the exact SIE model, given by [119]:

$$\alpha_{\text{NIE}}(x_1), \alpha_{\text{NIE}}(x_2) = \frac{q}{\sqrt{1-q^2}} \left(\operatorname{arcsin}\left(\frac{\sqrt{1-q^2}x_1}{\sqrt{x_1^2+x_2^2+x_c^2}}\right), \operatorname{arsinh}\left(\frac{\sqrt{1-q^2}}{q}\frac{x_2}{\sqrt{x_1^2+x_2^2+x_c^2}}\right)\right). \tag{2.55}$$

The left panel of Figure 2.10 illustrates an NIE lens with the same axis ratio and GW source positions as the SIE lensing system shown in Figure 2.7. Since the central singularity is removed, the radial cut is replaced by a radial caustic, and correspondingly, the radial critical line changes from a point-like singularity to an extended closed curve. Consequently, a GW source (blue star) that is split into two images in the SIE case is no longer multiply imaged in the NIE case, because it lies outside all caustics.

Second, galaxies are not isolated systems. Their gravitational potentials are perturbed by tidal fields from nearby mass concentrations, such as neighboring galaxies, galaxy groups, or the surrounding large-scale structure [103]. These environmental effects are commonly modelled by introducing an external shear field, which modifies the lensing potential. The external shear is parameterised by its magnitude γ_{ext} and position angle ϕ_{ext}. The corresponding Cartesian shear components are defined as $\gamma_1 \equiv \gamma_{\text{ext}} \cos(2\phi_{\text{ext}})$ and $\gamma_2 \equiv \gamma_{\text{ext}} \sin(2\phi_{\text{ext}})$, where γ_1 describes stretching along the coordinate axes, while γ_2 corresponds to stretching along axes rotated by $45°$. The lensing potential associated with the external shear is given by [119]:

$$\psi_\gamma(x_1, x_2) = \frac{\gamma_1}{2}(x_1^2 - x_2^2) + \gamma_2 x_1 x_2. \tag{2.56}$$

Since lensing potentials combine linearly, the total deflection field of a sheared NIE lens is obtained by superposing the external shear contribution onto the NIE deflection, as follows:

$$\boldsymbol{\alpha}_{\text{tot}}(x_1, x_2) = \boldsymbol{\alpha}_{\text{NIE}}(x_1, x_2) + \begin{pmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}. \tag{2.57}$$

The right panel of Figure 2.10 shows the same NIE lensing system with external shear. The addition of external shear modifies both the shape and orientation of the critical curves and caustics. For the adopted values of $\gamma_1 = 0.2$ and $\gamma_2 = 0.3$, the nonzero γ_2 component indicates that the shear principal axis is rotated with respect to the coordinate axes, leading to asymmetric distortions of the lensing configuration. In particular, a positive (negative) γ_2 corresponds to a counterclockwise (clockwise) rotation of the shear principal axis by an angle $\phi_{\text{ext}} = \frac{1}{2}\tan^{-1}(\gamma_2/\gamma_1)$. This rotated perturbation induces distinct responses in the radial and tangential lensing structures: the radial caustic and its associated critical curve are both distorted and rotated clockwise, acquiring a peanut-like morphology, whereas the tangential caustic and its corresponding critical curve are stretched and rotated counterclockwise along the shear principal direction.

2.3 Lensing of Electromagnetic Waves

Before discussing the GW lensing in detail, it is instructive to first establish the observational context of traditional EM lensing. Currently confirmed GW sources, specifically CBCs, are transient and effectively point-like. In contrast, the background sources in EM lensing span a wide range of spatial extents and temporal behaviours. In addition to transient point-like sources, such as quasars and supernovae[3], EM lensing also involves extended sources, including distant galaxies and galaxy clusters, as well as diffuse cosmological fields, such as the CMB. In this section, we provide a brief overview of several representative EM lensing sources.

The wavelengths of observed EM waves are negligibly small compared to the Schwarzschild radii of typical astrophysical lenses. Consequently, EM lensing is accurately described by the geometric-optics approximation, ensuring that the theoretical formalism outlined in Section 2.1 remains applicable[4].

2.3.1 Extended Sources

When the background source is not a point-like object but an extended object, the analysis of the lensing effects becomes more complex than that for point-like sources. Instead of solving the lens equations for a single discrete source position β, one must map an entire extended surface brightness distribution.

According to Liouville's theorem, gravitational lensing conserves the surface brightness of the background source. The observed surface brightness at an image position $\boldsymbol{\theta}$ therefore equals the intrinsic surface brightness at the corresponding source position $\boldsymbol{\beta}$, i.e., $I(\boldsymbol{\theta}) = I_s(\boldsymbol{\beta})$. The total observed flux S_{obs} is defined as the integral of $I(\boldsymbol{\theta})$ over the solid angle subtended by the image. Consequently, the magnification of the observed flux originates entirely from the stretching of this apparent image area. Using the Jacobian matrix $\mathcal{A}$ defined in Eq. (2.18), the total magnified flux is given by [106]:

$$S_{\text{obs}} = \int I(\boldsymbol{\theta}) d^2\boldsymbol{\theta} = \int I_s(\boldsymbol{\beta}) |\det \mathcal{A}^{-1}| d^2\boldsymbol{\beta} = \int I_s(\boldsymbol{\beta}) \mu(\boldsymbol{\beta}) d^2\boldsymbol{\beta}. \tag{2.58}$$

3. Stars can also serve as point-like background sources, for example, in microlensing events induced by exoplanets or other substellar objects [93]. However, such observations probe distances that are substantially smaller than the typical detection horizon of GW observation, which is not discussed in this thesis.
4. While the geometric optics approximation holds for macroscopic lenses, wave-optics effects in EM lensing may arise in the extreme microscopic regime, such as femtolensing induced by primordial black holes (PBHs) or ultra-light DM clumps with masses $M \lesssim 10^{-13}\, M_\odot$. This regime is beyond the scope of this thesis; see [128, 129, 130] for details.

In galaxy surveys, the magnification factor modifies the observed flux of the host galaxy, which is conventionally reported in terms of an apparent magnitude. Consequently, the apparent magnitude of the i^{th} lensed source-galaxy image, m_i' can be written as:

$$m_i' = m - 2.5\log_{10}(\mu_i), \tag{2.59}$$

where m denotes the unlensed apparent magnitude of the background source.

In the strong lensing regime, deflection across an extended background source produces highly stretched, macroscopic image configurations, such as giant luminous arcs or Einstein rings surrounding the foreground lens [103]. Consequently, the observational identification of individual strong-lensing systems is relatively straightforward. Furthermore, the presence of multiple images enables robust measurements of the relative magnifications, μ_{rel}, between the images, which provides constraints for lens reconstruction. By combining ray-tracing techniques with parametric lens models, the lens equation can be solved numerically, and the lens mass distribution can be constrained from the observed lensed images [103].

Despite these geometric advantages, extended sources typically exhibit steady, continuous emission over human timescales. As a result, the arrival-time delays between multiple lensed images cannot usually be measured directly. Nevertheless, once a lens model is constructed by solving the lens equation, the corresponding time-delay surface on the image plane can be computed. If the mass distribution of the lens is well constrained, approximate time delays between the images can be inferred.

Conversely, in the weak lensing regime, where the background source lies well outside the caustics or the lens potential is insufficiently deep, multiple images are not produced. Instead, the main observable effect is a subtle, coherent distortion of the source shape, parameterised by the local gravitational shear γ. Because this induced stretching ($\gamma \sim 0.01$) is typically much smaller than the intrinsic ellipticity of the background galaxies ($\epsilon \equiv 1-q \sim 0.3$), the weak-lensing signature for any individual galaxy is indistinguishable from its intrinsic shape [106]. Consequently, γ cannot be measured from a single object and must instead be inferred statistically. The analysis requires averaging the measured ellipticities of a large ensemble of background galaxies within a localised region of the sky, under the assumption that the intrinsic galaxy orientations are randomly distributed. The precision of the measurement is limited by intrinsic shape noise and by the finite number of galaxies available for averaging, which sets the statistical uncertainty of the shear estimate [131].

Furthermore, although the general magnification μ depends on both convergence κ and shear γ, the weak-lensing limit ($\kappa \ll 1$, $|\gamma| \ll 1$) allows the expression to be Taylor-expanded. Under this limit, the second-order shear terms become negligible, and the magnification is dominated by the local convergence [131]:

$$\mu = \frac{1}{(1-\kappa)^2 - |\gamma|^2} \approx 1 + 2\kappa. \tag{2.60}$$

On the largest cosmological scales, gravitational lensing of the CMB operates entirely in the weak lensing regime [101]. The intrinsic, unlensed CMB temperature and polarisation fields are statistically isotropic, meaning their primordial fluctuations exhibit no preferred directional alignment. CMB lensing analyses (e.g. [132]), therefore, measure the spatial coherence of the deflection field, that is, the degree to which neighboring photon trajectories are correlated by the same intervening large-scale structure. Observationally, this cumulative deflection manifests as a subtle smoothing of the acoustic peaks, the characteristic oscillatory features in the CMB power spectrum that trace primordial sound waves in the early Universe [133].

2.3.2 Transient Sources

In contrast to extended sources, transient sources, including gamma-ray bursts (GRBs), kilonovae from BNS mergers, and CCSNe, can provide time-dependent lensing signatures. These sources exhibit specific emission profiles, encompassing rapid luminosity evolution, distinct spectral transitions, and multi-wavelength components. For instance, GRBs and BNS mergers generate a prompt high-energy burst in gamma-rays or X-rays, followed by a fading afterglow spanning optical to infrared frequencies [134, 135]. Similarly, CCSNe display characteristic light-curve declines accompanied by temperature-dependent spectral line evolution as the stellar ejecta expands and cools [40, 136].

Transient sources can undergo microlensing by low-mass objects, such as stars, stellar remnants, or star clusters, which induces a time-dependent magnification superimposed on the intrinsic light curve [137, 138, 139]. Microlensing introduces chromatic variations due to the physical expansion of the transient source. Over time, the effective source size increases and becomes comparable to the Einstein radius of the microlens. Since emission at different wavelengths originates from distinct regions within the expanding ejecta–with high-energy emission concentrated in a compact inner core and lower-energy emission originating from a larger, cooler envelope–the differential magnification across the source becomes wavelength-dependent. Therefore, observed light-curve deviations that exhibit spectral distortions or colour-dependent magnification confirm the presence of microlensing [140, 141].

In the strong lensing regime, by contrast, the signals are achromatically magnified or demagnified. Strongly lensed multiple images therefore exhibit identical light-curve shapes across all spectral bands, differing only in relative time delays and magnifications, so that repeated *echoes* of the same transient event may be observed. The short duration and rapid variability of transient events enable direct cross-correlation of the light curves, yielding measurements of the relative arrival-time delays (Δt) between images[5].

Furthermore, EM transient events typically occur within galaxies, and because the transient emission fades rapidly, obtaining a spectroscopic redshift directly from the transient is observationally challenging. However, once transients or their afterglows are spatially localised, follow-up observations of the galaxies hosting them provide their redshifts (z_s) (e.g, [17]). The combination of measurable time delays and accurate redshifts allows strongly lensed transients to serve as powerful cosmological probes. Incorporating Δt and z_s into a lens model enables inference of the system's time-delay distance, $D_{\Delta t}$, which is inversely proportional to the Hubble constant, H_0. Lensed transients thus provide an independent method for cosmography and for testing the physics of the expanding universe.

2.4 Lensing of Gravitational Waves

The typical wavelengths of GWs ($\lambda_{\rm gw}$) are longer than those of EM waves. For lensing objects with masses below the galactic scale, these wavelengths are comparable to or exceed the lens's Schwarzschild radius ($R_{\rm sch}$). Specifically, for GWs from CBCs in the frequency band of ground-based detectors, the condition $\lambda_{\rm gw} \gtrsim R_{\rm sch}$ is satisfied for lens masses $\lesssim 10^5\, M_\odot$, under which the geometric-optics approximation breaks down. Consequently, wave optics is required for the lensing analysis in this mass regime [145, 146, 147].

In the frequency domain, the lensing effect is characterised by the amplification factor, $F(f)$. The lensed GW strain, $h_l(f)$, is then given by [145, 146, 147]:

$$h_l(f) = F(f) h_u(f), \tag{2.61}$$

where f is the frequency and $h_u(f)$ denotes the intrinsic unlensed waveform. This section outlines the theoretical framework for GW lensing. We describe how GWs are affected by lensing, derive the amplification factor, and detail the observable features of the resulting lensed signals.

5. Although a quasar is not a transient source, its central engine, a supermassive black hole accreting matter, produces stochastic luminosity variations across broad spectral bands over months to years [142]. When a quasar is strongly lensed, this same stochastic light curve is replicated in each image. By conducting long-term monitoring campaigns and cross-correlating the light curves, observers can measure the relative time delays between the multiple image trajectories (see e.g., [143, 144]).

2.4.1 GW in Weak Gravitational Field

Assuming a locally flat spacetime perturbed by a weak Newtonian potential $U(\boldsymbol{r}) \ll 1$ due to the lens, such that the system remains in the weak-field limit (i.e., far from the strong-gravity regime of compact objects like BHs), the background metric $g_{\mu\nu}^{(B)}$ is expressed as:

$$\mathrm{d}s^2 = -(1+2U)\mathrm{d}t^2 + (1-2U)\mathrm{d}\boldsymbol{r}^2 \equiv g_{\mu\nu}^{(B)}\mathrm{d}x^\mu \mathrm{d}x^\nu. \tag{2.62}$$

The propagation of a GW through this spacetime introduces a small metric perturbation. The total metric tensor $g_{\mu\nu}$ is given by:

$$g_{\mu\nu} = g_{\mu\nu}^{(B)} + h_{\mu\nu}, \tag{2.63}$$

where $h_{\mu\nu}$ represents the GW strain. In the weak-field limit, since the curvature radius of the background spacetime is much larger than the GW wavelength, the polarisation tensor $e_{\mu\nu}$ is parallel-transported along the background null geodesics, which allows the tensor perturbation to be decoupled into a constant polarisation state and a propagating scalar amplitude ϕ, such that $h_{\mu\nu} = \phi e_{\mu\nu}$ [148].

The propagation equation $\partial_\mu(\sqrt{-g^{(B)}} g^{\mu\nu}_{(B)} \partial_\nu \phi) = 0$ is satisfied in the curved background. Substituting the components of the background metric (Eq. (2.62)) and retaining terms up to first order in the weak potential, U, the equation reduces:

$$\left(\nabla^2 - \partial_t^2\right)\phi(\boldsymbol{r},t) = 4U\partial_t^2\phi(\boldsymbol{r},t). \tag{2.64}$$

Transforming to the frequency domain via the Fourier transform $\tilde{\phi}(\boldsymbol{r},\omega) = \int \phi(\boldsymbol{r},t)e^{i\omega t}dt$, such that ∂_t map to $-i\omega$, one obtains the Helmholtz equation:

$$\left(\nabla^2 + \omega^2\right)\tilde{\phi} = 4\omega^2 U\tilde{\phi}, \tag{2.65}$$

where $\omega = 2\pi f$ is the angular frequency of the GW.

To determine the effect of lensing on a GW, Eq. (2.65) must be solved with appropriate boundary conditions specifying the locations of the source and the observer. For a GW source at distance D_S, the incident GW is naturally modelled as a spherical wave, and the solution is evaluated at the observer located at distance D_L. As mentioned in Section 2.1, the thin lens approximation is valid for GW lensing. Under this thin-lens approximation, the three-dimensional potential $U(\boldsymbol{r})$ is projected onto a two-dimensional lens plane, effectively acting as a discrete phase screen. Based on the Huygens-Fresnel principle, the lensed wave at the observer can be interpreted as the superposition of secondary spherical wavelets originating from all points on the lens plane [147, 149].

Following the formalism of [147], Eq. (2.65) can be solved using Green's function, which provides the response at the observer to a point source on the lens plane. For a wave propagating over a distance r, the Green's function is defined as:

$$G_\omega(r) = -\frac{1}{4\pi}\frac{e^{i\omega r}}{r}. \tag{2.66}$$

Applying Green's theorem over the lens plane, the lensed wave amplitude $\tilde{\phi}_l$ is then expressed as a surface integral, known as the Kirchhoff-Fresnel diffraction integral:

$$\tilde{\phi}_L = \frac{1}{4\pi}\int_S \mathrm{d}^2\boldsymbol{\xi}\left[\tilde{\phi}\frac{\partial}{\partial n}\left(\frac{e^{i\omega r}}{r}\right) - \frac{e^{i\omega r}}{r}\frac{\partial}{\partial n}\tilde{\phi}\right], \tag{2.67}$$

where S denotes the lens plane, $\boldsymbol{\xi}$ are its physical coordinates, n is the unit normal vector to the surface, and $\tilde{\phi}$ is the incident scalar field evaluated at the lens plane. $\tilde{\phi}_L$ explicitly accounts for the diffraction of the wave by the lens, and serves as the starting point for constructing the lensed GW signal at the observer.

2.4.2 Amplification Factor

Amplification factors are defined as the ratio between the scalar amplitudes of the lensed ($\tilde{\phi}_L$) in Eq. (2.67) and unlensed ($\tilde{\phi}$) waves [147]:

$$F(\omega) = \frac{\tilde{\phi}_L(\omega)}{\tilde{\phi}(\omega)}. \tag{2.68}$$

Under the configuration described in the previous section (see Figure 2.2), and accounting for cosmological effects in an expanding universe–where distances correspond to angular diameter distances, and the observed frequency is redshifted–the amplification factors in Eq. (2.68) are given by [147]:

$$F(\omega, \boldsymbol{y}) = \frac{D_S\xi_0^2\,(1+z_L)\,\omega}{2\pi i D_{LS}D_L}\int d^2x\left[\exp(i\frac{D_S}{D_{LS}D_L}\xi_0^2\,(1+z_L)\,\omega T(\boldsymbol{x},\boldsymbol{y}))\right], \tag{2.69}$$

where $T(\boldsymbol{x},\boldsymbol{y})$ is the dimensionless Fermat potential in Eq. (2.25). It is convenient to simplify Eq. (2.69) by introducing the dimensionless frequency, defined as:

$$w \equiv \frac{D_S}{D_{LS}D_L}\xi_0^2\,(1+z_L)\,\omega. \tag{2.70}$$

Finally, the amplification factors reduce to [146, 147]:

$$F(w, \boldsymbol{y}) = \frac{w}{2\pi i}\int d^2x\,[\exp(iwT(\boldsymbol{x},\boldsymbol{y}))], \tag{2.71}$$

In general, the integral in Eq. (2.69) cannot be evaluated analytically due to the nontrivial structure of $T(\boldsymbol{x}, \boldsymbol{y})$, necessitating numerical evaluation. However, the integrand contains a rapidly oscillating phase factor ($\exp(iwT(\boldsymbol{x}, \boldsymbol{y}))$), which makes the numerical evaluation computationally expensive, particularly at high frequencies (large w).

For axially symmetric lens models, however, the two-dimensional integral can be reduced to a one-dimensional form [147]:

$$F(w, y) = -iwe^{iwy^2/2} \int_0^\infty dx\, xJ_0(wxy) \exp\left[iw\left(\frac{1}{2}x^2 - \psi(x) + \phi_m(y)\right)\right], \tag{2.72}$$

where J_0 is the zeroth-order Bessel function, $\psi(x)$ is the dimensionless lensing potential, and $\phi_m(y)$ denotes the minimum arrival time of the unlensed signal. Analytical solutions are available only in limited cases. In particular, for a PM lens, the amplification factor is given by [147]:

$$F(w, y) = \exp\left[\frac{\pi w}{4} + i\frac{w}{2}\left\{\ln\left(\frac{w}{2}\right) - 2\phi_m(y)\right\}\right] \times \Gamma\left(1 - \frac{i}{2}w\right) {}_1F_1\left(\frac{i}{2}w, 1; \frac{i}{2}wy^2\right), \tag{2.73}$$

where $\phi_m(y) = \frac{(x_m - y)^2}{2} - \ln x_m$ with $x_m = \frac{y + \sqrt{y^2+4}}{2}$ and ${}_1F_1$ is the confluent hypergeometric function of the first kind[6] [150]. Eqs. (2.71)-(2.73) can be substituted into Eq. (2.61) to compute the lensed GW waveform from a given unlensed waveform.

When GWs are lensed in the wave-optics regime ($w \lesssim 1$), the phenomenon is referred to as *microlensing* [145, 146, 147, 151, 152, 153, 154], typically induced by lenses that can be well approximated as the PM model. In the time domain, microlensed GW signals exhibit beating patterns. Figure 2.11 shows representative examples of microlensed waveforms. The waveform morphology is increasingly modified, and the beating pattern becomes more pronounced as the microlensing effect is more significant (i.e., for higher M_L^z and smaller y). The patterns arise from the interference between multiple signals with microsecond-scale time delays, which is fully encoded in the $F(w)$ computed in the wave-optics formalism. These morphological distortions encode the physical parameters of the microlens. Consequently, by analysing the interference pattern within the strain data, one can directly extract the redshifted lens mass and the impact parameter without solving the lens equation.

When the wavelength of a GW is much smaller than the Schwarzschild radius of the lens, corresponding to the regime $w \gg 1$, the geometric-optics approximation becomes valid. In this limit, wave-optics effects such as diffraction are negligible, and the standard formalism of EM lensing (Section 2.1) can be directly applied. In the diffraction integral (Eq. (2.71)), the phase factor $\exp(iwT(\boldsymbol{x}, \boldsymbol{y}))$ becomes highly oscillatory at large w. As a

6. See also https://dlmf.nist.gov/13.2E2 for the mathematical definition.

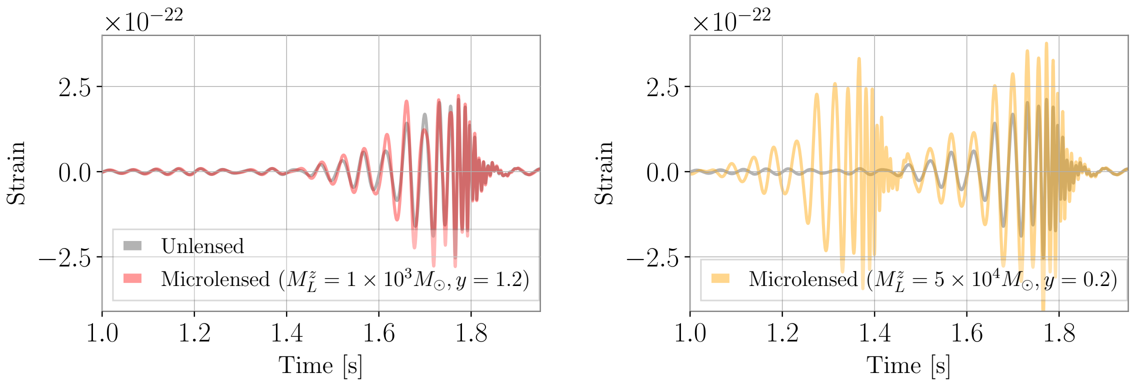

Figure 2.11: Unlensed waveform from a BBH event, shown as a reference, together with microlensed waveforms in the cases of mild (red, left panel) and strong (orange, right panel) microlensing.

result, contributions from different regions of the lens plane interfere destructively, except at the stationary points of the Fermat potential, defined by ($\nabla_{\boldsymbol{x}} T = 0$). These stationary points correspond to the discrete image positions determined by the lens equation in Eq. (2.17).

By evaluating the diffraction integral in the high-frequency limit using the stationary-phase approximation [147], the dominant contributions arise only from the stationary points and the amplification factor reduces to a discrete sum over the geometric images:

$$F(w, \boldsymbol{y}) \simeq \sum_j |\mu_j|^{1/2} \exp\left[iwT_j - i\pi n_j\right], \tag{2.74}$$

where the index j labels the j-th image, μ_j is the magnification factor, and T_j is the dimensionless Fermat potential evaluated at the image position. The integer parameter n_j is the Morse index, which encodes a frequency-independent phase shift associated with the image type. Specifically, $n_j = 0, 1, 2$ correspond to Type I, II, and III images, respectively (see Section 2.1.3 for details about image types). In the geometric-optics limit, the lensed waveform is described as a superposition of discrete signals arriving at different arrival times. Each contribution is magnified by $|\mu_j|^{1/2}$ and acquires a phase shift determined by the corresponding time delay and Morse index [155, 156, 157].

If the time delay between multiple signals is sufficiently large compared to the duration of the GW signal, the individual signals arrive at the detector as temporally separated events and can be identified as distinct signals. In this strong lensing regime, the lensed GW events generally exhibit large magnifications consistent with EM strong lensing [158, 159, 160, 161, 162, 163, 164, 165]. In contrast, smaller-scale objects, such as DM subhalos or intermediate-mass black holes (IMBHs), induce a regime known as *millilensing* or *phenomenological lensing* [166]. In this scenario, the geometric-optics approximation remains valid ($w \gg 1$), while the time delay between signals can be shorter than the signal duration.

As a result, the individual signals overlap in time and cannot be resolved separately, leading to interference. The detector therefore records a single, superimposed waveform characterised by specific amplitude and phase modulations, which can be described by Eq. (2.74), rather than multiple distinguishable events.

2.5 Comparisons between EM and GW Lensing

The field of multimessenger astronomy has now become firmly established. As both EM and GW observations continue to advance, multimessenger lensing is transitioning into an observationally accessible regime [102]. Understanding the key differences between EM and GW lensing, and how they can be jointly exploited, is therefore important.

As noted in the previous sections, the differences between EM and GW lensing for typical astrophysical compact objects (e.g., stars, stellar remnants, star clusters) and DM subhalos ($M_L < 10^5\, M_\odot$) are straightforward. GW lensing generally occurs in the wave-optics regime, resulting in frequency-dependent diffraction patterns, whereas EM lensing remains in the geometric-optics regime, producing achromatic flux magnification. Similarly, the distinction in detector response is also trivial. EM instruments measure the time-averaged energy flux, scaling with the magnification μ, while GW interferometers record the strain amplitude, scaling as $|\mu|^{1/2}$. These basic differences arise directly from the distinct wave properties of EM and GW signals. We now turn to the observable signatures, where their physical distinctions manifest in measurable quantities and can have significant implications for multi-messenger lensing studies.

In the strong-lensing regime, both EM and GW lensing effects are well described by geometrical optics. Consequently, the magnification factors, time delays, and image types (or, equivalently, phase shifts in the GW case) arise identically. However, the observational inference of these quantities differs between the two messengers. In EM lensing, multiple images exhibit different apparent magnitudes and spatial parities. In contrast, for GW signals, magnification manifests as a modification of the inferred luminosity distance, while the image parity introduces a topological phase shift (the Morse phase) in the coalescence phase.

As an illustrative example, consider a galaxy-galaxy strong-lensing system in which the source galaxy hosts a GW source. The magnification factor associated with the i^{th} lensed image modifies both the observed EM flux and the GW amplitude, leading to a shifted apparent magnitude, m_i' in Eq. (2.59) and an effective luminosity distance ($d_{\text{L},i}^{\text{eff}}$), as follows:

$$d_{\text{L},i}^{\text{eff}} = \frac{d_{\text{L}}}{\sqrt{\mu_i}}, \tag{2.75}$$

where d_{L} denotes the true luminosity distance to the GW source. Note that, although a galaxy is an extended source and therefore experiences spatially varying magnification across its surface, these variations are typically not significant on arcsecond scales. As a result, the magnification at the specific location of the GW source is expected to be well approximated by the average magnification inferred from the lensed galaxy image. If $\mu > 1$, lensing increases the observed flux, making the lensed EM image brighter and thus reducing its apparent magnitude. For GW signals, the strain amplitude scales inversely with luminosity distance, so magnification leads to a correspondingly smaller inferred distance.

In addition, the relative magnification factor ($\mu_{\mathrm{rel,EM}}$) between the i^{th} and j^{th} lensed galaxy images in EM observation is given by:

$$\mu_{\mathrm{rel,EM}} \equiv \frac{\mu_j}{\mu_i} = 10^{-0.4(m'_j - m'_i)} \tag{2.76}$$

while for GW signals, the relative magnification factor ($\mu_{\mathrm{rel,GW}}$) can be expressed as:

$$\mu_{\mathrm{rel,GW}} \equiv \frac{\mu_j}{\mu_i} = \left(\frac{d_{\mathrm{L},i}^{\mathrm{eff}}}{d_{\mathrm{L},j}^{\mathrm{eff}}}\right)^2 . \tag{2.77}$$

These two quantities are expected to be consistent within uncertainties.

Let us move to another observable, the relative time delays between multiple lensed images or signals. For transient sources, including GW signals from CBCs and EM transients such as SNe or GRBs introduced in Section 2.3.2, these delays can be measured directly from the arrival times of each signal. In GR framework, the time delay between the i^{th} and j^{th} lensed signals is given by:

$$\Delta t_{\mathrm{EM}} = \Delta t_{\mathrm{GW}} \equiv |t_i - t_j|, \tag{2.78}$$

where t_i and t_j are the arrival times of i^{th} and j^{th} signal, respectively.

For galaxy-galaxy lensing systems, however, the relative time delays of EM images are not directly observable. Instead, they are inferred from the Fermat potentials, T_i and T_j, evaluated at the image positions under an assumed lens model. The corresponding time delays are then given by:

$$\Delta t_{\mathrm{EM}} \equiv |t_i - t_j| = \frac{1 + z_L}{c} \frac{D_L D_S}{D_{LS}} (|T_i - T_j|). \tag{2.79}$$

In contrast to the magnification case, GW observations can measure relative time delays with very high precision, often at the millisecond level. Given the large mass scale of galaxy lenses, the variation of time delay across different image positions within an extended lensed galaxy can be significant. Therefore, with sufficiently high angular resolution, one can construct a detailed time-delay map from EM lens modelling. Comparing this map with the precisely measured Δt_{GW} enables a more accurate localisation of the GW source within the host galaxy.

Beyond enabling precise source localisation, multi-messenger strong lensing provides a unique laboratory for probing the fundamental physics of GR with exceptional precision. Within GR, GWs and EM waves are predicted to propagate at the same speed in vacuum ($c_{\mathrm{GW}} = c_{\mathrm{EM}}$), as implied by Eq. (1.13). If a multi-messenger transient, such as a BNS merger, is strongly lensed, the multiple discrete GW images and their EM counterparts will arrive at the observer with characteristic lensing-induced time delays [167, 168, 169]. Comparing the relative time delays measured independently in the GW and EM channels, therefore, provides a direct and highly sensitive test of whether the two messengers propagate identically over cosmological distances. Since the propagation path spans billions of light-years, even an extremely small deviation between c_{GW} and c_{EM} would accumulate over cosmic distances and manifest as a measurable offset between the GW and EM lensing time delays. Therefore, a statistically significant discrepancy between the two independent measurements would serve as a compelling signature of physics beyond GR.

In addition to possible deviations in propagation speed, several modified gravity theories predict that the amplitude evolution of GWs can differ from the GR prediction during cosmological propagation [170, 171, 172]. As a consequence, the luminosity distance inferred from GW observations may differ from that inferred from EM observations of the same source. Strong lensing is particularly valuable in this context because lens magnification enables the detection of intrinsically faint GW and EM sources at substantially higher redshifts than would otherwise be observable. Such high-z systems are especially powerful probes of modified gravity, as even small deviations from the GR prediction can accumulate over cosmological distances and become observationally measurable.

A commonly adopted phenomenological parametrisation of this effect is given by [173]:

$$\begin{aligned} \frac{d_L^{\mathrm{GW}}(z)}{d_L^{\mathrm{EM}}(z)} &\equiv \Xi(z) \\ &= \Xi_0 + \frac{1-\Xi_0}{(1+z)^n}, \end{aligned} \tag{2.80}$$

where Ξ_0 and n parameterise deviations from the GR prediction ($d_L^{\mathrm{GW}} = d_L^{\mathrm{EM}}$). Robust constraints on these parameters require accurate cosmological distance measurements.

Strong lensing itself provides an independent method for constraining cosmological distances through time-delay cosmography with both EM and GW observations, yielding complementary cosmological information from the same lensed system. The measured time delays between multiple lensed images constrain the time-delay distance ($D_{\Delta t}$) introduced in Eq. (2.25), which depends on the source and lens redshifts as well as the underlying cosmological parameters. Since cosmological distances become increasingly sensitive to the background cosmology at higher z, strongly lensed systems can place particularly precise constraints on the distance-redshift relation (see Section 3.1). Therefore, comparing the cosmological distance relation inferred from multi-messenger lensing observables with the luminosity distance directly obtained from GW signals enables a direct test of modified GW propagation on cosmological scales.

Overall, strongly lensed multi-messenger observations constitute a uniquely powerful laboratory for simultaneously probing cosmology, fundamental physics, and the astrophysical environments or formation channels of compact binary mergers.

Chapter 3

Gravitational-Wave Cosmology

Modern cosmology is primarily built upon the ΛCDM model, named after its two key constituents: Λ (Lambda), representing the cosmological constant associated with dark energy (DE) introduced in Eq. (1.6), and CDM, referring to cold dark matter. This framework is grounded in GR and the cosmological principle, which assumes that the Universe is homogeneous and isotropic on large scales. Within this model, the energy content of the Universe is dominated by two non-luminous components: CDM, which consists of non-relativistic matter that governs structure formation through gravitational clustering, and DE, which provides the repulsive pressure driving the observed accelerated expansion of the Universe.

To quantify the relative contributions of these components, multiple observational probes are combined. The CMB provides a measurement of the total energy density of the Universe, indicating that spacetime is consistent with spatial flatness [174, 175]. Baryon Acoustic Oscillations (BAO), which arise from sound waves propagating in the photon-baryon plasma before recombination, imprint a characteristic length scale in the large-scale distribution of matter [176]. The length scale serves as a standard ruler for tracking the expansion history of the Universe [177]. When combined with luminosity distance measurements from Type Ia supernovae [178, 179], these probes constrain cosmological parameters in a complementary manner. Their joint constraints converge on a narrow region of parameter space.

While the joint EM observations successfully established the ΛCDM model as a standard model, recent observational tensions suggest that it may be incomplete. The most prominent example is the *Hubble tension* [180], a statistically significant discrepancy between measurements of the Hubble constant, H_0, inferred from early-Universe observations, such as the CMB [174], and those obtained from late-time distance ladder methods based on e.g. Cepheid variables and Type Ia supernovae [181]. This discrepancy motivates the need for independent cosmological probes.

GW cosmography provides such an independent avenue to probe the expansion history of the Universe by directly measuring the distance-redshift relation [182]. Observable CBC quantities such as the luminosity distance, $d_{\mathrm{L}}(z)$, encode the kinematic evolution of cosmic expansion, providing a model-independent description that can be interpreted within spe-

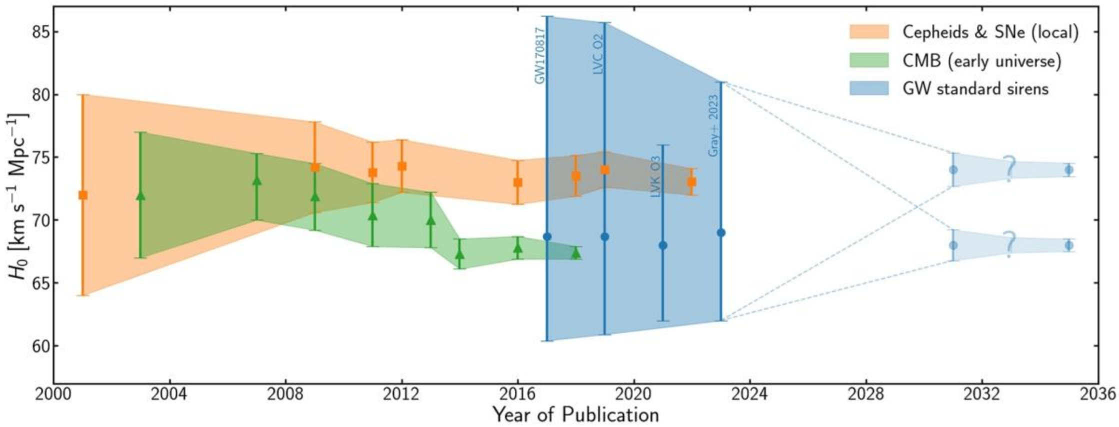

Figure 3.1: Constraints on the Hubble constant (H_0) from CMB, local distance-ladder measurements, and GW observations, highlighting the Hubble tension and the potential of future GW observations to resolve it. Figure is from [183].

cific cosmological models, such as ΛCDM. Cosmological parameters in the ΛCDM model are then inferred by fitting the reconstructed distance-redshift relation. Figure 3.1 illustrates that GW cosmography can provide a powerful probe of the Hubble tension identified by EM observations.

In this chapter, we review the distance-redshift relation in an expanding Universe and briefly introduce the ΛCDM framework, along with its current limitation, namely the Hubble tension. We then describe how GW observations provide an independent avenue for measuring key cosmological parameters, such as H_0, and discuss their associated uncertainties and limitations.

3.1 The Cosmographic Distance-Redshift Relation

The standard cosmological paradigm posits that the Universe is expanding from an initial state of high density and temperature (the Big Bang), and is homogeneous and isotropic on large scales. The spacetime geometry satisfying these conditions is described by the Friedmann–Lemaître–Robertson–Walker (FLRW) metric. Combined with the EFEs introduced in Eq. (1.6), this metric provides the theoretical foundation of the standard model of cosmology, upon which the ΛCDM framework is built [1]. In this section, we derive the relation between cosmological distance and the observable redshift in the expanding spacetime.

3.1.1 Friedmann-Lemaître-Robertson-Walker Model

In the FLRW metric, the spacetime interval is expressed as [1]:

$$ds^2 = -c^2dt^2 + a^2(t)\left[\frac{dr^2}{1-kr^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2)\right], \tag{3.1}$$

where t is the cosmic time, r, θ, ϕ are the reduced-circumference polar coordinates, k represents the spatial curvature, and $a(t)$ is the time-dependent scale factor. EM waves and GWs propagate along null geodesics ($ds^2 = 0$). For a signal travelling radially ($d\theta = d\phi = 0$) from a source at comoving coordinate $r = R$ to an observer at the origin $r = 0$, Eq. (3.1) reduces to:

$$\frac{c\,dt}{a(t)} = \frac{dr}{\sqrt{1-kr^2}}. \tag{3.2}$$

Integrating Eq. (3.2) from the time of emission t_e to the time of observation t_0, i.e., along the path between the source and the observer, yields the *comoving distance*:

$$\begin{aligned} d_C &= \int_{t_e}^{t_0} \frac{c\,dt}{a(t)} \\ &= \int_0^R \frac{dr}{\sqrt{1-kr^2}}. \end{aligned} \tag{3.3}$$

Here, k determines the mapping between the comoving distance d_C and the radial comoving coordinate r. A value of $k > 0$ corresponds to a closed Universe with positive curvature, $k = 0$ to a spatially flat Universe, and $k < 0$ to an open Universe with negative curvature. Eq. (3.3) provides the exact relation in each case [1]:

$$d_C = \begin{cases} \frac{1}{\sqrt{|k|}}\sin^{-1}(\sqrt{|k|}r) & (k > 0) \\ r & (k = 0) \\ \frac{1}{\sqrt{|k|}}\sinh^{-1}(\sqrt{|k|}r) & (k < 0) \end{cases} \tag{3.4}$$

In the flat case, the radial coordinate coincides with the comoving distance. In contrast, for nonzero curvature, the relation becomes non-linear, exhibiting trigonometric (closed) or hyperbolic (open) behaviour at large distances.

Consider two successive wave crests emitted at t_e and $t_e + \delta t_e$, and observed at t_0 and $t_0 + \delta t_0$, respectively. Here δt_e and δt_0 represent the emitted and observed wave periods. Since the comoving coordinate R of the source is fixed, the spatial integral is identical for both crests, leading to:

$$\int_{t_e}^{t_0} \frac{c\,dt}{a(t)} = \int_{t_e+\delta t_e}^{t_0+\delta t_0} \frac{c\,dt}{a(t)}. \tag{3.5}$$

Subtracting the common interval $\int_{t_e+\delta t_e}^{t_0} c\,dt/a(t)$ from both sides gives:

$$\int_{t_e}^{t_e+\delta t_e} \frac{dt}{a(t)} = \int_{t_0}^{t_0+\delta t_0} \frac{dt}{a(t)}. \tag{3.6}$$

The periods δt_e and δt_0 are extremely short compared to the cosmological timescale over which $a(t)$ varies. Therefore, $a(t)$ can be treated as constant over each infinitesimal interval, and Eq. (3.6) simplifies to:

$$\frac{\delta t_e}{a(t_e)} = \frac{\delta t_0}{a(t_0)}. \tag{3.7}$$

The time interval between successive wave crests is stretched by $a(t)$, which leads to the observed cosmological redshift, defined as the fractional change in wavelength [184]:

$$z = \frac{\lambda_0 - \lambda_e}{\lambda_e}, \tag{3.8}$$

where λ_0 and λ_e are the observed and emitted wavelengths, respectively. Since wavelength is proportional to the wave period ($\lambda = c\delta t$), Eq. (3.7) leads to the fundamental relation between redshift and scale factor [1, 184]:

$$\begin{aligned} z &= \frac{\lambda_0}{\lambda_e} - 1 \\ &= \frac{\delta t_0}{\delta t_e} - 1 \\ &= \frac{a(t_0)}{a(t_e)} - 1. \end{aligned} \tag{3.9}$$

While the comoving distance d_C remains fixed, the *proper distance* $d_P(t)$ between the source and observer evolves with time, reflecting the expansion of the Universe. At any given cosmic time t, the proper distance is simply given by:

$$d_P(t) = a(t)d_C, \tag{3.10}$$

which evolves with time through the scale factor $a(t)$. To determine the evolution, the dynamics of $a(t)$ must be specified, which are governed by the EFEs in Eq. (1.6). The large-scale matter content of the homogeneous Universe is well approximated as a perfect fluid, with stress–energy tensor [21]:

$$T_{\mu\nu} = \left(\rho + \frac{p}{c^2}\right) u_\mu u_\nu + p g_{\mu\nu}, \tag{3.11}$$

where ρ represents the total energy density, p is the isotropic pressure, and $u_\mu = (-c, 0, 0, 0)$ is the four-velocity of a comoving observer. In this frame, the time–time component reduces to $T_{00} = \rho c^2$. For the FLRW metric, the time-time component of the Einstein tensor (Eq. (1.5)) is given by:

$$\begin{aligned} G_{00} &= R_{00} - \frac{1}{2} g_{00} R \\ &= -\frac{3\ddot{a}}{ac^2} + \frac{3}{c^2}\left[\frac{\ddot{a}}{a} + \left(\frac{\dot{a}}{a}\right)^2 + \frac{kc^2}{a^2}\right] \\ &= \frac{3}{c^2}\left[\left(\frac{\dot{a}}{a}\right)^2 + \frac{kc^2}{a^2}\right], \end{aligned} \tag{3.12}$$

where overdots denote derivatives with respect to cosmic time, t. Substituting the T_{00} and G_{00} components into the time–time component of the EFEs yields the *first Friedmann equation* [1]:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3}, \tag{3.13}$$

which defines the square of the *Hubble parameter*, $H^2(t) \equiv (\dot{a}/a)^2$, as determined by the total energy density, spatial curvature, and the cosmological constant.

To express Eq. (3.13) in dimensionless form, each term is normalised by the expansion rate $H^2(t)$, which naturally introduces the density parameters associated with each component. This procedure is equivalent to expressing each term in the Friedmann equation in units of the critical density [1, 184]:

$$\rho_c(t) = \frac{3H^2(t)}{8\pi G}, \tag{3.14}$$

defined as the energy density required for a spatially flat Universe ($k = 0$).

For matter (m) and radiation (γ), whose energy densities evolve with cosmic expansion[1], their corresponding density parameters are defined as [186]:

$$\Omega_i(t) \equiv \frac{8\pi G \rho_i(t)}{3H^2(t)}, \quad i = m, \gamma. \tag{3.15}$$

For k and Λ, the associated density parameters are defined as [186]:

$$\Omega_k(t) \equiv -\frac{kc^2}{a^2(t)H^2(t)}, \quad \Omega_\Lambda(t) \equiv \frac{\Lambda c^2}{3H^2(t)}. \tag{3.16}$$

1. See [185], however, for a broader classification of possible energy components, some of which remain speculative.

With these definitions, the Friedmann equation in Eq. (3.13) can be written in a dimensionless form:

$$\Omega_\gamma(t) + \Omega_m(t) + \Omega_\Lambda(t) + \Omega_k(t) = 1, \tag{3.17}$$

which holds at all times.

In practice, cosmological observations constrain these parameters at present t_0, defining $\Omega_{i,0} \equiv \rho_{i,0}/\rho_c$. $\Omega_{k,0}$ and $\Omega_{\Lambda,0}$, are likewise obtained by evaluating $H(t)$ and $a(t)$ at t_0. The evolution of each component is governed by its equation of state. More specifically, the relation between pressure (p) and density (ρ) in Eq. (3.11) can be expressed through the equation of state parameter, w [1, 184]:

$$w \equiv \frac{p}{\rho c^2}, \tag{3.18}$$

where w is a dimensionless parameter, analogous to the relation found in the ideal gas law ($P = nRT/V$); however, in cosmology, it is generalised to include relativistic effects and vacuum energy.

For a perfect fluid in the FLRW model, w determines how the energy density responds to the expansion of the Universe. The time evolution of $\rho(t)$ is governed by the continuity equation, which is derived from the conservation of the energy-momentum tensor ($\nabla_\mu T^{\mu\nu} = 0$) in the FLRW metric [1]:

$$\dot{\rho} + 3\frac{\dot{a}}{a}\left(\rho + \frac{p}{c^2}\right) = 0. \tag{3.19}$$

Substituting Eq. (3.18) into Eq. (3.19) yields:

$$\frac{\dot{\rho}}{\rho} = -3(1+w)\frac{\dot{a}}{a}. \tag{3.20}$$

Integrating both sides with respect to time yields:

$$\rho \propto a^{-3(1+w)}. \tag{3.21}$$

Accordingly, the dimensionless density parameters in Eqs. (3.15) and (3.16) evolve with the scale factor as determined by their respective equations of state [1, 187]:

Radiation (Ω_r): For relativistic particles such as photons, the pressure satisfies $p = \frac{1}{3}\rho c^2$ ($w = 1/3$). Their energy density decreases due to both volumetric dilution, scaling as a^{-3}, and the redshifting of photon energies, which introduces an additional factor of a^{-1}. Consequently, the radiation energy density scales as a^{-4}, or $(1+z)^4$.

Matter (Ω_m): For non-relativistic matter, including baryons and dark matter, the pressure is negligible ($w \simeq 0$). As the rest mass is conserved, the energy density decreases solely due to the expansion of comoving volume, leading to a scaling of a^{-3}, or $(1+z)^3$.

Curvature (Ω_k): The spatial curvature term originates geometrically from the Ricci scalar of the FLRW metric. Its mathematical form, k/a^2, implies an inverse-square dependence on the scale factor, i.e., a^{-2}, or $(1+z)^2$.

Cosmological Constant (Ω_Λ): The cosmological constant is modelled as a constant energy density, representing the intrinsic energy of the quantum vacuum. This corresponds to an equation-of-state parameter $w = -1$ in Eq. (3.21). As a result, the vacuum energy density remains invariant under cosmic expansion, yielding a scaling of a^0, or $(1+z)^0$.

Using these scaling relations, the Friedmann equation can be expressed entirely in terms of present-day density parameters and redshift. Dividing Eq. (3.13) by H_0^2 yields:

$$\begin{aligned}\left(\frac{H(z)}{H_0}\right)^2 &= \Omega_{\gamma,0}a^{-4} + \Omega_{m,0}a^{-3} + \Omega_{k,0}a^{-2} + \Omega_{\Lambda,0} \\ &= \Omega_{\gamma,0}(1+z)^4 + \Omega_{m,0}(1+z)^3 + \Omega_{k,0}(1+z)^2 + \Omega_{\Lambda,0}.\end{aligned} \tag{3.22}$$

It is often convenient to introduce the *dimensionless Hubble parameter*[2],

$$\begin{aligned}E(z) &\equiv \frac{H(z)}{H_0} \\ &= \sqrt{\Omega_{\gamma,0}(1+z)^4 + \Omega_{m,0}(1+z)^3 + \Omega_{k,0}(1+z)^2 + \Omega_{\Lambda,0}},\end{aligned} \tag{3.23}$$

which encapsulates the expansion history independently of H_0.

3.1.2 Luminosity Distance

The luminosity distance to a source, d_L, is conventionally defined through the inverse-square law:

$$F = L/(4\pi d_L^2), \tag{3.24}$$

where F is the observed flux and L is the intrinsic luminosity of the source. We now reinterpret Eq. (3.24) within an expanding cosmological framework.

The luminosity L is defined as the total energy ΔE_e emitted over a proper time interval Δt_e ($L = \Delta E_e/\Delta t_e$). In the FLRW spacetime, photons emitted at t_e propagate outward and are distributed over a spherical surface of radius $d_P(t_0)$ at the time of observation t_0. The proper surface area of this sphere is given by:

$$\begin{aligned}A_P &= 4\pi d_P^2 \\ &= 4\pi a^2(t_0) d_C^2.\end{aligned} \tag{3.25}$$

Cosmic expansion induces two independent effects on the observed radiation [184, 187]:

2. Note that this differs from the dimensionless Hubble constant, defined as $h \equiv H_0/100$ km s^{-1} Mpc^{-1}

1. **Energy redshift**: The energy of each photon is reduced by a factor $(1+z)$, such that $\Delta E_0 = \Delta E_e/(1+z)$.
2. **Time dilation**: The arrival interval between successive photons is stretched, $\Delta t_0 = \Delta t_e(1+z)$.

Since the flux is defined as energy received per unit area per unit time, the observed flux becomes:

$$\begin{aligned} F &= \frac{\Delta E_0}{\Delta t_0 A_P} \\ &= \frac{\Delta E_e/(1+z)}{\Delta t_e(1+z)[4\pi a^2(t_0) d_C^2]} \\ &= \frac{L}{4\pi a^2(t_0) d_C^2 (1+z)^2}. \end{aligned} \tag{3.26}$$

Comparing Eq. (3.26) with Eq. (3.24) defines the luminosity distance in an expanding Universe:

$$d_L = a(t_0) d_C (1+z). \tag{3.27}$$

Hence, d_L encodes both the geometric separation and the effect of cosmic expansion.

To express d_L explicitly as a function of z, we relate cosmic time to redshift. From Eq. (3.9), we have $1+z = a(t_0)/a(t) = 1/a(t)$, where the present-day scale factor is unity. Differentiating both sides with respect to cosmic time yields:

$$\begin{aligned} \frac{dz}{dt} &= -\frac{\dot{a}}{a^2} \\ &= -(1+z)\frac{\dot{a}}{a} \\ &= -(1+z)H(z). \end{aligned} \tag{3.28}$$

This relation can be rewritten as:

$$dt = -\frac{dz}{H(z)(1+z)}. \tag{3.29}$$

Substituting Eq. (3.29) into Eq. (3.3), and noting that t_e and t_0 correspond to redshift z and $z=0$, respectively, the time integral can be recast as a redshift integral as follows:

$$d_C(z) = c\int_0^z \frac{dz'}{H(z')}. \tag{3.30}$$

Finally, using Eq. (3.27), the *luminosity distance–redshift relation* is given by [184, 188]:

$$d_L(z) = c(1+z)\int_0^z \frac{dz'}{H(z')}. \tag{3.31}$$

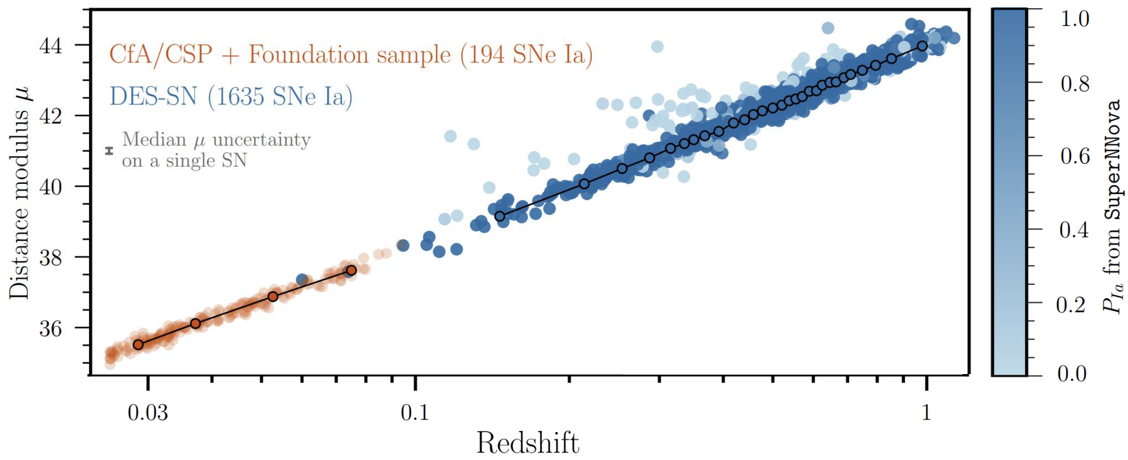

Figure 3.2: Hubble diagram constructed from supernovae over a range of redshifts. The distance modulus is defined as $\mu = 5 \log_{10} (d_{\mathrm{L}}/\mathrm{pc}) - 5$. The colour bar indicates the probability that a supernova is of Type Ia. Adapted from [191].

3.1.3 Hubble-Lemaître Law

The validity of Eq. (3.31) was first established observationally by Georges Lemaître and later by Edwin Hubble in the local Universe, where measurements were restricted to the low-redshift regime ($z < 1$) [189, 190]. In this limit, the cosmic expansion rate can be approximated as constant, $H(z) \approx H_0$, since the dimensionless Hubble parameter satisfies $E(z) \simeq 1$. Expanding Eq. (3.31) to first order in z and neglecting $\mathcal{O}(z^2)$ terms reduces d_L to a linear function of z as follows:

$$\begin{aligned} d_L &\approx c(1+z)\left(\frac{z}{H_0}\right) \\ &\approx \frac{cz}{H_0}. \end{aligned} \tag{3.32}$$

For local astrophysical sources, the cosmological redshift reduces to the non-relativistic Doppler shift, $z \approx v_r/c$, where v_r is the recession velocity of the sources. Substituting this relation into Eq. (3.32) recovers the standard kinematic form of the *Hubble–Lemaître law* [189, 190]:

$$v_r = H_0 d_L. \tag{3.33}$$

This result demonstrates that the general relativistic description of luminosity distance for EM or GW sources reduces, in the low-redshift limit, to the empirically established Hubble–Lemaître law. As shown in Figure 3.2, the Hubble–Lemaître law provides a good description of the recent data in the regime $z < 1$. It is worth noting that the Hubble–Lemaître law is formally expressed in terms of the proper distance, $v_r = H_0 d_P$. In the low-redshift limit, for a spatially flat Universe, the luminosity distance satisfies $d_L \approx d_C = d_P$.

3.1.4 The ΛCDM Framework and Hubble Tension

In Section 3.1.1, we reviewed the FLRW models, which are based on the FLRW metric and the Friedmann equations. These models provide the theoretical basis for the standard cosmological paradigm, namely the ΛCDM model. This model describes the dynamical content of the Universe by assuming a spatially flat universe ($\Omega_k = 0$) dominated today by two components: cold dark matter (Ω_c) and cosmological constant (Ω_Λ), together with a small baryonic contribution (Ω_b) and negligible radiation ($\Omega_\gamma \approx 0$) at late times. Consequently, the model implies the constraint $\Omega_{m,0} + \Omega_{\Lambda,0} = 1$, where $\Omega_{m,0} = \Omega_{b,0} + \Omega_{c,0}$. Under the ΛCDM framework, the Hubble parameter is often rewritten as:

$$H(z) = H_0\sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{\Lambda,0}}. \tag{3.34}$$

A central goal of observational cosmology is to determine the precise values of H_0, $\Omega_{m,0}$, and $\Omega_{\Lambda,0}$ in Eq. (3.34). Current estimates of the density parameters are empirically established through the intersection of independent EM observational probes. Each methodology imposes a distinct geometric constraint on the $(\Omega_{m,0}, \Omega_{\Lambda,0})$ plane.

For example, measurements of CMB anisotropies tightly constrain the angular scale of the sound horizon[3] at recombination, fixing the ratio between the physical sound horizon and the angular diameter distance to the last-scattering surface [175, 192]. However, this constraint is inherently degenerate, as different combinations of cosmological parameters (e.g., H_0, Ω_m, and Ω_Λ) can yield the same observed scale [193]. BAO provide a complementary probe by using the pre-recombination sound horizon as a standard ruler at lower redshifts, allowing the expansion history to be constrained through measurements of the angular diameter distance and $H(z)$ [194, 195, 196]. When combined, CMB and BAO observations break the geometric degeneracies by anchoring early-Universe physics while mapping the later-time expansion history, thereby yielding tight and robust constraints on cosmological parameters. Within the ΛCDM framework, the latest DESI DR 2 (Dark Energy Spectroscopic Instrument Data Release 2) [197] results yield $\Omega_{m,0} = 0.3027 \pm 0.0036$ from the joint CMB+BAO analysis, implying $\Omega_{\Lambda,0} \simeq 0.7$.

Despite its overall success, the ΛCDM model faces several challenges. Among these, the Hubble tension is the $> 4\sigma$ discrepancy between the H_0 value inferred from early-Universe data (CMB and BAO) and that measured directly in the late Universe using distance-ladder techniques [180][4]. Under the ΛCDM framework, early-Universe probes, anchored by the sound-horizon scale at recombination, consistently prefer a lower Hubble constant; for in-

3. The sound horizon is the comoving distance travelled by acoustic waves in the photon–baryon fluid before recombination.
4. The discrepancy is approximately $\sim 5.3\sigma$ when comparing with CMB measurements and $\sim 4.5\sigma$ when comparing with recent DESI BAO results [197]. The CMB and BAO measurements themselves also exhibit a discrepancy at the $> 2\sigma$ level.

stance, the DESI DR 2 results show $H_0 \simeq 67$–68 km s^{-1} Mpc^{-1} [197]. In contrast, local distance-ladder measurements from the SH$_0$ES program find $H_0 \simeq 73$ km s^{-1} Mpc^{-1} [181]. This discrepancy is not merely numerical. Since H_0 sets the absolute scale of the distance-redshift relation, a few percent difference suggests that either unknown systematics remain in the measurement techniques or the ΛCDM model is incomplete.

In summary, while ΛCDM provides an excellent phenomenological description of a wide range of observations, the Hubble tension highlights the need for models beyond ΛCDM or independent, complementary probes to test its validity.

3.2 GWs as Probes of Cosmic Expansion Rate

As discussed in the previous section, the expansion history of the Universe is encoded in the distance–redshift relation (Eq. (3.31)), with the luminosity distance d_L serving as a key observable. GW observations provide a direct means of measuring this quantity and have therefore emerged as a robust and independent cosmological probe. Unlike traditional standard candles, which rely on the hierarchical cosmic distance ladder and its associated calibration uncertainties, GWs enable a direct measurement of d_L through the self-calibrating nature of the waveform, as encoded in the phase and amplitude evolution predicted by GR.

To date, ground-based interferometers have observed CBCs across a wide range of distances. BBH mergers have been observed over a broad range of distances, spanning from a few hundred megaparsecs to several gigaparsecs ($z \sim 0.1$–1) [20], whereas BNS mergers are typically observed within a few hundred megaparsecs [198, 199]. This observational reach enables GW observations to probe cosmological models and constrain fundamental parameters [200, 201, 202]. In the low-redshift regime, GW distance measurements from CBCs can be directly compared with Type Ia SNe, providing an independent cross-check of the local distance scale. Looking ahead, next-generation ground-based detectors and space-based observatories will extend GW observations to $z \gtrsim 1$, probing a redshift regime complementary to those of SNe and the CMB, and thereby offering new insights into the origin of the Hubble tension.

The efficacy of GWs as probes of cosmic expansion depends on two key factors: the precision and accuracy of luminosity-distance measurements and the determination of the source redshift. The latter is most robustly obtained through the identification of an EM counterpart or a host galaxy. However, GW-based cosmography is subject to several sources of uncertainty affecting both distance and redshift inference, including detector calibration errors, the degeneracy between d_L and the binary inclination angle (ι), gravitational lensing effects, and peculiar velocities of host galaxies, which perturb the redshift–distance mapping.

3.2.1 Standard Sirens

The LVK Collaboration has detected numerous CBC signals and inferred their luminosity distances from the observed GW strain. Recall that, as shown in Eq. (1.23), the strain amplitude $h(t)$ of a CBC in-spiral is inversely proportional to the luminosity distance:

$$h \propto \frac{\mathcal{M}_c^{5/3}}{d_L} f^{2/3}. \tag{3.35}$$

Since $\mathcal{M}_c$ governs the frequency evolution $\dot{f}(t)$ at leading post-Newtonian order, it can be tightly constrained from the phase evolution of the signal. Once $\mathcal{M}_c$ is determined, the intrinsic amplitude of the waveform is correspondingly fixed. Consequently, d_L can be extracted directly from the observed strain amplitude without reference to external cosmic distance ladders. For this reason, CBCs are referred to as *standard sirens* [203, 204], providing a self-calibrating measure of distance.

When combined with complementary redshift information, standard sirens enable a direct measurement of H_0 and provide an independent cross-check of the expansion history of the Universe relative to constraints from the CMB, BAO, and Type Ia SNe, thereby offering a potential avenue to probe the Hubble tension. Standard sirens are generally classified into two categories: bright and dark sirens.

Bright sirens are events accompanied by identifiable EM counterparts. A representative example is the BNS merger, which can produce short GRB [205, 206] and kilonova emission [207, 208], as introduced in Sections 1.2.1 and 2.3.2. NSBH mergers can also produce EM counterparts, provided that the NS is tidally disrupted. Due to the typically higher masses of astrophysical BHs compared to NSs, NSBH systems can be detected to greater distances [209]. For those events, the redshift can be determined by identifying the host galaxy through EM follow-up observations and measuring its spectroscopic redshift[5]. The BNS merger GW170817 [16] remains the prototypical example, providing the first H_0 measurement from a GW event with EM counterparts [213].

In contrast, *dark sirens*, such as BBH mergers, lack EM counterparts, and their redshifts must be inferred statistically. Two primary approaches are commonly employed. The first is the *catalogue method*, which utilises galaxy catalogues (e.g., GLADE+ [214]) within the GW localisation volume, assigning probabilistic weights to potential host galaxies [215, 216, 217, 218, 219, 220]. The second is the *spectral siren method*, which exploits features of the compact-object mass distribution (e.g., the mass gap or characteristic mass scales) to partially break the degeneracy between intrinsic source mass and cosmological redshift [221, 222, 223]. By comparing the observed distribution of detector-frame mass

5. As a special case, BNS events can also provide redshift information through tidal deformability (characterised by the Love number [210], leading to so-called *love sirens*. Since tidal effects depend on the source-frame mass, the observed detector-frame mass can be combined with tidal information to partially break the mass–redshift degeneracy even in the absence of an EM counterpart [211, 212].

($m_{\rm d} \equiv m_{\rm s}(1+z)$) with an assumed source-frame mass (m_s) distribution, the redshift (z) can be inferred in a population-statistical sense. More recently, alternative approaches based on cross-correlation techniques have been proposed, in which constraints on H_0 are obtained by comparing the large-scale angular distribution of galaxies with the luminosity-distance distribution of GW events [224, 225, 226].

By measuring sets of (d_L, z) from an ensemble of GW events (or from a single event in the case of GW170817), one can perform hierarchical Bayesian inference to constrain H_0, as will be discussed in Section 4.4. Figure 3.3 shows the inferred H_0 values from various standard-siren methods and their combination with the Gravitational-Wave Transient Catalog-4.0 (GWTC-4.0) data [202]. Among these, the bright-siren measurement based on GW170817 provides the most precise constraint. However, the current samples are not yet sufficiently large to reach a precision comparable to that of Planck or SH$_0$ES, and therefore yield comparatively broader inferences.

With an increasing number of standard-siren detections, low-redshift events ($z \lesssim 1$) primarily constrain H_0 in the linear Hubble–Lemaître regime. At higher redshifts—potentially accessible with next-generation detectors or through strong lensing—standard sirens become increasingly sensitive to $H(z)$ in Eq. (3.34), and thereby to other cosmological parameters, such as Ω_m. As such, they probe a redshift regime complementary to that of Type Ia SNe and CMB observations, helping to break parameter degeneracies and improve the inference of H_0.

3.2.2 Uncertainties and Biases

For precision cosmology, precise and accurate measurements of the luminosity distance d_L are essential. For individual events, the typical fractional uncertainty in d_L is at the level of a few tens of percent, primarily limited by the SNR and the intrinsic degeneracy between d_L and the binary inclination angle (ι; since both parameters modulate the observed GW amplitude).

Breaking this degeneracy requires sensitivity to the signal's polarisation, achieved via a network of at least three detectors with non-aligned orientations. In such a network, the relative responses to the two GW polarisations ($+$ and $\times$) enable constraints on the inclination angle, thereby partially disentangling it from d_L (see Eq. (1.23)). Accurate sky localisation further aids this process by constraining the detectors' antenna response functions in Eq. (1.29). In addition, calibration uncertainties in the detector response introduce systematic errors in the inferred strain amplitude. These arise from imperfect knowledge of the instrument's frequency-dependent amplitude and phase response, effectively leading to biases in the reconstructed waveform and, consequently, in the inferred luminosity distance. Such uncertainties are typically modelled as Gaussian perturbations and incorporated into the likelihood through marginalisation (see Section 4.3.3 for a more detail).

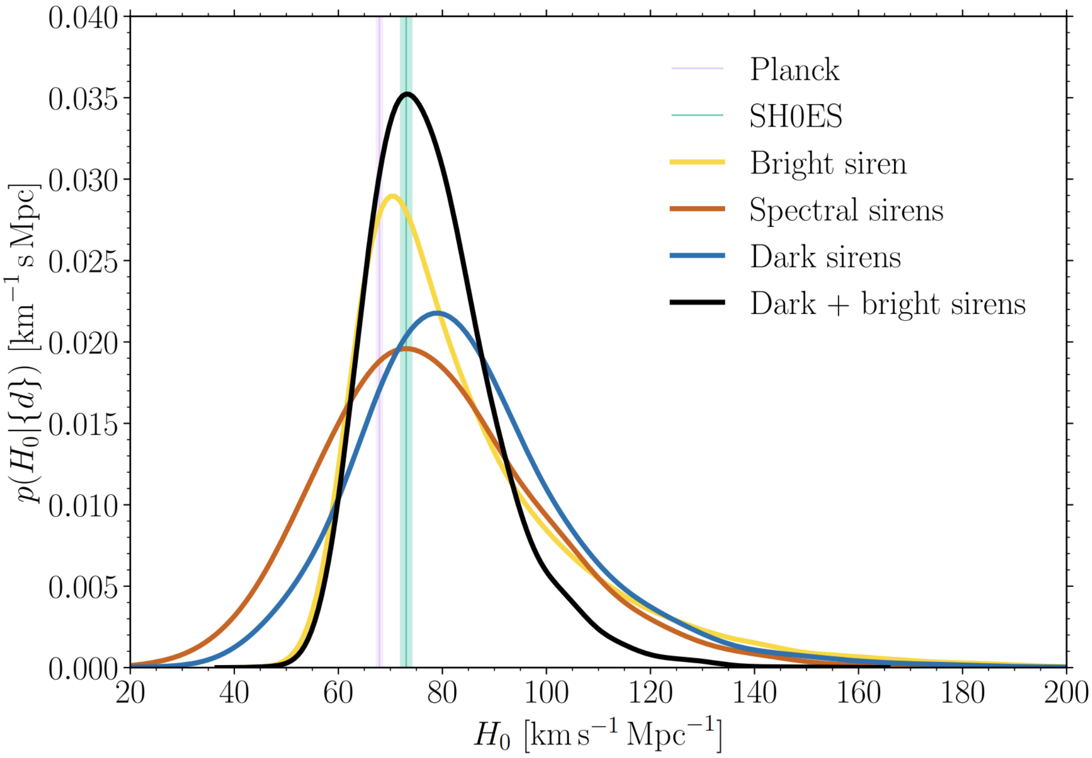

Figure 3.3: Probability density distributions of H_0 inferred from various standard-siren methods and their combination with the GWTC-4.0 catalogue data. Figure is from [202].

Beyond these instrumental and intrinsic limitations, several astrophysical and cosmological systematics must be considered. One important effect is gravitational lensing. Although the rate of strongly lensed GW events is expected to be low for current detectors, weak lensing is more common and cannot be identified on an event-by-event basis. Unlike strong lensing or microlensing—which can produce distinctive signatures such as multiple images or waveform distortions (see Section 2.4)—weak lensing induces stochastic magnification or demagnification of the signal, thereby biasing the inferred d_L.

As detector sensitivity improves, GW observations will probe progressively higher redshifts, where the probability and impact of lensing increase. At higher redshifts ($z \gtrsim$ 1–2), this effect becomes increasingly significant, with the dispersion in magnification commonly approximated as $\sigma_{\text{lens}} \simeq 0.066z$ [204]. This additional scatter can dominate the distance uncertainty budget and obscure the subtle redshift dependence of $d_L(z)$ arising from cosmological parameters such as the dark energy equation of state $w(z)$ or the spatial curvature Ω_k. Conversely, in the very low-redshift regime ($z \ll 1$), peculiar velocities of host galaxies introduce additional uncertainty in the redshift measurement through Doppler shifts [227]. This effect perturbs the redshift–distance relation and is particularly relevant for bright sirens, which are typically observed in the nearby Universe [228, 229].

Overall, robust cosmological inference requires statistical analyses over large event samples, along with mitigation strategies, potentially enabled by overlapping large-scale structure surveys.

Chapter 4

Bayesian Inference in Gravitational-Wave Astronomy

Extracting physical information from GW signals is fundamentally a problem of statistical inference. Since GW signals are typically buried deep within instrumental noise and depend on a multi-dimensional parameter space, simple point-estimation techniques are entirely inadequate. Instead, Bayesian inference is strictly required to rigorously quantify uncertainties and correlations between parameters [230, 231, 232, 233, 234, 235, 236]. Parameter estimation of GW signals enables the inference of source properties from observational data, including component masses, spins, and orbital dynamics, as well as the statistical properties of the underlying population and interactions with the local environment [20, 237]. Furthermore, such analyses allow us to identify the presence of intervening massive objects between GW sources and detectors and constrain fundamental cosmological parameters [202, 238].

Considering the CBC signals, Bayesian inference is performed over a high-dimensional parameter space (typically $\sim$15 dimensions, and higher when including spin precession, eccentricity, or lensing effects), which presents a severe computational challenge [236]. Traditional grid-based numerical integration scales exponentially with dimensionality, making it computationally intractable due to the so-called *curse of dimensionality* [239]. Consequently, evaluating the likelihood requires pipelines equipped with advanced stochastic sampling techniques (e.g., [236, 240]). More recently, the need to evaluate millions of waveform likelihoods has motivated the development of machine-learning-based approaches that can substantially accelerate GW parameter inference (e.g., [241, 242])

In this chapter, we present the theoretical framework of Bayesian inference for GW signals, along with representative examples and advanced applications. We begin by outlining the mathematical framework of Bayesian inference. We then describe the principles of stochastic sampling algorithms, focusing on Markov Chain Monte Carlo [243, 244, 245] and nested sampling [246, 247] methods, and discuss recent developments employing machine learning techniques, such as normalising flows, to accelerate inference. Subsequently, we present a detailed treatment of parameter estimation for CBC signals, with particular emphasis on BBH systems. Finally, we review the framework of hierarchical Bayesian analysis, highlighting applications to BBH population inference and measurements of the Hubble constant using GW observations.

4.1 Fundamentals of Bayesian Statistics

Bayesian inference provides a consistent logical framework for updating scientific knowledge in light of new experimental data. Unlike the frequentist paradigm, which defines probability as the long-run relative frequency of repeatable events and treats model parameters as fixed quantities, Bayesian statistics interprets probability as a measure of plausibility [248, 249]. This interpretation allows probability distributions to be assigned directly to physical parameters, thereby enabling rigorous quantification of parameter uncertainties from a single observation. In the context of GW astronomy, where GW signals are often deeply buried within detector noise, the Bayesian framework is indispensable for extracting physical parameters of GW sources from the strain data and for distinguishing between competing astrophysical and cosmological models.

4.1.1 Bayes' Theorem

The state of knowledge is updated through *Bayes' theorem*, whereby the *priors* (π), encoding pre-existing knowledge about the parameters, is combined with the *likelihood* ($\mathcal{L}$), which quantifies the probability of the observed data (d) given a set of parameters ($\boldsymbol{\theta}$). The resulting product is normalised to yield the *posteriors* ($p(\boldsymbol{\theta}|d)$), which encapsulate the updated constraints on the parameters inferred from the data as follows [248, 250]:

$$p(\boldsymbol{\theta} \mid d) = \frac{\mathcal{L}(d \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\mathcal{Z}(d)}, \tag{4.1}$$

where $\mathcal{Z}(d) \equiv \int \mathcal{L}(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}$ is the Bayesian *evidence* or the likelihood marginalised over the full parameter space.

In the context of GW astronomy, d can correspond to the strain time series measured by detectors, while $\boldsymbol{\theta}$ denotes the set of physical parameters describing the GW source [251]. The prior $\pi(\boldsymbol{\theta})$ encodes pre-measurement knowledge or physical constraints on these parameters; for example, component masses must be positive, spin magnitudes are bounded, and the secondary-to-primary mass ratio lies within the interval $[0, 1]$. The likelihood $\mathcal{L}(d \mid \boldsymbol{\theta})$ incorporates both the properties of the detector noise and the theoretical waveform model. Under the standard assumption of stationary Gaussian noise, it can be derived using

Eq. (1.40) introduced in Section 1.3.2. The evidence $\mathcal{Z}(d)$ serves not merely as a normalisation constant for parameter estimation, but also plays a central role in model selection [232, 252]. In particular, it enables quantitative comparison between competing hypotheses, such as a noise-only model versus a signal model, or between unlensed and lensed GW scenarios.

While the evidence $\mathcal{Z}(d)$ is obtained by marginalising over the full parameter space, in practice one is often interested in a subset of parameters. In such cases, the posterior distribution for parameters of interest can be obtained by marginalising the joint posterior over the remaining nuisance parameters [248].

For parameters $\boldsymbol{\theta}_k$, the marginalised posterior is given by:

$$\begin{aligned} p(\boldsymbol{\theta}_k \mid d) &= \int p(\boldsymbol{\theta} \mid d) d\boldsymbol{\theta}_{\setminus k} \\ &= \frac{1}{\mathcal{Z}(d)} \int \mathcal{L}(d \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\setminus k}, \end{aligned} \tag{4.2}$$

where $\boldsymbol{\theta}_{\setminus k}$ denotes the set of all parameters except θ_k. Correspondingly, the marginalised likelihood is defined as:

$$\mathcal{L}_{\text{marg}}(d \mid \boldsymbol{\theta}_k) \equiv \int \mathcal{L}(d \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\setminus k}. \tag{4.3}$$

Such marginalisation reduces the effective dimensionality of the parameter space, which can improve computational efficiency and accelerate inference. More importantly, the marginalised posterior for $\boldsymbol{\theta}_k$ consistently incorporates the uncertainties in all other parameters. As a result, it is generally broader than the corresponding conditional distribution obtained by fixing the remaining parameters to specific values.

Once the posterior distributions are obtained, the inferred parameters and their uncertainties can be summarised using suitable estimators. The median of the 1-dimensional marginalised posterior is commonly adopted as a point estimate for a given parameter. The associated uncertainty is quantified using a *credible interval* (C.I.), defined as the region of parameter space that contains a specified fraction of the posterior probability mass [248, 250]. Let $F(\theta)$ denote the cumulative distribution function (CDF) of the marginalised posterior $p(\theta \mid d)$:

$$F(\theta) \equiv \int_{-\infty}^{\theta} p(\theta' \mid d)\, d\theta'. \tag{4.4}$$

Then, a central C.I. containing a fraction q of the posterior probability is defined by the bounds θ_{low} and θ_{high} satisfying:

$$F(\theta_{\text{low}}) = \frac{1-q}{2}, \qquad F(\theta_{\text{high}}) = \frac{1+q}{2}. \tag{4.5}$$

For example, setting $q = 0.9$ yields the 90% C.I.. This framework allocates equal posterior probability to the lower and upper tails, ensuring a symmetric coverage in probability, although the resulting interval need not be symmetric in parameter values if the posterior is skewed.

Furthermore, the relative precision of a parameter can be quantified using the fractional uncertainty, defined as the width of the central C.I. divided by the median of the marginalised posterior:

$$\delta_\theta \equiv \frac{\theta_{\text{high}} - \theta_{\text{low}}}{\theta_{\text{median}}}, \tag{4.6}$$

where θ_{median} is the median of the posterior. This dimensionless quantity provides a standardised measure of parameter uncertainty, allowing direct comparison of the relative precision across parameters with different physical units or scales.

4.1.2 Hypothesis Testing

While parameter estimation infers the parameters associated with a given hypothesis, scientific inference more broadly requires assessing the relative plausibility of competing hypotheses, a procedure commonly referred to as *hypothesis testing* [253]. Consider two hypotheses, $\mathcal{H}_1$ (e.g., the presence of a GW signal) and $\mathcal{H}_0$ (e.g., the absence of a signal in the data). The relative support for $\mathcal{H}_1$ over $\mathcal{H}_0$ is quantified by the ratio of their posterior probabilities, known as the *odds ratio*:

$$\begin{aligned} \mathcal{O}_0^1 &\equiv \frac{p(\mathcal{H}_1|d)}{p(\mathcal{H}_0|d)} \\ &= \frac{\mathcal{Z}(d|\mathcal{H}_1)}{\mathcal{Z}(d|\mathcal{H}_0)} \frac{\pi(\mathcal{H}_1)}{\pi(\mathcal{H}_0)} \\ &= \mathcal{B}_0^1 \Pi_0^1, \end{aligned} \tag{4.7}$$

where $\mathcal{B}_0^1$ denotes the *Bayes Factor*, defined as the ratio of evidences for the same dataset d under the two competing hypotheses, and Π_0^1 represents the *prior odds*, encoding the prior belief assigned to each hypothesis.

Although $\mathcal{O}_0^1$ in Eq. (4.7) is the final statistic to select the favoured hypothesis, it is common practice to report $\mathcal{B}_0^1$ alone. In the absence of compelling theoretical or observational constraints, competing hypotheses are typically assigned equal prior probabilities ($\Pi_0^1 = 1$). For example, when testing for a lensing signature in data without prior population information, the lensed and unlensed hypotheses are often treated with equal prior weight. An additional motivation for reporting the Bayes factor is that it isolates the empirical information contained in the data [236]. The evidence $\mathcal{Z}$ encapsulates the full data-driven support for a given hypothesis, but the prior odds Π are based on external assumptions independent of the dataset. By quoting $\mathcal{B}_0^1$, one provides a result that can be consistently reinterpreted under alternative prior choices. For instance, if the GW lensing rate is known to be $\sim 10^{-3}$ per event, the corresponding prior odds ($\Pi \sim 10^{-3}$) can be applied a Bayes factor to recover the full odds ratio.

A value of $\mathcal{B}_0^1 > 1$ indicates that the data favour $\mathcal{H}_1$ over $\mathcal{H}_0$, with larger values corresponding to stronger levels of support. In practice, robust inference requires that such conclusions remain stable under reasonable variations of the prior [253]. If modest changes in prior width qualitatively alter the inference, the result cannot be considered robust. In addition, the Bayesian evidence provides a quantitative formulation of Occam's razor [254]. Hypotheses with higher-dimensional parameter spaces distribute their prior probability over larger volumes, thereby reducing the evidence unless the additional parameters yield a substantial improvement in likelihood. Consequently, the Bayes factor penalises unnecessary model complexity, favouring models that provide a better balance between goodness of fit and model complexity.

Despite its utility, the Bayes factor alone is not sufficient to establish a confident detection in GW data analysis. In typical applications, the likelihood used to compute $\mathcal{B}_0^1$ is constructed under the assumption of stationary Gaussian detector noise. However, real GW detector data contain non-Gaussian and non-stationary transient artifacts, commonly referred to as *glitches*, which can mimic astrophysical signals [72, 252]. As a result, instrumental noise can occasionally produce data that are spuriously favoured by a signal model, leading to false positives.

To assess the statistical significance of a candidate GW event, an empirical background distribution of the detection statistic must therefore be constructed [81, 255]. This background represents the distribution of $\mathcal{B}_0^1$ (or an equivalent ranking statistic) obtained purely from noise. The statistical significance of a candidate is then quantified using the FAR, defined as the expected rate at which noise alone produces events with a detection statistic equal to or exceeding that of the candidate. Within this framework, a robust GW detection requires satisfying two complementary criteria: the Bayes factor must strongly favour the signal hypothesis over the Gaussian noise model, and the FAR must be sufficiently low to demonstrate that the observed event is highly unlikely to arise from instrumental or environmental sources.

4.2 Stochastic Sampling

To obtain the posteriors of model parameters, one must evaluate the likelihood of the data given a hypothesis described by a set of parameters. In GW astronomy for CBCs, the parameter space spans at least 15 dimensions. Evaluating the posterior distribution over such a high-dimensional parameter space via direct numerical integration or grid-based quadrature is computationally prohibitive.

In a grid-based approach, the likelihood is evaluated on a predefined lattice spanning each parameter dimension. For instance, even a coarse resolution of 10 points per parameter in a 15-dimensional space requires 10^{15} likelihood evaluations. This exponential scaling renders grid-based methods impractical for CBC parameter estimation. Consequently, GW data analysis relies on stochastic sampling algorithms that generate samples distributed according to the posterior without requiring likelihood evaluation over the entire parameter space.

4.2.1 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) algorithms construct a sequence of samples whose empirical distribution converges to a target distribution $p(\boldsymbol{\theta}|d)$. These methods are rooted in Monte Carlo integration, where an expectation value with respect to $p(\boldsymbol{\theta}|d)$ is approximated using random samples. For any integrable function $f(\boldsymbol{\theta})$, the posterior expectation is defined as:

$$\mathbb{E}[f] = \int f(\boldsymbol{\theta}) p(\boldsymbol{\theta}|d)\, d\boldsymbol{\theta}. \tag{4.8}$$

Rather than evaluating this integral directly, one can approximate it using samples $\{\boldsymbol{\theta}_i\}$ drawn from the posterior in the limit of large N [256]:

$$\mathbb{E}[f] \approx \frac{1}{N} \sum_{i=1}^{N} f(\boldsymbol{\theta}_i). \tag{4.9}$$

In practice, however, direct sampling from $p(\boldsymbol{\theta}|d)$ is not feasible in a high-dimensional space. MCMC addresses this issue by constructing a *Markov chain* $\{\boldsymbol{\theta}_i\}$ whose stationary distribution is the target posterior. Here, ‘stationary‘ means that once the chain has converged, the probability of finding the chain at a point $\boldsymbol{\theta}$ is given by $p(\boldsymbol{\theta}|d)$, and this distribution remains invariant under further evolution of the chain. The defining property of a Markov chain is that the next sample depends only on the current one:

$$p(\boldsymbol{\theta}_{i+1}|\boldsymbol{\theta}_i, \boldsymbol{\theta}_{i-1}, \dots) = p(\boldsymbol{\theta}_{i+1}|\boldsymbol{\theta}_i). \tag{4.10}$$

By constructing transition rules that satisfy conditions such as detailed balance with respect to $p(\boldsymbol{\theta} \mid d)$, the Markov chain is guaranteed to converge to the posterior distribution in the long-run limit. The resulting correlated samples $\{\boldsymbol{\theta}_i\}$ can then be used in Eq. (4.9).

One of the most widely used MCMC algorithms is the *Metropolis-Hastings* (MH) algorithm [243, 244], which is utilised because it enforces the detailed balance condition (reversibility) while allowing asymmetric proposal distributions. The algorithm operates through the following iterative procedure:

1. Initialise at a parameter point $\boldsymbol{\theta}_i$.
2. Draw a candidate point $\boldsymbol{\theta}_{\text{prop}}$ from a proposal distribution $Q(\boldsymbol{\theta}_{\text{prop}}|\boldsymbol{\theta}_i)$.
3. Compute the *acceptance ratio* $\alpha(\boldsymbol{\theta}_i \to \boldsymbol{\theta}_{\text{prop}})$ [256]. The MH algorithm constructs α such that the resulting Markov chain satisfies the *detailed balance* condition with respect to $p(\boldsymbol{\theta} \mid d)$, which is given by:

$$p(\boldsymbol{\theta}_i \mid d)\, T(\boldsymbol{\theta}_{\text{prop}} \mid \boldsymbol{\theta}_i) = p(\boldsymbol{\theta}_{\text{prop}} \mid d)\, T(\boldsymbol{\theta}_i \mid \boldsymbol{\theta}_{\text{prop}}), \tag{4.11}$$

where T denotes the transition kernel, which is defined as:

$$T(\boldsymbol{\theta}_{\text{prop}} \mid \boldsymbol{\theta}_i) = Q(\boldsymbol{\theta}_{\text{prop}} \mid \boldsymbol{\theta}_i)\, \alpha(\boldsymbol{\theta}_i \to \boldsymbol{\theta}_{\text{prop}}). \tag{4.12}$$

Eq. (4.11) then yields the constraint:

$$\frac{\alpha(\boldsymbol{\theta}_i \to \boldsymbol{\theta}_{\text{prop}})}{\alpha(\boldsymbol{\theta}_{\text{prop}} \to \boldsymbol{\theta}_i)} = \frac{p(\boldsymbol{\theta}_{\text{prop}} \mid d)\, Q(\boldsymbol{\theta}_i \mid \boldsymbol{\theta}_{\text{prop}})}{p(\boldsymbol{\theta}_i \mid d)\, Q(\boldsymbol{\theta}_{\text{prop}} \mid \boldsymbol{\theta}_i)}. \tag{4.13}$$

A sufficient choice that satisfies this condition while ensuring $0 \leq \alpha \leq 1$ is:

$$\alpha(\boldsymbol{\theta}_i \to \boldsymbol{\theta}_{\text{prop}}) = \min\left(1, \frac{p(\boldsymbol{\theta}_{\text{prop}} \mid d)\, Q(\boldsymbol{\theta}_i \mid \boldsymbol{\theta}_{\text{prop}})}{p(\boldsymbol{\theta}_i \mid d)\, Q(\boldsymbol{\theta}_{\text{prop}} \mid \boldsymbol{\theta}_i)}\right). \tag{4.14}$$

Eq. (4.14) guarantees that $p(\boldsymbol{\theta} \mid d)$ is a stationary distribution of the Markov chain.
4. Draw a value u from a Uniform distribution $\mathcal{U}(0, 1)$.
5. If $u \leq \alpha$, accept the candidate and set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_{\text{prop}}$; otherwise retain the current point, $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$.

This mechanism enables the Markov chain to preferentially explore high-probability regions while still allowing occasional transitions to lower-probability states, thereby ensuring global exploration of the parameter space. Note that the Bayesian evidence cancels in the acceptance ratio in Eq. (4.14), such that only the likelihood and prior need to be evaluated. Consequently, standard MCMC algorithms sample from an unnormalised target distribution and do not directly provide the evidence. Estimating the evidence, therefore, requires additional techniques, such as thermodynamic integration, which introduces a sequence of tempered likelihoods that interpolate between the prior and the posterior [257].

MCMC methods efficiently explore high-dimensional spaces and are guaranteed to converge to the target distribution in the limit of infinite samples. In practice, however, their efficiency can degrade for multi-modal distributions, where the chain may remain confined to a local maximum. This limitation is particularly relevant in GW data analysis, where degeneracies in the detector response can produce distinct modes, such as the ring-like structures in sky localisation. Furthermore, practical implementations require careful assessment of convergence and the choice of an appropriate termination criterion, as strong autocorrelations within the chain can substantially reduce the effective sample size and lead to biased or unreliable inference [258, 259]. Moreover, the generated Markov chain requires post-processing. An initial *burn-in* phase is discarded to remove samples that are strongly influenced by the initial condition $\boldsymbol{\theta}_0$, ensuring that the retained samples are representative of the stationary distribution [245, 260].

In astrophysical applications, the standard MH algorithm is often replaced by ensemble-based samplers, most notably the affine-invariant implementation provided in the Python package `emcee` [261]. Rather than evolving a single Markov chain, `emcee` simultaneously evolves an ensemble of *walkers* that explore the parameter space in parallel. At each iteration, a walker proposes a new position by moving along the line connecting it to another randomly selected walker, allowing the ensemble to collectively adapt to the global geometry of the target distribution.

4.2.2 Nested Sampling

As mentioned above, standard MCMC methods do not directly provide the Bayesian evidence, and additional post-processing is required to estimate it. An alternative class of Monte Carlo methods is nested sampling [246, 247], which is specifically designed to compute the Bayesian evidence $\mathcal{Z}$ while simultaneously producing posterior samples of the model parameters as a byproduct. For this reason, nested sampling is widely adopted in GW data analysis, where both parameter estimation and hypothesis testing are required.

Unlike MCMC methods, which walk through the parameter space to sample the posterior directly, nested sampling reformulates the multidimensional evidence integral into a one-dimensional integral. Specifically, the Bayesian evidence for a model defined over a multidimensional parameter space $\boldsymbol{\theta}$ is given by:

$$\mathcal{Z} = \int \mathcal{L}(\boldsymbol{\theta})\, \pi(\boldsymbol{\theta})\, d\boldsymbol{\theta}, \tag{4.15}$$

where $\mathcal{L}(\boldsymbol{\theta})$ and $\pi(\boldsymbol{\theta})$ denote the likelihood and prior, respectively. Nested sampling proceeds by introducing the *prior volume* X, defined as the fraction of the prior volume where $\mathcal{L}(\boldsymbol{\theta})$ exceeds a certain threshold λ:

$$X(\lambda) = \int_{\mathcal{L}(\boldsymbol{\theta})>\lambda} \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}. \tag{4.16}$$

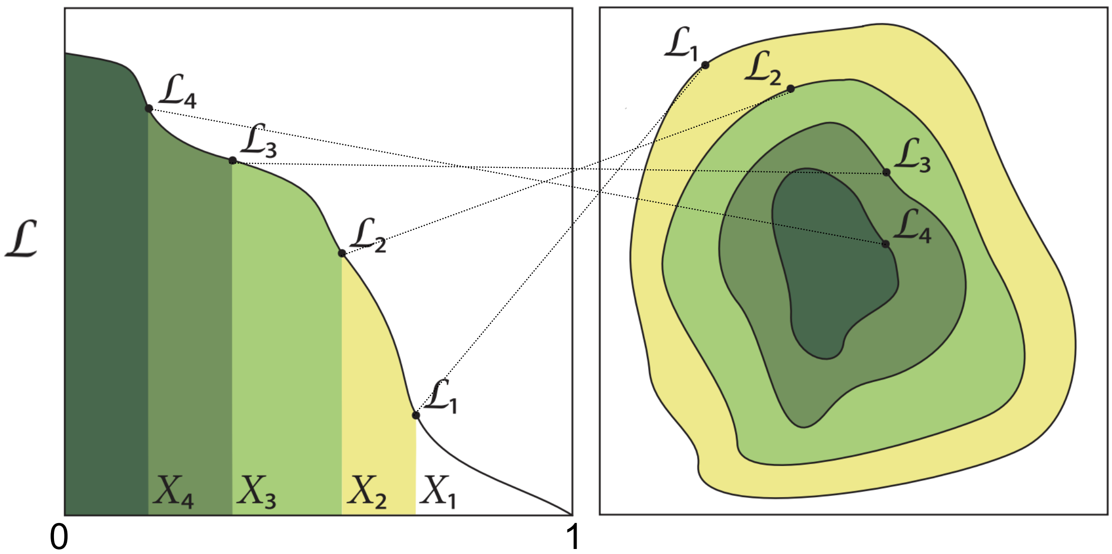

Figure 4.1: Illustration of partitioning the multi-dimensional parameter space into nested subsets based on $\mathcal{L}_i$. The left panel shows the progressive reduction of the prior mass (X_i) as λ increases, while the right panel depicts the corresponding iso-likelihood contours in the parameter space. Adapted from [262].

As the likelihood threshold λ increases, the allowed region $\mathcal{L} > \lambda$ shrinks, and the corresponding prior volume X decreases monotonically from 1 to 0, which allows the likelihood to be reparameterised as a function of prior volume, $\mathcal{L}(X)$. With this change of variables, Eq. (4.15) can be rewritten as a one-dimensional integral as follows:

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)\, dX, \tag{4.17}$$

which corresponds to integrating the likelihood over nested *shells* of decreasing prior volume (see Figure 4.1).

The algorithm is initialised with a set of N_{live} samples (live points) drawn from the prior. At each iteration i, the live point with the lowest likelihood, $\mathcal{L}_i$, is removed and replaced by a new sample drawn from the prior subject to the constraint $\mathcal{L} > \mathcal{L}_i$. This constraint progressively restricts the sampling region to smaller prior volumes, generating a nested sequence $X_0 > X_1 > X_2 > \cdots$. Each discarded sample is assigned a weight given by the shrinkage in prior volume, $w_i \approx X_{i-1} - X_i$, such that the evidence is accumulated as:

$$\mathcal{Z} \approx \sum_i \mathcal{L}_i\, w_i. \tag{4.18}$$

Meanwhile, after iteration i, the remaining N_{live} live points occupy a region of prior volume X_i, corresponding to likelihood values greater than the current minimum $\mathcal{L}_i$. This region can contribute only a limited additional amount to the total evidence. Among the live points, the maximum likelihood is denoted by $\mathcal{L}_{\max}$, which provides an upper bound on the remaining contribution as follows:

$$\Delta\mathcal{Z}_i \lesssim \mathcal{L}_{\max} X_i. \tag{4.19}$$

Once this upper bound becomes sufficiently small compared to the accumulated evidence, further iterations will not significantly change the result, and the algorithm can be terminated. Upon termination, a discrete representation of the posterior distribution is reconstructed from the same set of weighted samples as a byproduct. The posterior weight assigned to each discarded point is:

$$p_i = \frac{\mathcal{L}_i w_i}{\mathcal{Z}}, \tag{4.20}$$

which can be used to compute expectation values of model parameters.

Nested sampling is generally more computationally expensive than standard MCMC methods since each iteration requires generating a new sample from the prior under a likelihood constraint. As the iterations proceed, the likelihood threshold λ increases, and the region satisfying this condition becomes progressively smaller. In high-dimensional problems, such regions occupy only a tiny fraction of the original prior volume, so most trial points drawn from the prior fail to satisfy the constraint and are rejected. The repeated rejection significantly increases the computational cost.

In addition, the algorithm maintains a set of live points throughout the run, and each iteration removes one point and replaces it with a new constrained sample. The overall computational cost, therefore, depends on both the number of live points and the efficiency of sampling within the restricted region. The latter is particularly sensitive to the choice of prior, as excessively broad priors enlarge the initial parameter space and prolong the contraction toward the high-likelihood region.

A number of publicly available inference packages implementing nested sampling are widely used in GW data analysis. Among these, `dynesty` [263] provides a flexible and pure-Python implementation, supporting both static and dynamic nested sampling, where the number of live points can be adaptively allocated to improve sampling efficiency. Another commonly used package is `PyMultiNest` [264], which interfaces with the `MultiNest` algorithm [265] and employs ellipsoidal decompositions of the live point set to efficiently sample from constrained prior volumes. Both approaches are designed to address the central challenge of nested sampling, namely, the efficient generation of new samples within increasingly restricted regions of parameter space.

More recently, machine learning-based approaches have been introduced to further improve sampling efficiency. In particular, `nessai` [241] incorporates normalising flows to construct flexible, data-driven proposals for nested sampling. A normalising flow is a sequence of invertible transformations that maps a simple base distribution (e.g., a multivariate Gaussian) to a complex target distribution. By learning this transformation from the current set of live points, the sampler can approximate the geometry of the constrained posterior region and generate new samples with significantly higher acceptance rates. This approach is particularly effective for problems with complex, multi-modal, or highly degenerate structures, where traditional sampling strategies may become inefficient.

4.3 Analysis of GWs from BBHs

To date, all GW signals identified as significant detections by the LVK Collaboration have been consistent with signals from CBCs, including BBHs, NSBHs, and BNSs [20]. Inference of source properties from GW signals provides direct constraints on the astrophysical nature of compact objects. For BBHs, parameter estimation extracts the mass and spin distributions, directly addressing the existence of IMBHs and the boundaries of theoretical mass gaps. For systems involving NSs, measurements of the tidal deformability constrain the nuclear equation of state at supra-nuclear densities. In NSBH systems, the inferred mass ratio and spin configuration determine whether the NS undergoes tidal disruption or direct plunge, thereby constraining their formation and evolutionary pathways.

Beyond individual events, the ensemble of inferred parameters enables population-level studies. In particular, the distributions of masses and spins probe the underlying formation channels, distinguishing between isolated binary evolution and dynamical capture. Furthermore, as discussed in Chapter 3, the direct encoding of luminosity distance in the waveform allows CBCs to serve as standard sirens for independent cosmological measurements. Moreover, precise inference enables stringent tests of GR in the strong-field, highly dynamical regime, including consistency checks between the in-spiral and ringdown phases [266].

In this section, we detail the application of the Bayesian inference framework to extract the physical parameters of the GW sources from the observed strain data. Given that the scope of this thesis is restricted to BBH mergers, the following discussion focuses on the physical properties and parameter estimation procedures specific to BBH systems.

4.3.1 Standard 15 BBH Parameters and Waveform Models

Under the assumption of quasi-circular orbits, i.e., the binary has radiated away any initial orbital eccentricity through GW emission before entering the detector band, a BBH signal is fully described by a 15-dimensional parameter vector, $\boldsymbol{\Theta}$, which can be decomposed into intrinsic and extrinsic parameters. The intrinsic parameters, $\boldsymbol{\theta}_{\rm int}$, characterise the physical properties of the binary and are independent of the observer. These include the component masses (m_1, m_2) and the spin vectors of each object $(\vec{\chi}_1, \vec{\chi}_2)$, each of which has three degrees of freedom, yielding six spin components in total. Figure 4.2 illustrates the intrinsic parameters of a BBH system.

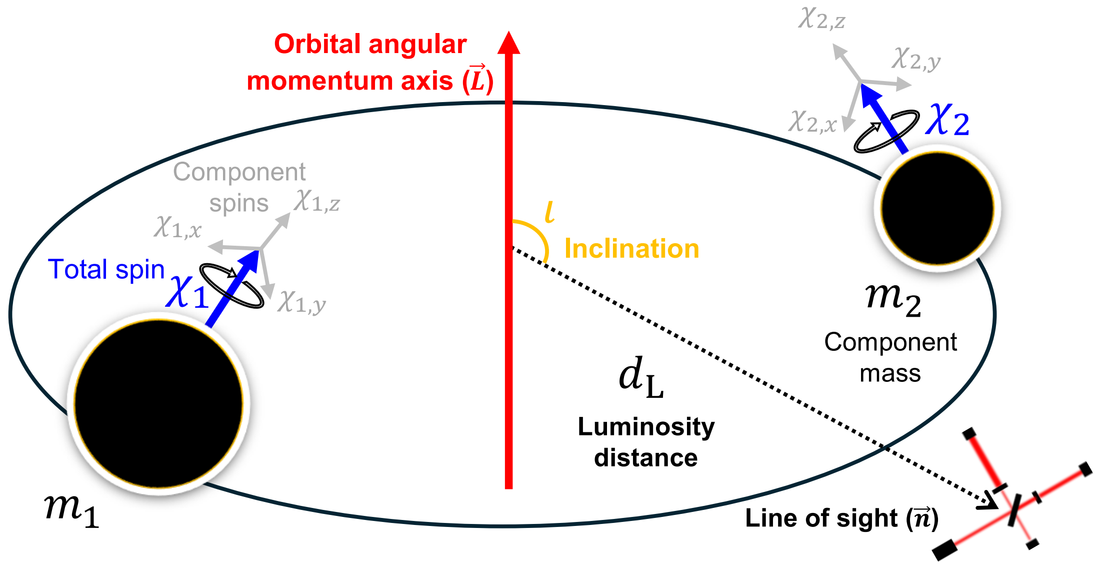

Figure 4.2: Cartoon illustration of a BBH system. The component masses, spins, luminosity distance, and inclination angle are indicated.

The extrinsic parameters, $\boldsymbol{\theta}_{\mathrm{ext}}$, describe the orientation and location of the source relative to the detector. The luminosity distance (d_L) and the inclination angle (ι) are shown in Figure 4.2. The sky position is specified by the right ascension (RA; α) and declination (Dec; δ). Additional extrinsic parameters include the polarisation angle (ψ), which specifies the orientation of the wave polarisation axes, and the coalescence time (t_c) and phase (ϕ_c), which define the temporal and phase reference of the signal. All fifteen BBH parameters are summarised in Table 4.1.

Parameter	Symbol	Unit
Intrinsic		
Component masses	m_1, m_2	$M_\odot$
Component spins	$\chi_{1x}, \chi_{1y}, \chi_{1z}, \chi_{2x}, \chi_{2y}, \chi_{2z}$	—
Extrinsic		
Luminosity distance	d_L	Mpc
Inclination	ι	rad
Right ascension	α	rad
Declination	δ	rad
Polarisation angle	ψ	rad
Coalescence time	t_c	s
Coalescence phase	ϕ_c	rad

Table 4.1: Summary of the 15 parameters of a quasi-circular BBH. The parameters are divided into intrinsic properties, which determine the binary's physical properties and waveform, and extrinsic properties, which describe its orientation and how it is observed.

As defined in Eq. (1.30), the response of a GW interferometer to a CBC signal is given by a linear combination of the two polarisation states (Eq. (1.23)), weighted by the antenna pattern functions (Eq. (1.29)). Since the matched-filtering and parameter estimation are typically performed in the frequency domain, it is convenient to work with the Fourier transform of the strain, $\tilde{h}(f)$. The Fourier-transformed GW strain is given by:

$$\begin{aligned}\tilde{h}(f;\mathbf{\Theta}) = \frac{1}{d_L}\left(F_+(\alpha,\delta,\psi)\frac{1+\cos^2\iota}{2}\,\tilde{h}_+\left(f;\boldsymbol{\theta}_{\text{int}}\right) + F_\times(\alpha,\delta,\psi)\cos\iota\,\tilde{h}_\times\left(f;\boldsymbol{\theta}_{\text{int}}\right)\right) \\ \times \exp\left(2\pi i f t_c + \phi_c\right).\end{aligned} \tag{4.21}$$

Here, the h_+ and $h_\times$ are described by only $\boldsymbol{\theta}_{\text{int}}$. The luminosity distance, d_L, sets the overall amplitude scale ($\propto d_L^{-1}$), while the inclination angle ι determines the relative contributions of the two polarisations.

In principle, an exact GW waveform in Eq. (4.21) can be obtained by solving the full EFE through numerical relativity (NR) simulations. However, such simulations are computationally expensive, requiring substantial computational resources for a single waveform. Hence, the direct use of NR is impractical for parameter estimation, which requires $\mathcal{O}(10^5 - 10^7)$ waveform evaluations. To overcome this limitation, GW data analysis relies on *waveform models* (or *approximants*), which provide fast and accurate representations of the signal.

The simplest class of waveform models is based on post-Newtonian (PN) theory, which provides an analytical expansion of the in-spiral waveform in powers of the orbital velocity [31]. PN-based waveform models, such as the `TaylorF2` approximant, provide an analytical description of the in-spiral signal in the frequency domain by expanding the phase as a truncated PN series [30, 31, 267, 268, 269, 270, 271]. These models are computationally efficient but are limited to the early in-spiral regime and fail to accurately capture the late in-spiral, merger, and ringdown phases. To achieve a more realistic description of the full coalescence, including the strong-field regime, more sophisticated waveform models are required. Such models combine analytical approximations with NR results. Specifically, analytical frameworks such as PN theory accurately describe the early in-spiral, but break down in the strong-field regime. To extend their validity, free coefficients appearing in these models are adjusted so that the resulting waveforms closely match high-accuracy NR simulations across the parameter space.

One widely used framework is the *effective one-body (EOB)* formalism, which maps the relativistic two-body problem onto an EOB system, in which a single effective particle moves in a deformed spacetime that encodes the influence of both masses [272, 273, 274, 275]. Rather than relying on a PN expansion at a given order (e.g., 3.5PN or 4PN), EOB reorganises the series into a form that extends its validity toward the strong-field regime and attaches a ringdown model based on QNMs. The remaining degrees of freedom are then calibrated to NR. However, EOB waveforms are generated by solving coupled ordinary differential equations in the time domain, which can be computationally expensive.

An alternative approach is the *phenomenological (IMRPhenom) waveform* family, which constructs the waveform directly in the frequency domain. These models divide the signal into three regimes: early in-spiral, intermediate, and merger-ringdown, and represent the waveform in terms of its amplitude and phase as:

$$\tilde{h}(f;\boldsymbol{\Theta}) = \mathcal{A}(f;\boldsymbol{\Theta})\, e^{i\Psi(f;\boldsymbol{\Theta})}, \tag{4.22}$$

where $\mathcal{A}(f;\boldsymbol{\Theta})$ and $\Psi(f;\boldsymbol{\Theta})$ denote the amplitude and phase, respectively. Particularly, the phase is modelled using piecewise analytical ansatzes of the form:

$$\Psi(f) = \begin{cases} \Psi_{\mathrm{PN}}(f) + \Psi^{\mathrm{phen}}_{\mathrm{insp}}(f), & f < f_{\mathrm{int}}, \\ \Psi^{\mathrm{phen}}_{\mathrm{int}}(f), & f_{\mathrm{int}} \leq f < f_{\mathrm{MR}}, \\ \Psi^{\mathrm{phen}}_{\mathrm{MR}}(f), & f \geq f_{\mathrm{MR}}, \end{cases} \tag{4.23}$$

where f_{int} and f_{MR} denote the transition frequencies between the in-spiral, intermediate, and merger-ringdown regimes. $\Psi_{\mathrm{PN}}(f)$ corresponds to the PN contribution valid in the early in-spiral, while $\Psi^{\mathrm{phen}}(f)$ represents phenomenological corrections calibrated to hybrid waveforms constructed from analytical approximations and NR. This procedure ensures that the model accurately reproduces both the weak-field and strong-field regimes. Since the models provide explicit analytical expressions for the waveform amplitude and phase in the frequency domain, the IMRPhenom waveform family is computationally efficient and widely used in large-scale analyses.

Lastly, *NR surrogate* models adopt a data-driven approach. Instead of solving the EFE or relying on analytical expansions, these models interpolate a precomputed set of NR waveforms using reduced-basis and empirical interpolation techniques. As a result, NR surrogates achieve near-NR accuracy while being orders of magnitude faster to evaluate. Their main limitation is that they are only valid within the region of parameter space covered by the NR training set, and may not be reliable for extreme mass ratios, high spins, or other configurations not included in the training data.

In practice, multiple waveform model families are employed for parameter estimation, as different models can yield systematically different waveforms even for the same set of source parameters[1]. Therefore, it is standard to analyse each GW event using multiple waveform models to assess systematic uncertainties and ensure the robustness of the inferred results.

1. However, we cannot determine which waveform model performs better in terms of accuracy [276, 277].

4.3.2 Advanced Morphological Features

Waveform models based on the standard 15 BBH-parameter description often rely on simplifying assumptions, such as aligned spins, negligible orbital eccentricity, and dominant quadrupole $(\ell, m) = (2, 2)$ emission. However, in more special astrophysical scenarios, BBHs can exhibit additional physical effects that generate complex morphological features in $\tilde{h}(f)$ (Eq. (4.21)). These features carry crucial information that can help break degeneracies between waveform parameters, enhance the identification of lensing signatures, and improve constraints on cosmological parameters. In the following, we consider three such effects.

Firstly, we consider *spin precession*, which arises when the component spin vectors $\vec{\chi}_1$ and $\vec{\chi}_2$ are misaligned with the orbital angular momentum $\vec{L}$. In such configurations, spin–orbit coupling generates torques that cause the orbital plane itself to precess during the in-spiral. More specifically, the orbital angular momentum and the individual spin vectors precess around the approximately conserved total angular momentum, given by:

$$\vec{J} = \vec{L} + \vec{S}_1 + \vec{S}_2, \tag{4.24}$$

where $\vec{S}_i$ denotes the spin angular momentum of the i-th compact object.

The dominant observable effect of precession is commonly characterised by the effective precession spin parameter χ_p [278], defined as:

$$\chi_p \equiv \frac{1}{A_1 m_1^2} \max\left(A_1 S_{1\perp}, A_2 S_{2\perp}\right), \tag{4.25}$$

where

$$A_1 = 2 + \frac{3q}{2}, \qquad A_2 = 2 + \frac{3}{2q}. \tag{4.26}$$

$S_{i\perp}$ denotes the component of the spin perpendicular to $\vec{L}$. Physically, χ_p quantifies the degree to which spin-induced precession modulates the GW signal. Geometrically, precession causes the orientation of the binary relative to the observer to vary with time. As a result, the inclination angle ι and polarisation angle ψ become time-dependent quantities, i.e., $\iota(t)$ and $\psi(t)$.

As shown in Figure 4.3, when comparing GWs from BBHs with identical parameters except for their spin configurations, the highly precessing system ($\chi_p \simeq 0.85$) exhibits characteristic amplitude and phase modulations in the observed waveform relative to the non-precessing case, particularly during the in-spiral phase. These modulations introduce additional structure into the signal that cannot be reproduced by non-precessing waveforms, thereby helping to break degeneracies among the mass ratio, spin, inclination angle, and luminosity distance.

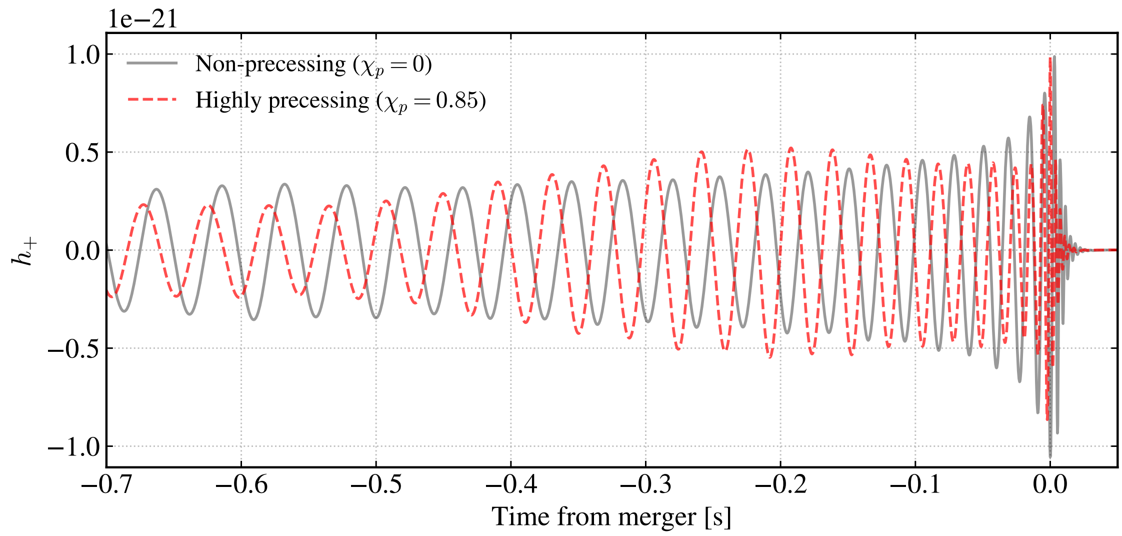

Figure 4.3: Waveform comparison between the time-domain h_+ strains from a non-precessing BBH and a highly precessing BBH ($\chi_p \simeq 0.85$), with all other source parameters kept identical.

Secondly, we consider *orbital eccentricity* (e). Standard CBC waveform templates generally assume quasi-circular orbits because GW radiation efficiently circularises binaries long before they enter the sensitive frequency band of ground-based detectors. However, binaries formed dynamically in dense stellar environments, such as globular clusters or galactic nuclei, may retain measurable residual eccentricity when entering the detector band [279, 280].

For eccentric binaries, the orbital separation and velocity vary periodically throughout the orbit, becoming largest near periastron and smallest near apastron. Consequently, the GW emission is no longer dominated by a single smoothly evolving chirp frequency. Instead, eccentricity introduces additional harmonics of the orbital frequency and modifies the in-spiral rate itself. As shown in Figure 4.4, the quasi-circular BBH exhibits a smooth, featureless amplitude spectrum in the frequency domain, following the characteristic $f^{-7/6}$ scaling during the in-spiral phase. In contrast, orbital eccentricity introduces additional harmonic structure, manifested as significant amplitude modulations in the low-frequency regime (20–100 Hz). As the binary evolves toward merger, GW emission progressively circularises the orbit, suppressing these modulations and causing the two spectra to converge at higher frequencies.

Lastly, we consider *higher-order modes* (HOMs). The dominant contribution to GW emission is typically described by the quadrupolar $(\ell, m) = (2, \pm 2)$ modes. For binaries with nearly equal masses observed close to face-on ($\iota \approx 0°$), these modes alone provide an excellent approximation to the full waveform [281]. However, when the binary possesses a substantially asymmetric mass ratio or is viewed close to edge-on ($\iota \sim 90°$), sub-dominant modes such as $(2, 1)$, $(3, 2)$, $(3, 3)$, and $(4, 4)$ become increasingly important, and the mis-

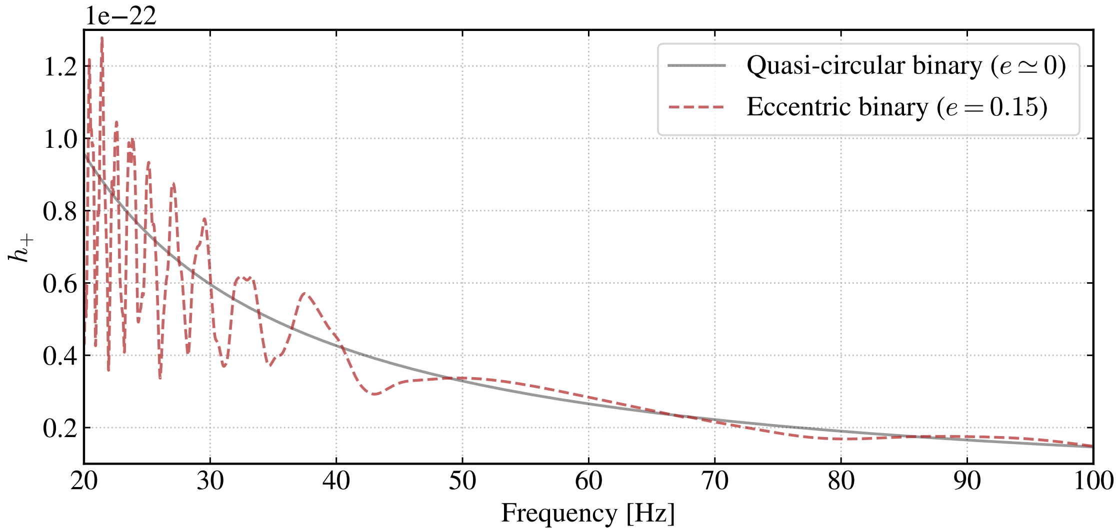

Figure 4.4: Waveform comparison between the frequency-domain h_+ strains from a quasi-circular BBH and a mildly eccentric BBH ($e = 0.15$), with all other source parameters kept identical.

match between the detected signal and waveforms including only the dominant $(2, 2)$ mode can become significant [282, 283, 284, 285]. Figure 4.5 shows the GW waveforms, including all modes, only the dominant $(2, 2)$ mode, and the sub-dominant modes for a BBH with mass ratio $q = 0.3$ and inclination angle $\iota \simeq 70°$.

These HOMs evolve at frequencies approximately proportional to integer multiples of the orbital frequency, i.e., $f_{\ell m} \sim m f_{\rm orb}$, producing additional frequency tracks in the GW spectrogram. The relative amplitudes of the HOMs depend strongly on both the mass ratio and the inclination, meaning that they encode geometric information not accessible from the dominant quadrupole mode alone. Consequently, incorporating HOMs into waveform models substantially improves constraints on the luminosity distance and inclination angle and helps reduce the well-known distance-inclination degeneracy, particularly for massive and asymmetric BBH systems [286, 287].

4.3.3 Marginalisation of Nuisance Parameters

Once the BBH parameters of interest ($\boldsymbol{\theta}$) and a waveform model are determined, the corresponding GW templates ($h_T(\boldsymbol{\theta})$) are generated for matched filtering. These templates allow us to establish the likelihood function, $\mathcal{L}_{\rm gw}$, which acts as the statistical bridge between the observed data and the physical model, as defined in Eq. (1.40).

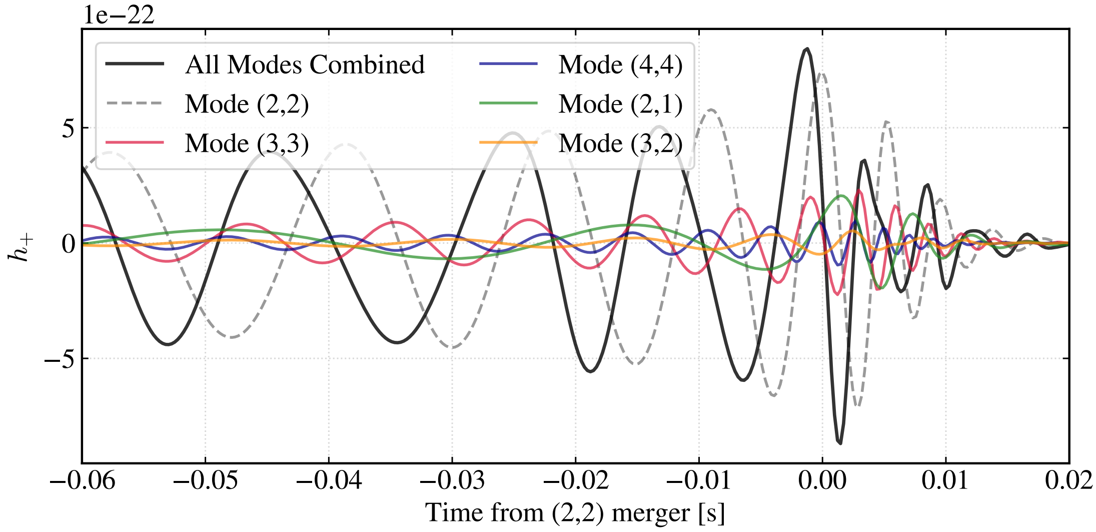

Figure 4.5: Time-domain GW waveforms including all modes, only the dominant $(2, 2)$ mode, and selected sub-dominant modes $(2, 1)$, $(3, 2)$, $(3, 3)$, and $(4, 4)$ for a BBH with $q = 0.3$ and $\iota \simeq 70°$. The x-axis denotes the time relative to the merger time of the $(2, 2)$ mode. The mismatch between the dominant $(2, 2)$-only waveform and the full waveform is clearly visible.

In practice, instead of sampling the full multidimensional parameter space, specific *nuisance* parameters are often marginalised to improve computational efficiency and alleviate degeneracies. This process involves integrating the likelihood over these parameters (see Eq. (4.3)). In BBH analyses, parameters such as the coalescence time, phase, luminosity distance, and instrumental calibration uncertainties are commonly treated in this manner. This strategy offers several methodological advantages:

- **Coalescence Time (t_c):** This parameter is marginalised using Fast Fourier Transforms (FFTs). This operation allows the likelihood to be evaluated across all possible time shifts simultaneously in the frequency domain, eliminating the need for the stochastic sampler to explore the highly correlated arrival-time parameter space and thereby significantly accelerating convergence.
- **Coalescence Phase (ϕ_c):** When the waveform model is restricted to the dominant $(\ell, |m|) = (2, 2)$ mode, the coalescence phase enters the frequency-domain waveform as a global phase factor $\propto e^{i2\phi_c}$, depending on the phase convention. This can remove the need to sample over the coalescence phase, reducing the dimensionality of the sampling problem.

- **Luminosity Distance (d_L):** As the distance enters as an overall amplitude scaling ($\propto d_L^{-1}$), it can be marginalised by integrating over the amplitude with an assumed distance prior. In practice, this marginalisation is implemented with precomputed lookup tables, allowing the sampler to efficiently recover the posterior on d_L in a post-processing step. This approach helps isolate the geometric degeneracy between distance and inclination, enabling more efficient exploration of the remaining parameters.
- **Calibration Uncertainties:** Detector systematics, including frequency-dependent amplitude and phase response errors, are incorporated by marginalising over calibration uncertainties assuming Gaussian-distributed deviations. This procedure accounts for instrumental systematics without introducing additional sampling dimensions, ensuring that inferred posteriors remain robust against calibration errors.

In the context of GW lensing, however, the standard marginalisation techniques described above must be applied with caution, as they can inadvertently suppress or bias key lensing signatures. In particular, marginalising over d_L is problematic due to the magnification–distance degeneracy. Since the observed amplitude scales as $\sqrt{\mu}/d_L$, treating d_L as a nuisance parameter can bias the inferred magnification μ, thereby affecting both lens and source parameter estimation.

Furthermore, lensing introduces additional phase modifications, manifesting as discrete Morse phase shifts ($n_j\pi/2$) for the j^{th} image in the strong-lensing regime, and as frequency-dependent phase shifts in the microlensing regime. Analytical marginalisation over the coalescence phase for each lensed signal can remove the relative phase information between strongly lensed images and obscure the frequency-dependent phase structure introduced by microlensing, making it difficult to identify image morphology and directly invalidating the lensing hypothesis test.

4.3.4 Parameter Estimation of a BBH Event

We now present a representative parameter-estimation analysis for a simulated nearly equal-mass BBH event using the methodology described in the previous sections. The injected source parameters used to generate the GW signal are summarised in Table 4.2, together with the priors adopted in the Bayesian inference analysis. The simulated signals are generated using the IMRPhenomXPHM-SpinTaylor [288, 289] waveform model with a sampling frequency of 4096 Hz.

We inject the signal into a three-detector network consisting of aLIGO Hanford, aLIGO Livingston, and AdV interferometers. To emulate realistic observing conditions, we assume the projected O4a PSDs for all detectors. The subsequent parameter-estimation analysis is performed using the `bilby` pipeline [240] with the nested sampler `dynesty` [263]. The likelihood is evaluated over the frequency range from $f_{\text{min}} = 20$ Hz to the Nyquist frequency $f_{\text{max}} = 2048$ Hz, corresponding to half of the adopted sampling frequency.

Table 4.2: Injected source parameters and prior ranges. The notation U[min, max] indicates a uniform prior bounded by the provided minimum and maximum values.

Parameter	Injected value	Prior
m_1 $[M_\odot]$	45.04	U$[5, 100]$
m_2 $[M_\odot]$	40.39	U$[5, 100]$
a_1	0.12	U$[0, 0.99]$
a_2	0.36	U$[0, 0.99]$
θ_1	2.82	Sine
θ_2	1.47	Sine
ι	2.50	Sine
ψ	1.00	U$[0, 2\pi]$
ϕ_{12}	0.83	U$[0, 2\pi]$
ϕ_{jl}	5.16	U$[0, 2\pi]$
RA	5.78	U$[0, 2\pi]$
DEC	0.88	Cosine
d_L [Mpc]	4778.85	UniformSourceFrame
t_c [s]	1126259642.41	-
ϕ_c	1.56	U$[0, 2\pi]$

We adopt the default prior choices in `bilby` commonly used for a precessing BBH parameter estimation (see Table 4.2). Particularly, uniform priors are imposed on the component masses, while the dimensionless spin magnitudes are sampled uniformly between 0 and 0.99, with the upper limit chosen to avoid numerical issues near extremal spins. The spin orientations are assumed to be isotropically distributed. For the luminosity distance, we adopt a uniform prior in the source-frame comoving volume, corresponding to a source population distributed uniformly in comoving volume, with an additional $(1+z)^{-1}$ factor accounting for cosmological time dilation between the source and detector frames.

The recovered posterior distributions for a representative subset of the intrinsic and extrinsic parameters are shown in Figure 4.6. For clarity, only the parameters most directly relevant to the binary masses, spins, orientation, sky localisation, and luminosity distance are displayed, while the full parameter estimation is performed in the complete 15-dimensional parameter space. The injected parameters are successfully recovered within the posterior credible regions for a network SNR of $\rho_{\rm net} \simeq 11.8$.

As shown in Figure 4.6, the recovered posterior distributions exhibit several characteristic parameter correlations commonly encountered in GW parameter estimation. Although the component masses (m_1, m_2) are both relatively well-constrained, their posteriors exhibit a noticeable correlation structure. This correlation arises because the in-spiral phase evolution is primarily governed by the chirp mass ($\mathcal{M}_c$), which is measured with particularly high precision from the signal's phase evolution. Consequently, different combinations of m_1 and m_2 that yield similar $\mathcal{M}_c$ values can produce waveforms with very similar in-spiral behaviour, leading to elongated correlated posterior distributions in the (m_1, m_2) plane.

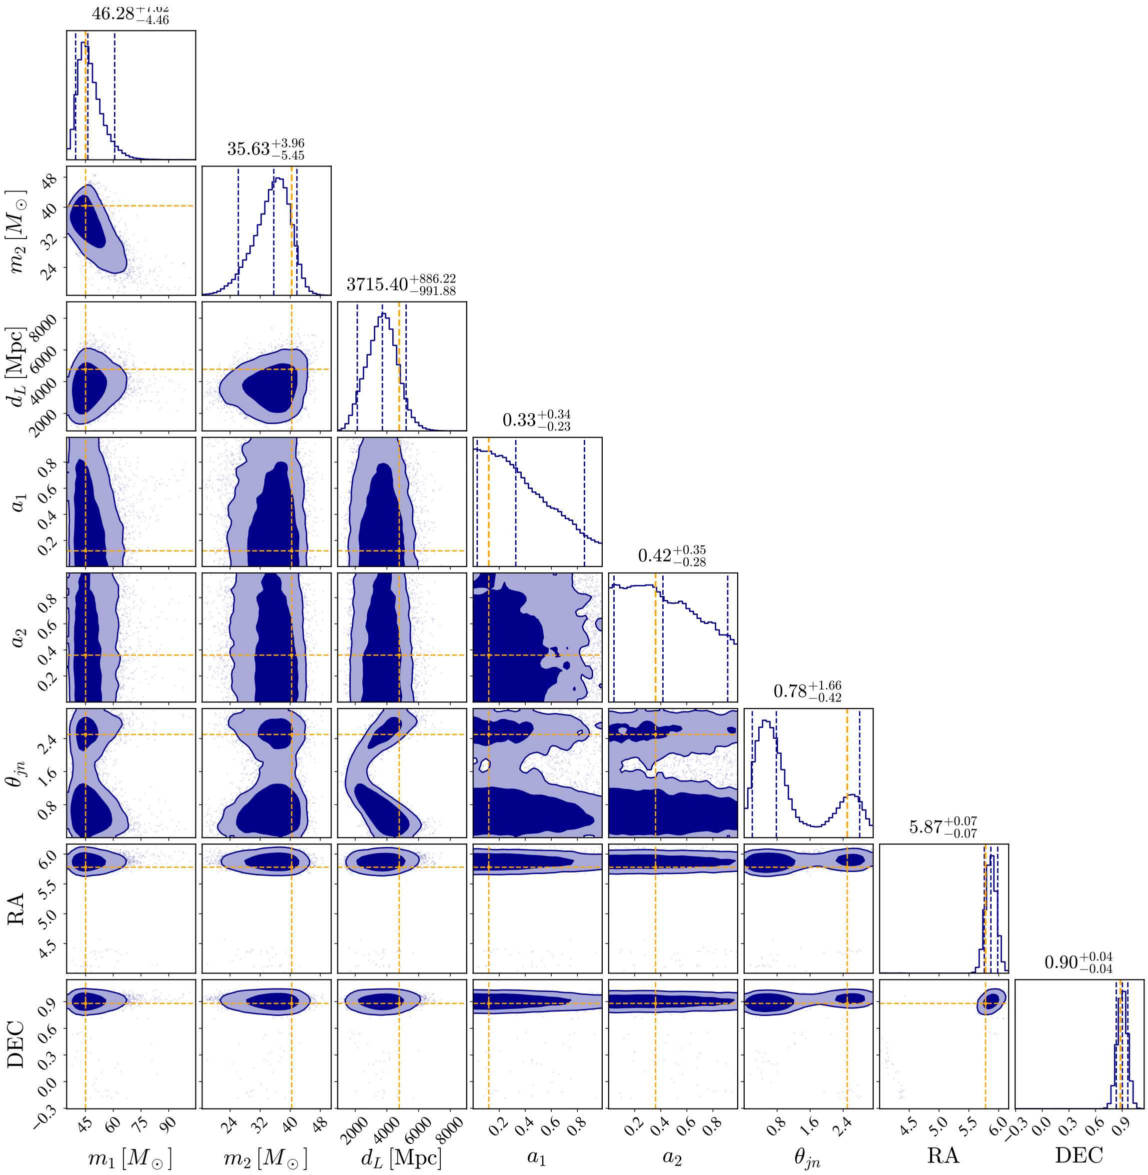

Figure 4.6: Corner plot showing the posterior distributions of a representative subset of BBH parameters inferred from the simulated event. Only the most relevant intrinsic and extrinsic parameters are shown for clarity. Orange dashed lines and markers indicate the injected values, while the navy dashed lines denote the 90% C.I..

The spin magnitudes remain comparatively weakly constrained, primarily because the injected system contains only moderate spins and limited spin-precession signatures. For such configurations, the spin-dependent waveform modulations contribute only subdominantly to the observed signal, making them more difficult to distinguish from changes induced by other parameters.

Correlations are also observed between the luminosity distance (d_{L}) and the inclination angle (ι). This degeneracy arises because binaries viewed closer to face-on at larger distances can produce strain amplitudes similar to those of more inclined binaries located nearer to the observer. Although HOMs, spin precession, and polarisation information from multiple detectors can partially break this degeneracy [209, 286, 287], it generally remains significant for BBH events with moderate network SNRs.

4.4 Hierarchical Bayesian Analysis

While standard parameter estimation described in Section 4.3 focuses on inferring the source properties of individual events based on observed data, hierarchical Bayesian analysis (HBA) extends this framework by modelling the population distribution of source parameters and combining information from multiple detections. Rather than treating each event independently, the HBA assumes that the event-level parameters are realisations drawn from a common population model, characterised by a set of *hyperparameters*, $\boldsymbol{\Lambda}$. These hyperparameters govern the shape and structure of the population distribution and are inferred jointly using the full set of observed events. By inferring $\boldsymbol{\Lambda}$, one can constrain population-level properties of GW sources. In addition, the same framework can be extended to probe cosmological parameters, including the expansion rate of the Universe, the matter density, and the dark energy equation of state.

4.4.1 Basic Framework

To infer the population properties of GW sources, described by source parameters ($\boldsymbol{\theta_s}$), is used to determine the posterior distribution of hyperparameters $p(\boldsymbol{\Lambda}|\{d\})$, given a set of observational data $\{d\} = \{d_1, d_2, \ldots, d_{N_{\mathrm{det}}}\}$ from N_{det} detected events. By Bayes' theorem,

$$p(\boldsymbol{\Lambda}|\{d\}) \propto \mathcal{L}(\{d\}|\boldsymbol{\Lambda})\pi(\boldsymbol{\Lambda}), \tag{4.27}$$

where $\mathcal{L}(\{d\}|\boldsymbol{\Lambda})$ denotes the likelihood of the data and $\pi(\boldsymbol{\Lambda})$ represents the hyper-prior encoding prior knowledge of the population.

Under the assumption that the occurrence of GW events follows an inhomogeneous Poisson point process, let $N(\boldsymbol{\Lambda})$ denote the expected total number of GW events, including both detected and undetected sources, for a given set of hyperparameters $\boldsymbol{\Lambda}$. The expected number of detected events among these N sources is then given by [290, 291]:

$$N_{\mathrm{exp}}(\boldsymbol{\Lambda}) = N(\boldsymbol{\Lambda}) \int p_{\mathrm{pop}}(\boldsymbol{\theta_s}|\boldsymbol{\Lambda}) P_{\mathrm{det}}(\boldsymbol{\theta_s}) d\boldsymbol{\theta_s}, \tag{4.28}$$

where $p_{\mathrm{pop}}(\boldsymbol{\theta_s}|\boldsymbol{\Lambda})$ denotes the population distribution of the source parameters, and $P_{\mathrm{det}}(\boldsymbol{\theta_s})$ is the detection indicator, which takes the value 1 if a source with parameters $\boldsymbol{\theta_s}$ exceeds the chosen detection threshold and 0 otherwise.

Consequently, the probability of detecting N_{det} events is given by a Poisson distribution with a mean of N_{exp}. The likelihood for the observed catalogue, accounting for selection effects arising from the finite detector sensitivity, can then be written as [291]:

$$\mathcal{L}(\{d\}|\boldsymbol{\Lambda}, N) \propto N(\boldsymbol{\Lambda})^{N_{\mathrm{det}}} e^{-N_{\mathrm{exp}}(\boldsymbol{\Lambda})} \prod_{i=1}^{N_{\mathrm{det}}} \int \mathcal{L}(d_i|\boldsymbol{\theta}_{s,i}) p_{\mathrm{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda}) d\boldsymbol{\theta}_{s,i}. \tag{4.29}$$

It is convenient to define the detection efficiency [291]:

$$\eta(\boldsymbol{\Lambda}) \equiv \frac{N_{\mathrm{exp}}}{N} = \int p_{\mathrm{pop}}(\boldsymbol{\theta_s}|\boldsymbol{\Lambda}) P_{\mathrm{det}}(\boldsymbol{\theta_s}) d\boldsymbol{\theta_s}. \tag{4.30}$$

Finally, N can be treated as a nuisance parameter and marginalised over. Adopting a scale-invariant prior, $\pi(N) \propto 1/N$, the marginal likelihood becomes [290]:

$$\begin{aligned}
\mathcal{L}(d \mid \boldsymbol{\Lambda}) &\propto \int dN\, (N)^{N_{\mathrm{det}}} e^{-N_{\mathrm{exp}}} \pi(N) \prod_{i=1}^{N_{\mathrm{det}}} \int \mathcal{L}(d_i|\boldsymbol{\theta}_{s,i}) p_{\mathrm{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda}) d\boldsymbol{\theta}_{s,i} \\
&= \int dN\, (N)^{N_{\mathrm{det}}-1} e^{-\eta N} \prod_{i=1}^{N_{\mathrm{det}}} \int \mathcal{L}(d_i|\boldsymbol{\theta}_{s,i}) p_{\mathrm{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda}) d\boldsymbol{\theta}_{s,i} \\
&= \Gamma(N_{\mathrm{det}}) \eta(\boldsymbol{\Lambda})^{-N_{\mathrm{det}}} \prod_{i=1}^{N_{\mathrm{det}}} \int \mathcal{L}(d_i|\boldsymbol{\theta}_{s,i}) p_{\mathrm{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda}) d\boldsymbol{\theta}_{s,i},
\end{aligned} \tag{4.31}$$

where $\Gamma(N_{\mathrm{det}})$ is independent of $\boldsymbol{\Lambda}$ and can be absorbed into the normalisation constant.

The detection efficiency $\eta(\boldsymbol{\Lambda})$ in Eq. (4.30) corresponds to an expectation value over the population distribution. Rather than directly sampling from p_{pop}, the efficiency can be efficiently evaluated using injection campaigns [292]. A large set of simulated mock GW signals is generated from a fiducial injection distribution π_{inj}, and passed through the detection pipeline to determine whether they are detectable. The efficiency can then be approximated via importance sampling as:

$$\begin{aligned}
\eta(\boldsymbol{\Lambda}) &= \int \frac{p_{\mathrm{pop}}(\boldsymbol{\theta}_s|\boldsymbol{\Lambda})}{\pi_{\mathrm{inj}}(\boldsymbol{\theta}_s)} \pi_{\mathrm{inj}}(\boldsymbol{\theta}_s) P_{\mathrm{det}}(\boldsymbol{\theta}_s) d\boldsymbol{\theta}_s \\
&\simeq \frac{1}{N_{\mathrm{inj}}} \sum_{i}^{N_{\mathrm{inj,det}}} \frac{p_{\mathrm{pop}}(\boldsymbol{\theta_i}|\boldsymbol{\Lambda})}{\pi_{\mathrm{inj}}(\boldsymbol{\theta}_{s,i})},
\end{aligned} \tag{4.32}$$

where N_{inj} is the total number of injected signals, and $N_{\text{inj,det}}$ is the number of detected injections. The factor $p_{\text{pop}}/\pi_{\text{inj}}$ accounts for the mismatch between the injection distribution and the target population model, ensuring that the importance sampling recovers the expectation value under p_{pop}.

The reliability of the estimate depends on the effective number of samples, which is given by:

$$N_{\text{eff}} = \frac{\left(\sum_i w_i\right)^2}{\sum_i w_i^2}, \tag{4.33}$$

where $w_i = \frac{p_{\text{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda})}{\pi_{\text{inj}}(\boldsymbol{\theta}_{s,i})}$. A small N_{eff} indicates that the estimate is dominated by a few high-weight samples, leading to large statistical uncertainty. To ensure a robust estimation, it is typically required that $N_{\text{eff}} > 4N_{\text{det}}$ [292].

To evaluate the total likelihood in Eq. (4.31), the remaining step is to compute the individual likelihood for d_i. In practice, parameter estimation for GW signals provides discrete posterior samples rather than analytic functions. Let $\mathcal{S}_i$ denote the number of posterior samples for the i^{th} event, which are typically obtained using a reference prior, $\pi_{\text{ref}}(\boldsymbol{\theta_s})$. The integral in Eq. (4.31) can then be approximated via Monte Carlo reweighting as [291]:

$$\int \mathcal{L}(d_i|\boldsymbol{\theta}_{s,i}) p_{\text{pop}}(\boldsymbol{\theta}_{s,i}|\boldsymbol{\Lambda}) d\boldsymbol{\theta}_{s,i} \approx \frac{1}{\mathcal{S}_i} \sum_{k=1}^{\mathcal{S}_i} \frac{p_{\text{pop}}(\boldsymbol{\theta}^k_{s,i}|\boldsymbol{\Lambda})}{\pi_{\text{ref}}(\boldsymbol{\theta}^k_{s,i})}, \tag{4.34}$$

where $\boldsymbol{\theta}^k_{s,i}$ denotes the k^{th} posterior samples of the i^{th} event. Finally, combining the above ingredients (Eqs. (4.31), (4.32), and (4.34)), the posterior distribution in Eq. (4.27) can be efficiently evaluated.

4.4.2 Parametric Models of BBH Populations and Cosmology

To fully specify Eq. (4.27), it is necessary to adopt parametric models for the GW source population, encoded in p_{pop} introduced in Section 4.4.1. In this thesis, we focus on the distributions of BBH component masses and their merger rate as a function of redshift. Recent analyses by the LVK Collaboration have employed a variety of parametric models to describe these distributions.

For the component masses, we introduce a fiducial model used in the GWTC-3 analysis [293], the Power Law+Peak model [290, 294]. This phenomenological model combines a truncated power-law component with a Gaussian peak, allowing it to capture the overall decrease in event rate with mass and a localised concentration of events around a preferred mass scale. Within the mass range $[m_{\text{min}}, m_{\text{max}}]$, the source-frame primary mass

(m_1) distribution is given by:

$$
\begin{aligned}
&p_{\text{pop}}\left(m_1 \mid \alpha, \lambda_{\mathrm{g}}, \mu_{\mathrm{g}}, \sigma_{\mathrm{g}}, m_{\min}, \delta_m\right) \\
&= \left[\left(1-\lambda_{\mathrm{g}}\right) m_1^{-\alpha} + \frac{\lambda_{\mathrm{g}}}{\sqrt{2\pi}\,\sigma_g} \exp\left(-\frac{(m_1-\mu_g)^2}{2\sigma_g^2}\right)\right] S(m_1|m_{\min}, \delta_m),
\end{aligned}
\tag{4.35}
$$

where α is the power-law index, and λ_g, μ_g, and σ_g denote the mixture fraction, mean, and standard deviation of the Gaussian component, respectively. $S(m_1|m_{\min}, \delta_m)$ represents a smoothing function introducing a gradual taper at the low-mass end, which is given by:

$$
\begin{aligned}
&S\left(m \mid m_{\min}, \delta_m\right) \\
&= \begin{cases}
0 & (m < m_{\min}), \\
\left[\exp\left(\frac{\delta_m}{m-m_{\min}} + \frac{\delta_m}{m-m_{\min}-\delta_m}\right)+1\right]^{-1} & (m_{\min} \le m < m_{\min}+\delta_m), \\
1 & (m \ge m_{\min}+\delta_m),
\end{cases}
\end{aligned}
\tag{4.36}
$$

where δ_m sets the width of the low-mass taper applied near $m_{\min}$. The function avoids an unphysical sharp cutoff at $m_{\min}$ and improves numerical stability when evaluating selection effects and likelihoods. For the secondary mass m_2, it is convenient to model the mass ratio $q_m \equiv m_2/m_1$ instead of directly sampling m_2. The q_m distribution is drawn from a truncated power-law distribution, with the low-end smoothing function defined in Eq. (4.36), which is given by:

$$
p_{\text{pop}}\left(q_m \mid m_{\min}, m_1, \beta\right) = (q_m)^{\beta} S(m_2|m_{\min}, \delta_m) \quad \text{when} \;\; m_{\min} < m_2 < m_1. \tag{4.37}
$$

where β is the power-law index.

In addition to the component mass distribution, the population model must specify the merger rate density as a function of redshift, $\mathcal{R}(z_s)$, which describes the redshift evolution of BBH mergers in the source frame. This is because it determines how GW sources are distributed over cosmic time and directly enters the calculation of the expected number of detections through the selection function.

A commonly used parametrisation is a simple power-law model:

$$
\mathcal{R}(z) = \mathcal{R}_0(1+z)^{\kappa}, \tag{4.38}
$$

where $\mathcal{R}_0$ denotes the merger rate at $z = 0$ and κ controls the redshift evolution. An alternative one is the Madau–Dickinson form [295, 296], which is motivated by the cosmic star formation history and its connection to compact binary formation:

$$
\mathcal{R}(z) = \mathcal{R}_0(1+z)^{\gamma} \frac{1+(1+z_p)^{-(\gamma+\kappa)}}{1+\left(\frac{1+z}{1+z_p}\right)^{\gamma+\kappa}}, \tag{4.39}
$$

where z_p is the redshift at which the merger rate peaks, and γ and κ denote the power-law indices of the redshift distribution before and after z_p, respectively.

More recent population analyses, such as the GWTC-4 study [237], adopt the Broken Power Law + Two Peaks model [297] as a fiducial model, which provides a better fit to the data than the Power Law + Peak model, reflecting its ability to capture the increasingly complex structure revealed by larger GW catalogues. In addition, weakly parametric approaches based on spline representations (e.g., B-splines) have also been explored. In this framework, the population distribution is modelled as a separable joint distribution, with each parameter (e.g., mass or redshift) described independently using one-dimensional basis splines. This approach allows the data to flexibly determine the shape of the distributions with reduced reliance on specific functional forms.

The hierarchical framework for the BBH population can be extended to incorporate cosmological parameters, which affect the inference through the conversion between detector-frame observables and source-frame quantities. In particular, the luminosity distance inferred from GW observations depends on z through the assumed cosmological model, while the conversion between detector-frame and source-frame masses also depends on z. Consequently, cosmological parameters enter the population inference, which can be incorporated by extending the population model as follows:

$$p_{\mathrm{pop}}(\boldsymbol{\theta}_s \mid \boldsymbol{\Lambda}); \rightarrow; p_{\mathrm{pop}}(\boldsymbol{\theta}'_s, z \mid \boldsymbol{\Lambda}_{\mathrm{BBH}}, \boldsymbol{\Lambda}_{\mathbf{cosmo}}), \tag{4.40}$$

where $\boldsymbol{\theta}'_s$ indicates the source parameters except for the luminosity distance, while $\boldsymbol{\Lambda}_{\mathrm{BBH}}$ and $\boldsymbol{\Lambda}_{\mathrm{cosmo}}$ denote the sets of BBH-population hyperparameters and cosmological parameters, respectively. In this formulation, the inference jointly constrains both the population properties and cosmology.

As previously shown in Figure 3.3, if independent redshift (z) information is available, one can directly compare the inferred luminosity distance with z, thereby enabling constraints on H_0.

Part II

Searching for Lensed Gravitational-Wave Signals

Chapter 5

Basic Lensing Search Techniques

As introduced in Section 2.4, GWs can experience a variety of lensing effects depending on the mass scale of the intervening lens. For relatively small astrophysical objects, GWs undergo microlensing in the wave-optics regime or phenomenological lensing under the geometrical-optics approximation. In contrast, massive structures such as galaxies or galaxy clusters can cause strong lensing, yielding clearly separated multiple images, or weak lensing, which manifests as a single but distorted signal.

Detecting these lensed GW signals requires dedicated search strategies beyond standard matched-filter analyses. In conventional GW searches, one correlates detector data with theoretically predicted unlensed waveforms (templates) to identify candidate events, as introduced in Section 1.3. For lensed signals, however, the templates must explicitly incorporate the modifications induced by the lens. Microlensing and phenomenological lensing produce single, morphologically distorted GW signals (see Section 2.4.2). Although multiple images are formally generated, the associated time delays are much shorter than the signal duration, causing the signals to overlap and appear as a single distorted waveform. Consequently, the waveform morphology is modified, and detection relies on matched-filtering-based Bayesian inference that employs lensed waveform templates explicitly accounting for the frequency-dependent amplification factor in Eq. (2.71) (e.g., [166, 298]).

Strong lensing, on the other hand, preserves the intrinsic waveform morphology. In the strong lensing case, the detection problem, therefore, reduces to identifying multiple events that share identical source parameters (e.g., masses, spins, sky location, etc.) and are statistically consistent with having originated from a common astrophysical source rather than from random coincidences[1] (e.g., [299, 300, 301, 302, 303, 304]). In addition, when HOMs in the GW signal are sufficiently strong, a single Type II image can provide a distinctive signature through its characteristic phase shift, providing a distinctive signature even without multiple detected images (e.g, [157, 305, 306, 307]).

1. Weak lensing produces signals whose morphology is indistinguishable from unlensed waveforms. The magnification is also subtle, making the signals effectively degenerate with unlensed ones in standard searches. As a result, they are challenging to detect with the current ground-based detectors.

In this chapter, we introduce the basic techniques employed to identify lensed GW signals in ground-based detector data. We consider three complementary search strategies: (1) a wave-optics search for single distorted signals that incorporates explicit lens models, targeting microlensed events; (2) a geometrical-optics, lens-model-independent search for single distorted signals, suitable for phenomenological lensing and type-II image searches; and (3) a multiple-image search for strong lensing, based on simultaneous Bayesian parameter estimation across candidate images. Finally, we briefly discuss more realistic astrophysical scenarios in which strongly lensed signals are further perturbed by small-scale structures embedded within the primary lens.

5.1 Single Distorted Lensed Signals

Lensing signatures can be identified from a single GW signal. The core search strategy is analogous to that used in standard unlensed GW analyses: instead of employing unlensed waveform templates, we use lensed waveforms that incorporate lensing effects. Specifically, we construct a lensing hypothesis and evaluate its likelihood within the Bayesian framework given by Eq. (4.1). We then compute the Bayes factor for the lensing hypothesis ($\mathcal{B}_{\mathrm{L}}$) and compare it to the corresponding Bayes factor under the unlensed hypothesis ($\mathcal{B}_{\mathrm{U}}$), yielding the relative Bayes factor, $\mathcal{B}^{\mathrm{L}}_{\mathrm{U}}$. If $\mathcal{B}^{\mathrm{L}}_{\mathrm{U}}$ is sufficiently high, the data favour the lensing hypothesis. This procedure corresponds to the Bayesian hypothesis testing framework described in Section 4.1.2.

5.1.1 Wave-optics Regime

Representative examples of single distorted lensed signals arise in the microlensing regime, corresponding to the wave-optics limit in which the GW wavelength is comparable to the lens scale. In this regime, the lens induces a frequency-dependent modulation of the waveform amplitude and phase. As a result, the microlensed waveform, h_{ML}, is obtained by multiplying the unlensed strain, h_{U}, by the complex amplification factor $F(f)$ given in Eq. (2.71).

As discussed in Section 2.4.2, the PM lens model provides a well-motivated description for microlensing by compact objects. In general, microlensing can also arise from more complex configurations involving multiple compact lenses. However, in this section, we restrict our discussion to the isolated PM lens case, while multiple-lens configurations are discussed separately in Section 5.2.2. The corresponding amplification factor for an isolated PM lens is given by Eq. (2.73) and is fully characterised by two lens parameters: redshifted lens mass (M_L^z) and dimensionless impact parameter (y). Within the Bayesian framework, the likelihood under the microlensing hypothesis is therefore:

$$\ln \mathcal{L}(d|\boldsymbol{\theta_s}, M_L^z, y) = -\frac{1}{2}\langle d - h_{\mathrm{ML}}(\boldsymbol{\theta_s}, M_L^z, y) | d - h_{\mathrm{ML}}(\boldsymbol{\theta_s}, M_L^z, y)\rangle, \tag{5.1}$$

	Microlensing effects		
	Strong	Moderate	Weak
ρ_{net}	11.56	11.87	11.37
$M_L^z\,[M_\odot]$	14110.08	1421.89	4187.91
y	0.14	0.25	1.92
$\log_{10}\mathcal{B}_{\text{U}}^{\text{ML}}$	6.23	3.51	0.26

Table 5.1: Network SNR, injected lens PM parameters, and $\log_{10}$ Bayes factors for the three mock microlensed signals.

where d denotes the detector data, $\boldsymbol{\theta_s}$ represents the GW source parameters.

To assess parameter recovery, we simulate three representative microlensed signals and inject them into a three-detector network consisting of aLIGO Livingston, aLIGO Hanford, and AdV, assuming projected O4a PSDs for each detector and Gaussian noise[2]. The three cases correspond to (i) strong microlensing effects (high M_L^z and low y), (ii) moderate effects (low M_L^z and low y), and (iii) weak effects (intermediate M_L^z and high y). All three injections are constructed to have similar network SNR, $\rho_{\text{net}} \simeq 11.5$. Each dataset is analysed under the microlensing hypothesis using `bilby` [240] with the `dynesty` [263] sampler. For comparison, we additionally perform parameter estimation under the unlensed hypothesis and compute the corresponding Bayes factor, $\mathcal{B}_{\text{U}}^{\text{ML}}$. Table 5.1 summarises the injected lensing parameters, network SNRs, and resulting Bayes factors for the three mock signals.

As shown in Figure 5.1, the posteriors of M_L^z and y are well constrained for the strong and moderate microlensing cases, exhibiting clear peaks consistent with the injected values. The posterior structure indicates that, even at moderate network SNRs, the waveform distortions induced by microlensing are sufficiently prominent to be distinguished from unlensed signals. Consistently, the corresponding Bayes factors, $\mathcal{B}_{\text{U}}^{\text{ML}}$, strongly favour the microlensing hypothesis over the unlensed hypothesis. In addition, although the moderate microlensing case involves a smaller lens mass, the low impact parameter still produces detectable effects, allowing recovery of the injected lens parameters. These results suggest that small impact parameters play a particularly important role in determining the observability of microlensing signatures.

In contrast, the weak microlensing case yields broad and weakly constrained posteriors without a clear peak near the injected values. Although the lens mass itself is not especially small in this case, the larger impact parameter suppresses the observable lensing features, making the microlensing-induced waveform distortions difficult to distinguish from statistical fluctuations or intrinsic parameter degeneracies at the given network SNR. Consistently, the corresponding Bayes factor is also substantially reduced.

2. The PSDs used in the simulations are publicly available from https://dcc.ligo.org/T2200043-v3/public.

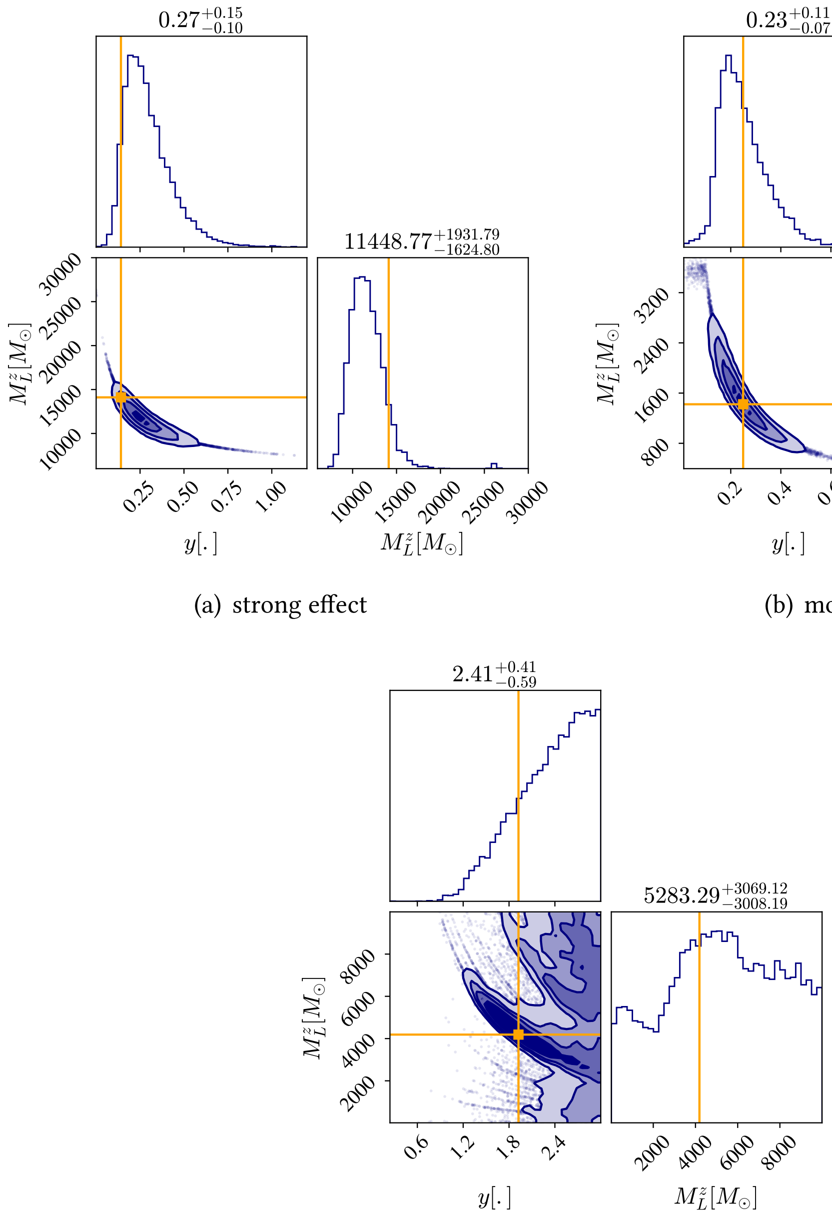

(a) strong effect

(b) moderate effect

(c) weak effect

Figure 5.1: Corner plots showing the posterior distributions of M_L^z and y for the (a) strong, (b) moderate, and (c) weak microlensing cases.

However, a Bayes factor alone is not sufficient to claim a microlensing detection. High values of $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}}$ may also arise from statistical fluctuations in the noise or from non-lensing effects. As discussed in Section 1.3.2, to assess the statistical significance, one needs to construct a background distribution using a large number of unlensed injections, each analysed under both the microlensing and unlensed hypotheses. The resulting distribution of $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}}$ under the unlensed background allows us to quantify the significance of the observed Bayes factor. The false-positive probability (FPP) of the observed Bayes factor, $\mathcal{B}_{\mathrm{obs}}$, is defined as:

$$\mathrm{FPP} = \frac{\text{number of background events with } \mathcal{B}_{\mathrm{U}}^{\mathrm{ML}} > \mathcal{B}_{\mathrm{obs}}}{\text{total number of background events}}. \tag{5.2}$$

If the FPP is sufficiently small, a microlensing detection may be claimed.

However, when multiple independent events are analysed, the probability of obtaining a large Bayes factor from noise or statistical fluctuations increases. This is commonly referred to as the *look-elsewhere effect*, for which the FPP must be corrected by accounting for the associated *trial factor* [308, 309]. Therefore, the significance inferred from a single event must be adjusted to account for the total number of independent events included in the search. In this context, each analysed GW event constitutes an independent trial. The relevant quantity is then the global FPP, defined as the probability of obtaining at least one event with $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}} > \mathcal{B}_{\mathrm{obs}}$ among all analysed events. Assuming statistical independence, the global FPP (equivalently, the global p-value) is given by:

$$\mathrm{FPP}_{\mathrm{global}} = 1 - (1 - \mathrm{FPP})^{N_{\mathrm{event}}}, \tag{5.3}$$

where N_{event} denotes the total number of independent events.

5.1.2 Geometrical-optics Regime

Similar to microlensing, phenomenologically lensed signals are characterised by a single distorted waveform arising from the coherent superposition of multiple lensed images [166]. In this framework, each image is described within the geometric-optics approximation, providing a lens-model-agnostic description. The lensed signal is thus modelled as a coherent sum of K image components rather than as discretely resolvable events, where each component shares the same intrinsic waveform as the unlensed signal.

As noted in Section 2.4, the amplification factor $F(f)$ for phenomenological lensing is basically the same as Eq. (2.74). However, since the main parameters to be estimated are the relative magnification μ_{rel}, the time delay Δt, and the total number of images K, it is more convenient to rewrite the lensed waveform as:

$$\begin{aligned} \tilde{h}_{\mathrm{Phenom}}(f) &= F(f)\tilde{h}_{\mathrm{U}}(f) \\ & \sqrt{\mu_1}\left[1 + \sum_{j=2}^{K} \sqrt{\mu_{\mathrm{rel},j}} \exp\left(2\pi i f \Delta t_j - i\pi n_j\right)\right] \tilde{h}_{\mathrm{U}}(f), \end{aligned} \tag{5.4}$$

where μ_1 denotes the magnification factor of the first-arriving signal, while $\mu_{\mathrm{rel},j}$ and Δt_j represent the magnification and time delay of the j^{th} image relative to the first, respectively.

Following the hypothesis-testing and background-construction framework established for the microlensing analysis in Section 5.1.1, the phenomenological analysis yields a Bayes factor ($\mathcal{B}_{\mathrm{U}}^{\mathrm{Phenom}}$), along with posterior distributions for the lensing observables. These results are subsequently evaluated using the global FPP defined in Eq. (5.3).

In the phenomenological lensing regime, the time delays between superposed images range from milliseconds to $\mathcal{O}(0.1)$ s [166]. Given the sensitive frequency band of current ground-based detectors, these delays correspond to lens masses in the range $M_L \sim 10^2$-$10^6\ M_\odot$ In this mass range, unlike microlensing, which is typically well-motivated by point-mass models, phenomenological lensing can arise from a variety of lens candidates, including IMBHs and dark-matter subhalos embedded within a larger galactic lens [166]. Consequently, for a robust astrophysical interpretation, the parameters inferred from this agnostic framework should be compared with the results of lens-model-dependent analyses based on representative physical models that produce such phenomenological signatures.

Shifting focus to the strong lensing regime, a single strongly lensed signal can show a lensing signature if it corresponds to a Type II image [157, 305, 306, 307]. For Type II images ($n_i = 0.5$), the phase factor in the amplification factor (Eq. (2.74)) becomes $\exp(-i\pi/2) = -i$, which corresponds to a constant $\pi/2$ phase shift across all frequencies. In the time domain, this operation is equivalent to a Hilbert transform of the waveform, introducing a frequency-dependent time delay for different frequency components.

For a signal dominated by the quadrupole mode ($\ell = |m| = 2$), this constant phase shift is completely degenerate with a shift in the coalescence phase (ϕ_c) of an unlensed signal, rendering it indistinguishable. However, this degeneracy can be broken if the signal contains non-negligible HOMs (e.g., $\ell = 3, m = 3$). Because the phase shift applies globally to all modes, while the physical phase evolution of each mode scales with its m-index ($\Phi_{\ell m} \approx m\Phi_{\rm orb}$), the relative phase between different harmonics is altered. The lensed waveform for a Type II image is given by [157]:

$$h_{\rm TypeII}(t) \propto \sum_{\ell,m} \mathcal{A}_{\ell m} e^{-i(\Phi_{\ell m}(t)+\pi/2)}, \tag{5.5}$$

where $\mathcal{A}_{\ell m}$ is the amplitude of each mode. In an unlensed signal, a shift in the orbital phase $\Phi_{\rm orb}$ by $\Delta\Phi$, a constant shift in the orbital phase at coalescence, would result in a phase shift of $m\Delta\Phi$ for each mode. Since the Type II transformation introduces a shift of $\pi/2$ for all m, no single value of $\Delta\Phi$ can mimic this effect across multiple harmonics (e.g., a $\pi/2$ shift looks like an orbital shift of $\pi/4$ for $m = 2$, but $\pi/6$ for $m = 3$). This inconsistency allows for the unambiguous identification of a Type II image.

Following the methodologies described above, the search for Type II images aims to calculate Bayes factors for observed GW events and compare them against a background distribution. Given that strong lensing can produce three distinct image types (Type I, II, and III), the analysis involves calculating the Bayes factors between Type II and Type I images ($\mathcal{B}^{II}_{I}$), as well as between Type II and Type III images ($\mathcal{B}^{II}_{III}$) [238]. If both Bayes factors are sufficiently high, corresponding to a low FPP, the evidence for lensing becomes compelling.

5.2 Multiple Repeated Lensed Signals

As discussed previously, searches for lensing signatures within a single distorted signal rely on identifying distinctive waveform modulations that are not expected in unlensed signals, such as interference-induced amplitude and phase variations, as well as the characteristic phase shifts associated with specific image types. In contrast, the search strategy for multiple lensed signals from the same source (i.e., strongly lensed GW (SLGW) events) is qualitatively different.

Because gravitational lensing in the geometric-optics limit is achromatic and does not alter the morphology of the GW waveform, intrinsic source parameters, such as the component masses (m_1, m_2), spin magnitudes (χ_1, χ_2), and source redshift (z_s), are invariant across all lensed images. In addition, the sky locations of different images are expected to be nearly identical, since the current GW detector network does not have sufficient angular resolution to distinguish them [20]. In contrast, the extrinsic parameters, such as the luminosity distance ($d_{\rm L}$), arrival time (t_c), and coalescence phase (ϕ_c), can differ between images [155, 156, 157]. Consequently, strong lensing searches rely on the expectation that multiple lensed images share the same intrinsic source properties, with the differences in observable extrinsic parameters being consistent with physically plausible lensing scenarios.

We discuss how to utilise the consistency in the intrinsic parameters to identify candidate pairs or multiplets of SLGW signals originating from a common source and lensed by a macrolens[3]. Furthermore, we extend our discussion to more complex scenarios in which both strong lensing and microlensing occur simultaneously. In these cases, the lens's complex mass distribution can produce both multiple discrete signals and subtle waveform distortions in individual images.

5.2.1 Joint Parameter Estimation

To quantitatively assess whether a pair of GW events originates from strong lensing of a common source, we analyse the events under two competing hypotheses: the strong-lensing hypothesis, $\mathcal{H}_{\rm SL}$, in which the event pair shares a common source, and the unlensed hypothesis, $\mathcal{H}_{\rm U}$, in which the two events are unrelated unlensed signals. We then compute the corresponding Bayes factor, $\mathcal{B}^{\rm SL}_{\rm U}$. This procedure is referred to as *joint parameter estimation* (JPE), where "joint" indicates that the two events are analysed in a correlated manner, either by conditioning one event on the parameters of the other or by estimating parameters for both events simultaneously [300, 301, 302].

3. A macrolens refers to a large-scale gravitational lens, such as a galaxy or galaxy cluster, that produces multiple, discretely resolvable signals and is well described by the geometric optics limit, in contrast to microlensing, which can induce wave-optics effects and interference patterns.

Let d_1 and d_2 denote the data streams containing the two GW signals. The Bayes factor is defined as [300, 301, 302]:

$$\begin{aligned}\mathcal{B}^{\rm SL}_{\rm U} &\equiv \frac{\mathcal{Z}(d_1, d_2 \mid \mathcal{H}_{\rm SL})}{\mathcal{Z}(d_1, d_2 \mid \mathcal{H}_{\rm U})} \\ &= \frac{\alpha^2}{\beta}\frac{\displaystyle\int d\boldsymbol{\theta}\, d\boldsymbol{\lambda}\, p(d_1, d_2 \mid \boldsymbol{\theta}, \boldsymbol{\lambda}, \mathcal{H}_{\rm SL})\, \pi(\boldsymbol{\theta}, \boldsymbol{\lambda} \mid \mathcal{H}_{\rm SL})}{\displaystyle\prod_{i=1}^{2}\int d\boldsymbol{\theta}_i\, p(d_i \mid \boldsymbol{\theta}_i, \mathcal{H}_{\rm U})\, \pi(\boldsymbol{\theta}_i \mid \mathcal{H}_{\rm U})},\end{aligned} \tag{5.6}$$

where $\boldsymbol{\theta}$ denotes the set of source parameters shared by the two lensed images, while $\boldsymbol{\lambda}$ represents lensing-related parameters and image-dependent extrinsic parameters. Here, $p(d_1, d_2 \mid \boldsymbol{\theta}, \boldsymbol{\lambda}, \mathcal{H}_{\rm SL})$ is the joint likelihood for the two signals under $\mathcal{H}_{\rm SL}$, and $\pi(\boldsymbol{\theta}, \boldsymbol{\lambda} \mid \mathcal{H}_{\rm SL})$ denotes the corresponding prior distribution for the parameters. In contrast, under $\mathcal{H}_{\rm U}$, each event is described by an independent parameter set, $\boldsymbol{\theta}_i$. Since the two events are assumed to be independent under $\mathcal{H}_{\rm U}$, the joint likelihood factorises into the product of the individual-event likelihoods. Note that α and β are normalisation constants associated with the strong-lensing and unlensed hypotheses, respectively, which account for selection effects arising from detector sensitivity and detection thresholds [302].

A high value of $\mathcal{B}^{\rm SL}_{\rm U}$ suggests that the event pair is more likely to originate from a strongly lensed common source. However, like the case of a single distorted signal discussed in Section 5.1, the Bayes factor alone is not sufficient to claim a decisive detection. The BBH population spans a broad and nearly continuous region of the physically allowed mass-spin parameter space, and therefore, unrelated events can coincidentally exhibit similar source parameters, resulting in false alarms. To quantify this possibility, it is necessary to construct the background distribution of $\mathcal{B}^{\rm SL}_{\rm U}$ expected from pairs of independent unlensed events and compute FPP or global FPP in Eqs. (5.2) and (5.3).

Up to this point, we have only considered pairs of GW events. However, as discussed in Section 2.2, a single GW source can produce more than two strongly lensed images. In such cases, the formalism of $\mathcal{B}^{\rm SL}_{\rm U}$ in Eq. (5.6) remains unchanged in principle, but the analysis becomes increasingly complicated because additional image-dependent parameters must be considered. Nevertheless, it is not always necessary to analyse all lensed images simultaneously for an initial search. More specifically, even if a quadruply lensed GW signal is observed, performing JPE on only a randomly selected image pair can still enable the identification of a strong-lensing candidate, since the two signals originate from the same underlying source. A more complete joint analysis incorporating all images would then be required as a follow-up study to fully characterise the lens system and improve constraints on the lensing configuration (e.g., [300, 310, 311]).

In practice, performing a full JPE for a pair of GW events is substantially more computationally expensive than standard parameter estimation for a single unlensed signal. This issue becomes increasingly important as the number of detected CBC events grows. For example, the GWTC-4 catalogue already contains more than 200 CBC events [20], and future observing runs are expected to detect significantly more signals. Since the number of possible event pairs scales as $N(N-1)/2$, the total number of candidate pairs increases rapidly with the catalogue size. Consequently, performing full JPEs for all possible pairs would require prohibitive computational resources and analysis time.

To address this challenge, recent strong-lensing searches [238] employed several low-latency ranking and filtering pipelines to identify promising candidate pairs before performing full JPEs. This hierarchical strategy restricts the computationally intensive follow-up analyses to only those event pairs flagged as potentially lensed, enabling a substantially more efficient use of computational resources.

5.2.2 Strong Lensing with Substructures

Astrophysical compact objects that produce microlensing signatures in the wave-optics regime are typically embedded within larger gravitational structures such as galaxies or galactic halos. Consequently, a GW propagating near such a system can experience both galaxy-scale strong lensing and additional perturbations from smaller embedded structures. In this scenario, the macrolens first generates multiple strongly lensed images (so-called *macroimages*), after which each image can undergo further lensing distortions due to a microlensing field composed of compact objects or substructures along its individual propagation path [312, 313, 314, 315, 316, 317, 318, 319, 320].

Figure 5.2 illustrates an example of such a complex lensing configuration, in which a BBH is lensed by microlenses embedded within a macrolens. Since the microlensing fields need not be identical across macroimages, reflecting different shear effects near each image as well as the non-uniform stellar distribution within the lens galaxy, each image may experience distinct waveform modulations. As a result, different macroimages do not necessarily share identical waveform morphologies, even though they originate from the same underlying source parameters. This situation poses a challenge for strong-lensing searches based on Eq. (5.6), because the corresponding JPE framework is lens-model-independent and primarily relies on the consistency of the inferred source parameters between the two events. Such approaches may not fully capture the signatures of complex lensing configurations. In particular, microlensing-induced waveform distortions can bias the inferred source parameters and potentially reduce the apparent consistency between strongly lensed images, thereby causing genuine lensed pairs to be missed [320, 321, 322].

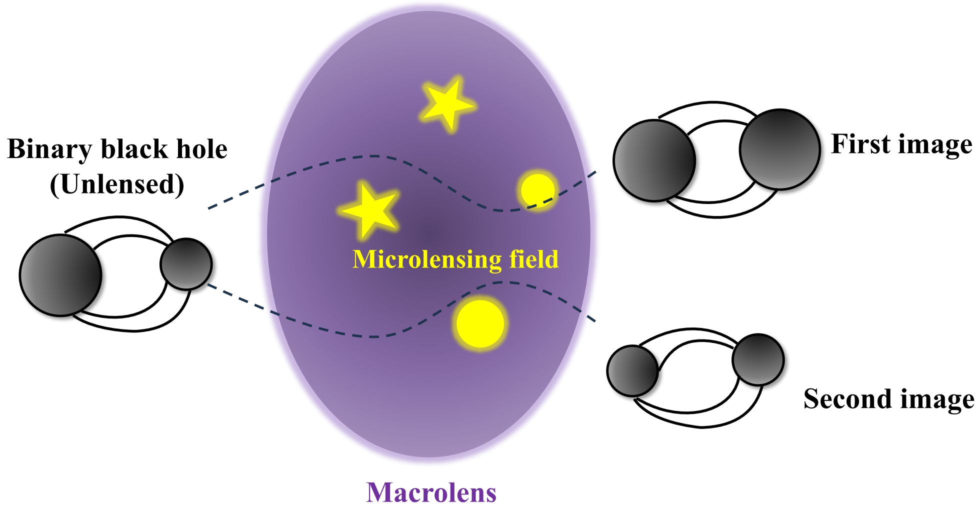

Figure 5.2: A complex lensing system in which a BBH signal is strongly lensed into multiple macroimages by a galaxy-scale macrolens, with additional microlensing effects arising from stellar fields or compact objects embedded within the macrolens.

A similar limitation also arises in searches for single distorted lensed signals. As discussed in Section 5.1.1, microlensing searches typically assume the isolated PM lens model. In realistic astrophysical environments, however, multiple compact lenses may contribute to the amplification factors. In addition, the macrolens itself can introduce external shear and convergence, further modifying the lensing effects [314, 318, 319].

As a simple example, the amplification factor in Eq. (2.71) for a system consisting of a macrolens and a PM microlens can be described by introducing an additional convergence (κ) and shear (γ), given by:

$$\begin{aligned} F(w,\boldsymbol{y}) &= \frac{w}{2\pi i}\int d^2x \left[\exp\left\{iw\left(\frac{1}{2}|\boldsymbol{x}-\boldsymbol{y}|^2-\psi(\boldsymbol{x})\right)\right\}\right] \\ &= \frac{w}{2\pi i}\int d^2x \left[\exp\left\{iw\left(\frac{1}{2}|\boldsymbol{x}-\boldsymbol{y}|^2-\ln\boldsymbol{x}+\frac{\kappa}{2}\left(|\boldsymbol{x}|^2\right)+\psi_\gamma(\boldsymbol{x})\right)\right\}\right], \end{aligned} \tag{5.7}$$

where $\psi_\gamma(\boldsymbol{x})$ is defined in Eq. (2.56).

Figure 5.3 compares three frequency-dependent amplification factors: those of an isolated PM microlens (see Eq. (2.73)), and an embedded-lens configuration in which an SIE macrolens contains the same PM microlens evaluated at two different macroimage positions, corresponding to Type I and Type II images (for the image type, see Section 2.1.3). The two embedded cases correspond to the amplification factors given by Eq. (5.7) for the same PM microlens evaluated at different macroimage locations within the macrolens po-

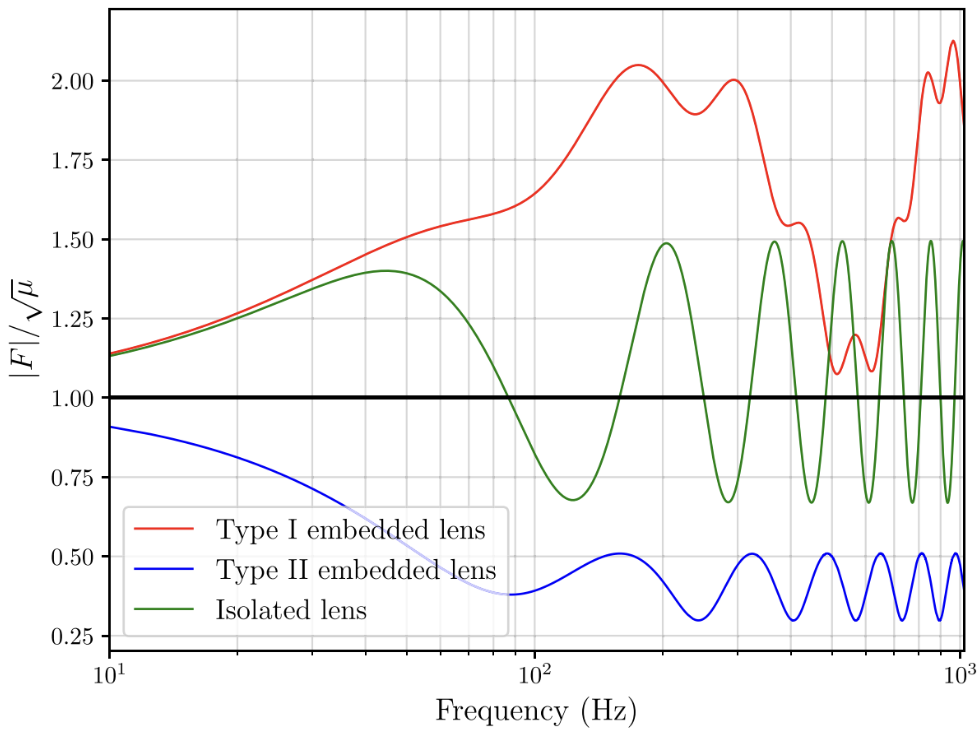

Figure 5.3: An example of amplification factors for an isolated PM lens and for the same PM perturber embedded in Type I and Type II macroimages produced by a common macrolens galaxy. Figure is from [318].

tential. This locational difference corresponds to the distinct local lensing environments, which can be effectively interpreted as different external shear conditions induced by the macrolens potential at each image position. Therefore, the two embedded amplification factors can differ substantially, and both deviate significantly from the isolated PM case.

Moreover, the example still represents an idealised configuration involving a single PM microlens. Realistic lens environments are expected to contain multiple compact objects or substructures, which produce even more complex interference patterns [315, 319, 323]. Consequently, the isolated PM approximation may fail to reproduce the true amplification factors, introducing non-negligible systematic mismatches between the assumed lens model and the actual lens. Such lens-model systematics can reduce the resulting Bayes factors and may weaken or even hinder the detection of microlensing signatures.

One of the most promising approaches for searching for such complexly lensed GW signals is to explicitly incorporate the effects of all relevant substructures into the lensed waveform model by parameterising the lensing perturbations and jointly inferring the lens parameters together with the CBC source parameters, analogous to the method introduced in Section 5.1. However, such analyses rapidly become computationally prohibitive. Bayesian inference for GW signals typically requires millions of likelihood evaluations, and the computational cost can increase dramatically with the complexity of the lens model. In addition, the integration in Eq. (5.7) itself is computationally expensive due to the highly oscillatory nature of the integrand. A more efficient approach for the rapid identification of both strongly lensed and microlensed signals will be discussed in Chapter 7.

Chapter 6

Search for Microlensing Signatures in O4a BBHs

Disclaimer

The following chapter is based on the paper *GWTC-4.0: Searches for Gravitational-Wave Lensing Signatures* [238], carried out by the LVK Lensing Group. The author contributed to this work both as a data analyst conducting microlensing searches for observed BBH events and as a member of the paper writing team, contributing to manuscript preparation and figure generation. This work has been submitted to the *Astrophysical Journal* for publication.

As discussed in Section 5.1, lensing signatures can be identified within a single distorted GW signal. However, from the first two observing runs (O1–O2) [324] through the third (O3) [325, 326], the LVK Collaboration reported no decisive evidence for lensing in GW observations.

We have extended these searches to BBH events observed during the first part of the fourth LVK observing run (O4a), with a specific focus on microlensing signatures, as detailed in Section 5.1.1. The author has contributed to microlensing searches since the first half of O3, continuing through analyses of the full O3 dataset through participation in data analysis and interpretation. Building on these earlier efforts, the analysis of O4a data adopts more advanced and dedicated frameworks: `LensingFlow`, an automated workflow manager designed for large-scale lensing analyses [327], and `Gravelamps`, a pipeline specialised in performing full wave-optics hypothesis testing employing a grid-based approach to calculate amplification factors [298]. These tools enable more systematic and computationally efficient exploration of microlensing effects.

In this chapter, we present the results of the O4a microlensing search and examine the event GW231123_135430 (hereafter GW231123) in more detail. To date, this special event exhibits the highest Bayes factor in favour of the microlensing hypothesis over the unlensed hypothesis among all analysed candidates. While the high Bayes factor does not yet provide definitive evidence for lensing, it makes the event a strong candidate. This motivates a more detailed investigation to distinguish possible lensing effects from intrinsic source properties or instrumental artifacts.

6.1 Analysis Setup

We analyse 84 BBH events observed during the O4a to assess the presence of microlensing signatures. As described in Section 5.1.1, we examine the properties of these events under the microlensing hypothesis and evaluate their Bayesian evidence. The analysis employs the isolated PM lens model, and for the waveform we adopt IMRPhenomXPHM-SpinTaylor [288, 289] (XPHM-ST) waveform model, one of the fiducial models used in the GWTC-4.0 analysis [20].

To estimate the two PM lens parameters, we define priors based on the expected physical characteristics of potential lenses. For the redshifted lens mass ($M_{\rm L}^z$), we employ a uniform prior between $1\,M_\odot$ and $10^4\,M_\odot$. This range spans the masses of stars and stellar remnants ($\sim 1\,M_\odot$) up to IMBHs ($\sim 10^4\,M_\odot$). Beyond this upper limit, the time delay between the two microimages typically becomes large enough that they are resolved as separate signals, rather than appearing as a single distorted waveform. Meanwhile, the prior for the dimensionless impact parameter (y) is set as $p(y) \propto y$ with $y \in [0.1, 3]$, maintaining consistency with previous microlensing searches [324, 325, 326]. A value of $y < 0.1$ is statistically rare because of the small cross-section, whereas for $y > 3.0$, the microlensing effect becomes negligible for current ground-based detectors.

It is important to note that a normalisation error was identified in the likelihood evaluation using `Bilby`, which serves as the foundational framework for both the unlensed analyses in this chapter and the `Gravelamps` pipeline (see Section 5.10 of Abac et al. [328] for more details). To address the likelihood issue, we employ a reweighting scheme applied directly to the nested sampling chains for all `Gravelamps` and complementary unlensed analyses. Unlike reweighting only the final posterior samples, reweighting the full nested sampling chains also consistently updates the Bayesian evidence. In practice, this is done by recalculating the weights of the nested sampling points using the corrected likelihood values, leading to a more robust estimate of the evidence.

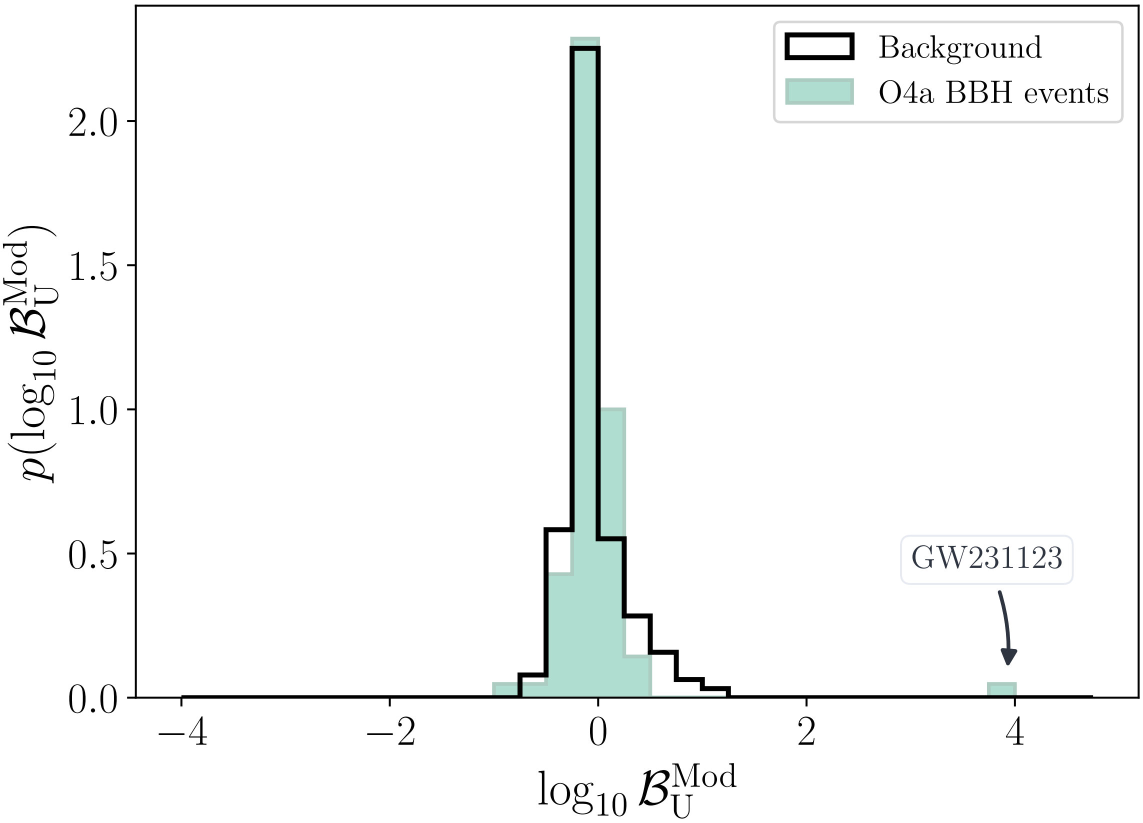

Figure 6.1: Histogram of $\log_{10}\mathcal{B}_{\mathrm{U}}^{\mathrm{Mod}}$ for O4a BBH events and background injections. GW231123 is clearly separated from the background distribution.

6.2 Point-Mass Microlensing Search

From the parameter estimation setup described above, we calculate the evidence for each event under the PM microlensing hypothesis ($\mathcal{B}^{\mathrm{ML}}$) and the unlensed hypothesis ($\mathcal{B}^{\mathrm{U}}$). The resulting Bayes factors, $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}}$, serve as our main statistic for the hypothesis test, as detailed in Section 5.1.1.

Figure 6.1 displays the distribution of $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}}$ for the 84 O4a candidates (green histogram). To robustly evaluate these values, we generated a background injection set comprising 254 detectable mock unlensed GW signals ($\rho_{\mathrm{net}} > 8$) (see Appendix B of [238] for more details). The distribution of $\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}}$ for the background is also shown in Figure 6.1 (black histogram). While most O4a events have low Bayes factors consistent with the background, GW231123 clearly stands out as an outlier, with $\log_{10}\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}} \simeq 3.8$, the most substantial value identified in our lensing searches to date [329].

To quantify the statistical significance of this result, we estimate the FPP defined in Eq. (5.2). In this context, the FPP represents the likelihood that an unlensed signal—due to noise fluctuations or intrinsic parameter degeneracies—could yield a Bayes factor as high as that of GW231123. Since the observed value lies outside the range of the background distribution, we cannot directly calculate the FPP from it. In particular, none of the $N_{\mathrm{bg}} = 254$

background samples produced a value larger than the observed Bayes factor. Rather than assigning FPP $= 0$, we place a conservative upper bound using the rule of succession [330], which accounts for the finite number of background samples:

$$\mathrm{FPP}_{\mathrm{upper}} = \frac{N(\mathcal{B}_{\mathrm{U}}^{\mathrm{ML}} > \mathcal{B}_{\mathrm{obs}}) + 1}{N_{\mathrm{bg}} + 1} = \frac{0+1}{254+1} \approx 0.0039, \tag{6.1}$$

where N_{bg} is the total number of the background events. An FPP of 0.39 % indicates that such a high Bayes factor is rare under the background distribution, making the signal stand out relative to typical events.

However, when accounting for the trial factor across the 84 independent events we analysed, the global FPP is significantly higher. Using Eq. (5.3), we obtain an upper bound for the global FPP of $\mathrm{FPP}_{\mathrm{global}} < 28\,\%$. The $\mathrm{FPP}_{\mathrm{global}}$ value is not sufficiently low to claim a definitive microlensing detection, which would typically require $\mathrm{FPP}_{\mathrm{global}} \ll 1\,\%$. This relatively high value is primarily due to the limited background size. With only 254 injections for an analysis of 84 events, the background is insufficient to resolve the far tail of the distribution. For example, to achieve a 3σ significance ($\mathrm{FPP}_{\mathrm{global}} \approx 0.0027$) for such a high Bayes factor, a background set of $\mathcal{O}(10^4)$ injections would be required to sample rare noise and degeneracy configurations.

Furthermore, the current background may not adequately represent events with properties similar to GW231123, which is characterised by a high total mass and evidence of non-negligible effective precessing spin [329]. In particular, such events are relatively rare in the background, and their absence can bias the FPP estimate. We consider two main effects:

1. **Precessing spin-lensing degeneracy**: The modulations induced by precession can mimic the interference patterns of lensing due to small mass objects [331]. If a background does not contain enough precessing BBHs, the risk of a false positive increases.
2. **Signal duration and non-Gaussianity**: Due to its high mass, GW231123 is a short-duration signal ($\sim 0.1\,$s), making it more susceptible to non-Gaussian instrumental noise, such as glitches. While our background is constructed to mitigate such non-Gaussian features, it may not fully capture their impact, which can be misinterpreted as either high spin [332] or a lensing signature.

Consequently, while the high Bayes factor of GW231123 is compelling, the current background limitations and potential degeneracies imply that more thorough follow-up is needed before we can claim microlensing.

6.3 GW231123: The Most Lensing-Favoured Event

6.3.1 Inferred Lens Properties

In Section 6.2, we found that GW231123 exhibits the highest Bayes factor ($\log_{10} \mathcal{B}^{\rm ML}_{\rm U}$) recorded to date in GW lensing searches. To assess the robustness of this result against waveform modelling uncertainties, we repeat the analysis using alternative waveform models IMRPhenomXO4a (XO4a) [333] and NRSur7dq4 (NRSur) [334]. These models incorporate different physical prescriptions and modelling systematics relative to the baseline XPHM-ST model. The recovered $\log_{10} \mathcal{B}^{\rm ML}_{\rm U}$ values for each waveform are shown in Table 6.1. Notably, the Bayes factors exhibit a strong dependence on the choice of waveform model. As all other aspects of the analysis, such as the data segment, prior assumptions, and noise treatment, are held fixed, the observed differences across waveform models indicate that the high Bayes factor obtained with XPHM-ST is likely driven by its specific modelling characteristics, rather than by external factors such as non-Gaussian noise transients.

	Waveform model		
	XPHM-ST	NRSur	XO4a
$\log_{10} \mathcal{B}^{\rm ML}_{\rm U}$	3.8	1.5	0.1
$M^z_L\,[M_\odot]$	911^{+504}_{-292}	803^{+489}_{-299}	687^{+13208}_{-356}
y	$0.61^{+0.27}_{-0.21}$	$0.69^{+0.37}_{-0.26}$	$2.13^{+0.79}_{-1.54}$

Table 6.1: Bayes factors and inferred point-mass lensing parameters for GW231123 across different waveform models.

We observe a clear hierarchical decrease in $\mathcal{B}^{\rm ML}_{\rm U}$ in the order of XPHM-ST, NRSur, and XO4a. This trend indicates an inverse relationship between $\mathcal{B}^{\rm ML}_{\rm U}$ and the quality of the fit (i.e., the Bayesian evidence) under the unlensed hypothesis. Specifically, among the three waveforms, XPHM-ST yields the lowest Bayesian evidence for the data, followed by NRSur, with XO4a providing the strongest evidence (see Appendix B in [329] for details). This suggests that waveform models that provide a relatively poorer fit to the data under the unlensed hypothesis may show a stronger preference for the additional flexibility introduced by the PM microlensing model. Consequently, the relatively high Bayes factor obtained with XPHM-ST may, at least in part, reflect the interplay between waveform modelling systematics and the added degrees of freedom in the microlensing model, rather than arising solely from a distinct microlensing signature.

For each microlensing analysis performed with a given waveform model, the inferred values of the two point-mass parameters (M^z_L, y), including their medians and 90 % C.I, are summarised in Table 6.1. The corresponding posterior distributions are shown in Figure 6.2. For the XPHM-ST and NRSur models, the posteriors are well constrained and exhibit clear peaks. In contrast, the XO4a results show broader posterior distributions, with support

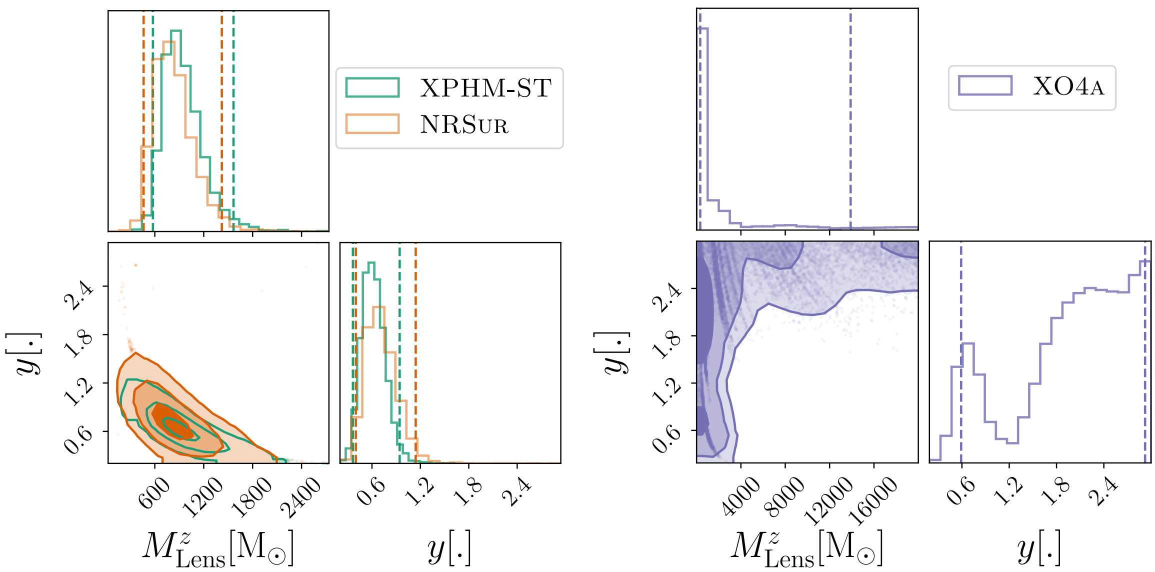

Figure 6.2: Posterior distributions of the redshifted lens mass and dimensionless impact parameter inferred from point-mass lensing analyses using the XPHM-ST (green), NRSur (orange), and XO4a (purple) waveform models.

extending over a wider region of parameter space, although some overlap with the peaks identified by XPHM-ST and NRSur is still apparent. This behaviour is consistent with the lower Bayes factor obtained for XO4a and indicates comparatively weaker support for the microlensing hypothesis.

The recovered redshifted lens mass ($M_L^z \sim 10^2$–10^4 $M_\odot$) and source redshift ($z_s \sim 0.2 - 0.7$) point to several plausible lens candidates. While an IMBH is a natural interpretation in this mass range, the lens could also correspond to a dense DM subhalo or a PBH, both of which are predicted to exist within galactic halos under various cosmological models. Although our analysis adopts an isolated PM model, a compact lens is more likely to be embedded within a complex environment rather than in isolation, as discussed in Section 5.2.2. Possible environments include the centers of globular clusters, where IMBHs may form, or galactic disks, where the lens could be embedded within a larger macrolensing field. The relatively high Bayes factor obtained for the isolated PM model may arise because the observed signal is primarily sensitive to the local gravitational potential within the Einstein radius of the compact object. In this local approximation, the influence of the surrounding stellar population or the host galaxy potential becomes subdominant to the effect of the primary deflector on the GW wavefront.

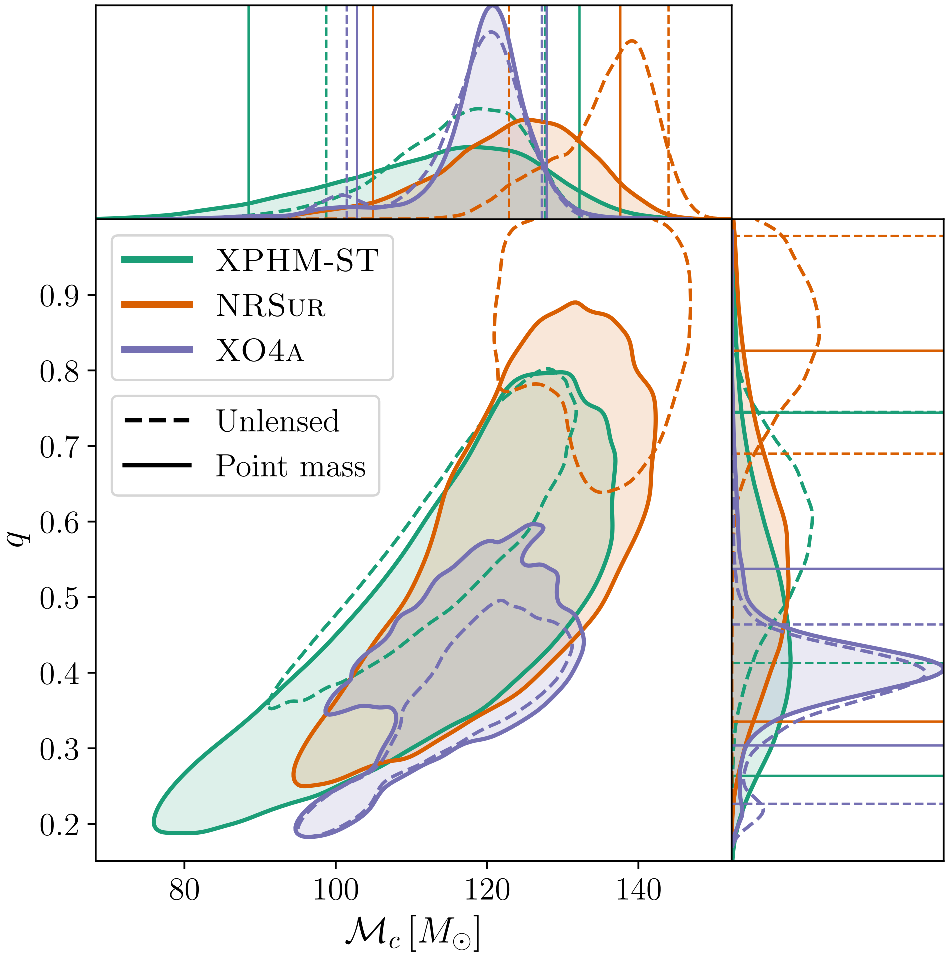

Figure 6.3: Comparison of the detector-frame chirp mass posteriors and mass-ratio under the PM microlensing (solid line) and unlensed (dashed line) hypotheses using XPHM-ST, NRSur, and XO4A waveform models.

6.3.2 Inferred Source Properties

In addition to the inferred lens properties, we compare the inferred source properties under the microlensing hypothesis to those obtained under the unlensed case.

Figure 6.3 presents the joint posterior distributions for the detector-frame chirp mass ($\mathcal{M}_c$) and mass ratio (q), with 90 % credible regions, along with the corresponding 1D marginal distributions for each waveform model, including their 90 % C.I.. Under the unlensed hypothesis, the inferred parameters exhibit noticeable discrepancies across the three waveform models. In particular, the NRSur model favours higher chirp masses and mass ratios than the IMRPhenom family. This behaviour likely reflects increased modelling uncertainty in the high-mass regime, where the signal contains relatively few in-spiral cycles and the inference becomes more sensitive to waveform systematics.

In contrast, under the microlensing hypothesis, the posterior distributions for the two parameters become broader across all models, mainly attributed to the increased degrees of freedom in the microlensing model. At the same time, the level of disagreement between waveform models is reduced. Notably, the NRSur posterior shifts toward lower masses and mass ratios, bringing it into closer agreement with the IMRPhenom-based results. This trend is also evident in the 2D credible regions, which show increased overlap compared to the unlensed case. Overall, while the broader posteriors indicate increased parameter degeneracy, the improved cross-model consistency suggests that the microlensing framework may partially mitigate waveform-dependent discrepancies present in the unlensed analysis.

We further investigate the impact of the microlensing hypothesis on the inferred source-frame masses ($m_{1,s}$ and $m_{2,s}$) and spin magnitudes (χ_1 and χ_2), focusing on the NRSur model, which provides the most physically rigorous representation by minimising the averaged mismatch with NR waveforms [329].

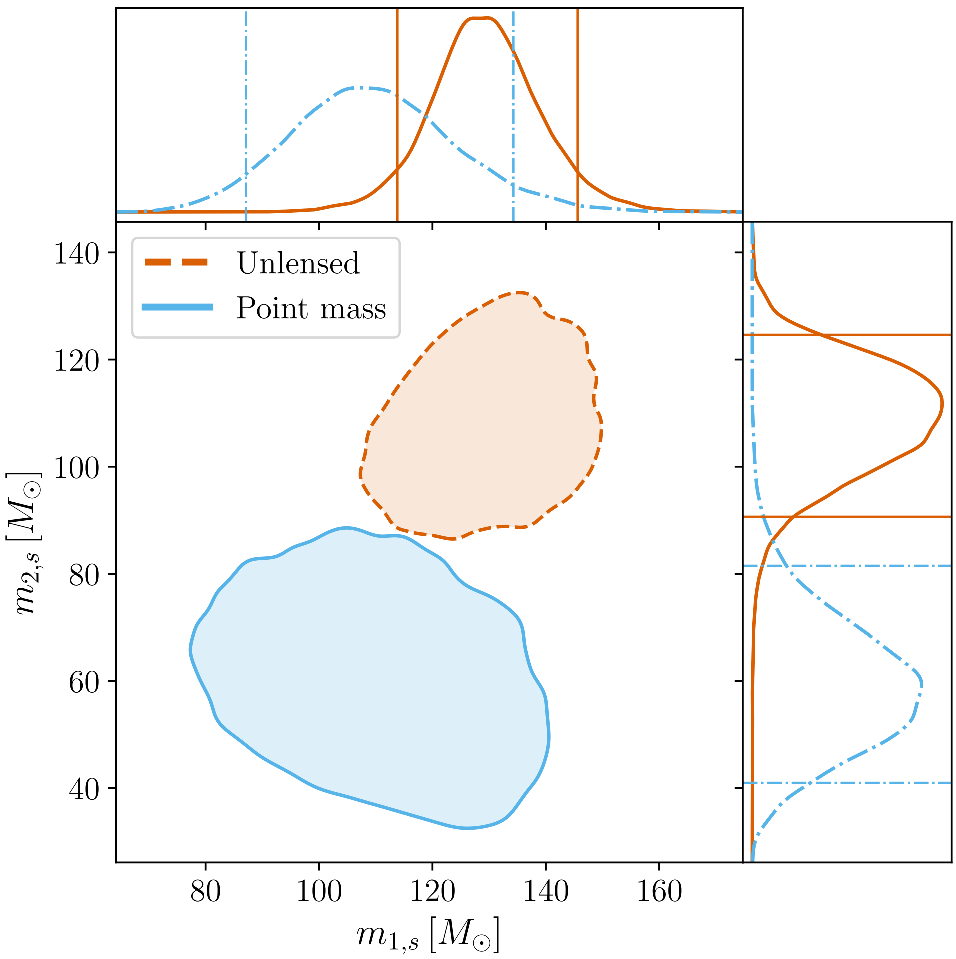

Figure 6.4: Comparison of the posterior distributions of the source-frame component masses for GW231123 inferred with the NRSur model under the unlensed (orange) and point-mass microlensing (sky blue) hypotheses.

A key feature of PM lensing is that it produces magnification only ($\mu \geq 1$). Under this interpretation, the observed GW signal corresponds to a magnified version of a more distant source. As a result, the inferred luminosity distance (d_L) increases, implying a higher source redshift (z_s). Since the source-frame masses scale as $m_s = m/(1+z_s)$, this leads to systematically lower intrinsic mass estimates. Figure 6.4 shows the posterior distributions of $m_{1,s}$ and $m_{2,s}$, under the microlensing and unlensed hypotheses. While the reduction in $m_{1,s}$ is relatively modest, $m_{2,s}$ decreases more significantly. Quantitatively, the median and 90 % C.I. shift from $129^{+16.4}_{-15.4}\ M_\odot$ and $106^{+15.2}_{-18.8}\ M_\odot$ in the unlensed case to $109^{+25.2}_{-21.9}\ M_\odot$ and $59^{+22.6}_{-18.0}\ M_\odot$ under the lensing hypothesis. This difference is consistent with the behaviour seen in Figure 6.3, where the NRSur analysis under the microlensing hypothesis favours smaller mass ratios. Consequently, whereas the unlensed interpretation places the event in an ambiguous outlier regime, the lensed scenario shifts it toward values more consistent with the known BBH population.

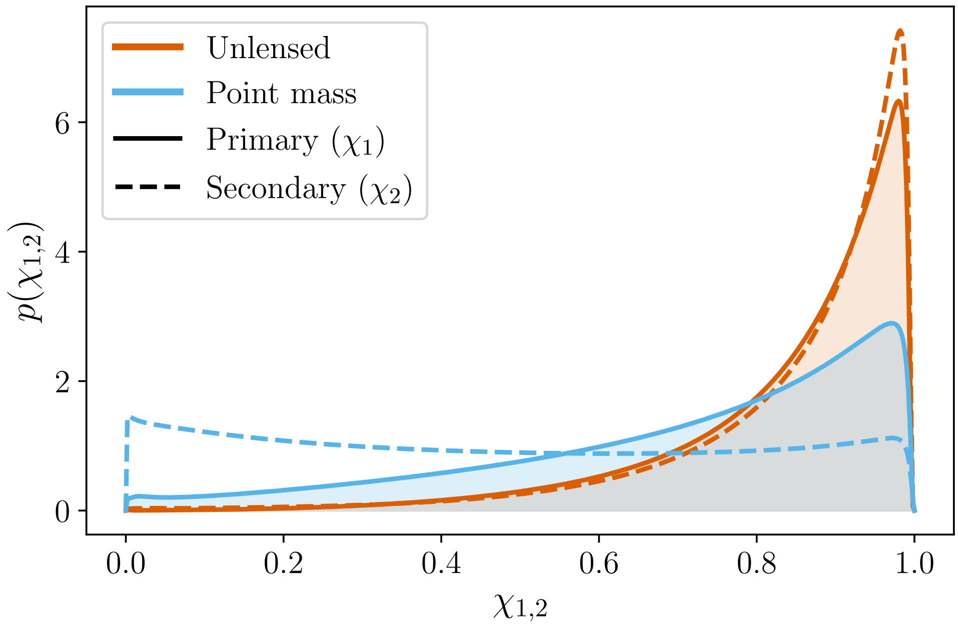

Figure 6.5: Comparison of the posterior distributions of the dimensionless spin magnitudes for GW231123 inferred with the NRSur model under the unlensed (orange) and point-mass microlensing (sky blue) hypotheses.

Furthermore, the inferred χ_1 and χ_2 shift toward more moderate values under the microlensing hypothesis across all waveform models, with χ_2 in particular exhibiting a broader and less peaked distribution. Figure 6.5 shows the corresponding posteriors for the NRSur cases. This behaviour can be understood in terms of a morphological degeneracy between spin-induced modulations and lensing-induced distortions. As discussed in [331], spin precession can introduce waveform modulations that resemble the interference-like features produced by a compact lens, particularly when the spin magnitudes are large. Within the microlensing framework, such features can instead be accommodated by lensing

Parameter	$\mathcal{M}_c$		χ_1		χ_2	
	KS	JSD	KS	JSD	KS	JSD
Unlensed	0.647	0.294	0.150	0.044	0.051	0.002
Lensed	0.342	0.089	0.097	0.009	0.026	4×10^{-4}

Table 6.2: KS statistics and JS divergence (in nats) computed between the chirp mass and spin magnitude posteriors inferred from individual-detector analyses under the unlensed and lensed hypotheses, using the NRSur waveform model

parameters, thereby reducing the preference for extremely high spin values. Similar to the trends observed in the mass parameters, this shift toward more moderate spin values improves the agreement across waveform models, leading to a more consistent interpretation of the signal.

6.3.3 Single-Detector Analysis and Consistency Checks

If a GW signal is genuinely affected by lensing, the associated distortions are expected to appear consistently across all detectors in the network. To test this, we perform independent microlensing analyses using data from the aLIGO Hanford (LHO) and aLIGO Livingston (LLO) detectors. As noted in the case of GW200208_130117 [335], a lensing-like feature present in only a single detector is more plausibly explained by local instrumental noise or a transient glitch that mimics the lensing morphology. In contrast, a consistent signature observed across independent detectors is less likely to have a terrestrial origin.

Using the NRSur model, we obtain $\log_{10} \mathcal{B}^{\rm ML}_{\rm U} = 0.62$ for LHO and 2.64 for LLO, both of which independently favour the microlensing hypothesis. The higher Bayes factor obtained for LLO can be attributed to its higher SNR relative to LHO, which enables a more precise resolution of the interference patterns. Part of the discrepancy between the two detectors may also arise from differences in data treatment: the LHO data were deglitched, whereas the LLO analysis used the original data [336]. Residual noise artifacts or variations in data conditioning can, therefore, introduce modest shifts in the inferred evidence.

To quantify the consistency of the intrinsic parameters, specifically the chirp mass, mass ratio, and spin magnitudes, inferred from each single-detector analysis, we compute the Kolmogorov–Smirnov (KS) statistic and the Jensen–Shannon (JS) divergence. These measures are applied to the posteriors obtained under both the unlensed and microlensing hypotheses. As shown in Table 6.2, both the KS and JS values are lower in the microlensing case than in the unlensed case, which indicates that the parameter estimates inferred from the two detectors are more consistent under the microlensing hypothesis, suggesting that it provides a more coherent description of the signal across the detector network.

To further assess whether the apparent preference for the lensing hypothesis reflects a genuine feature of the data or arises from modelling degeneracies, we compare the inferred waveforms with a model-independent reconstruction obtained from an unmodelled search pipeline [337]. This reconstruction is performed using cWB [44], which employs a wavelet basis to reconstruct coherent signal power across detectors, providing a data-driven baseline that does not rely on specific waveform assumptions.

Given the high $\mathcal{B}^{\mathrm{ML}}_{\mathrm{U}}$ values, which motivate the lensing interpretation, several scenarios can explain the observed agreement:

- **Consistent lensing signature**: The reconstructed waveform exhibits features consistent with the lensed templates, supporting the interpretation that the data contain lensing-induced distortions.
- **Degeneracy-driven agreement**: The reconstructed waveform does not exhibit clear lensing-specific features, whereas the lensed templates achieve a better fit by exploiting degeneracies (e.g., between spin-induced modulations and lensing effects), suggesting that the apparent support for lensing may be driven by model flexibility.
- **Ambiguous or noise-dominated structure**: The reconstructed waveform does not clearly favour either scenario, suggesting that the apparent agreement may be influenced by noise fluctuations or limited information content in the signal.

Figure 6.6 shows a time-domain comparison between the reconstructed waveform (h_{rec}) and the waveforms generated using the parameters inferred under the unlensed and microlensing hypotheses with the NRSur model (h_{model}). The shaded regions correspond to the 90 % C.I. of the waveforms obtained from parameter-estimation analyses. Visual inspection alone does not clearly indicate which waveform better matches the cWB reconstruction. We therefore quantify the agreement by computing the mismatch (MM), which measures the dissimilarity between two signals [76, 338]:

$$\mathrm{MM} = 1 - \max_{t_0, \phi_0} \frac{\langle h_{\mathrm{rec}} \mid h_{\mathrm{model}} \rangle}{\sqrt{\langle h_{\mathrm{rec}} \mid h_{\mathrm{rec}} \rangle \langle h_{\mathrm{model}} \mid h_{\mathrm{model}} \rangle}}, \tag{6.2}$$

where $\langle \cdot \mid \cdot \rangle$ denotes the noise-weighted inner product defined in Eq. (1.41). The maximisation is performed over relative time and phase shifts (t_0, ϕ_0) to account for trivial offsets in the coalescence time and phase, ensuring that the comparison reflects only intrinsic waveform differences rather than arbitrary alignments.

The term MM in Eq. (6.2) is computed using a restricted data segment spanning -0.15 s to $+0.1$ s around the trigger time (as illustrated in Figure 6.6), rather than the full 8 s data stream. Outside this interval, the signal is subdominant, and the data are effectively noise-dominated, so those regions do not contribute meaningfully to the comparison. Figure 6.7 shows the distributions of MM values between the cWB reconstruction and waveform realisations drawn from the unlensed and microlensed posteriors, restricted to samples within the 90 % C.I., for each LIGO detector. The mismatch between the cWB reconstruction and

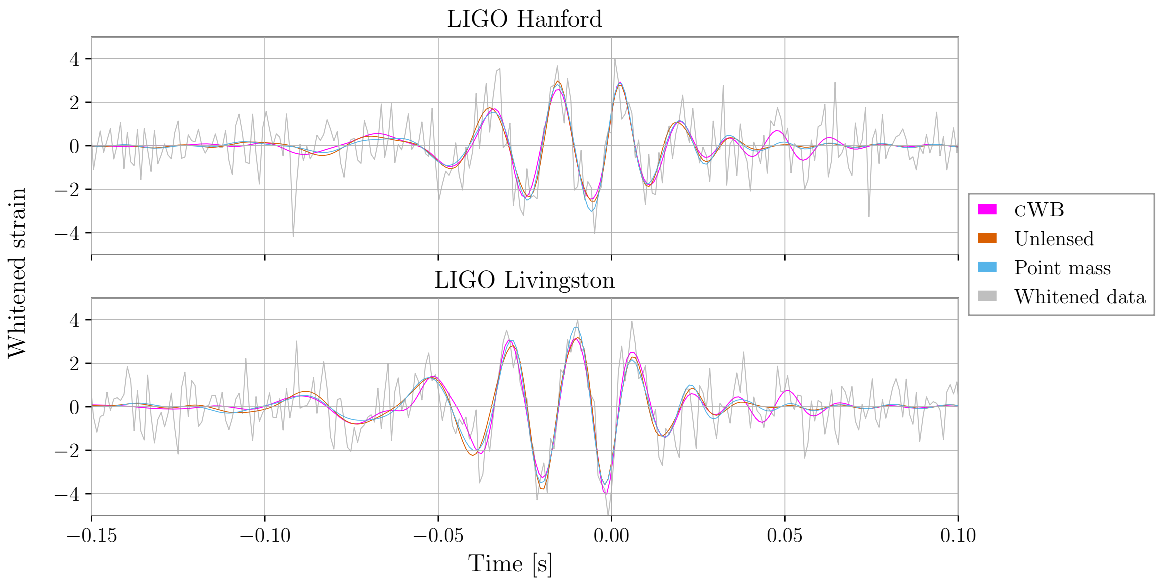

Figure 6.6: Comparison between the cWB reconstruction (magenta) and the 90 % credible waveforms inferred under the lensed (sky blue) and unlensed hypotheses (orange). The time represents the time relative to 2023 November 23 at 13:54:30.619 UTC.

the mean waveform is 0.029 (0.033) for the unlensed case and 0.022 (0.025) for the microlensed case in LHO (LLO), respectively, indicating that the microlensed reconstruction provides a closer match to the model-agnostic baseline. This trend is also reflected in the full mismatch distributions: for the unlensed model, the median and 90 % C.I. are $0.040^{+0.012}_{-0.009}$ and $0.045^{+0.012}_{-0.015}$ for LHO and LLO, while for the microlensed model they are $0.029^{+0.008}_{-0.006}$ and $0.033^{+0.010}_{-0.007}$. The difference between the mismatch distributions for the microlensed (sky-blue) and unlensed (orange) cases is further quantified using the KS statistic, yielding values of 0.75 for LHO and 0.68 for LLO. These results indicate that, although the two distributions partially overlap, they remain statistically distinguishable.

While the results consistently favour the microlensing hypothesis, it remains unclear whether this preference reflects genuine lensing-induced features or arises from degeneracies in the waveform modelling, particularly given the short duration of the signal, which limits the ability to distinguish between lensing effects and modelling degeneracies.

6.4 Conclusion

We reanalysed 84 BBH candidates detected during the first half of the LVK fourth observing run under the PM microlensing hypothesis, focusing on frequency-dependent lensing distortions. Among these, 83 events yield Bayes factors ($\mathcal{B}^{\mathrm{ML}}_{\mathrm{U}}$) that are either negligible or fully consistent with the expected unlensed background distribution. GW231123 stands out as the sole outlier, exhibiting a Bayes factor that exceeds the background and represents the highest value reported in GW lensing searches to date.

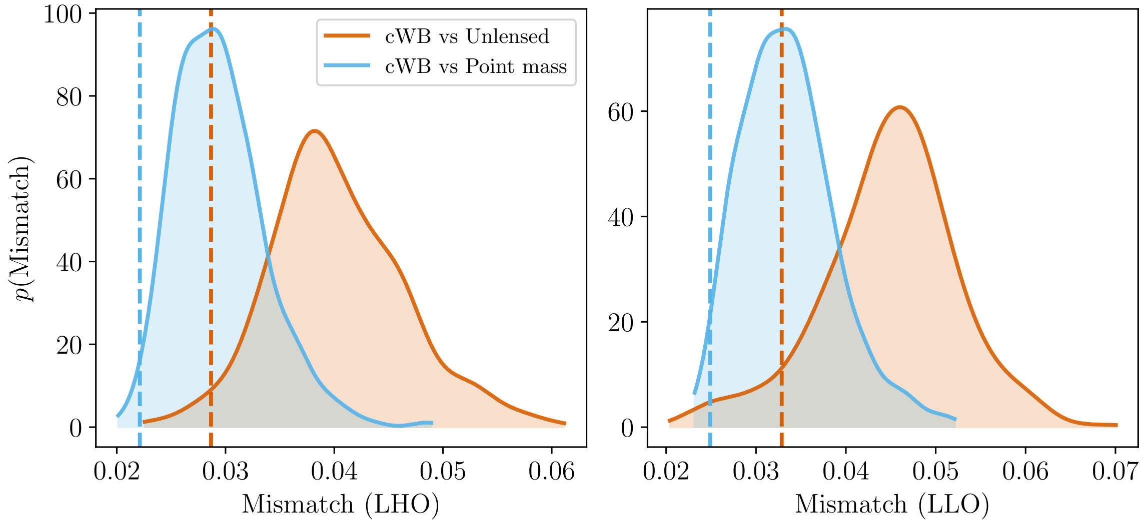

Figure 6.7: Distribution of the mismatch between the cWB reconstruction and the recovered lensed (sky blue) and unlensed (orange) waveforms for GW231123 within the 90 % C.I.. The vertical dashed lines indicate the mismatch values between the cWB reconstruction and the mean unlensed and mean lensed waveforms, respectively.

Due to the limited size of the current background injection set, we obtain an upper bound of 28 % for the global FAP, indicating that statistical fluctuations cannot yet be excluded with high confidence. Furthermore, we must consider the possibility that missing physics in current unlensed waveform models may artificially enhance the Bayes factor. Specifically, if the unlensed templates (e.g., XPHM-ST) fail to perfectly capture complex physical features present in the data, such as advanced spin precession or subtle NR corrections, the Bayesian evidence may compensate for these modelling residuals by favouring the additional degrees of freedom provided by the lensing parameters. In this case, the preference for lensing may reflect modelling degeneracies rather than a genuine lensing signal. Despite a more detailed investigation of the inferred lens properties, source parameters under the microlensing hypothesis, and consistency across individual detectors, we find that the evidence remains inconclusive. Although GW231123 represents the most compelling candidate identified so far, a robust lensing claim requires a larger background sample and improved control over waveform systematics.

Finally, it is worth considering an alternative and plausible astrophysical interpretation accessible to the current detector network: the overlapping-signals scenario [339]. In this picture, two or more independent GW signals from unrelated sources coincidentally arrive at the detectors within a very short time window. Such a superposition can produce a distorted waveform that mimics the interference patterns expected from micro-images separated by millisecond time delays. Under this scenario, [340] demonstrated that the Bayes factor favouring a multi-signal model over a single unlensed signal can also be significantly large, comparable to the value obtained for GW231123. Distinguishing between microlensing and overlapping signals remains challenging, as both can produce similar phenomenological signatures. Disentangling these possibilities, whether the signal arises from a single lensed source or from multiple coincident sources, will be essential for a definitive interpretation and constitutes an important direction for future work.

Chapter 7

Residual Test to Search for Microlensing Signatures in Strongly Lensed Gravitational-Wave Signals

Disclaimer

The following chapter is based on the paper *Residual Test to Search for Microlensing Signatures in Strongly Lensed Gravitational Wave Signals* [104]. The author is the first author of this paper and contributed to the development of methodology, data analysis, interpretation of the results, scientific discussion, manuscript preparation, and generation of all figures. This paper was published in the *Astrophysical Journal.*

Roughly, $\mathcal{O}(1)$ SLGW signals per year are expected for the LVK detectors at their design sensitivities, while third-generation (3G) detectors are anticipated to observe approximately $\mathcal{O}(10)$ such events per year [158, 159, 341]. To date, however, dedicated searches conducted by the LVK Collaboration across all BBH events from the O1 through O4a observing runs have yielded no conclusive evidence for either strong-lensing or microlensing signatures [238, 324, 325, 326, 335]. Current search schemes employ different methodologies depending on the lensing regime. Strong-lensing searches are typically agnostic to the lens model and rely on the geometrical-optics approximation [299, 300, 301, 302, 303, 304, 342, 343, 344, 345, 346, 347, 348], whereas the microlensing pipeline assumes a specific physical configuration, such as a single, isolated microlens [298].

In a more realistic scenario, as noted in Section 5.2.2, a GW signal strongly lensed by a galaxy or galaxy cluster is likely to experience further perturbation from stellar fields or DM substructures embedded within the lens. These effects can imprint complex signatures on individual SLGW signals. Consequently, the BBH parameters inferred from individual images can exhibit apparent discrepancies. Neglecting such microlensing effects in strong-lensing searches may lead to missed detections. Incorporating these additional complexities into current search pipelines, however, would substantially increase the computational cost. In the context of the rapidly growing number of detected events, such an approach is likely to become computationally prohibitive [349, 350, 351].

To circumvent this limitation, we propose a residual-based approach rather than performing full parameter estimation with template families that explicitly model both microlensing and strong-lensing effects. Specifically, if a standard strong-lensing analysis—which assumes that multiple images from the same source share an identical waveform morphology—is applied to strongly lensed signals that also exhibit microlensing, the resulting best-fit waveform will fail to fully capture the observed data. We then exploit the resulting residuals as a diagnostic to identify these additional lensing-induced distortions. This strategy is conceptually analogous to testing GR, which aims to identify deviations from GR-predicted templates within the observed signals [352, 353, 354, 355].

7.1 Methodology

7.1.1 Amplification Factors for Microlensing Fields within a Macrolens

The total lensing effect on a GW signal induced by a microlensing field embedded within a macrolens can be computed using Eq. (2.71), with an appropriate Fermat potential $T(\boldsymbol{x}, \boldsymbol{y})$ that describes the full lens configuration. In particular, the Fermat potential must incorporate both the macrolens contribution[1] and the microlensing effects, the former of which can be parametrised by the convergence (κ) and shear (γ) [103]. Under this setup, the Fermat potential for N point masses in the vicinity of a macroimage located at $\boldsymbol{x} = (x_1, x_2)$ is given by [319, 356, 357]:

$$
\begin{aligned}
T\left(\boldsymbol{x}, \boldsymbol{x}^m, \boldsymbol{y}_{\text{macro}} = 0\right) = & \frac{k}{2}\left((1-\kappa+\gamma)x_1^2 + (1-\kappa-\gamma)x_2^2\right) \\
& - \left[\frac{k}{2}\sum_m^N \ln\left(\boldsymbol{x}^m - \boldsymbol{x}\right)^2 + k\phi_-(\boldsymbol{x})\right],
\end{aligned}
\tag{7.1}
$$

where $\boldsymbol{x}^m$ denotes the position of the m^{th} microlens, and $k = 4Gm_{\text{L}}^z/c^3$, with m_{L} representing the average microlens mass. The term ϕ_- corresponds to a negative mass sheet introduced to compensate for the contribution of the microlenses, such that the total convergence remains unchanged. In Eq. (7.1), the macrolensing contribution is obtained through a local expansion of the lensing potential, where the expansion center can be chosen arbitrarily. For convenience, we set the macroimage position as the expansion center, i.e., $\boldsymbol{y}_{\text{macro}} = 0$.

Substituting Eq. (7.1) into Eq. (2.71), the amplification factor as a function of the dimensionless frequency w can be written as [323]:

$$
F(w) = \frac{w}{2\pi i}\int d^2\boldsymbol{x}\left[\exp\left\{iwT\left(\boldsymbol{x}, \boldsymbol{x}^m, \boldsymbol{y}_{\text{macro}} = 0\right)\right\}\right],
\tag{7.2}
$$

1. We assume an SIE profile (see Section 2.2.4) for the macrolens in this work.

7.1.2 Mock Data Generation and Joint Parameter Estimation

For mock lensed GW signals, we first simulate 100 BBH systems using PyCBC [358]. Source parameters are drawn via a Monte Carlo procedure based on a merger-rate density that traces the star formation rate (SFR) following [295], assuming a characteristic time delay of 50 Myr between star formation and the BBH merger (see Appendix B of [359] for details). We adopt a PowerLaw + Peak mass distribution with hyperparameters and parameter ranges as defined in [360] to sample the component masses (m_1, m_2). The inclination angle is drawn from $p(\iota) \propto \sin\iota$, while the remaining parameters—component spins $(\vec{a}_1, \vec{a}_2)$, polarisation angle ψ, right ascension α, declination δ, coalescence time t_c, and coalescence phase ϕ_c—are sampled from uniform distributions over their respective physically allowed ranges. The waveforms are then modelled with IMRPhenomPv2 [278].

For each BBH, we construct two pairs of lensed signals under two distinct lensing hypotheses:

- $\mathcal{H}_{\rm SL}$: the BBH is strongly lensed by a single macrolens, producing two macroimages.
- $\mathcal{H}_{\rm MLSL}$: the BBH is strongly lensed by the same macrolens as in $\mathcal{H}_{\rm SL}$, and each macroimage is further microlensed by stellar fields embedded within the macrolens.

Under $\mathcal{H}_{\rm SL}$, we generate 100 pairs of SLGWs using the previously made 100 BBH systems. For each BBH system, we assign strong-lensing parameters following the SIE model, including the axis ratio (q), velocity dispersion (σ_v), and impact parameter (y). The σ_v values are sampled from SDSS galaxy distributions [361], while the q values are drawn from a Rayleigh distribution with the scaling parameter corrected following [341]. The y values are sampled from a uniform distribution. Given these SIE lens parameters, the corresponding κ and γ are computed using Eqs. (??) and (2.21). For lensing systems that produce quadruple images, we select the first two arriving images to ensure that each system yields exactly two detectable GW signals.

Next, under $\mathcal{H}_{\rm MLSL}$, to generate signals that exhibit both strong-lensing and microlensing effects, we simulate microlensing fields by assuming a stellar mass distribution that follows the late Salpeter initial mass function (IMF) [362] and a Sérsic radial profile [363]. For consistency with [319, 364], we configure the IMF to include remnant objects, assuming their mass density constitutes 10 % of the total stellar mass density. Stellar microlens masses are sampled in the range $0.1–1.5\ M_\odot$, while compact remnant masses span $0.1–28\ M_\odot$. For each lensing system, the number of microlenses (N) is determined by the microlensing convergence (κ_*) as inferred from a Sérsic profile and the microlens masses. Table 7.1 summarises the main parameter distributions and ranges used to generate the lensed waveforms.

Finally, we construct a mock signal set consisting of 100 pairs of SLGW signals ($h^{\rm inj}_{\rm SL,1}, h^{\rm inj}_{\rm SL,2}$) and 100 pairs of signals that are both strongly lensed and microlensed ($h^{\rm inj}_{\rm MLSL,1}, h^{\rm inj}_{\rm MLSL,2}$) to serve as injections. Note that the lensing effect in $\mathcal{H}_{\rm SL}$ is described by $F_{\rm geo}(w)$ in Eq. (2.74), whereas for $\mathcal{H}_{\rm MLSL}$, the amplification is governed by the frequency-dependent factor $F(w)$

in Eq. (7.2). We inject all lensed signals into simulated stationary Gaussian noise n for the aLIGO-Hanford, aLIGO-Livingston, and AdV detectors, using the O4 design sensitivity PSDs [365][2]. To assess future prospects, we also employ the anticipated O5 sensitivity curves [365][3]. We consider only those signals with a network SNR exceeding the threshold of $\rho_{\text{net}} = 12$, following [319].

Parameter	Distribution or Model	Range
m_1	`Power Law + peak`	[5, 85]$M_\odot$
m_2	`Uniform`	[5, m_1]$M_\odot$
m_{L} (stellar)	`Salpeter IMF`	[0.1, 1.5]$M_\odot$
m_{L} (remnant)	Spera et al. (2015) [362]	[0.1, 28]$M_\odot$
κ_*	`Sérsic profile`	-
$N(\kappa_*, m_{\text{L}})$	-	[10^3, 10^6]
κ, γ	`SIE` (q, σ_v, y)	-

Table 7.1: The assumptions for component masses, macrolens, and microlens parameters adopted to simulate lensed GW signals.

To conduct JPEs for the 200 pairs of lensed GW signals, we employ the `Golum` pipeline [300, 303]. The joint likelihood function in this pipeline assumes that the two input signals originate from the same source and are purely strongly lensed, thereby neglecting any potential microlensing effects.

First, from the JPEs performed on the 100 pairs of SLGW signals, we obtain the maximum-likelihood waveforms for the first lensed signals ($h^{\text{maxL}}_{\text{SL},1}$). As a complementary step, we also extract the maximum-likelihood values of the lensing observables, including relative magnification factors (μ_{rel}), differences between arrival times (Δt), and differences between Morse factors ($\Delta\phi_n$). One can then determine the maximum-likelihood waveforms of the second lensed signals ($h^{\text{maxL}}_{\text{SL},2}$) by transforming the luminosity distances, coalescence times, and Morse factors according to:

$$\begin{aligned} d_L^{\text{maxL}}|_{\text{SL},2} &= d_L^{\text{maxL}}|_{\text{SL},1} &&\times\sqrt{\mu_{\text{rel}}^{\text{maxL}}}, \\ t_c^{\text{maxL}}|_{\text{SL},2} &= t_c^{\text{maxL}}|_{\text{SL},1} &&+\Delta t^{\text{maxL}}, \\ \phi_n^{\text{maxL}}|_{\text{SL},2} &= \phi_n^{\text{maxL}}|_{\text{SL},1} &&+\Delta \phi_n^{\text{maxL}}. \end{aligned} \tag{7.3}$$

2. The PSD data were obtained from the LIGO Document Control Center (DCC) under document LIGO-T2000012, available at `https://dcc.ligo.org/LIGO-T2000012/public`.
3. The projected O5 PSDs for each detector are sourced from the Pycbc.psd package [366].

As discussed in Section 4.3.2, waveform morphology can be significantly altered in the presence of non-negligible eccentricity, precessing spins, or HOMs. Nevertheless, Eq. (7.3) remains valid even for lensed waveforms exhibiting such phase modulations, provided that the JPE utilises waveform templates incorporating these features [157, 306, 367]. We underline that a JPE analyses two lensed signals simultaneously, enabling accurate inference of lensing observables, such as the Morse factors and their difference [367]. Consequently, $h^{\rm maxL}_{\rm SL,2}$ can be obtained from $h^{\rm maxL}_{\rm SL,1}$ in the frequency domain.

Similarly, we obtain the maximum-likelihood waveforms for the pairs of microlensed signals ($h^{\rm maxL}_{\rm MLSL,1}$, $h^{\rm maxL}_{\rm MLSL,2}$). The morphologies of $h^{\rm maxL}_{\rm MLSL}$ differ from those of $h^{\rm inj}_{\rm MLSL}$, as the templates used in the pipeline account only for strong lensing. This mismatch is expected to manifest as non-negligible residuals.

7.1.3 Residual Test

Since the adopted pipeline does not account for microlensing, the discrepancies between $h^{\rm maxL}_{\rm MLSL}$ and $h^{\rm inj}_{\rm MLSL}$ are expected to be significantly larger than those between $h^{\rm maxL}_{\rm SL}$ and $h^{\rm inj}_{\rm SL}$. To statistically quantify these differences between the maximum-likelihood waveforms retrieved under the two hypotheses, we investigate the residuals ($r \equiv d - h$), which represent the strain remaining after subtracting the signal model (h) from the observed data (d). If the true GW signal embedded in the data perfectly matches the model, the residual will be consistent with pure detector noise[4].

We assume that the data in the $i^{\rm th}$ detector, containing a lensed GW signal, can be expressed as:

$$d^i = h^i_{\rm SL} + {\rm d}h^i + n^i, \tag{7.4}$$

where $h^i_{\rm SL}$ is the purely strongly lensed signal excluding microlensing effects, ${\rm d}h^i$ represents the frequency-dependent distortion induced by microlensing, and n^i is the detector noise.

Under $\mathcal{H}_{\rm SL}$, the lensing effects are governed by $F_{\rm geo}$ and the waveform morphology remains unchanged, meaning the ${\rm d}h$ terms vanish. Consequently, the residual in the $i^{\rm th}$ detector is given by:

$$\begin{aligned} r^i|\mathcal{H}_{\rm SL} &\equiv d^i - h^{i,{\rm maxL}} \\ &= h^i_{\rm SL} + n^i - h^{i,{\rm maxL}} \\ &= \delta h^i_{\rm SL} + n^i, \end{aligned} \tag{7.5}$$

4. Non-Gaussian features, such as glitches, can occasionally overlap with a signal. In such cases, the residuals would not align with the stationary detector noise characterised at other times.

where $\delta h^i_{\rm SL}$ denotes the difference between the true lensed signal and the maximum-likelihood waveform recovered under $\mathcal{H}_{\rm SL}$. Since $h^{i,\rm maxL}$ is obtained by fitting the noisy data, rather than the noiseless signal, $\delta h^i_{\rm SL}$ does not generally vanish. In the idealised limit where the maximum-likelihood waveform exactly reproduces the observed data, $h^{i,\rm maxL} \to h^i_{\rm SL} + n^i$, and hence $\delta h^i_{\rm SL} \to -n^i$, giving $r^i \to 0$. In practice, for a correctly specified hypothesis and a waveform model, the residual should instead be statistically consistent with the instrumental noise.

Conversely, $\mathrm{d}h$ is non-zero under $\mathcal{H}_{\rm MLSL}$. In this case, the residual is given by:

$$\begin{aligned} r^i|\mathcal{H}_{\rm MLSL} &\equiv d^i - h^{i,\rm maxL} \\ &= h^i_{\rm SL} + \mathrm{d}h^i + n^i - h^{i,\rm maxL} \\ &= \delta h^i_{\rm MLSL} + n^i, \end{aligned} \tag{7.6}$$

where $\delta h^i_{\rm MLSL}$ denotes the difference between the true microlensed signal, $h^i_{\rm SL} + \mathrm{d}h^i$, and the maximum-likelihood waveform inferred under $\mathcal{H}_{\rm SL}$. Since the analysis pipeline uses templates that do not account for microlensing, this difference generally contains a contribution from the microlensing effect in addition to any waveform mismatch. As with $\delta h^i_{\rm SL}$, $\delta h^i_{\rm MLSL}$ does not necessarily remain non-zero, since the maximum-likelihood waveform is obtained by fitting the noisy data. However, if the microlensing contribution cannot be adequately absorbed by the $\mathcal{H}_{\rm SL}$ waveform model, it can leave a coherent structure in the residual that is inconsistent with the expected instrumental-noise distribution. A schematic representation of the expected residuals under each hypothesis is shown in Figure 7.1, illustrating that the residual under $\mathcal{H}_{\rm MLSL}$ is more pronounced than that under $\mathcal{H}_{\rm SL}$.

To evaluate the extent to which residuals deviate from pure detector noise (denoted by $\mathcal{H}_{\rm N}$) under a given signal hypothesis $\mathcal{H}_{\rm X}$, where X is either SL or MLSL, we analyse the cross-correlation between residuals from different detectors. We introduce a statistic, Λ, to represent this quantity. In [368], the cross-correlation of residuals—defined as the deviation between observed microlensed strain and the best-fit unlensed templates—was utilised to search for microlensing signatures in otherwise unlensed GW observations. Following the methodology in [347], these cross-correlations are computed by introducing a time shift τ to one residual relative to another. This approach assesses the coherence of the residual power across the detector network. We adapt this technique by calculating the inner product of the residuals from each detector to obtain Λ, which quantifies the alignment between the two signals.

For strain data from the $i^{\rm th}$ and $j^{\rm th}$ detectors with corresponding residuals r^i and r^j, Λ is defined as:

$$\Lambda_{\mathcal{H}_{\rm X}} = \exp\left[\frac{1}{2}\left\langle r^i(t)_{\mathcal{H}_{\rm X}}, r^j(t+\tau)_{\mathcal{H}_{\rm X}}\right\rangle\right], \tag{7.7}$$

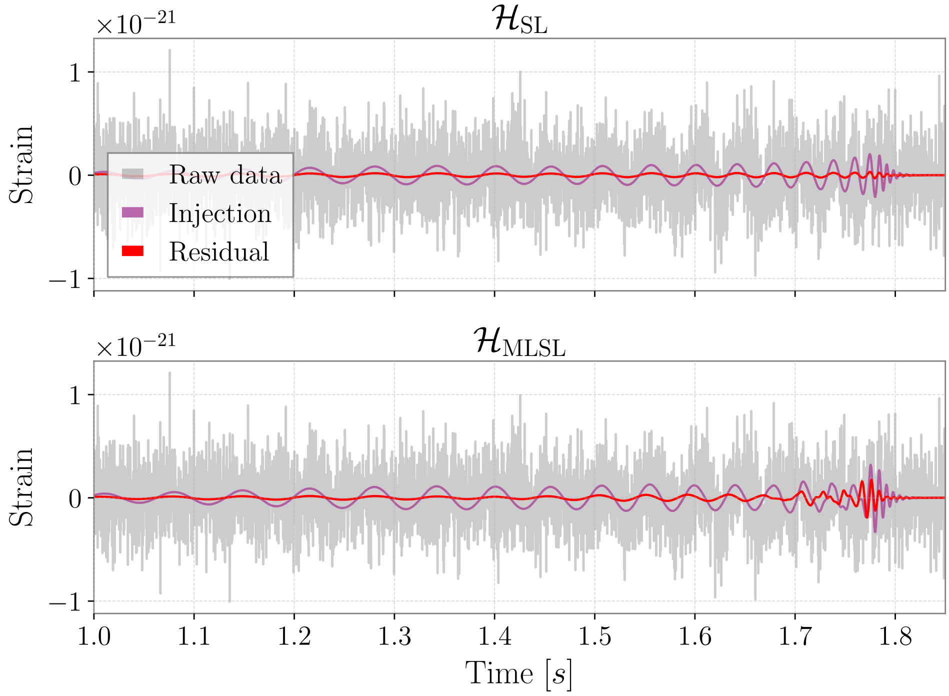

Figure 7.1: Time domain strains of raw data (grey solid lines), injected waveforms (purple solid lines), and residuals (red solid lines) for a representative BBH event in the LIGO-Hanford detector, assuming $\mathcal{H}_{\mathrm{SL}}$ and $\mathcal{H}_{\mathrm{MLSL}}$ in the top and bottom panel, respectively. We assume the same BBH and macrolens parameters for the two injected waveforms.

where τ is the relative time delay between the i^{th} and j^{th} detectors, and the inner product $\langle \cdot, \cdot \rangle$ is defined as in Eq. (1.41). When calculating Λ, we do not use a time window spanning the entire signal duration. Instead, we adopt a targeted time window of the form:

$$t \in [t_c - \alpha t_s + \Delta t, t_c + \Delta t], \tag{7.8}$$

where t_c is the coalescence time, t_s is the signal duration inferred from the JPE, and α is a scaling constant used to tune the window size. The term Δt ensures sufficient coverage of the ringdown phase; we set $\Delta t = 0.2$ s. In the subsequent analysis, we use $\alpha = 0.25$, as this value yields optimal performance for our residual-based diagnostic. A detailed discussion on the efficiency of the time window will be presented in Section 7.2.3.

Recall that the residual r is decomposed into a model-mismatch component δh and a detector noise component n. When evaluating Eq. (7.7), δh contributes directly to the inner product because it remains coherent across detector segments when aligned with the correct time shift τ. In contrast, the noise components n are uncorrelated between detectors due to their stochastic nature, and thus, their cross-contributions tend to average to zero. Consequently, a higher value of δh results in a higher value of Λ.

The distributions of Λ obtained under each hypothesis can be differentiated if their corresponding residuals are sufficiently distinct. In this study, the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution effectively serves as a background distribution, allowing us to determine which hypothesis is favoured for a given event by evaluating where its computed Λ value falls relative to this background. We consider an event to favour $\mathcal{H}_{\rm MLSL}$ if its $\Lambda_{\mathcal{H}_{\rm MLSL}}$ value exceeds the upper bound of the 2σ C.I.[5] of the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution.

It is important to note that, while the AdV detector is included in the initial parameter estimation, it is excluded from the calculation of Λ in Eq. (7.7). This is because AdV's noise level is generally higher than that of the LIGO detectors. Although a larger detector network typically improves the total network SNR, the inclusion of a high-noise detector can degrade the residual test performance by broadening the Λ distribution, thereby reducing sensitivity to microlensing signatures.

7.2 Results

7.2.1 Analysis using Realistic Population-based Events

From the JPEs conducted on the two independent sets of 100 pairs, simulated under the scenarios of pure strong lensing and superimposed microlensing (see Sec. 7.1.2), we obtain the maximum-likelihood waveforms (i.e., 200 $h_{\rm SL}^{\rm maxL}$ and 200 $h_{\rm MLSL}^{\rm maxL}$). Using Eqs. (7.5)-(7.7), we compute the values of $\Lambda_{\mathcal{H}_{\rm SL}}$ from the 200 $r_{\rm SL}$ and $\Lambda_{\mathcal{H}_{\rm MLSL}}$ from the 200 $r_{\rm MLSL}$. Note that the instrumental noise realisations n^i used for both $r^i_{\rm SL}$ and $r^i_{\rm MLSL}$ are identical. Additionally, we generate a set of 200 pure Gaussian noise segments (N) for each detector, based on the same PSDs used for the data generation, to compute the corresponding $\Lambda_{\mathcal{H}_{\rm N}}$ values.

Since the detector noise is assumed to be stationary and Gaussian, and the JPE framework accurately recovers $h_{\rm SL}$, the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution is expected to closely follow the $\Lambda_{\mathcal{H}_{\rm N}}$ distribution, which typically behaves as a normal distribution with a mean of 1 and unit variance. In contrast, the distribution of $\Lambda_{\mathcal{H}_{\rm MLSL}}$ is expected to deviate from $\Lambda_{\mathcal{H}_{\rm N}}$, indicating that the microlensing-induced terms $\delta h_{\rm MLSL}$ in Eq. (7.6) contribute more significantly to the residual than $\delta h_{\rm SL}$.

Figure 7.2 shows the Λ distributions for the residuals and Gaussian noises, along with their corresponding Gaussian kernel density estimates (KDEs). As previously noted, an event is identified as a candidate that is both strongly lensed and microlensed if its Λ value exceeds the upper bound of the 95 % (2σ) C.I. of the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution. For O4 sensitivity, 11 % of the 200 microlensed events favour $\mathcal{H}_{\rm MLSL}$. A primary factor contributing to this relatively low percentage is that most events in our realistic population exhibit small mismatch

5. All one-dimensional credible intervals computed in this work are central, meaning they contain equal probability mass above and below the bounds.

values (MM ≤ 0.03) relative to $h_{\rm SL}$ (see Eq. (6.2))[6]. In this chapter, the MM is defined as follows:

$$\begin{aligned}\mathrm{MM} &\equiv 1-\mathcal{M}\left(h_{\rm MLSL}, h_{\rm SL}\right) \\ &= 1-\max_{\phi_0, t_0} \frac{\langle h_{\rm MLSL}, h_{\rm SL}\rangle}{\sqrt{\langle h_{\rm MLSL}, h_{\rm MLSL}\rangle \langle h_{\rm SL}, h_{\rm SL}\rangle}}.\end{aligned} \tag{7.9}$$

Such events are intrinsically difficult to distinguish, even when using specialised microlensing pipelines like GRAVELAMPS [298]. In Section 7.2.2, we will discuss how the detection efficiency improves for high-mismatch events, which are expected to be more readily detectable.

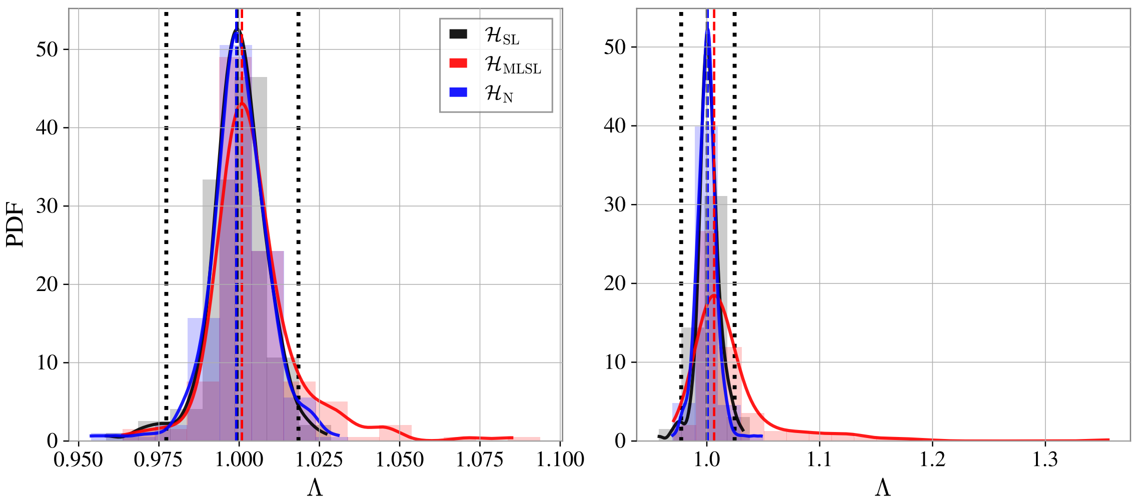

Figure 7.2: *Left panel*: The Λ distributions and their corresponding Gaussian KDEs for each of 200 **N** (blue), $\boldsymbol{r}_{\rm SL}$(black), and $\boldsymbol{r}_{\rm MLSL}$ (red) samples obtained using the O4 sensitivity. $\boldsymbol{r}_{\rm SL}$ and $\boldsymbol{r}_{\rm MLSL}$ have the same underlying noise strain. *Right panel*: The results obtained by employing the O5 sensitivity.

In addition to the O4 analysis, we repeat our study using the projected LIGO O5 sensitivity to quantify the improvement in our residual-based test. We anticipate enhanced performance if: (1) the $\delta h_{\rm SL}$ term in Eq. (7.5) is minimised through more precise parameter estimation, and (2) the reduced noise floor maximises the relative contribution of dh to the Λ statistic in Eq. (7.6). As expected, for the same population of strongly lensed and microlensed events, the percentage of signals favouring $\mathcal{H}_{\rm MLSL}$ increases to 21.5 % at O5 sensitivity.

6. MM values are computed by averaging the two MM values obtained using $h_{\rm SL}$ and $h_{\rm MLSL}$ detected in the aLIGO-Hanford and aLIGO-Livingston detectors.

To verify that more significantly microlensed events yield larger Λ values, we calculate the difference $\Delta\Lambda \equiv \Lambda_{\mathcal{H}_{\rm MLSL}} - \Lambda_{\mathcal{H}_{\rm SL}}$ for individual BBH events. This allows us to investigate the correlation between $\Delta\Lambda$ and both the MM and the maximum strain of $\delta h_{\rm MLSL}$—the latter representing the residual amplitude remaining after subtracting the best-fit strongly lensed template. A positive $\Delta\Lambda$ value indicates that $r_{\rm MLSL}$ is more significant than $r_{\rm SL}$.

As shown in Figure 7.3, microlensed events with larger δh are concentrated in the high $\Delta\Lambda$ region, consistent with the definition of Λ in Eq. (7.7). Similarly, events with relatively high MM values tend to exhibit higher $\Delta\Lambda$, although the correlation between $\Delta\Lambda$ and δh is stronger than that between $\Delta\Lambda$ and MM. In the O4 case, most high-MM events are located within the 2σ C.I. of the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution (black dotted lines), as the noise power is relatively high compared to the residual power. Consequently, the statistical fluctuations of $\Lambda_{\mathcal{H}_{\rm SL}}$ tend to be broader than those of $\Delta\Lambda$. In contrast, the reduced noise levels in the O5 case result in most high-MM events above the 2σ threshold.

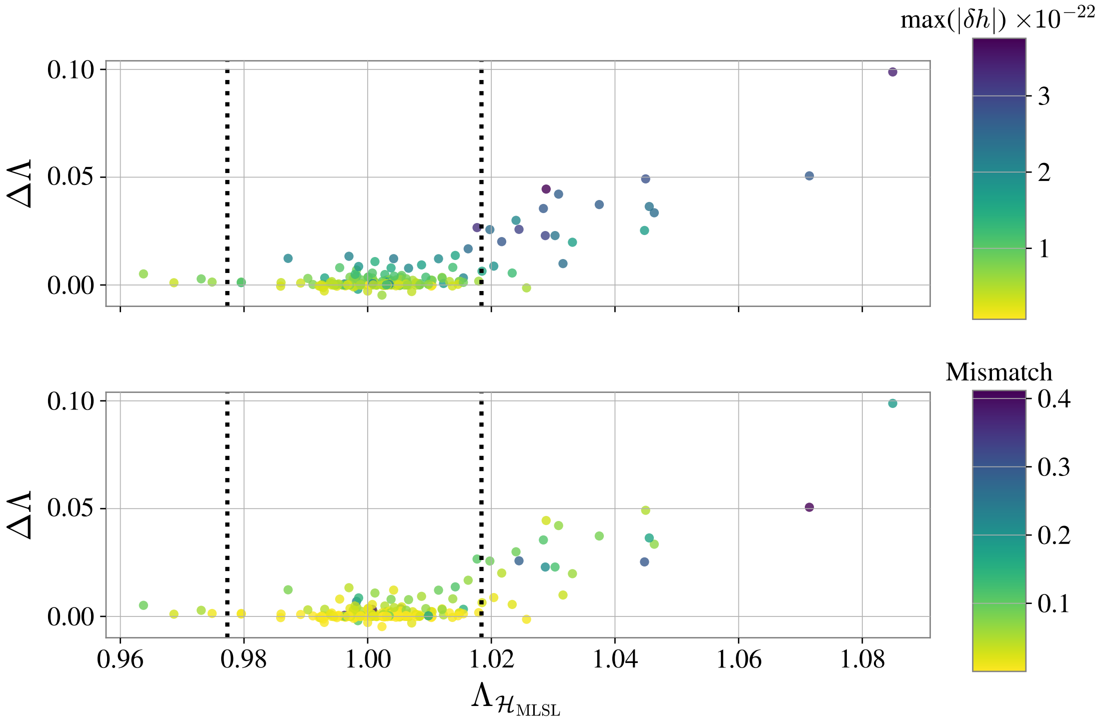

(a) O4 sensitivity

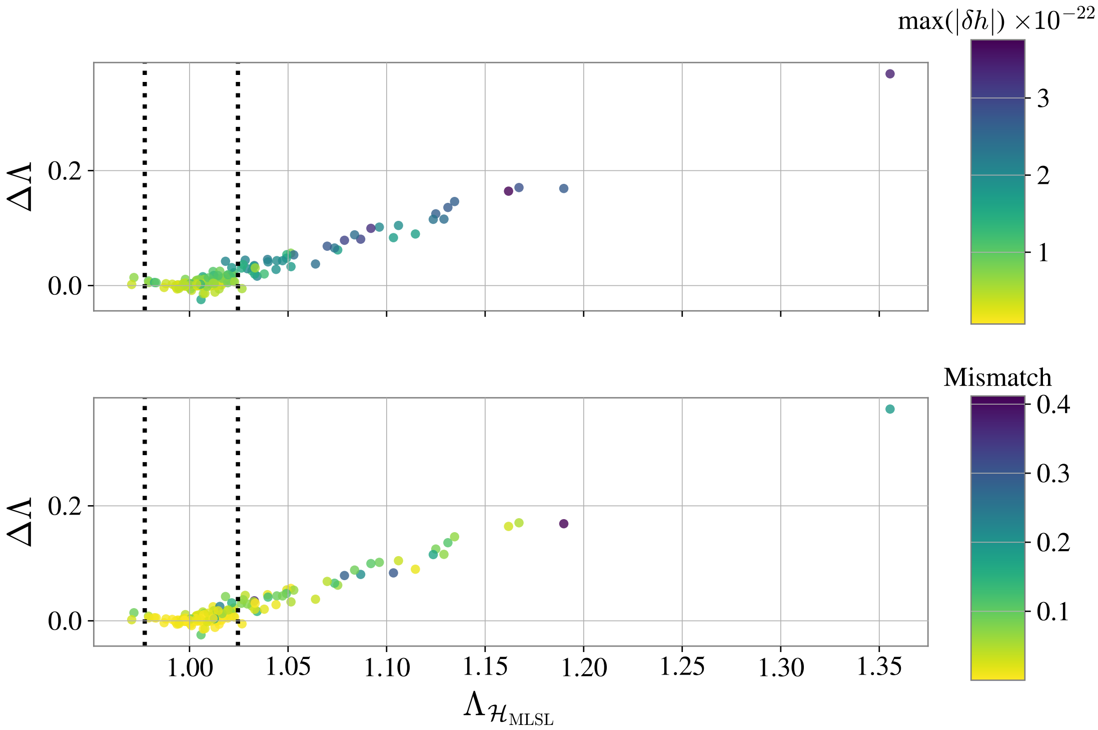

(b) O5 sensitivity

Figure 7.3: Scatter plots showing the correlation between $\Lambda_{\mathcal{H}_{\rm MLSL}}$ and $\Delta\Lambda$ for (a) O4 and (b) O5 sensitivities, colour-coded based on max($|\delta h|$) and mismatch, denoting the maximum strain of the pure residual and mismatch between $h_{\rm SL}$ and $h_{\rm MLSL}$ computed by averaging the values obtained from aLIGO-Hanford and aLIGO-Livingston, respectively.

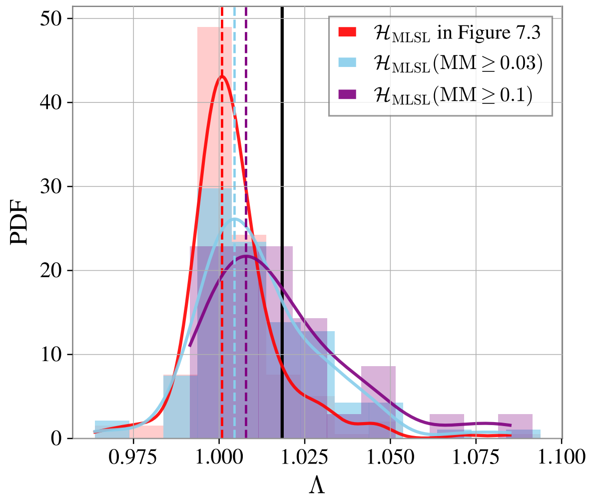

Figure 7.4: The same configuration as Figure 7.2, but now showing Λ distributions for realistic population-based microlensed events (red), events with MM ≥ 0.03 (cyan), and events with MM ≥ 0.1 (purple).

It is important to note that a high MM does not necessarily indicate a large δh. If a significantly distorted signal has a low SNR, the absolute value of δh remains low despite the high MM. Conversely, even for weakly distorted signals, if the time delay between microlensed images is very short, the resulting repeated merger part of the waveform may not be fully captured by the recovered best-fit waveform. This leads to a high δh and, consequently, a high $\Delta\Lambda$.

7.2.2 Analysis using High-mismatch Events

The majority of the simulated $h_{\rm MLSL}$ exhibit small MM values, which is consistent with expectations for a realistic microlensing population [315]. Therefore, we also focus on microlensed events with MM ≥ 0.03 and MM ≥ 0.1, which are more likely to be detectable with current pipelines, to better evaluate the effectiveness of our residual-based test. To obtain more robust results for the high-MM regime, we created an additional set of events with MM ≥ 0.3 by artificially enhancing the microlensing effects. Specifically, we increased the macrolens magnification to $\mu = 100$ to produce significantly microlensed signals that do not follow the realistic population distribution.

Figure 7.4 compares the Λ distributions for events with MM ≥ 0.03 and MM ≥ 0.1 against the realistic population-based distributions shown in Figure 7.2. For events with higher MM values, there is a clear increase in the proportion of signals identified as microlensed (i.e., those with Λ values exceeding the upper bound of the 2σ C.I. of the $\Lambda_{\mathcal{H}_{\rm SL}}$ distribution). This confirms that the residual test is a valid diagnostic for identifying microlensed events in the high-MM regime. A summary of the results is tabulated in Table 7.2.

O4 sensitivity			
Condition	N_{total}	N_{ML}	R_{ML}
Realistic population	200	22	0.110
Events with MM ≥ 0.03	94	26	0.277
Events with MM ≥ 0.1	35	12	0.343

O5 sensitivity			
Condition	N_{total}	N_{ML}	R_{ML}
Realistic population	200	43	0.215
Events with MM ≥ 0.03	94	49	0.521
Events with MM ≥ 0.1	35	23	0.657

Table 7.2: Summary of the results of residual test depicted in Figures 7.2 and 7.4. N_{total}, N_{ML}, and R_{ML} denote the total number of events satisfying the specified condition, the number of events discerned as a microlensed signal, and the ratio between the two numbers.

Furthermore, we constructed receiver operating characteristic (ROC) curves to assess the overall performance of our residual-based approach. As shown in Figure 7.5, all area under the curve (AUC) values exceed 0.5, indicating that the residual test effectively distinguishes between microlensed and purely strongly lensed GW signals. The AUC values improve with higher detector sensitivity and larger mismatch values, consistent with the trends presented in Table 7.2.

7.2.3 Effect of the Time Window on Λ

In Sections 7.2.1 and 7.2.2, Λ values were calculated using a specific time window of the form shown in Eq. (7.8), with $\alpha = 0.25$ and $\Delta t = 0.2\,\text{s}$, effectively excluding the early inspiral portion of the time-domain waveform. For BBH sources, the GW strain is lower during the early inspiral stage and reaches its maximum during the late inspiral and merger phases. Consequently, the early parts of the signal are more susceptible to noise fluctuations compared to segments closer to the merger. Our approach mirrors those adopted in tests of GR [353, 354, 355], where narrower time windows are generally employed to minimise noise contributions that are unrelated to potential non-GR features. However, a narrow window must be chosen carefully for microlensing, as the time delays between microimages vary. A window that is too restrictive risks excluding data segments, particularly in the post-merger phase, that may contain significant microlensing signatures. To address this, we investigate various time windows covering different portions of the GW signal to evaluate the impact of noise fluctuations on the Λ distributions.

A representation of the $\Lambda_{\mathcal{H}_{\text{MLSL}}}$ distributions obtained using various time windows is shown in Figure 7.6. The coloured segments (red, cyan, and purple) of the curves represent events correctly identified as microlensed, while the gray segments correspond to events identified as unmicrolensed. As shown, smaller values of α result in a higher proportion of

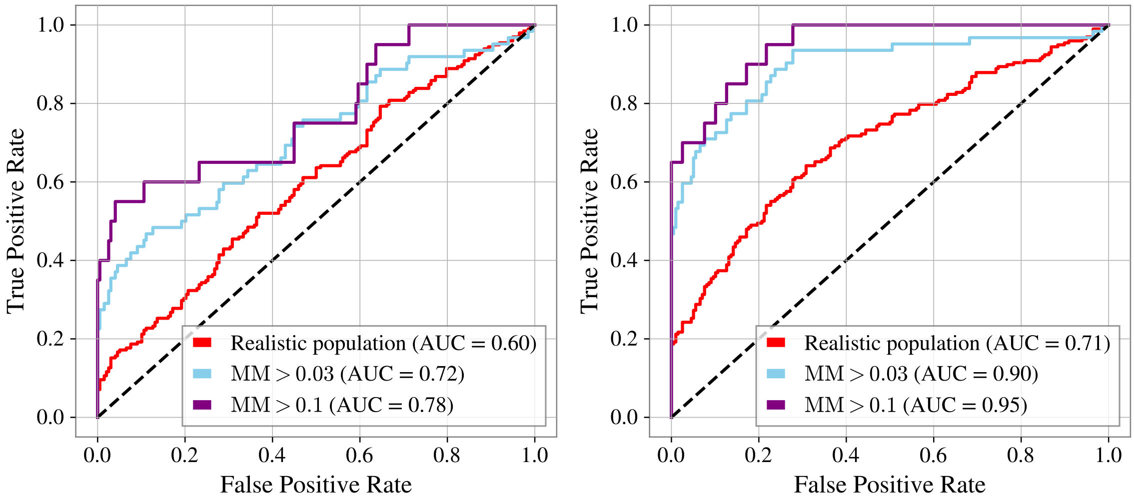

Figure 7.5: ROC curves for the residual test constructed by comparing each of the three event groups—$\Lambda_{\mathcal{H}_{\rm MLSL}}$ values of events drawn from a realistic population (red), events with MM ≥ 0.03 (cyan), and events with MM ≥ 0.1 (purple)—against the background distribution given by the $\Lambda_{\mathcal{H}_{\rm SL}}$ values shown in Figure 7.2 (*Left*: O4 sensitivity, *Right*: O5 sensitivity). The corresponding AUC values have also been estimated.

the $\Lambda_{\mathcal{H}_{\rm MLSL}}$ distribution being identified as microlensed. This trend is expected because GW microlensing is most prominent when the time delays between images are shorter than the chirp time, and the resulting signatures are more distinctive in the high-frequency regime. Since the merger and ringdown phases predominantly consist of high-frequency components, using a targeted time window centered on the merger rather than the entire signal duration enhances the efficiency of the residual test. The detection ratios ($R_{\rm ML}$) for each time window and mismatch condition are summarised in Table 7.3. Empirically, we find that a time window with $\alpha = 0.25$ yields the optimal efficiency for microlensing detectability.

7.3 Conclusion

Strongly lensed GW signals likely experience microlensing effects due to stellar fields embedded within the lens galaxy, which imprints complex patterns onto the GW waveforms. Current strong-lensing Bayesian inference pipelines are not designed to detect these distortions, as they typically assume repeated signals share an identical time-frequency evolution. However, incorporating both strong lensing and microlensing into a joint Bayesian analysis is computationally challenging due to the expanded parameter space and the complexity of calculating amplification factors for dense stellar fields.

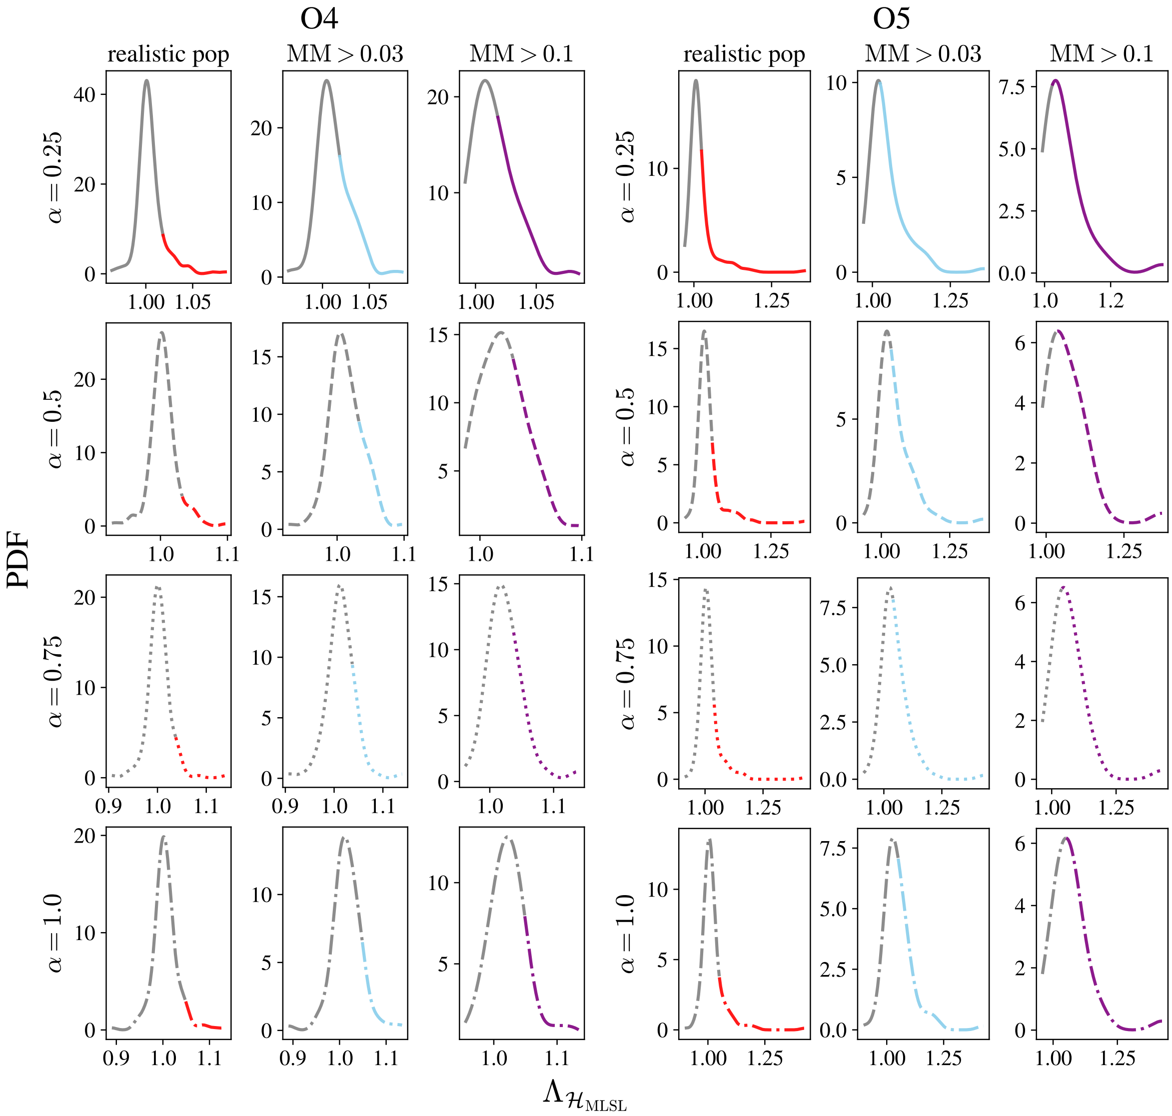

Figure 7.6: $\Lambda_{\mathcal{H}_{\rm MLSL}}$ distributions obtained using various time windows determined by a constant α and mismatch conditions (*Left*: O4 sensitivity, *Right*: O5 sensitivity). Coloured segments in the distributions represent microlensed events with $\Lambda_{\mathcal{H}_{\rm MLSL}}$ falling within the range identified as microlensed, whereas grey segments indicate microlensed events misclassified as unmicrolensed.

Our results highlight the potential of residual tests as a robust diagnostic for identifying microlensing signatures in SLGW signals. This approach can serve as a low-latency pre-analysis pipeline, enabling us to efficiently prioritise events for follow-up with more intensive, full-parameter-estimation microlensing pipelines. By focusing on the residuals—the difference between the recovered best-fit purely strongly lensed waveform and the observed data—we can uncover events with significant microlensing signatures that might otherwise remain overlooked in traditional analyses. The distribution of cross-correlation values for both strongly lensed and microlensed signals, compared with that for purely strongly lensed ones, suggests that microlensing effects can create statistically measurable deviations in residual tests.

O4 sensitivity				
Condition	$R_{\rm ML}$	$R_{\rm ML}^{\alpha=0.5}$	$R_{\rm ML}^{\alpha=0.75}$	$R_{\rm ML}^{\alpha=1.0}$
Realistic population	0.110	0.075	0.065	0.050
Events with MM ≥ 0.03	0.277	0.213	0.170	0.106
Events with MM ≥ 0.1	0.343	0.314	0.257	0.114

O5 sensitivity				
Condition	$R_{\rm ML}$	$R_{\rm ML}^{\alpha=0.5}$	$R_{\rm ML}^{\alpha=0.75}$	$R_{\rm ML}^{\alpha=1.0}$
Realistic population	0.215	0.160	0.155	0.140
Events with MM ≥ 0.03	0.521	0.447	0.447	0.404
Events with MM ≥ 0.1	0.657	0.657	0.629	0.571

Table 7.3: The same configuration as in Table 7.2, but showing the ratio of events identified as microlensed to the total number of events.

As discussed in earlier studies, however, certain waveform features such as spin precession and orbital eccentricity can mimic microlensing by introducing similar modulations [321, 331, 369], potentially leading to residuals being mistakenly attributed to microlensing. While we have neglected these effects in this study, as lensing of BBHs with extreme precession or eccentricity is expected to be rarer than the standard population [370, 371, 372], distinguishing between residuals arising from microlensing versus those from intrinsic source properties should be addressed in future work.

In our current cross-correlation framework, we utilised maximum-likelihood waveforms to obtain a single statistic per event. A more exhaustive approach would involve using the full posterior distribution of waveforms to generate a distribution of cross-correlation values for each event. One could then assess the degree of overlap between this distribution and the full cross-correlation distribution obtained from all events. Furthermore, while we used Gaussian noise realisations to establish a baseline, real detector noise often includes non-Gaussian features, such as glitches. Therefore, robust deglitching procedures must be applied before the residual test, although microlensing features may be inadvertently removed if the lens model is not incorporated into the cleaning.

It is also important to note that while we employed one of the state-of-the-art microlensing models and realistic parameter ranges as a baseline, the results could vary with different microlensing population models and macromodel magnifications. In general, higher macrolens magnification and higher density of microlenses lead to larger mismatches, thereby increasing the probability of detection. For further investigation into the role of the macromodel across various parameter spaces, we refer the reader to [312, 315, 323]. Beyond our baseline estimates for a realistic population, we also quantify our detection capability using selected high-mismatch waveforms (Table 7.2), allowing our results to be interpreted within the context of these specific distortions. A comprehensive investigation of diverse microlensing fields in conjunction with complex galaxy lens models is reserved for future work.

When focusing specifically on high-mismatch events, which are typically detectable by existing microlensing pipelines, the residual test demonstrates improved performance at both O4 and O5 sensitivities. However, a caveat remains: excessively distorted microlensed signals may suffer substantial SNR loss during the initial matched-filtering process if the templates do not account for microlensing effects [322]. In such cases, a dedicated search pipeline incorporating microlensed templates may be required to avoid missing these events entirely.

When considering the entire dataset containing microlensed events in both low- and high-mismatch regimes, our results suggest that aLIGO O4 sensitivity may not be sufficient to reliably detect most realistic microlensing effects. However, the projected improvements for O5 are expected to significantly enhance our ability to discern these signatures. In the era of next-generation detectors, where the detection rate of lensed GW signals will be significantly higher, such low-latency methods will become increasingly important. We anticipate our approach to be even more efficient for next-generation detectors due to the longer signal durations, higher SNRs, and improved sensitivity in the high-frequency regime. A detailed quantitative study of the efficiency in these future scenarios will be conducted in future work.

Finally, the computational advantage of our method is substantial. A single JPE followed by the residual test for a pair of lensed events takes approximately half a day, as a single likelihood evaluation in our approach requires only a few milliseconds. In contrast, using a frequency-domain source model that explicitly incorporates both strong and microlensing effects in a standard search pipeline (e.g., within a nested sampling framework) results in likelihood evaluations taking approximately 10 seconds per waveform. This makes the full analysis roughly 10^4 times slower than our method. Consequently, the residual test presented in this chapter has the potential to serve as a high-speed, low-latency pre-analysis tool, alerting us to potential microlensing candidates before committing to more computationally intensive investigations.

Part III

Cosmological and Population Inference

Chapter 8

Lens Reconstruction for Delensing

Disclaimer

The following chapter is based on two papers: *Gravitational Lensing aided Luminosity Distance Estimation for Compact Binary Coalescences* [373], published in *Physical Review D*, and *Inferring Properties of Dark Galactic Halos using Strongly Lensed Gravitational Waves* [311], published in the *Astrophysical Journal*. In the former work, the author served as a co-first author and contributed to the development of the analysis framework, data analysis, manuscript preparation, and generation of all figures. In the latter work, the author was the first author and contributed to the development of the methodology, data analysis, interpretation of the results, scientific discussion, manuscript preparation, and generation of all figures.

In microlensing analysis, a lens model is typically assumed a priori, and the lens parameters are inferred simultaneously during parameter estimation. While some correlations exist, the source parameters are not completely degenerate with the lens parameters. In contrast, the analysis of detected SLGW signals, which often relies on model-independent comparisons of the relative magnification factors, time delays, and phase shifts between multiple images, faces a fundamental challenge: the inferred luminosity distance of each image is completely degenerate with the lensing magnification. Consequently, the measured distances do not represent the *true* luminosity distance to the source.

This degeneracy poses a critical issue for GW cosmography. As discussed in Section 3.2.2, lensing introduces a systematic bias in the inferred luminosity distance. Consequently, any cosmological inference based on this biased distance, such as measurements of the Hubble constant (H_0), will also be biased, even if the source redshift is precisely determined from EM counterparts. Therefore, in the study of SLGWs, it is imperative to perform a lens reconstruction to *delens* the signals and recover the true luminosity distance. By accurately modelling the mass distribution of the lens responsible for multiple imaging, the distance–magnification degeneracy can be broken or mitigated. This approach enables simultaneous constraints on lens properties and the recovery of unbiased luminosity distances required for precision cosmology.

In this chapter, we first present a proof-of-principle study of the delensing procedure using analytic lens models, including the PM and SIS models described in Section 2.2.1 and Section 2.2.2, respectively. In these simplified cases, lens reconstruction and the recovery of the true luminosity distance can be performed using closed-form analytical relations. We specifically examine how SLGW detection can improve the precision of luminosity distance measurements compared to unlensed signals, thereby strengthening constraints for GW cosmography. We then extend the analysis to more realistic scenarios. We explore numerical approaches to lens reconstruction under more general lens models in Sections 2.2.3–2.2.6, where the mapping between lens parameters and lensing observables is no longer one-to-one, precluding closed-form solutions.

8.1 Analytic Delensing with PM and SIS Models

For the PM and SIS models, within the geometric-optics approximation, the relevant lens parameters are the redshifted lens mass and the dimensionless impact parameter, $\theta_l = (M_l^z, y)$. The SIS model provides a simple and widely used description of galaxy-scale lenses in the strong-lensing regime, and is therefore well-motivated for studies of SLGWs. In contrast, the PM model is not generally representative of galactic-scale lenses, although it can serve as an effective approximation in simplified configurations where the mass distribution is highly centrally concentrated. Despite this, we adopt the PM model alongside the SIS model for its analytical simplicity and its utility as a baseline model, as both models admit closed-form relations between lensing observables and lens parameters.

In this section, we present an analytic delensing procedure and quantify the reduction in luminosity distance uncertainty by using mock BBHs lensed by the two lens models, constructed to resemble real BBHs selected from GWTC-3 [374].

8.1.1 Effective Luminosity Distance

As described in Eqs. (2.35) and (2.43), the PM lens produces two images for any non-zero y, whereas the SIS lens yields two images only for $y < 1$.

For both the PM and SIS models, the magnification of each image is uniquely determined by the impact parameter y. As a result, the effective luminosity distances defined in Eq. (2.75) can be directly expressed as a function of y. For the PM lens, this relation is given by:

$$d_{\mathrm{L},\pm}^{\mathrm{PM}} = \frac{d_{\mathrm{L}}}{\sqrt{\frac{1}{2} \pm \frac{y^2+2}{2y\sqrt{y^2+4}}}}, \tag{8.1}$$

Table 8.1: Injection parameters and SNR values of three BBHs similar to GW200208_222617, GW200216_220804, and GW200219_094415.

Event	m_1 $[M_\odot]$	m_2 $[M_\odot]$	χ_{eff}	d_{L} [Gpc]	SNR
GW200208_222617	51.0	12.3	0.45	4.1	7.94
GW200216_220804	51.0	30.0	0.10	3.8	10.35
GW200219_094415	37.5	27.9	−0.08	3.4	9.86

while for the SIS lens it takes the form:

$$d_{\text{L},\pm}^{\text{SIS}} = \frac{d_{\text{L}}}{\sqrt{1 \pm \frac{1}{y}}}, \tag{8.2}$$

where the $+$ and $-$ signs denote the first-arriving (positive-parity) and second-arriving (negative-parity) images, respectively.

8.1.2 Parameter Estimation Framework

We simulate BBHs based on selected interesting candidates from the GWTC-3 catalogue [374], including GW200208_222617 (a low mass-ratio candidate; $q_m = m_2/m_1$), GW200216_220804 (an eccentric candidate), and GW200219_094415 (a high-spin candidate). These systems are chosen to span different regions of the binary parameter space, and to contrast with the equal-mass BBH case typically used as a baseline. In particular, such systems tend to exhibit more complex waveform morphology (see Section 4.3.2), which can make parameter estimation more challenging and may lead to larger uncertainties in inferred source properties, including the luminosity distance. We adopt the median values of the component masses (m_1, m_2), effective spin χ_{eff} and luminosity distance d_{L}, which are summarised in Table 8.1.

We generate both injection and template GW waveforms using the IMRPhenomXPHM model [288] to avoid systematic biases arising from waveform mismatches. For each simulated unlensed signal $(h_u(f))$, we perform parameter estimation using the `bilby` library [240] and the `dynesty` [263]. This setup is labeled as *normal PE* in Figure 8.1. We assume a detector network consisting of aLIGO-Hanford (H), aLIGO-Livingston (L), and AdV (V), referred to as the HLV network. To achieve optimal distance estimation under ideal sensitivity conditions and to assess the impact of detector noise and performance, we consider two configurations:

- **Case A**: Design PSD [365] with zero-noise realisations.
- **Case B**: O3a PSDs with Gaussian noise, where the O3a PSDs correspond to the first three months of the third observing run[1].

1. The corresponding PSD data for H, L, and V can be found at `https://dcc.ligo.org/LIGO-T2000012/public`

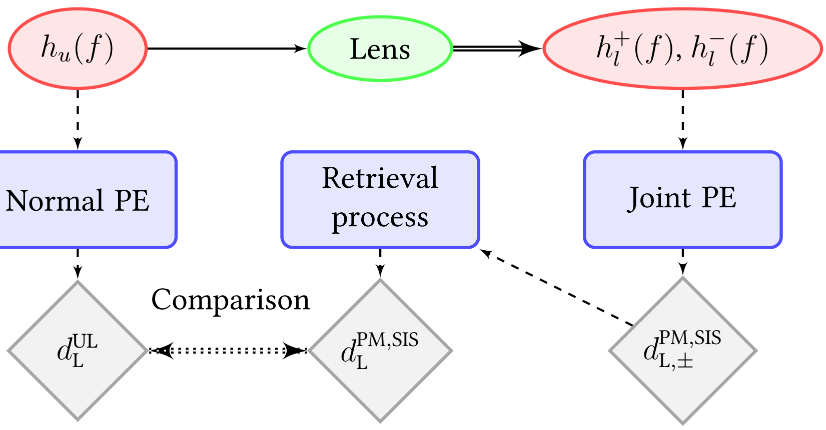

Figure 8.1: Flowchart of parameter estimation for an unlensed signal and its corresponding strongly lensed pair constructed from it. The posterior of $d_{\mathrm{L}}^{\mathrm{UL}}$ is first inferred from the unlensed signal. The posteriors of $d_{\mathrm{L}}^{\mathrm{PM,SIS}}$ for the two lensed signals are then obtained via joint parameter estimation and retrieval. Finally, the inferred distance posteriors are compared to assess the precision of the posterior on the true luminosity distance.

For the luminosity distance prior, we assume a uniform distribution in Euclidean volume, corresponding to $p(d_{\mathrm{L}}) \propto d_{\mathrm{L}}^2$, i.e., a constant source density in Euclidean space (the low-redshift limit of a comoving volume prior), over the range $[0.1, 10]$ Gpc. For other source parameters, we adopt the default precessing BBH priors implemented in `Bilby` [240, 375].

Next, we simulate mock lensed signals by assuming that each unlensed signal is lensed by both a PM and an SIS lens, producing two lensed signals per model (i.e., three unlensed signals produce a total of six lensed signals).

As discussed earlier, we restrict the dimensionless impact parameter to the range $y \in [0.1,\ 1.0)$, ensuring the formation of at least two images. The lower bound is motivated by the expected occurrence rate of SLGWs [301, 324, 325, 326, 376]. As an optimistic scenario for strong lensing, we consider the regime of maximum amplification, which occurs at $y = 0.1$ for the lens models considered here; accordingly, we set $y = 0.1$ in our injections. We also adopt a lens mass of $M_{lz} = 10^{11.5}\, M_{\odot}$, corresponding to typical galactic-scale lenses. This choice yields time delays between the two lensed images ranging from weeks to months over the considered range of y, consistent with expectations for galaxy-scale strong lensing.

To construct lensed signal pairs, we use the same intrinsic source parameters as in the unlensed case, except for the luminosity distance, which is modified by lensing magnification. For example, the chirp mass is assumed to be identical for all images, i.e., $\mathcal{M}_{c,+} = \mathcal{M}_{c,-} = \mathcal{M}_c$. The sky locations of the lensed images are also taken to be identical, as current detector sensitivities do not enable us to resolve small angular separations between

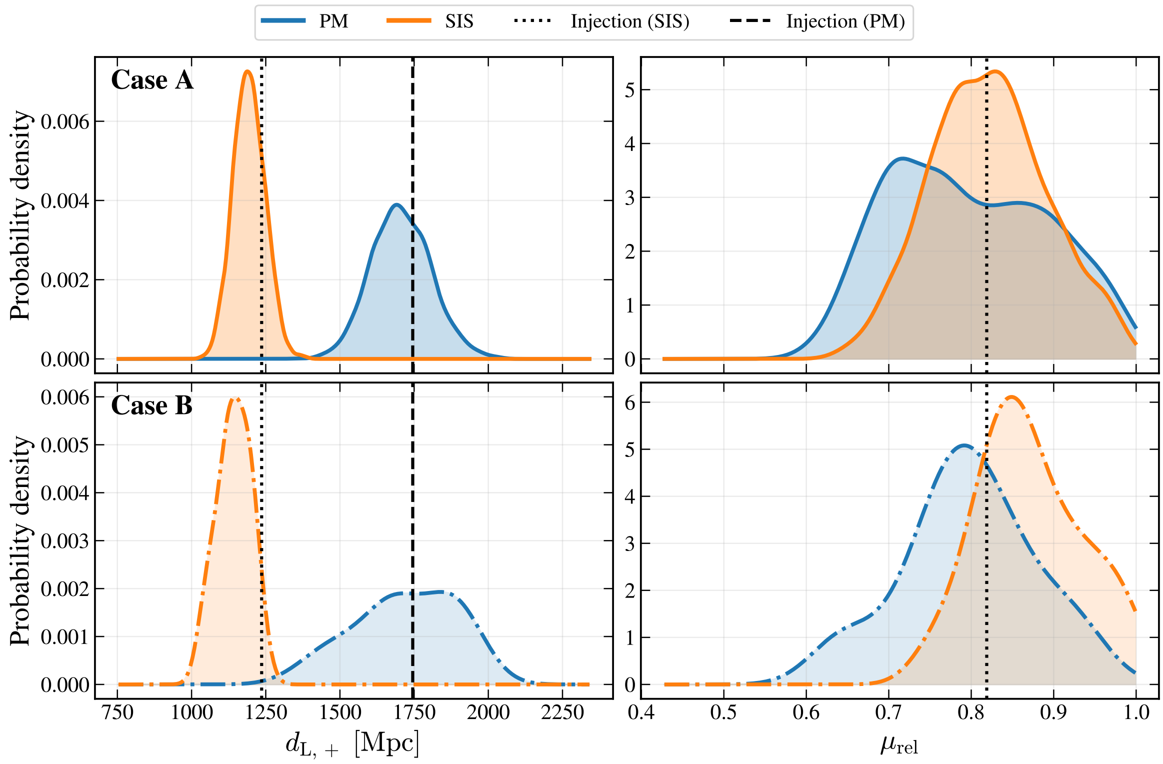

Figure 8.2: Posteriors of $d_{L,+}$ and μ_{rel} obtained from JPEs of a pair of signals from GW200208_222617 lensed by SIS (orange) and PM (blue) models. Black dotted and dashed lines indicate the injected values of $d_{L,+}$ for the SIS and PM cases, respectively. The injected values of μ_{rel} are nearly identical for the two lens models ($\mu_{rel}^{PM} \simeq \mu_{rel}^{SIS} \simeq 0.818$); therefore, only a single vertical line is shown for μ_{rel}.

images (e.g., [374]). We do not include time delays or phase shifts in the analysis, as they do not directly affect luminosity distance estimation. The prior assumptions for the lensed signals are taken to be the same as those used for the unlensed case, except that we adopt a uniform prior on the relative magnification, μ_{rel} [300].

For Cases A and B, we perform JPEs on the lensed signal pairs for each lens configuration using the method described in Section 5.2.1, implemented via the `Golum` pipeline [300, 303] (labeled as *joint PE* in Figure 8.1). While the full analysis infers both source parameters and lensing observables, we focus on the relative magnifications and the luminosity distances, which are the quantities of interest in this work.

8.1.3 Results

From the JPE of the lensed signal pairs, we obtain posteriors for the relative magnification factors and effective luminosity distances for both the PM and SIS models under Cases A and B. Figure 8.2 shows the resulting posteriors of μ_{rel} and $d_{L,\pm}^{PM,SIS}$ for each pair of lensed signals from the GW200208_222617-like event as a representative example. The top and bottom panels correspond to Cases A and B, respectively. The posteriors $p(\mu_{rel})$ are consistent with the injected values for both lens models and are nearly identical, as $y = 0.1$ yields $\mu_{rel} \simeq$

0.8 for both models. In contrast, the distance posteriors differ between the models. The distributions $p(d_{\mathrm{L},+})$ inferred under the SIS model are narrower than those obtained with the PM model within the 99 % C.I., indicating improved precision in parameter estimation. This difference arises from the higher μ_+, and consequently higher SNR, in the SIS model ($\mu_+^{\mathrm{SIS}} = 11$, $\mathrm{SNR}_{\mathrm{SIS}} \simeq 26.34$) compared to the PM model ($\mu_+^{\mathrm{PM}} \simeq 5.5$, $\mathrm{SNR}_{\mathrm{PM}} \simeq 18.66$). In addition, the presence of noise in Case B leads to a noticeable reduction in the precision of the luminosity-distance measurement, particularly for the PM case. For the SIS case, however, the effect appears less significant due to its substantially higher SNR.

Following [300], we improve parameter constraints by reweighting the posterior samples of the individual lensed signals. The reweighted posterior is constructed by combining the posteriors of the lensing and apparent source parameters (see Eq. (15) of [300]). For both models, the first-arriving (+) image is more strongly magnified and is therefore used as the reference. In this procedure, the posterior $p(d_{\mathrm{L},+})$ is refined using information from $p(d_{\mathrm{L},-})$.

To recover the true luminosity distances, we convert the reweighted effective distances using μ_+ posteriors, which are obtained analytically from the reweighted μ_{rel} posteriors (denoted as *retrieval process* in Figure 8.1). Using Eqs. (2.36) and (2.44), we obtain the posterior of y, which is then mapped to μ_+ via Eqs. (2.35) and (2.43). To illustrate an optimal scenario in which only the uncertainty in $d_{\mathrm{L},+}$ is taken into account, we estimate the true luminosity distance using a fixed μ_+ corresponding to its maximum-likelihood value, without propagating the uncertainty in μ_+. For Cases A and B, the retrieved luminosity distance distributions for the lensed pair from the GW200208_222617-like event are shown in Figure 8.3, together with the posterior of the true distance obtained from parameter estimation of the corresponding unlensed signal ($d_{\mathrm{L}}^{\mathrm{UL}}$) for comparison.

To quantify how much the recovered luminosity distances are improved in precision relative to the unlensed case, we compare the widths of the 99 % C.I. of the recovered distances posteriors, denoted by $\mathcal{W}_{99}$. We define the ratio $\mathcal{R}_{99} \equiv \mathcal{W}_{99}^{\mathrm{UL}}/\mathcal{W}_{99}^{\mathrm{X}}$, where X denotes one of the adopted O3b-like BBHs, labeled here using the shorthand event identifiers 208, 216, and 219. For convenience, we use abbreviated event labels obtained by dropping the first three and last six digits of the original event IDs in the superscripts of $\mathcal{W}_{99}$ and $\mathcal{R}_{99}$. Table 8.2 summarises the resulting $\mathcal{W}_{99}$ and $\mathcal{R}_{99}$ values for the three mock BBHs considered in this section.

The results are broadly consistent across all BBHs: the luminosity-distance posteriors inferred from lensed signals are more tightly constrained than those obtained from the corresponding unlensed signals. Differences in $\mathcal{W}_{99}$ and $\mathcal{R}_{99}$ values for the three BBHs are mainly due to the SNRs of the original unlensed events since identical lens parameters are adopted for each lens model. For example, the GW200208_222617-like event has the lowest

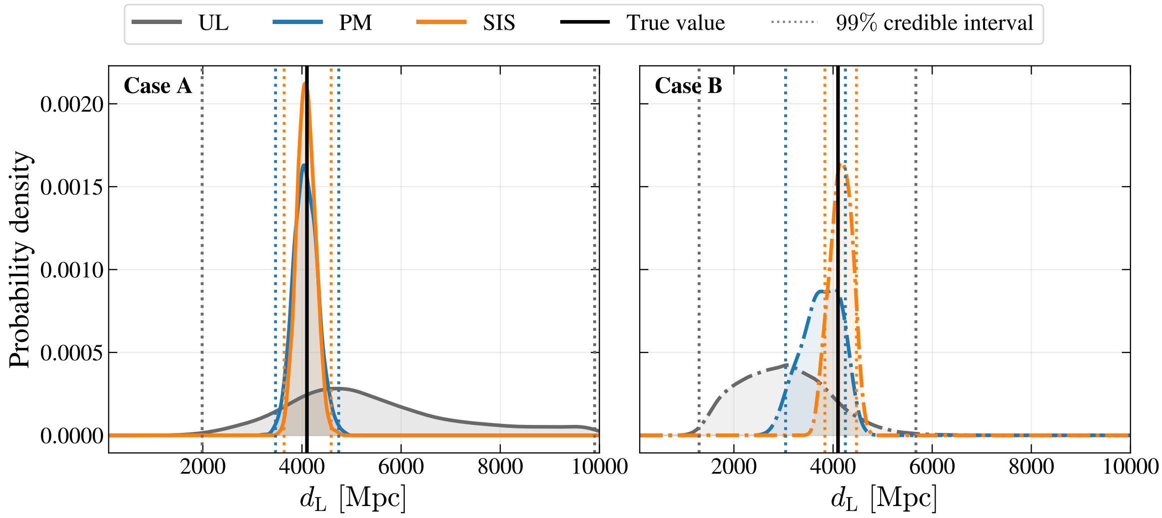

Figure 8.3: Posteriors of the recovered true luminosity distance (d_L) of the GW200208_222617-like event for the PM (blue) and SIS (orange) cases. For comparison, the distance posterior obtained from the corresponding unlensed signal, $p(d_L^{UL})$ (grey), is also shown. The left and right panels correspond to the results from Cases A and B, respectively. The black solid line indicates the true value. Coloured vertical dashed lines denote the lower and upper bounds of the 99 % C.I for each posterior distribution in the corresponding colour.

unlensed network SNR among the three systems considered in Case A. Correspondingly, its unlensed distance posterior exhibits the broadest C.I., with $\mathcal{W}_{99} \simeq 8$ Gpc. As a result, the impact of strong lensing appears most pronounced for this event, yielding $\mathcal{R}_{99}^{208} \sim 8$, compared to $\mathcal{R}_{99} \sim 4$ for the other BBHs.

As expected, the presence of noise tends to reduce the improvement from strong lensing, since the SNR enhancement from magnification and from combining multiple repeated signals becomes less effective. Nevertheless, even in the presence of Gaussian noise, the recovered luminosity-distance precision remains improved by a factor of approximately two to seven relative to the unlensed case (see the Case B results in Table 8.2). Overall, these proof-of-principle results demonstrate the potential of strong lensing to improve luminosity-distance measurements by providing multiple magnified observations of the same GW source, leading to tighter distance constraints

8.2 Beyond Simple Lens Models: Numerical Reconstruction

As shown in the previous section, delensing based on successful lens reconstruction can significantly enhance the precision of luminosity distance estimates. However, for more realistic lens models, the lens parameters cannot be determined analytically from lensing observables. In this section, we present a numerical approach to infer lens parameters from

Case A: HLV design PSDs with zero-noise realisation

Signal	$\mathcal{W}_{99}^{208}$ [Mpc]	$\mathcal{R}_{99}^{208}$	$\mathcal{W}_{99}^{216}$ [Mpc]	$\mathcal{R}_{99}^{216}$	$\mathcal{W}_{99}^{219}$ [Mpc]	$\mathcal{R}_{99}^{219}$
Unlensed	7,918	1.00	4,806	1.00	4,167	1.00
Lensed (PM)	1,274	6.21	1,786	2.69	1,146	2.68
Lensed (SIS)	948	8.35	1282	3.75	1,147	3.63

Case B: HLV O3a PSDs with a Gaussian-noise realisation

Signal	$\mathcal{W}_{99}^{208}$ [Mpc]	$\mathcal{R}_{99}^{208}$	$\mathcal{W}_{99}^{216}$ [Mpc]	$\mathcal{R}_{99}^{216}$	$\mathcal{W}_{99}^{219}$ [Mpc]	$\mathcal{R}_{99}^{219}$
Unlensed	4,373	1.00	7,675	1.00	3,961	1.00
Lensed (PM)	1,206	3.63	1,882	4.08	2,133	1.86
Lensed (SIS)	637	6.86	1,693	4.53	1,437	2.76

Table 8.2: The $\mathcal{W}_{99}$ and $\mathcal{R}_{99}$ values that are obtained from GW200208_222617-, GW200216_220804-, and GW200219_094415-like BBHs. Results from Cases A and B are compared for each injection.

GW signals using the rejection sampling technique. This method can be applied to a wide range of lens models once the relevant lens parameters are defined. We also explore how the reconstruction can be further improved by incorporating multi-messenger information, such as the identification of the host galaxy of the GW source and the lens galaxy.

8.2.1 Analysis Setup

We simulate nine distinct GW lensing systems based on the lens models described in Section 2.2. The parameters are randomly chosen from a pool of $\boldsymbol{\theta}_l$ that can generate the specified detectable multiple images. To maintain consistency, we use a single intrinsic unlensed signal from a GW150914-like event, generated with the IMRPhenomXPHM waveform model [288]. For most lensing configurations, we assume a luminosity distance of $d_\mathrm{L} = 3$ Gpc and a redshifted lens mass of $M_l^z = 10^{12}\ M_\odot$. However, for the NFW and PNFW models, the strong lensing cross-sections for sources at low redshifts are significantly smaller than the other models (see Figure 2 in [311]). To ensure a non-negligible lensing probability for the two models, we instead inject the signals at $d_\mathrm{L} = 10$ Gpc with a higher redshifted lens mass of $M_l^z = 10^{13}\ M_\odot$.

The resulting lensing configurations span from double-image to quintuple-image systems, with the image multiplicity determined by model-dependent lens parameters. Simulated lensed GW signals are injected into Gaussian noise characterised by the projected LIGO-Virgo O4 PSD [365]. Consistent with the projected detectability of the detector network [365], we assume that lensed GW signals ($h_l(f)$) are detected with a network SNR $\rho_\mathrm{net} > 8$ except for NFW and PNFW cases. An exception arises for Type III images in NFW and PNFW models, which are strongly demagnified by the lens geometry and often yield lower SNRs. For each system, the injected lens parameters ($\boldsymbol{\theta}_l$) are listed in Table 8.3. The

ID	Model	Injected $\boldsymbol{\theta}_l$	No. of images
1	PM	$y = 0.22$	Double
2	SIS	$y = 0.22$	Double
3	NFW	$(y,\ \kappa_s) = (0.004,\ 0.167)$	Triple
4	SIE	$(q,\ y) = (0.85,\ 0.22)$	Double
5	SIE	$(q,\ y) = (0.63,\ 0.22)$	Quadruple
6	SPEMD	$(q,\ y,\ \gamma) = (0.85,\ 0.22,\ 1.5)$	Double
7	SPEMD	$(q,\ y,\ \gamma) = (0.58,\ 0.17,\ 0.75)$	Quadruple
8	PNFW	$(q,\ y,\ \kappa_s) = (0.6,\ 0.011,\ 0.167)$	Triple
9	PNFW	$(q,\ y,\ \kappa_s) = (0.82,\ 0.008,\ 0.167)$	Quintuple

Table 8.3: Injected $\boldsymbol{\theta}_l$ values used to simulate the nine GW lensing systems considered in this section.

recovered SNRs vary across systems: Systems 1, 2, and 4–7 yield $\rho_{\text{net}} \in [11, 16]$, while Systems 3, 8, and 9 fall within $\rho_{\text{net}} \in [4, 10]$. For highly demagnified images, we adopt an SNR threshold of $\rho_{\text{net}} = 4$, assuming such signals can be identified with dedicated sub-threshold search pipelines [325, 326, 377].

To conduct JPEs for each lensing system, we again use the `Golum` pipeline [300, 303]. As this pipeline is optimised to analyse two signals, the estimations are performed on pairs of the first lensed signal ($h_{l,1}$) and the j-th lensed signal ($h_{l,j}$), which means that we can get posteriors of $\mu_{\text{rel},1j}$ and Δt_{1j}, where $j = 1, 2, ..., n$ if the system has n lensed GW signals. We assume uniform prior distributions for $\mu_{\text{rel},1j}$ and Δt_{1j}. For the source parameters, we adopt the prior distributions for precessing BBHs in [375]. Priors for the main parameters used in this work are shown in Table 8.4.

Parameter	Prior	Range
m_1, m_2	Uniform	[5, 100]$M_\odot$
d_L	UniformComovingVolume	[0.1, 5]Gpc OR [0.1, 10]Gpc
RA	Uniform	[0, 2π]
DEC	Isotropic	-
μ_{rel}	Uniform	[0.1,10]
Δt	Uniform	-

Table 8.4: Priors on the main source parameters, including component masses (m_1, m_2), luminosity distance (d_{L}), and sky location (RA, DEC), as well as lensing observables (μ_{rel}, Δt) employed in this work. A broader prior on d_{L} ([0.1, 10] Gpc) is adopted for NFW and PNFW systems to account for injected effective distances, $d_{\text{L}}^{\text{eff}} > 5\text{Gpc}$.

8.2.2 Lensing Observable Catalogue and Rejection sampling

A key challenge in lens reconstruction is that, for most lens models beyond simple cases such as PM and SIS, the mapping between the lensing observables and the lens parameters is not one-to-one. In addition, the observables are inferred as posterior distributions rather than precise point estimates, and their associated uncertainties further hinder accurate recovery of the underlying lens properties.

To address this, we employ a rejection sampling technique [254] to infer $\boldsymbol{\theta_l}$. This approach is particularly well-suited for our framework because, while the mapping from observables to parameters is not analytically invertible, forward modelling allows us to efficiently compute $\mu_{\rm rel}$ and Δt for any given set of $\boldsymbol{\theta_l}$. By comparing these forward-modelled predictions with the likelihood implied by the inferred posteriors of the observables, we effectively sample the posterior distribution of $\boldsymbol{\theta_l}$.

For each aforementioned lens model, we select sets of lens parameters $\boldsymbol{\theta_l}$ that produce multiple image solutions (i.e., at least two images). For each $\boldsymbol{\theta_l}$, we calculate the corresponding relative magnification factors ($\mu_{{\rm rel},1j}$) and Fermat potential (ΔT_{1j}) using LENSTRONOMY [378]. Note that the physical time delay is given by $\Delta t = D_{\Delta t} \Delta T$, where $D_{\Delta t}$ does not affect the image configuration in the geometrical optics limit. The resulting parameter–observable combinations are stored as ordered triplets, ($\Theta = \{\boldsymbol{\theta_l}, \mu_{\rm rel}, \Delta T\}$) forming *lensing observable catalogues.*

We use the constructed catalogues for the detected lensed signals in each system listed in Table 8.3. For simplicity, we assume that the correct lens model is known for each system. Figure 8.4 illustrates an example using the posterior distribution $P(\mu_{\rm rel})$ from lensing system ID 2 (see Table 8.3). In the rejection sampling procedure, candidate values of $\mu_{\rm rel}$, drawn from the catalogue triplets Θ, are first proposed according to a distribution $Q(\mu_{\rm rel})$ (solid blue line). Each sample is then accepted with probability proportional to the ratio $P(\mu_{\rm rel})/MQ(\mu_{\rm rel})$, where M is chosen such that $P(\mu_{\rm rel}) \leq MQ(\mu_{\rm rel})$ (solid yellow line) over the domain, ensuring that the accepted samples populate the region under the target distribution $P(\mu_{\rm rel})$ (solid black line). Green and magenta dots indicate the accepted and rejected Θ, respectively.

Finally, the accepted values of $\boldsymbol{\theta_l}$ effectively represent the posterior probability distributions of the lens parameters. Furthermore, the associated time-delay samples, ΔT, can be converted into the corresponding redshifted lens mass samples, M_l^z, by applying the analytical relationship defined in Eq. (2.44).

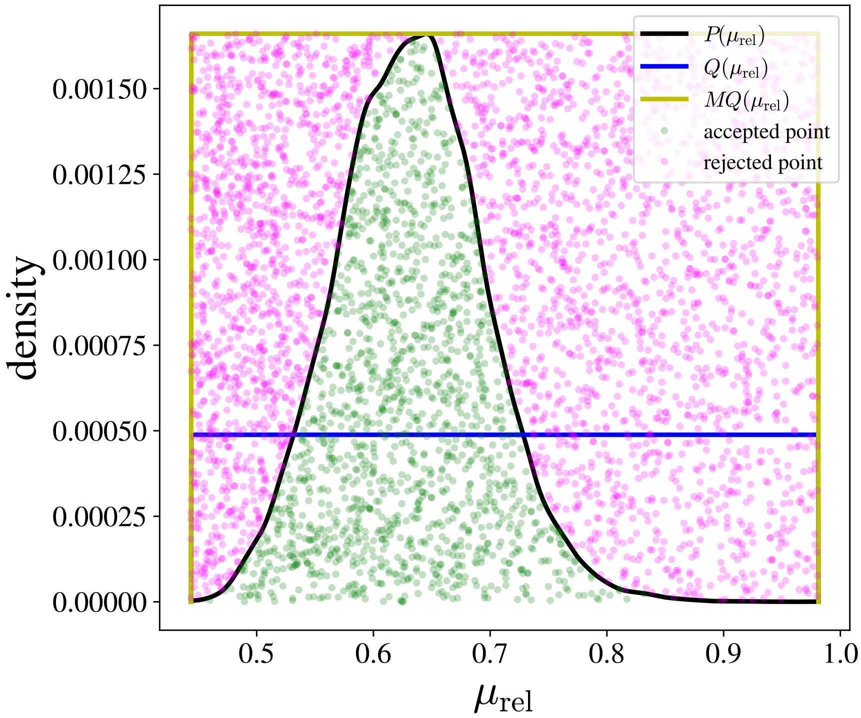

Figure 8.4: Illustration of rejection sampling.

In terms of computational performance, the rejection sampling technique is efficient. For a double-image system, the total computation time is typically less than one CPU hour. For more complex systems with triple to quintuple images, the required time increases to a few CPU hours. This is due to the sequential nature of the sampling process, where the rejection sampling for the $j^{\rm th}$ image pair (first and $j^{\rm th}$) is conducted only after the posterior results for the preceding $i^{\rm th}$ pair (first and $i^{\rm th}$, where $i < j$) have been obtained.

8.2.3 Axially Symmetric Lenses

To confirm the validity of the rejection sampling technique, we first compare the $\boldsymbol{\theta}_l$ posteriors obtained via rejection sampling against the analytical posteriors derived through univariate transformations of the $\mu_{\rm rel}$ posteriors using Eqs. (2.36) and (2.44) for the PM and SIS models, respectively. As shown in Fig. 8.5, the recovered posteriors for M_l^z and y for System 1 (PM Double) and System 2 (SIS Double) peak near the injected values. The results from the two methods are nearly identical: the analytical posteriors are represented by dashed gray and black lines, while the rejection sampling results are shown in solid colours. The solid black vertical lines indicate the true injected values, and the vertical dashed lines denote the 90 % C.I. for each parameter.

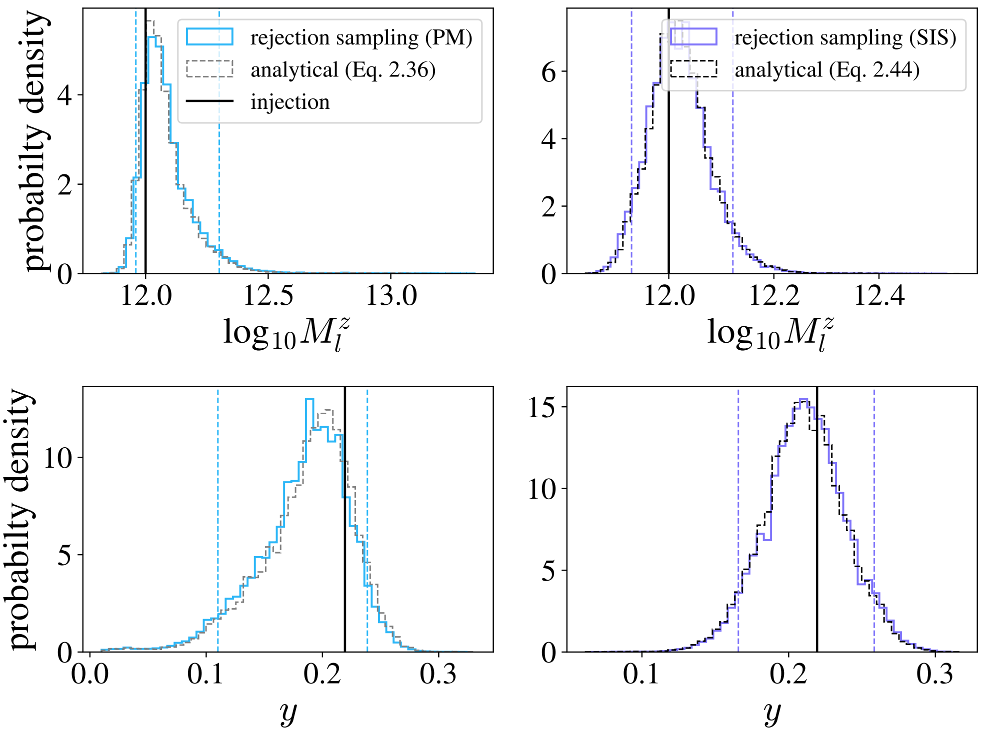

Figure 8.5: Posteriors of $\boldsymbol{\theta}_l$ and redshifted lens mass M_l^z for System IDs 1 and 2. The coloured distributions are obtained using the rejection sampling technique, while the grey distributions are derived analytically from $\mu_{\rm rel}$ posteriors using Eqs. (2.36) and (2.44).

The high degree of agreement between the analytical and numerical distributions successfully demonstrates the validity of the rejection sampling technique as a tool for mapping lensing observables back to lens parameters. However, it is important to note that while this validates the methodology, rejection sampling does not inherently bypass all physical degeneracies. The ability to successfully recover $\boldsymbol{\theta}_l$ remains subject to intrinsic degeneracies among parameters that may exist within more complex lens models.

As introduced in Section 2.2.3, among the spherically symmetric lens models considered, the NFW model does not have an analytic expression for the lensing observables. Thus, we use the rejection sampling technique for System ID 3 (NFW triple). After identifying the estimated $\mu_{\rm rel}$ posteriors with the detected triple images as a target distribution, we draw $\mu_{\rm rel}$ samples from the lensing observable bank of the NFW to conduct rejection sampling.

Figure 8.6 shows the recovered posteriors for the lens parameters $\boldsymbol{\theta}_l$ and the redshifted lens mass M_l^z for the NFW case. Although the injected values are successfully recovered within the 90 % C.I., the posteriors are poorly constrained. In particular, κ_s remains largely consistent with the prior, indicating that it is weakly constrained. This behaviour reflects the intrinsic nature of galaxy-scale lensing: for an NFW profile, the relatively low surface density implies that strong lensing, and hence the formation of multiple images, occurs only for very small values of y. Within this narrow range, many of the (y, κ_s) combinations

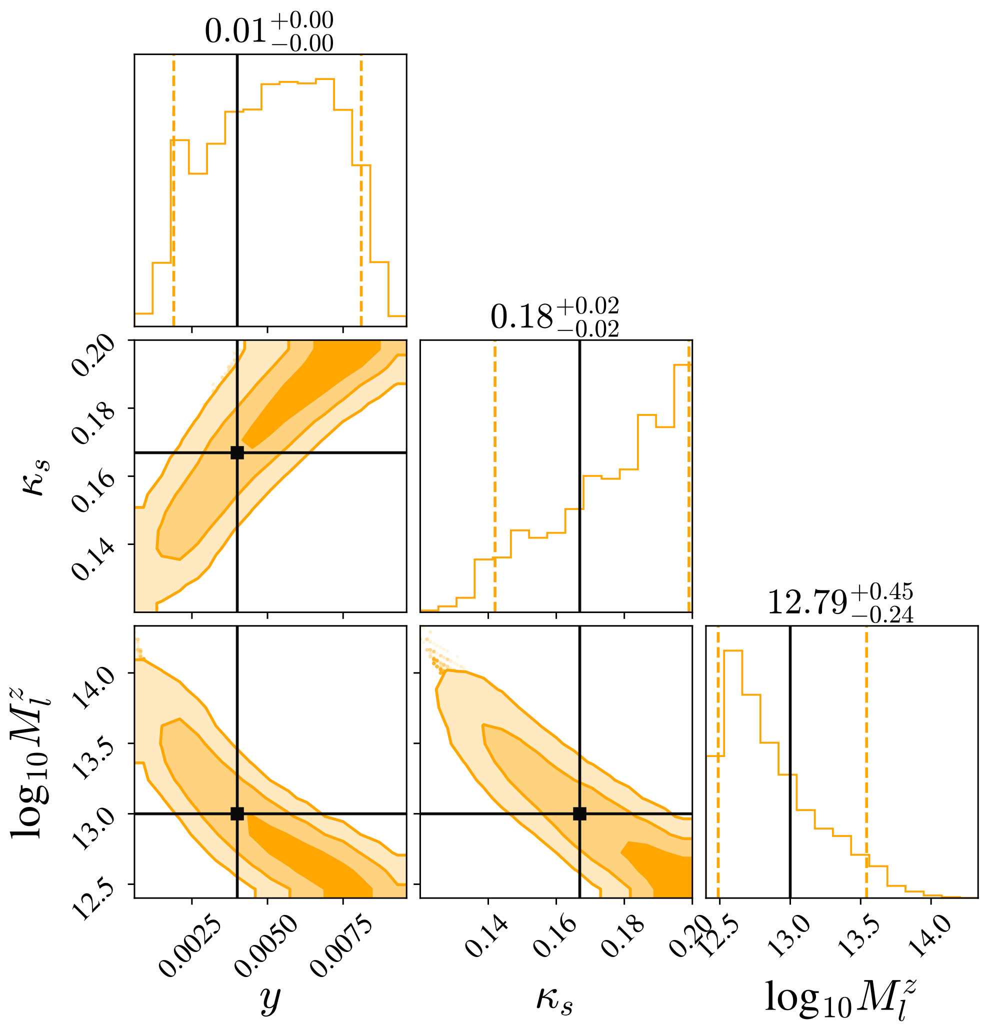

Figure 8.6: Corner plot showing the posteriors of the lens parameters $\boldsymbol{\theta}_l = (y, \kappa_s)$ and redshifted lens mass M_l^z for System ID 3.

from the prior produce nearly identical $\mu_{\rm rel}$ values, leading to a strong degeneracy between these parameters. Accordingly, the M_l^z does not peak at the true value. The precision of the obtained $\mu_{\rm rel}$ posterior is currently insufficient to break this fundamental degeneracy, resulting in the broad, uninformative posteriors.

8.2.4 Non-Axially Symmetric Lenses

In more realistic astrophysical scenarios, strong lens candidates such as galaxies are not perfectly spherical but exhibit ellipticity. The introduction of ellipticity complicates the lens model and introduces stronger degeneracies between $\boldsymbol{\theta}_l$. Consequently, we expect the posteriors for $\boldsymbol{\theta}_l$ and M_l^z in non-axially symmetric models to be more poorly constrained than those in the axially symmetric cases. In addition, the inclusion of extra parameters to describe the lens profile, such as the profile density slope γ, further increases the complexity of the parameter space and exacerbates these degeneracies.

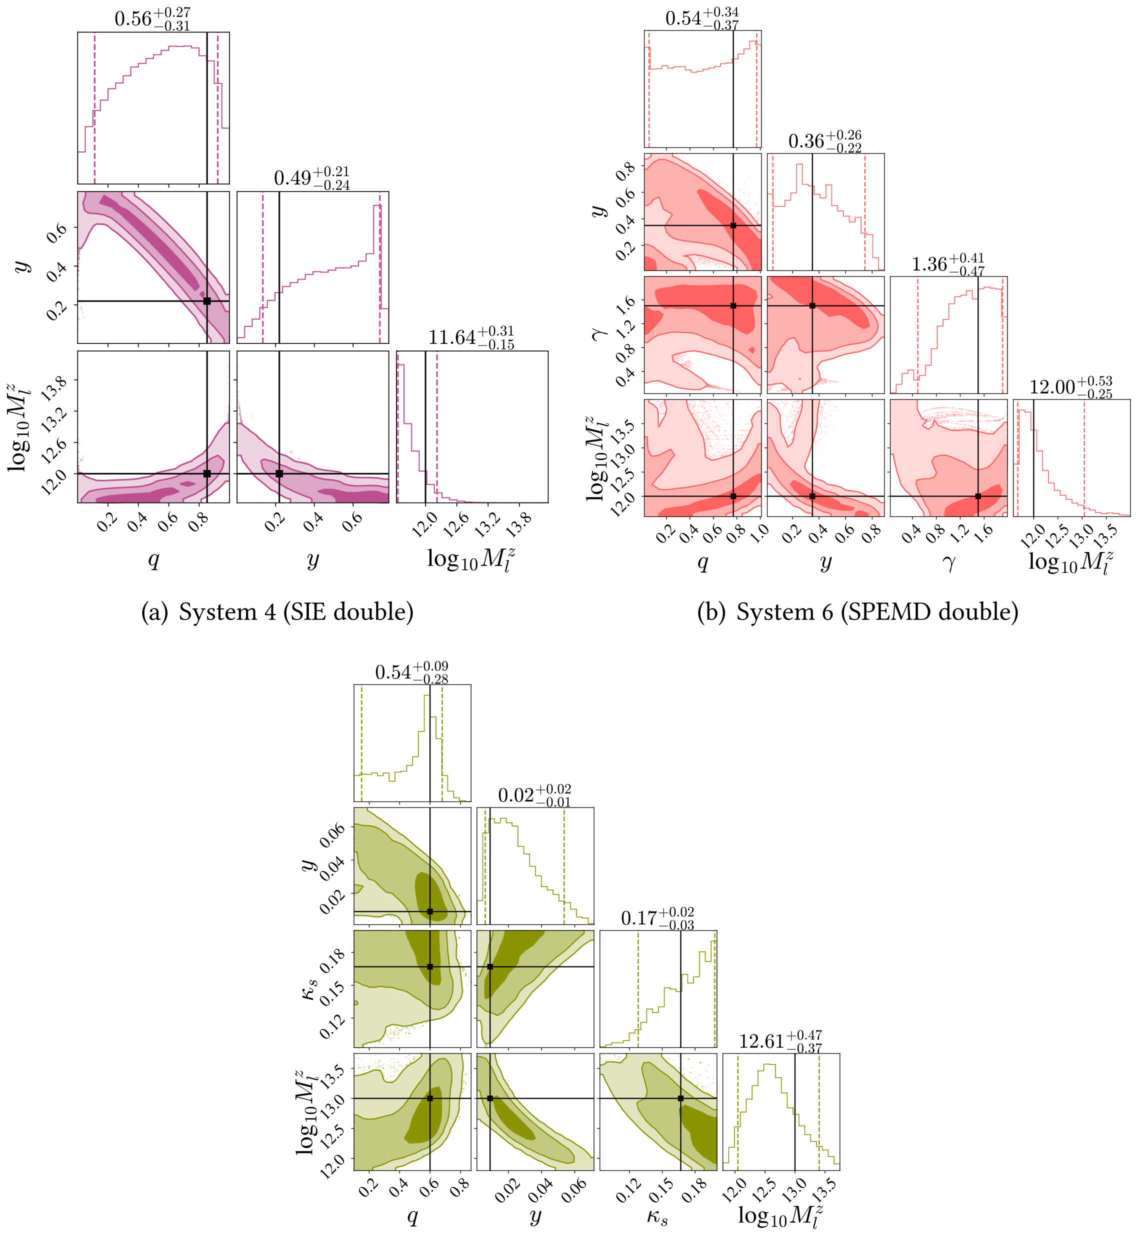

(a) System 4 (SIE double)

(b) System 6 (SPEMD double)

(c) System 8 (PNFW triple)

Figure 8.7: Corner plots showing the posteriors of $\boldsymbol{\theta_l}$ and M_l^z for System IDs 4, 6, and 8.

Figure 8.7 shows the inferred posteriors for System IDs 4, 6, and 8, corresponding to double-image configurations (SIE and SPEMD) and a triple-image configuration (PNFW). While the true injected values lie within the 2σ credible regions of the joint 2D posteriors for all three models, the marginal 1D posteriors often do not peak at the true values and remain poorly constrained. Notably, the M_l^z posteriors are better constrained than the structural parameters. This is because M_l^z is primarily constrained by the time delays between images (Δt), which are typically better constrained than $\mu_{\rm rel}$ with GW observations. These results indicate that the properties of an elliptical lens cannot be robustly recovered when the lensing configuration produces only double or triple images under our modelling assumptions.

However, the detection of more than two images provides additional independent constraints that help break degeneracies between lens parameters. By obtaining more lensing observables ($\mu_{\mathrm{rel,j}}$ and Δt_j), the parameter constraints become tighter, leading to better-constrained posteriors compared to double-image systems. Figure 8.8 shows the results for System IDs 5, 7, and 9, which produce quadruple or quintuple images. To account for the possibility that only a subset of images may be observed in practice, we consider two scenarios: (1) the detection of only a pair of GW images (gray) and (2) the detection of all available GW images (coloured).

Table 8.5 quantifies the improvement in lens parameter recovery when these additional constraints from more images are included. We define the ratio, $\mathcal{R}^{\theta_l}_{\mathrm{A/B}}$, as the relative width of the 90 % C.I. between the *pair-only* case (A) and the *all-images* case (B). Except for the SIE case, the C.I. widths of $\boldsymbol{\theta_l}$ narrow down by tens of percent when all images are included. This additional information improves the agreement between the inferred posteriors and the injected values, particularly for the SPEMD case, where all posteriors converge well to the injected values. In contrast, the PNFW case continues to exhibit a significant degeneracy between y and κ_s, similar to the NFW results. Thus, for non-axially symmetric models, the posteriors of $\boldsymbol{\theta_l}$ and M_l^z remain only weakly constrained by GW information alone. Despite these limitations, the posterior width of q consistently decreases across all models. This improved constraint on the lens structure is particularly valuable for multi-messenger astronomy, as it can aid in localising and identifying candidate lens galaxies associated with the GW host galaxy within the inferred sky area.

8.2.5 Non-Axially Symmetric Lenses with EM Observation

While GW information alone is often insufficient to recover lens parameters with high precision, incorporating EM observations offers a powerful means of breaking existing parameter degeneracies. By leveraging extra constraints from EM data, we can significantly refine the estimation of structural parameters and, consequently, obtain a more robust constraint on the redshifted lens mass M_l^z. For instance, if a lens galaxy is identified within the GW sky localisation region, its structural properties, such as the axis ratio and position angle, can be directly measured from high-resolution imaging. In addition, the image positions of the lensed host galaxy can provide further spatial constraints.

However, identifying the correct lens galaxy remains challenging, as the GW localisation uncertainty is still large enough to contain thousands of candidate galaxies. In this context, several studies [310, 379, 380] have proposed localising the host–lens system using wide-field surveys such as the Legacy Survey of Space and Time (LSST) [381]. This approach generally assumes the joint detection of lensed GW and EM signals from the same

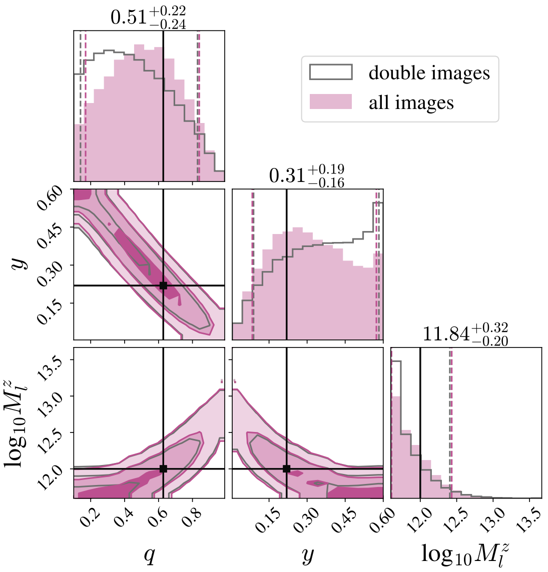

(a) System 5 (SIE quadruple)

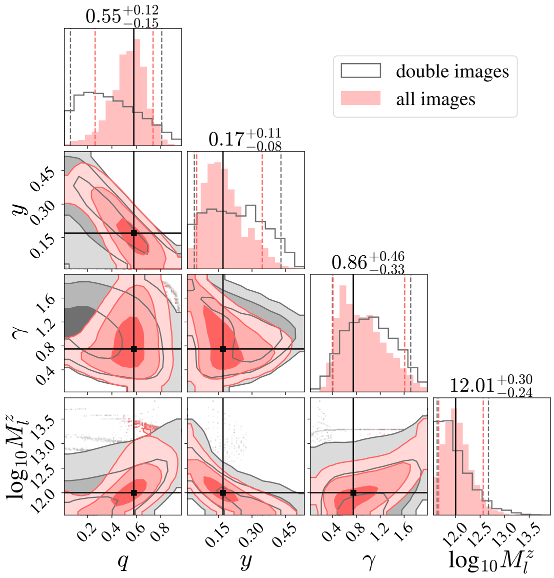

(b) System 7 (SPEMD quadruple)

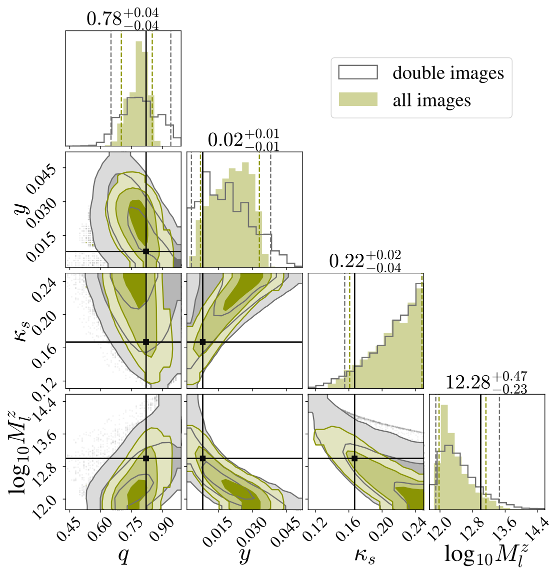

(c) System 9 (PNFW quintuple)

Figure 8.8: Same configurations as in Figure 8.7, but showing the posteriors of $\boldsymbol{\theta_l}$ and M_l^z for System IDs 5, 7, and 9.

host galaxy[2]. By combining lensing observables from both messengers, specifically through the joint application of Eqs. (2.76)–(2.79), the number of candidate systems within the GW localisation region can be significantly reduced, enabling identification of the most probable lens-source configuration.

2. In general, the host galaxy of a GW source is also lensed by the intervening galactic halo.

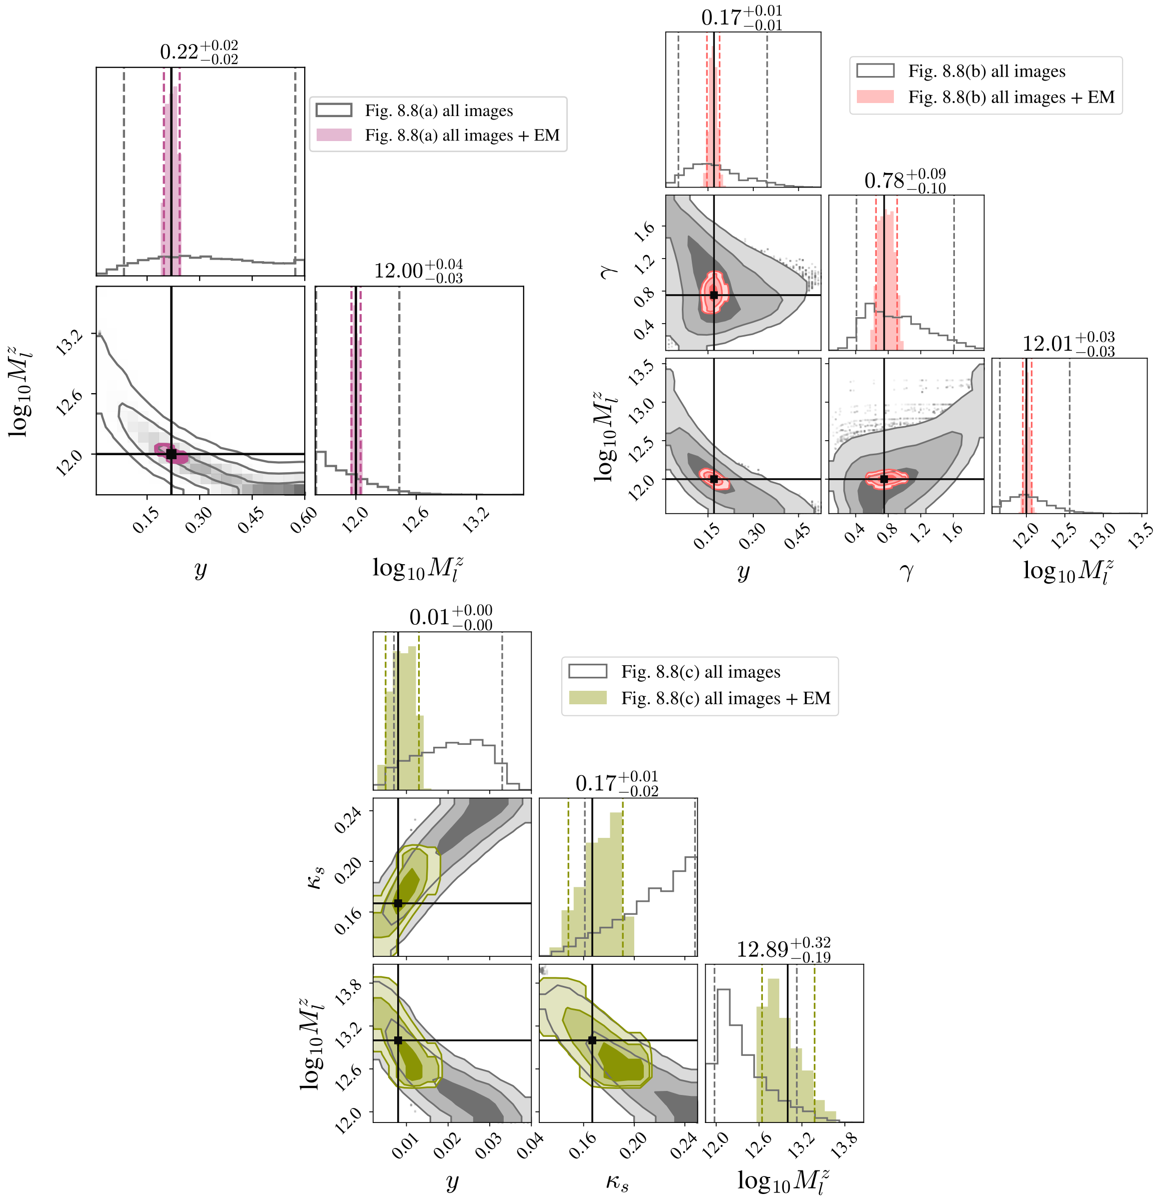

Figure 8.9: Same configurations as in Figure 8.8, but with the assumption that quadruple (quintuple) image systems are detected with the additional EM observations and thereby, axis ratios and image positions are confined with $\pm$ 5 % and $0.1''$ error ranges.

In particular, [310] argue that the detection of quadruple GW images can significantly reduce the sky localisation area with current LIGO–Virgo detectors. Furthermore, as shown in Figure 8.8, the axis ratio of a lens galaxy can be more precisely constrained when four or more GW images are identified, providing a critical starting point for identifying the lens galaxy in EM surveys.

To quantify the impact of multi-messenger constraints, we consider a scenario in which the lens galaxy has been successfully identified. Motivated by typical observational uncertainties, we assume that the axis ratio and EM image positions are constrained to within ± 5 % and 0.1 arcseconds, respectively. Figure 8.9 shows the recovered posteriors for $\boldsymbol{\theta_l}$ and M_l^z for the non-axially symmetric lens systems under these assumptions. All $\boldsymbol{\theta}_l$ posteriors

$P_{\text{dbl}}(\boldsymbol{\theta}_l)$ vs $P_{\text{all}}(\boldsymbol{\theta}_l)$ in Figure 8.8					
Model	$\mathcal{R}^{q}_{\text{dbl/all}}$	$\mathcal{R}^{y}_{\text{dbl/all}}$	$\mathcal{R}^{\log_{10} M_l^z}_{\text{dbl/all}}$	$\mathcal{R}^{\gamma}_{\text{dbl/all}}$	$\mathcal{R}^{\kappa_s}_{\text{dbl/all}}$
SIE	1.03	1.01	0.97	-	-
SPEMD	1.55	1.33	1.22	1.11	-
PNFW	1.89	1.36	1.30	-	1.06

$P_{\text{all}}(\boldsymbol{\theta}_l)$ vs $P_{\text{all+EM}}(\boldsymbol{\theta}_l)$ in Figure 8.9				
Model	$\mathcal{R}^{y}_{\text{all/all+EM}}$	$\mathcal{R}^{\log_{10} M_l^z}_{\text{all/all+EM}}$	$\mathcal{R}^{\gamma}_{\text{all/all+EM}}$	$\mathcal{R}^{\kappa_s}_{\text{all/all+EM}}$
SIE	9.87	8.17	-	-
SPEMD	6.94	7.90	4.21	-
PNFW	3.07	1.51	-	1.85

Table 8.5: Quantitative evaluation of the improvement in lens parameter recovery from additional constraints, such as multiple GW images and indicative EM observations.

become significantly tighter and peak near the injected values. We quantify this improvement in Table 8.5 using the ratio $\mathcal{R}^{\theta_l}_{\text{all/all+EM}}$, defined as the relative width of the 90 % C.I of the *all-GW images* to the *all-GW + EM* scenarios. Most importantly, M_l^z is recovered with high precision. By combining the inferred M_l^z with the lens redshift z_l obtained from EM spectroscopy, the intrinsic lens mass M_l can be determined, completing the reconstruction of the lensing system.

8.3 Conclusion

In this chapter, we first demonstrate the delensing process through analytic lens reconstruction. Our proof-of-concept study shows that we can retrieve the true luminosity distance with significantly improved precision by detecting SLGWs from a single BBH. Even in the presence of simulated detector noise, an optimal strong-lensing configuration enhances the distance constraint by a factor of several relative to the unlensed case. Although the initial demonstration is restricted to simple PM and SIS models with double images, more complex lens models, such as those introduced in Section 8.2.4, can yield higher magnifications and increased image multiplicity, further improving the precision of luminosity distance estimates.

As the lens models become more complex, the reconstruction procedure becomes increasingly challenging. An inference framework based on rejection sampling is implemented to address this challenge, further demonstrating that the precision of lens reconstruction is significantly improved when joint GW and EM observations are available. In addition to the geometric information provided by EM imaging, spectroscopic measurements of both the lens galaxy and the lensed host galaxy enable precise redshift measurement, supplying an additional and critical constraint on the lensing configuration, such as the Einstein radius.

Successful lens reconstruction directly enables more precise and accurate luminosity distance estimates, providing critical input for studies of the formation and evolution of compact binary populations. Furthermore, these improved distance measurements offer a powerful avenue for constraining the Hubble constant through the Hubble–Lemaître law discussed in Section 3.1.3. By mitigating the intrinsic uncertainties in standard GW distance inference, SLGWs may play a key role in resolving the current Hubble tension in the coming decades.

However, real-world detection scenarios are expected to be significantly more challenging, and several key issues remain to be addressed. We highlight three primary directions for future investigation.

First, observational completeness. In practice, only a subset of the full image ensemble may be detected (e.g., two images from a quadruple-image SIE system), either because additional images have not yet arrived due to long time delays or because their SNRs fall below the detection threshold. A natural extension is to incorporate rejection sampling to compare different image-multiplicity scenarios and identify the statistically preferred configuration.

Second, model selection and hypothesis testing. Because the lens density profile is not known a priori, rejection sampling should be embedded within a JPE framework to enable rigorous hypothesis testing. Such an approach can not only identify the most probable lens model but also improve search pipelines by reducing FAR, as suggested by [382, 383].

Finally, complex lensing environments (macro–micro lensing). As discussed in Chapter 7, galactic substructures, such as star clusters and DM subhalos, can induce additional microlensing effects on top of the macrolensing signal. Neglecting these contributions may bias lens parameter estimates. Developing parameter estimation frameworks that incorporate the associated wave-optics effects remains an important goal for future work, particularly in the context of space-based interferometers, where higher SNR and improved sky localisation will help disentangle these complex degeneracies [384, 385, 386].

Chapter 9

The Impact of Strong Lensing on H_0 Measurements with Gravitational-Wave Dark Sirens

Disclaimer

The following chapter is based on the paper *The Impact of Strong Lensing on Hubble Constant Measurements with Gravitational-Wave Dark Sirens* [387]. The author was the first author and contributed to every stage of this work, including mathematical derivations, framework development, data generation and analysis, interpretation of the results, scientific discussion, figure production, and manuscript preparation. This paper was published in the *Astrophysical Journal Supplement Series*.

As discussed in Section 3.1.4, the precise value of H_0 remains under debate due to the Hubble tension, between early-Universe measurements inferred from the cosmic microwave background [175], which report $H_0 = 67.4 \pm 0.5$ km s^{-1} Mpc^{-1}, and late-Universe measurements based on the Cepheid-supernova distance ladder [181], which find $H_0 = 72.3 \pm 1.4$ km s^{-1} Mpc^{-1}. This discrepancy has motivated growing interest in independent approaches to measuring H_0. In this context, as introduced in Section 3.2, GW standard sirens provide a promising alternative, particularly given the rapid expansion of GW astronomy over the past decade, with more than 200 CBCs detected by the LVK network [20].

In this chapter, we focus on the galaxy-catalogue-based dark siren method [219, 220], in which each GW event is statistically associated with multiple candidate host galaxies within its sky localisation volume. By weighting these galaxies according to their spatial probabilities and luminosities, one can construct a probabilistic redshift distribution for the GW source and subsequently perform hierarchical inference of H_0. For the galaxy-catalogue-based method, improving constraints on H_0 primarily depends on two factors: accurate sky localisation to reduce the number of potential host galaxies, and precise measurements of the luminosity distance. In this context, we consider SLGW events, which, as demonstrated in Section 8.1, are expected to yield improved luminosity-distance precision.

In addition, SLGWs can provide more informative constraints than unlensed events for sky localisation [310, 388], both of which are essential for robust host-galaxy identification in the dark-siren framework. Furthermore, the additional lensing information encoded in these systems can provide complementary cosmological constraints on H_0.

Although no unambiguous detection of SLGWs has yet been reported in the latest GWTC catalogue [238], the LVK detector network is expected to observe GW strong lensing at a rate of $\mathcal{O}(1)$ Gpc^{-3}, yr^{-1} at design sensitivity [158, 159, 160, 326, 341, 359, 389]. Such systems are expected to provide cosmological information analogous to that obtained from time-delay measurements of strongly lensed quasars [111, 390, 391]. As the number of detected SLGWs increases in future LVK observing runs [102, 326] and with next-generation observatories such as the Einstein Telescope [200, 392, 393], joint analyses incorporating both source and lens population models will become increasingly important. At the same time, unlensed dark sirens should not be excluded from analyses, as combining lensed and unlensed GW events within a unified hierarchical framework can substantially improve the precision and robustness of the inferred H_0. In this regard, tools such as `gwcosmo` [219, 220, 394] provide a flexible Bayesian framework for incorporating population information and jointly inferring cosmological parameters from both lensed and unlensed GW events.

In this chapter, we investigate the impact of SLGWs on H_0 inference, highlighting how lensing information can strengthen cosmological constraints relative to conventional unlensed dark sirens. Two observational scenarios are considered. In the first scenario, multiple lensed signals originating from a single dark siren event are detected. Once the host galaxy is uniquely identified through lensing observables and cross-correlation with galaxy catalogues, the Hubble–Lemaître law (Eq. (3.33)) can be directly used to infer H_0. In the second scenario, a population of lensed dark siren sources is detected, requiring a hierarchical framework that consistently combines all events. This framework incorporates comparisons with galaxy-galaxy lensing catalogues to construct more informative redshift maps and thereby improve the precision of the inferred H_0. We additionally quantify potential systematic biases in H_0 inference arising from neglected or incorrectly modelled lensing effects. Throughout this chapter, we assume a fiducial flat ΛCDM cosmology with $\Omega_m = 0.25$, $\Omega_b = 0.044$, $\Omega_\Lambda = 0.7$, and $H_0 = 70$ km, s^{-1}, Mpc^{-1}.

9.1 Mock Data

9.1.1 Galaxy catalogue

We adopt the MICE-Grand Challenge light-cone halo and galaxy catalogue (`MICECATv2`) [395, 396, 397, 398, 399] for our galaxy distribution, while separately investigating how catalogue incompleteness affects the inference of H_0. The simulation volume covers approximately 5000 deg^2 of the sky over a redshift range of $0.07 < z < 1.42$. The catalogue provides absolute and apparent magnitudes across multiple photometric bands, alongside both true and observer-frame redshifts. We specifically select the r-band for our analysis, as MICECATv2 is calibrated using the r-band Schechter function from [400], with galaxy evolution models applied at higher redshifts [401]. Although completeness estimates were originally derived using i-band data, we use the r-band data to ensure a direct and consistent application of the Schechter parameters for incompleteness corrections. Here, completeness is defined as the fraction of observable galaxies successfully included in the catalogue.

To maintain computational efficiency within gwcosmo, we utilise a randomly selected one-eighth subset of the full MICECATv2 catalogue, comprising approximately 6.2×10^7 galaxies[1]. We select galaxies with r-band absolute magnitudes $M_r < -20$ and assume this subset is 100 % complete down to an r-band apparent magnitude $m_r < 24$ at $z = 1.42$ [401]. The absolute-magnitude cut is imposed to achieve an approximately uniform number density in comoving volume, thereby constructing a volume-limited sample. This selection prevents an overabundance of faint, low-redshift galaxies while avoiding a significant bias toward only the brightest sources. We exclude galaxy clusters to focus exclusively on galaxy-galaxy strong lensing. Hereafter, this complete catalogue is referred to as the *Unlensed Galaxy catalogue* (UGC0).

Within the landscape of galaxy surveys, certain catalogues focus exclusively on lensed systems; for example, the SLACS survey [402, 403] targets confirmed strong-lensing systems. Complementary efforts have produced dedicated catalogues of strong gravitational lenses, such as lenscat [404], which facilitate the identification of host galaxies associated with lensed transient events. In this context, we simulate a catalogue of lensed galaxies derived from UGC0. Compared to the unlensed case, a lensed galaxy catalogue contains significantly fewer potential host-galaxy candidates per unit sky area within a given GW localisation region. This sparsity effectively mitigates source-identification uncertainty in the statistical H_0 inference performed with gwcosmo. To simulate lensing systems, we assume that all galaxies in our mock catalogue follow SIE density profiles described in Section 2.2.4.

Since MICECATv2 does not directly provide the SIE lens parameters, we model their underlying distributions. We adopt the functional forms described in [359]; however, rather than fixing the galaxy number density at a specific redshift within the Schechter function, we compute it independently at each redshift using the MICECATv2 data and sample velocity dispersion σ_v accordingly. For the shape parameters, we assume a uniform distribution for the orientation angle ϕ_e and a Rayleigh distribution for the axis ratio q, extending the lower limit to $q_{\text{min}} = 0.1$ to include highly elliptical galaxies.

1. The resulting number density is 0.00146 Mpc^{-3} (reduced from the original 0.0117 Mpc^{-3}), which more closely approximates the number densities of observational catalogues like GLADE+ [214] used in dark siren GW cosmology.

The number of strongly lensed galaxies, N_{Lgal}, in a given catalogue is determined by the optical depth, τ_{gal}, as follows:

$$
\begin{aligned}
N_{\mathrm{Lgal}} &= N_{\mathrm{gal}} \times \tau_{\mathrm{gal}} \\
&= N_{\mathrm{gal}} \times \sum_{i}^{N_{\mathrm{gal}}} \int \tau(z_i) P(z_i) dz_i,
\end{aligned}
\tag{9.1}
$$

where $\tau(z_i)$ is the optical depth for the i^{th} galaxy at redshift z_i, and $P(z_i)$ represents the probability that a galaxy from the catalogue resides at redshift z_i. Thus, τ_{gal} corresponds to the combined optical depth across the entire galaxy population in the given universe (see Eq. (2.31)).

Since we adopt an SIE lens model, only two or four images of a source galaxy can be produced. The resulting lensed galaxy catalogue contains $\sim 4.6 \times 10^4$ double-image systems and $\sim 1.0 \times 10^4$ quadruple-image systems. Hereafter, this complete lensed galaxy catalogue, comprising all double- and quadruple-image systems generated from the UGC0, is referred to as the *Lensed Galaxy catalogue* (LGC0). Figure 9.1 shows the redshift distribution of galaxies in the UGC0 alongside the distributions of double- and quadruple-image systems in the LGC0. As expected, the distribution of lensed galaxies peaks at higher redshifts, as galaxies at higher redshift have a higher probability of encountering a foreground lens.

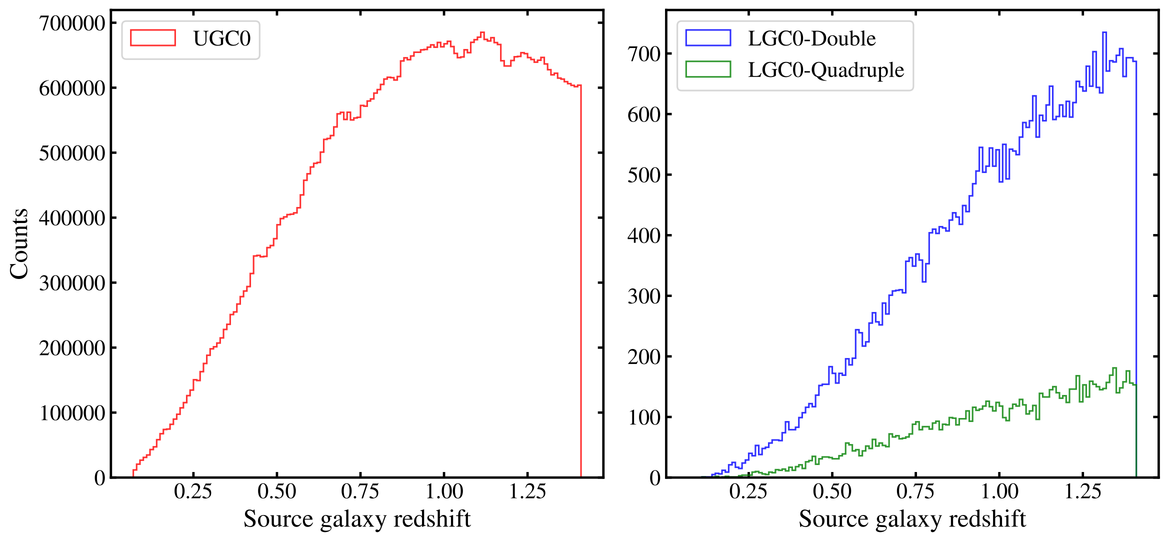

Figure 9.1: The redshift distributions of galaxies in the unlensed catalogue (UGC0; red), and of background galaxies in double-image (blue) and quadruple-image (green) systems in the lensed catalogue (LGC0).

In practice, observational galaxy catalogues are incomplete and do not capture every potential host galaxy of a GW source. This incompleteness must be accounted for when inferring H_0 using dark sirens and galaxy catalogues. To mimic these observational limitations, we impose an apparent magnitude cut of $m_{\rm th} = 21.87$, removing galaxies fainter than this threshold from both the UGC0 and LGC0. The resulting incomplete catalogues are hereafter referred to as UGC1 and LGC1, respectively; their application in the H_0 measurement will be discussed in Section 9.3.

9.1.2 Dark Sirens

For our dark-siren simulations, we adopt the parametric models of BBH populations described in Section 4.4.2, along with the source parameter ranges and hyperparameters inferred from the GWTC-3 analysis [293]. We assume a fixed observation time of $T_{\rm obs} = 2.26$yrs, corresponding to the duration of the O4 observing run. For the remaining BBH parameters, including spins, we assume uniform distributions.

Each simulated BBH is assigned to a host galaxy in the UGC0 catalogue following a physically motivated weighting scheme. A galaxy's probability of hosting a BBH depends on its SFR, stellar mass, metallicity, and the delay time between binary formation and merger [405, 406]. While explicitly incorporating all these factors can improve H_0 convergence by accurately reflecting the true host distribution, inaccurate or oversimplified prescriptions risk introducing systematic biases into the H_0 posterior [407, 408]. In this study, rather than modelling the full complexity of host-galaxy properties, we adopt a luminosity-based weighting scheme as a practical approximation, assigning higher hosting probabilities to more luminous galaxies. Unlensed GW signals are generated using the IMRPhenomXPHM-SpinTaylor waveform model [288, 289] and injected into a detector network consisting of aLIGO–Livingston, aLIGO–Hanford, and AdV at their expected O4 sensitivities [365]. We consider only those signals detected in all three detectors with a network SNR above 8 ($\rho_{\rm net} > 8$).

In addition to the unlensed GW signals, SLGW signals are generated by applying the lensing effects of galaxies in the LGC0 to the previously simulated unlensed signals using the amplification factors (see Eq. (2.74)). To determine the number of lensed GW signals ($N_{\rm lgw}$), we follow a procedure analogous to that used in constructing the LGC0: the optical depth for BBHs is computed based on the galaxies in the UGC0, such that:

$$N_{\rm lgw} = N_{\rm gw} \times \tau_{\rm BBH}. \tag{9.2}$$

Note that $\tau_{\rm BBH}$ is the combined optical depth for the entire BBH population, computed using the complete unlensed catalogue, UGC0.

In practice, it is computationally prohibitive to calculate the dimensionless impact parameters (y) for every potential galaxy-galaxy pair while accounting for sky positions and angular diameter distances between the source, lens, and observer. Therefore, we randomly select N_{lgw} source-lens pairs from the LGC0, where each source galaxy is assumed to host a BBH that is strongly lensed by its associated foreground lens galaxy. Consequently, y values that ensure strong lensing are assigned, which are drawn from a uniform distribution over the range $y \in [0.01, 1.2]$. The upper bound of 1.2 corresponds to the maximum for which strong lensing can occur within the axis-ratio range adopted for our SIE model. Table 9.1 summarises the assumed distributions for the representative source and lens parameters.

Parameter	**Distribution model**	**Range**
m_1	`Power Law + Peak`	$[5, 112]M_{\odot}$
m_2	`Truncated Power Law`	$[5, m_1]M_{\odot}$
q	`Rayleigh`	[0.1, 1.0]
σ_v	`Schechter`	[50,500]km/s
y	`Uniform`	[0.01, 1.2]

Table 9.1: Assumed probability distributions for component masses (m_1, m_2) of the source of dark sirens and the main lens parameters, including the axis-ratio (q) and the velocity dispersion (σ_v), and dimensionless impact parameter (y), employed in simulating strongly lensed GW signals.

Among the simulated SLGW signals, only those with $\rho_{\text{net}} > 8$, evaluated using the same detector configuration as in the unlensed case, are considered detected[2]. In a realistic scenario, it is not immediately known whether a detected signal is lensed. Candidate lensed pairs (or higher-order multiplets) must first be identified by low-latency search pipelines. For instance, posterior-overlap tests identify events with mutually consistent intrinsic parameters—such as component masses and spins—and overlapping sky localisations, while accounting for differences in arrival times and strain amplitudes [299]. For high-ranking candidates, JPE is subsequently performed. During this step, as discussed in Section 5.2.1, the lensing evidence is computed by jointly analysing the data segments under the lensing hypothesis, thereby assessing whether the candidate multiplet is coherently described as multiple lensed signals of the same GW source. In this work, the true associations between lensed signals are known by construction, allowing us to perform the JPE directly on the correct signal sets. Only those systems for which the lensing hypothesis is statistically supported by the JPE analysis are retained for the subsequent cosmological inference.

2. In actual observations, thresholds for the astrophysical probability (p_{astro}) and FAR are applied to ensure signals are of astrophysical origin. For this study, we assume all signals satisfying the SNR threshold also meet these criteria.

9.2 H_0 Measurement with a Strongly Lensed Dark Siren

Assuming the detection of a lensed dark siren, significantly improved sky localisation can be achieved relative to the unlensed case [310, 388], along with more precise measurements of the luminosity distance (d_L) [373], as described in Chapters 3 and 8. Within this more tightly constrained sky area, candidate lens galaxies and host galaxies can be identified, enabling the use of the host redshift in combination with the improved luminosity distance to infer H_0 with higher precision. However, the detectability of the lensing system remains dependent on the lens mass, luminosity, and redshift.

9.2.1 Method

If a source-lens system can be uniquely specified, the redshift of the source galaxy (z_s) can be obtained via EM observations, either through photometric or spectroscopic methods. The obtained z_s can then be used to apply the Hubble-Lemaître law ($H_0 = v_r/d_p$), where

$$v_r(z) = \int_0^z \frac{c\, dz'}{\sqrt{\Omega_m(1+z')^3 + \Omega_\Lambda}} \tag{9.3}$$

is the recessional velocity of the source galaxy, and d_p is its proper distance from the observer. In a flat-ΛCDM universe, where the scale factor at $z = 0$ is set to $a_0 = 1$, the proper distance is equivalent to the comoving distance to the source galaxy ($d_C = d_L/(1 + z_s)$). In this study, we ignore the peculiar velocity of the source galaxy. This assumption can introduce biases in H_0 measurement using bright sirens, such as binary-neutron stars, but it is reasonable for dark-siren analyses, where host galaxies are typically not located in the nearby Universe ($d_\mathrm{L} < 200$ Mpc) [228, 229].

Since the redshift of a galaxy cannot be determined perfectly, we assign uncertainties such that the measured redshift follows a Gaussian distribution, where the the standard deviation is σ_z. For the photometric and spectroscopic methods, we adopt $\sigma_z = 0.033(1 + z_s)$ and $\sigma_z = 0.001(1 + z_s)$, respectively [409]. Using these estimated redshift distributions, the recessional velocity of the source can be determined, and the measured luminosity distance can be converted into the corresponding comoving distance. Note that the measured luminosity distance from the j^th lensed signal is not the true distance to the source, but is rescaled by the lensing magnification factor, as described in Eq. (2.75). Recovering the true distance, therefore, requires delensing, which involves estimating the individual magnification factors μ_j from lens reconstruction.

From the lens-model-independent JPE for the lensed signals described in Section 5.2.1, one can obtain not only source parameters, but also lensing observables for each pair of signals. For a double-image system, this corresponds to a single triplet, ($\mu_{\mathrm{rel},12}, \Delta t_{12}, \Delta\phi_{n,12}$). For a quadruple-image system, three such triplets are theoretically available, ($\mu_{\mathrm{rel,1j}}, \Delta t_{1j}, \Delta\phi_{n,1j}$)

with $j = 2, 3, 4$, although fewer triplets may be obtained if some images are missing or undetectable. These lensing observables can then be converted into model-dependent lens parameters either analytically [410] or numerically using a rejection-sampling framework introduced in Section 8.2.

Once the lens parameters are constrained, one can further utilise shape parameters from EM observations, such as the axis ratio (q) and elliptical angle (ϕ_e) of the lens galaxies, as well as the inferred redshifted lens mass, to aid in uniquely identifying the source-lens system within the sky localisation. This approach relies on the assumption that both the lens galaxy and the lensed source images are included in the catalogue. Therefore, we use LGC0 for this case. When accounting for catalogue incompleteness by imposing $m_{\rm th}$, both the lens galaxy and the lensed source galaxy images must be contained within LGC1.

With the lens parameters determined, we can reconstruct the lens galaxy under a given lens model to obtain individual magnification factors, which are used to delens the apparent luminosity distance and retrieve the true distance to the source. However, it is important to note that the rejection-sampling approach does not always yield accurate lens reconstructions, owing to degeneracies between lensing observables and the underlying lens properties, as well as from the limited availability of EM information due to survey magnitude limits. This issue will be addressed in Section 9.4.1.

9.2.2 Results

Among the simulated double- and quadruple-image systems, we present four representative detected cases. For each system type, we identify: (1) the system with the highest mean SNR ($\rho_{\rm mean}$), defined as the average network SNR of its lensed signals, and (2) the system with the lowest $\rho_{\rm mean}$.

For a given detector network and comparable $\rho_{\rm mean}$, the sky localisation of a double-image system is more precise than that of an unlensed event. A quadruple-image system yields further improvement since more independent measurements of the same event are available. Figure 9.2 displays the skymap posteriors for the four representative systems. When cross-referencing these localisations with the complete lensed galaxy catalogue (LGC0), the inferred 90 % credible sky areas ($\Delta\Omega_{90\,\%}$) contain $\mathcal{O}(1)$ and $\mathcal{O}(2\text{–}3)$ galaxies with properties compatible with the GW observables for the double- and quadruple-image cases, respectively. The exact counts of identified galaxies for each system are summarised in Table 9.2.

To further constrain the potential lensing candidates, we utilise lensing observables, specifically the relative magnification factors and the time delays between images. Since these quantities can be directly obtained from the lensed galaxy catalogue, we can compare the EM-derived properties of a candidate source–lens system with the parameters inferred from the lensed GW signals.

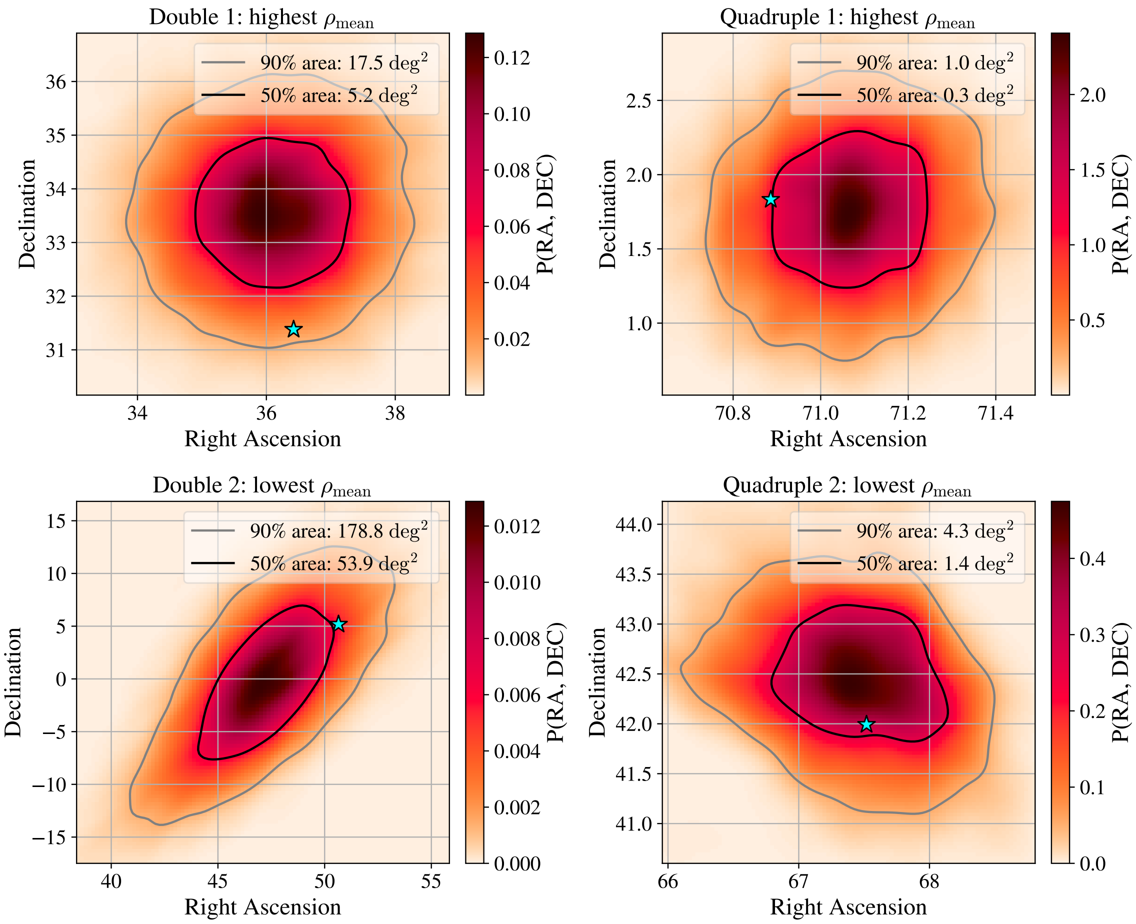

Figure 9.2: Sky localisation maps for the selected double- and quadruple-image systems. The star symbols mark the true sky positions of the corresponding GW lensing systems.

As shown in Figure 9.3, for the double-image systems in the LGC0, the number of candidates consistent with the inferred lensing observables is reduced to a few. For GW signals, the time delay is measured with high precision, providing a significantly more stringent constraint than the relative magnification factors, which are inferred from the less precise luminosity distances. In practice, uniquely identifying the source-lens pair in a double-image configuration, a prerequisite for measuring the source-galaxy redshift, remains challenging. Such identification is typically feasible only for high-SNR events that

Table 9.2: The mean SNR values ($\rho_{\rm mean}$), 90 % credible sky area ($\Delta\Omega_{90\,\%}$), numbers of galaxies within these area ($N_{\rm G-G}$), and reduced numbers of galaxies whose properties are consistent with those inferred from the lensed GW signals for each system ($N'_{\rm G-G}$).

System	$\rho_{\rm mean}$	$\Delta\Omega_{90\,\%}[\rm deg^2]$	$N_{\rm G-G}$	$N'_{\rm G-G}$
Double 1	22	17.5	306	5
Double 2	8.1	178.8	1971	1
Quadruple 1	30.0	0.97	1	1
Quadruple 2	9.0	4.35	13	1

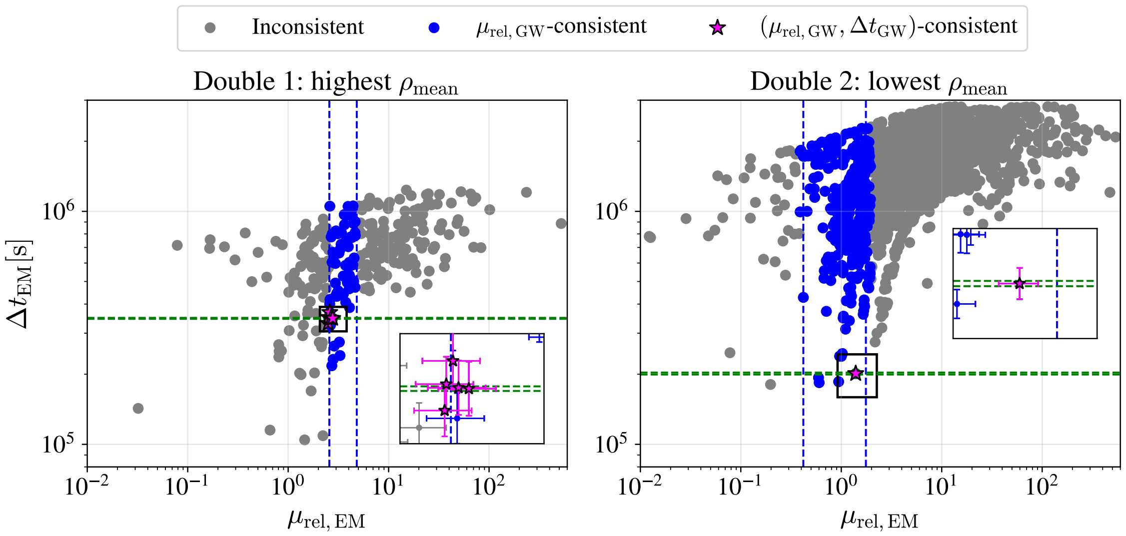

Figure 9.3: Distributions of the median relative magnification factors and time delays obtained from EM observations for the source–lens galaxy systems within the localised sky area of each double-image case. In the inner box, the error bars of $\mu_{\rm rel,EM}$ and $\Delta t_{\rm EM}$ are shown.

yield exceptionally tight sky localisations (e.g., $\Delta\Omega_{90\,\%} \lesssim 20$ deg^2) [310, 380]. In contrast, quadruple-image systems provide much tighter constraints. The additional images allow for the unique identification of the source-lens system among only a handful of candidates, as shown in Table 9.2.

Despite the inherent difficulties in double-image cases, we proceed under the assumption that the source-lens pair can be uniquely identified for each system. Once the lensing systems are correctly identified, as discussed in Section 8.2.5, we can further exploit the observed image positions of the lensed source galaxy and the shape parameters of the lens galaxy to more tightly constrain the remaining lens parameters, especially the impact parameter y. This enables a more accurate determination of the individual magnification factors[3]. Following [311], we assume that the shape parameters (q, ϕ_e) can be determined from EM observations with an uncertainty of $\pm 5\,\%$, and that image positions can be measured with a precision of $0.1''$.

The inferred individual magnification factors are presented in Figure 9.4. These distributions are utilised to retrieve the true luminosity distances from the previously obtained apparent distances following Eq. (2.75). Specifically, we perform a stochastic delensing procedure by drawing random samples from the posterior distribution of the magnification factor and multiplying them by samples drawn from the apparent-distance posterior. This yields a corresponding posterior realisation of the true luminosity distance. As expected,

3. In general, a galaxy, as an extended source, does not share the same source-plane position y as the point-like GW event embedded within it (i.e., $y_{\rm gal} \neq y_{\rm gw}$). Since the magnification factor is highly sensitive to y, it is essential to utilise $y_{\rm gw}$ rather than $y_{\rm gal}$ for cosmological inference.

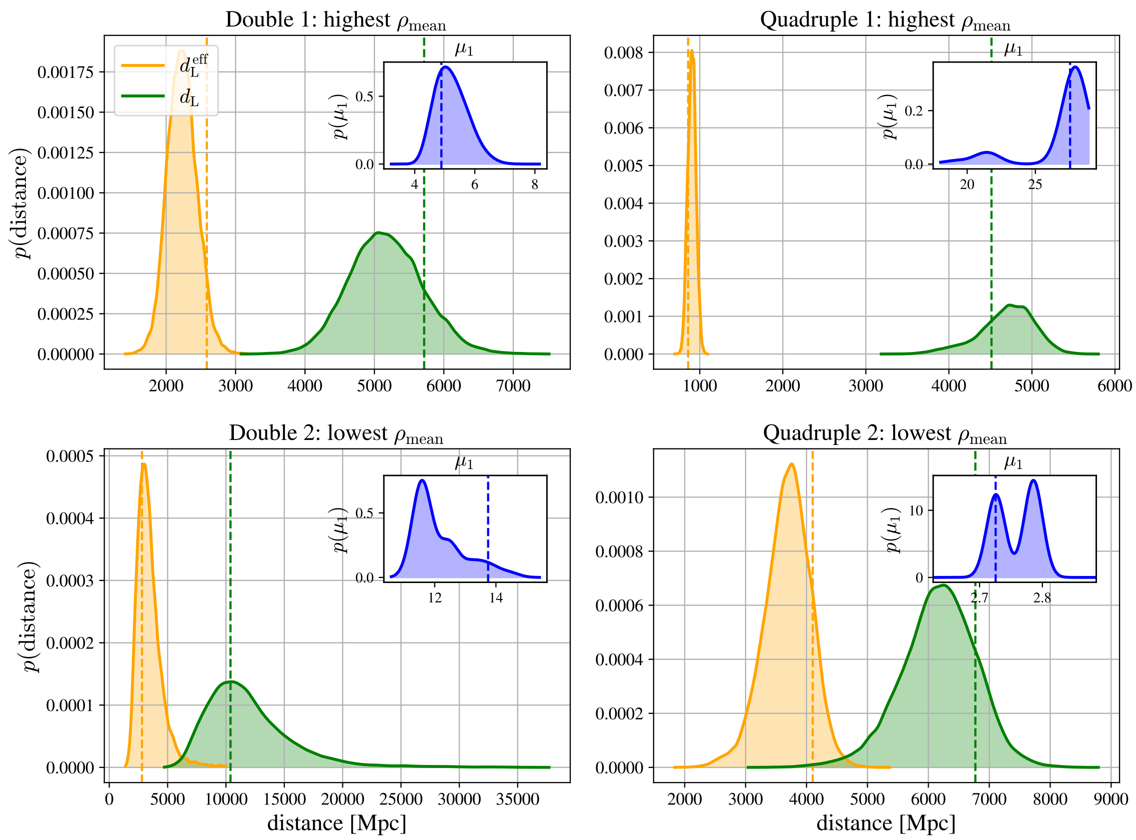

Figure 9.4: Posteriors of the apparent (orange) and retrieved true (green) luminosity distances for each case, along with the recovered magnification factor (blue), μ_1, for the first-arriving lensed GW signal obtained via rejection sampling.

lens reconstruction, the inference of individual magnification factors, is less accurate for double-image systems, where only two observables (μ_{rel} and Δt) are available. In contrast, quadruple-image systems provide much higher accuracy due to the four additional constraints provided by the extra images.

Using the recovered true luminosity distances along with the photometric or spectroscopic redshifts of the identified host galaxies in the lensing systems, we compute the corresponding comoving distances and recessional velocities. The Hubble–Lemaître law is then applied to derive H_0. Figure 9.5 shows the H_0 posteriors obtained for the four representative cases. The 90 % C.I. for the quadruple-image systems, when combined with spectroscopic redshifts, are significantly narrower than those in the other cases.

To quantify the degree of constraint on the H_0 posteriors, we define the relative uncertainty, δ_{H_0}, as:

$$\delta_{H_0} \equiv \frac{90\,\%\ \text{C.I.}}{H_{0,\text{med}}}, \tag{9.4}$$

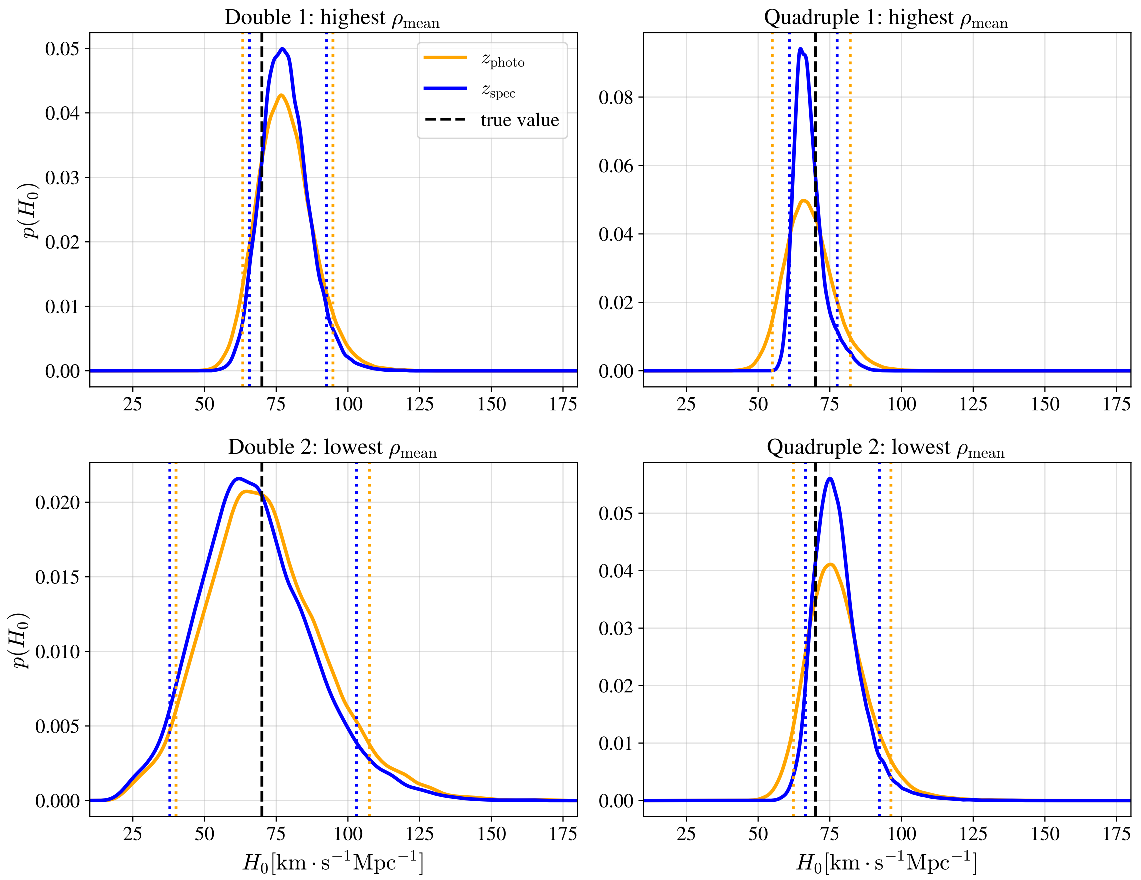

Figure 9.5: The posteriors of the Hubble constant for the four representative lensing systems.

where $H_{0,\text{med}}$ denotes the median of the posterior. Normalising the C.I. by the median, rather than the mean, which is inherently more sensitive to skewed or long-tailed distributions, provides a dimensionless measure of the fractional uncertainty, analogous to the coefficient of variation. A lower δ_{H_0} value thus indicates a more tightly constrained posterior.

For each representative system, the estimated H_0 with its 90 % C.I. and the corresponding δ_{H_0} are summarised in Table 9.3. Among systems with the same number of lensed images, a higher ρ_{mean} leads to tighter H_0 posteriors and, consequently, a lower δ_{H_0}, as expected, since higher-SNR events yield narrower posteriors for the inferred luminosity distance. Furthermore, at a comparable SNR, quadruple-image systems yield narrower posteriors than double-image systems, resulting in systematically lower δ_{H_0} values. This improvement arises because the four lensed signals provide more independent constraints, enhancing the precision of the delensed luminosity distance.

The low SNR of lensed GW signals contributes significantly to the large uncertainties, as the inferred apparent luminosity distances are often poorly constrained. However, for sources located at high redshift, the uncertainty can remain substantial regardless of the SNR. Such high-z detections generally fall into two representative cases: (1) light BBH events that become observable only through significant magnification, and (2) intrinsic-

Table 9.3: Relative uncertainties and the corresponding Hubble constant posteriors, along with their 90 % C.I., inferred using spectroscopic redshifts for the double- and quadruple-image lensing systems.

System	H_0 [km · s^{-1} · Mpc^{-1}]	δ_{H_0}
Double 1	$77.7^{+14.6}_{-11.8}$	0.34
Double 2	$67.4^{+48.4}_{-33.6}$	1.22
Quadruple 1	$66.3^{+7.2}_{-5.7}$	0.19
Quadruple 2	$76.5^{+16.3}_{-10.0}$	0.34

ally heavy BBH events that require only moderate magnification to exceed the detection threshold. In the first case, the apparent luminosity distances may be relatively well-constrained initially; however, after the delensing procedure is applied to recover the true distances, the variance increases in proportion to the magnification factor. In the second case, the apparent distances are inherently broad, as the luminosity distance-redshift relation flattens at high z, leading to increased uncertainty in the inferred distances. In both scenarios, the resulting H_0 posteriors are expected to be poorly constrained.

Figure 9.6 shows how δ_{H_0} varies across the ρ_{mean}-z_s plane for all detected lensed signals. For both doubly and quadruply lensed systems, δ_{H_0} increases towards lower ρ_{mean} and higher z_s. At high redshifts, double-image systems generally exhibit uniformly higher δ_{H_0} values. In contrast, for quadruple-image systems, the uncertainty is dominated by low ρ_{mean} rather than high z_s. This indicates that detecting quadruple-image systems is particularly beneficial for constraining H_0 from distant sources, as the additional image constraints help mitigate the uncertainties inherent to high-z observations.

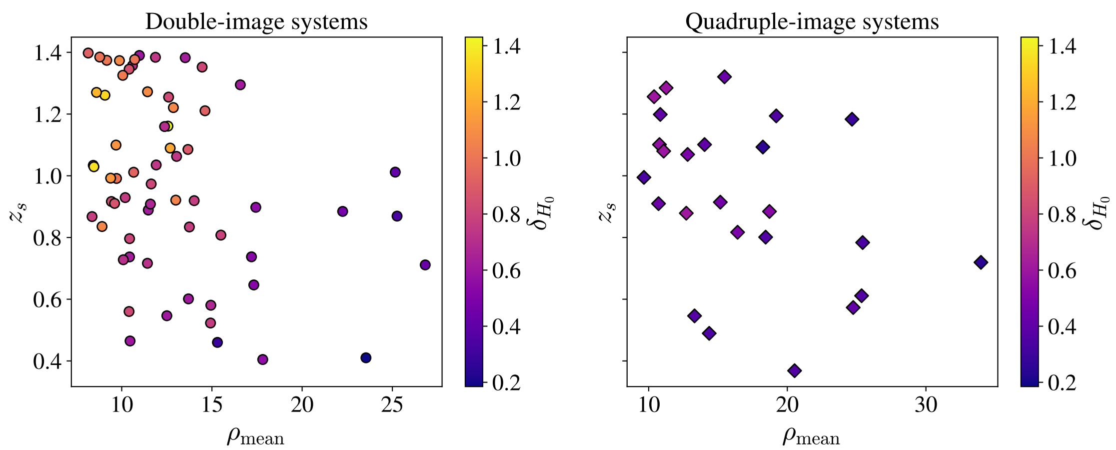

Figure 9.6: Scatter plots showing the distributions of ρ_{mean} and z_s for all detected doubly- (circle markers) and quadruply-lensed (diamond markers) GW systems in our sample pool, colour-coded by δ_{H_0}.

Beyond statistical uncertainties related to detector sensitivity, inaccurate delensing can introduce additional errors even when the apparent luminosity distance is well constrained, and the source lies in a relatively low-z universe. Such errors may arise during the rejection sampling process due to degeneracies between lensing observables and the underlying lens parameters. If the individual magnification factors obtained through delensing, which are used to recover the true luminosity distance, are not sufficiently constrained, the resulting distance posterior will broaden, losing the sharp peak required for an accurate measurement. Consequently, achieving a tightly constrained delensed luminosity distance is essential for inference of precise H_0.

In this section, we assume a complete galaxy-galaxy lensing catalogue and the unique identification of each lensing system. Under these idealised conditions, delensing is expected to be highly accurate, allowing us to neglect the associated systematic uncertainties. In practice, however, observational limitations often lead to multiple candidate host galaxies within the GW sky-localisation region, or to missing hosts that are fainter than the catalogue's magnitude threshold. These factors introduce additional uncertainty into the H_0 inference for individual events. The complexity further increases when combining information from multiple dark sirens across different lensing systems. In the following section, we describe the framework for combining multiple lensed dark sirens while accounting for catalogue incompleteness, the redshift distributions of potential hosts, and the effects of imperfect delensing.

9.3 H_0 Measurement with Multiple Strongly Lensed Dark Sirens

When multiple strongly lensed dark sirens from different lensing systems are detected, the individual H_0 posteriors discussed in Section 9.2 cannot be combined through simple multiplication. Instead, a joint likelihood must be constructed and then multiplied by the H_0 prior distribution to obtain the overall posterior distribution. In this section, we adapt `gwcosmo` [219, 220, 394][4], a Bayesian inference pipeline specifically designed for cosmological studies and population analysis using GW events, to estimate the H_0 posterior for lensed GW observations.

---

4. The `gwcosmo` code is available at `https://git.ligo.org/lscsoft/gwcosmo`.

9.3.1 Method

Given the detection of $N_{\mathrm{lens,det}}$ lensed GW systems, each consisting of multiple signals denoted collectively as $\{x_{\mathrm{lgw}}\}$ with corresponding detection flags $\{D_{\mathrm{lgw}}\}$, the posterior for H_0 is expressed as:

$$\begin{aligned} p(H_0 \mid \{x_{\mathrm{lgw}}\}, \{D_{\mathrm{lgw}}\}) \propto\, & p(H_0)\, p(N_{\mathrm{lens,det}} \mid H_0) \\ & \times \prod_{i=1}^{N_{\mathrm{lens,det}}} \mathcal{L}(x_{\mathrm{lgw},i} \mid D_{\mathrm{lgw},i}, H_0), \end{aligned} \tag{9.5}$$

where $p(H_0)$ is the prior on H_0 and $p(N_{\mathrm{lens,det}} \mid H_0)$ is the probability of detecting $N_{\mathrm{lens,det}}$ lensed systems for a given H_0. In the unlensed case, the detection count term $p(N_{\mathrm{det}} \mid H_0)$ is typically treated as independent of H_0 by adopting a rate prior of $p(R) \propto 1/R$ [219, 220, 290]. However, in the lensed case, the detection probability is a function of magnification, which is H_0-dependent. Consequently, $p(N_{\mathrm{lens,det}} \mid H_0)$ retains a dependence on H_0 through the magnification-dependent selection effects. The last term in Eq. (9.5) represents the joint likelihood of the $N_{\mathrm{lens,det}}$ systems. The likelihood for a single lensed system is defined as:

$$\mathcal{L}(x_{\mathrm{lgw}} \mid D_{\mathrm{lgw}}, H_0) = \frac{p(x_{\mathrm{lgw}} \mid H_0)}{p(D_{\mathrm{lgw}} \mid H_0)}, \tag{9.6}$$

where

$$p(D_{\mathrm{lgw}} \mid H_0) = \int p(D_{\mathrm{lgw}} \mid x_{\mathrm{lgw}}, H_0) \times p(x_{\mathrm{lgw}} \mid H_0) \mathrm{d}x_{\mathrm{lgw}}. \tag{9.7}$$

The term $p(D_{\mathrm{lgw}} \mid x_{\mathrm{lgw}}, H_0)$ takes a value of unity if the data $\{x_{\mathrm{lgw}}\}$ is successfully identified as a GW lensing system, and 0 otherwise. As discussed in Section 5.2, a system is classified as strongly lensed when multiple signals satisfy the detection criterion (in this work, $\rho_{\mathrm{net}} > 8$), and exhibit mutually consistent intrinsic parameters and sky localisations, supported by a sufficiently high Bayes factor favouring the strong-lensing hypothesis. In practice, however, the finite duty cycle of GW detectors implies that one or more signals in a lensed multiplet may be missed if the detectors are offline at the time of arrival. In such cases, only a subset of the lensed signals, potentially just a single signal, may be observed. The signal would likely be misclassified as unlensed and analysed accordingly, thereby introducing a potential systematic bias in the inferred H_0. We discuss this effect and its implications in Section 9.4.2.

The term $p(x_{\rm lgw} \mid H_0)$ represents the probability of obtaining the observed lensed GW data for a given value of H_0. In practice, $p(x_{\rm lgw} \mid H_0)$ depends on the full set of parameters that determine the SNR of the GW event, which includes not only H_0, but also the lens and source parameters (θ_l, θ_s)[5], the redshifts (z_s, z_l), and the specific sky patches (Ω) that constitute the sky-localisation map. Hereafter, all dependencies on $x_{\rm lgw}$ are written explicitly. Following this convention, Eq. (9.6) can be reformulated as:

$$
\begin{aligned}
&\mathcal{L}(x_{\rm lgw}|D_{\rm lgw}, H_0) \\
&= \left[\iiiint \sum_j^{N_{\rm pix}} p(D_{\rm lgw}|\Omega_j, \theta_s, \theta_l, z_s, z_l, H_0) p(\theta_s, \theta_l, z_s, z_l, |\Omega_j, H_0) p(\Omega_j) d\theta_s d\theta_l dz_s dz_l\right]^{-1} \\
&\times \iiiint \sum_j^{N_{\rm pix}} p(x_{\rm lgw}|\Omega_j, \theta_s, \theta_l, z_s, z_l, H_0) p(\theta_s, \theta_l, z_s, z_l, |\Omega_j, H_0) p(\Omega_j) d\theta_s d\theta_l dz_s dz_l,
\end{aligned}
\tag{9.8}
$$

where the term $p(x_{\rm lgw} \mid \dots)$ is represented in a pixelated form by partitioning the sky-localisation map into discrete patches on a HEALPix grid [411, 412][6]. In this scheme, each sky patch is assigned the subset of likelihood samples whose inferred sky locations fall within that specific region [394]. This pixelation approach is adopted to account for the sky-varying incompleteness characteristic of galaxy catalogues. Under this discretisation, the integral over the sky position is effectively replaced by a summation over the pixels.

The term $p(\theta_s, \theta_l, z_s, z_l \mid \Omega_j, H_0)$ represents the joint prior distribution of the source and lens parameters, along with their respective redshifts, for a given H_0. In contrast to the analysis of individual dark-siren events described in Section 9.2, where a unique host is identified, the source and lens redshifts are inferred statistically by cross-matching the sky localisation of the detected lensed signals with a galaxy-galaxy lensing catalogue, such as the LGC0.

For a specific sky patch Ω_j, the redshift distributions of the source and lens galaxies are determined by selecting all catalogue galaxies located within that patch. These patch-specific distributions define the *line-of-sight (LoS) redshift prior*, which is independent of the intrinsic distributions of the GW source parameters. Consequently, the joint prior can be decomposed as:

$$
p(\theta_s, \theta_l, z_s, z_l, |\Omega_j, H_0) = p(\theta_s|H_0) p(z_s, z_l, \theta_l|\Omega_j, H_0). \tag{9.9}
$$

5. θ_s denotes the intrinsic source parameters, excluding the source redshift z_s. Likewise, θ_l refers to the lens parameters other than the lens redshift z_l.
6. `http://healpix.sf.net`.

Because the intrinsic properties of the GW source—such as masses, spins, and orbital parameters—are independent of the sky location, the first term $p(\theta_s \mid H_0)$ carries no explicit dependence on Ω_j.

The denominator in Eq. (9.8) accounts for selection effects arising from the finite sensitivity of GW detectors, which introduces a bias toward louder signals in the detected GW population. Accurate modelling of these effects is required to avoid selection biases, such as Malmquist bias, and to obtain unbiased estimates of population-level parameters [292]. In the context of lensing, the selection effect is defined as the fraction of all GW sources whose lensed signals are detected with a network SNR exceeding the detection threshold, $\rho_{\rm thr} = 8$. For doubly lensed systems, we require both signals to individually exceed $\rho_{\rm thr}$. For quadruply lensed systems, at least two signals must exceed the threshold, as a minimum of two distinct detections is necessary to confidently identify a strongly lensed event.

In principle, one could perform a dedicated sub-threshold search to identify counterparts of a detected super-threshold signal that fall below the detection threshold due to insufficient magnification or demagnification [377, 413]. Furthermore, a lensed signal could potentially be identified from a single image by searching for waveform distortions in Type-II images, such as those induced by the phase shift of HOM [157, 367] as described in Section 5.1. However, these effects are neither directly relevant to the high-magnification regime nor contribute to the improvements in sky localisation. Consequently, we do not consider these scenarios in this work.

Regarding EM selection effects, although a complete catalogue was assumed in Section 9.2, realistic galaxy surveys may fail to include source-galaxy images or lens galaxies whose apparent magnitudes fall below the survey's detection threshold. The resulting incompleteness introduces additional selection effects as a part of the redshift information required for the LoS redshift prior remains missing. To account for these effects, we construct the LoS redshift prior by marginalising over the absolute and apparent magnitudes of both the source and lens galaxies, following the methodology in [220]. Since the LoS prior is defined for galaxies that are potential hosts within the GW sky-localisation region, we introduce binary parameters s and l, denoting, respectively, whether a given source galaxy hosts a GW source and whether the detected signals are genuine counterparts corresponding to multiple lensed images of a single underlying GW source.

As described in Section 9.1.1, we impose an apparent-magnitude threshold $m_{\rm th} = 21.87$ to construct incomplete catalogues from the UGC0 and LGC0. At this threshold, the resulting UGC1 contains approximately half of the galaxies in the UGC0 ($\sim 3.1 \times 10^7$ galaxies). Under the same cut, the LGC1 retains approximately 34 % of the galaxy–galaxy lensing systems from the LGC0. Using these four simulated catalogues, we compute the LoS redshift priors on a HEALPix grid. We adopt a resolution of $n_{\rm side} = 128$ for the UGCs to ensure consistency with the framework in [220]. For the LGCs, we use $n_{\rm side} = 64$ because the number of objects decreases significantly due to the small strong-lensing optical depth, making a coarser grid more appropriate for maintaining sufficient statistical coverage per pixel.

By substituting the joint likelihood of the lensed GW signals from Eq. (9.8) into Eq. (9.5), and incorporating the joint prior expressions from Eq. (9.9), the posterior for H_0 can be written explicitly as:

$$
\begin{aligned}
&p(H_0|\{x_{\rm lgw}\},\{D_{\rm lgw}\}) \propto p(H_0)p(N_{\rm lens,det}|H_0) \\
&\times \left[\iiiint p(D_{\rm lgw}|\theta_s,\theta_l,z_s,z_l,H_0)p(\theta_s|H_0)\sum_j^{N_{\rm pix}} p(z_s,z_l,\theta_l|\Omega_j,H_0,s,l)d\theta_l d\theta_s dz_l dz_s\right]^{-N_{\rm lens,det}} \\
&\times \prod_i^{N_{\rm lens,det}}\left[\iiiint \sum_j^{N_{\rm pix}} p(x_{{\rm lgw},i}|\Omega_j,\theta_s,\theta_l,z_s,z_l,H_0)p(\theta_s|H_0)p(z_s,z_l,\theta_l|\Omega_j,H_0,s,l)d\theta_l d\theta_s dz_l dz_s\right].
\end{aligned}
\tag{9.10}
$$

A comprehensive derivation of Eq. (9.10), along with a detailed explanation of the contribution of each constituent term, is provided in Appendix A.

9.3.2 Results

To estimate H_0, we construct two separate analyses using (i) unlensed signals with UGCs and (ii) lensed signals with LGCs. In each case, we draw random sub-samples from the simulated population of detectable events to emulate a realistic observing run. The first sample pool comprises 250 unlensed signals, consistent with the number of events detected until the LVK fourth observing run. Under an optimistic scenario, the second pool consists of six doubly lensed and two quadruply lensed systems. Although this number of lensed events exceeds conservative expectations[7], it is adopted to ensure sufficient statistics for reconstructing smooth posterior distributions while maintaining the expected ratio between double- and quadruple-image systems.

Figure 9.7 shows the H_0 posteriors for four analysis scenarios: unlensed and lensed GW signals combined with either complete or incomplete galaxy catalogues. Note that we employ spectroscopic redshift measurements from the galaxy catalogues. In all cases, the posteriors converge well to the true value; however, the 90 % C.I. is noticeably narrower for the lensed signals with complete catalogues. Specifically, among the lensed systems, doubly-lensed signals exhibit broader sky localisation (as shown in Figure 9.2), resulting in numerous potential host galaxies along the line of sight and complicating the identification of the true lensing system. Furthermore, the uncertainty in their luminosity dis-

7. In a realistic observing scenario, a fraction of lensed GW signals may be missed due to the finite detector duty cycle and the limited efficiency of search pipelines.

tance remains relatively large, yielding only modest constraints on H_0. In contrast, while quadruply lensed signals are less frequent, they provide substantially improved sky localisation, which significantly reduces the number of candidate host galaxies and narrows the luminosity-distance uncertainty, resulting in considerably tighter constraints on H_0.

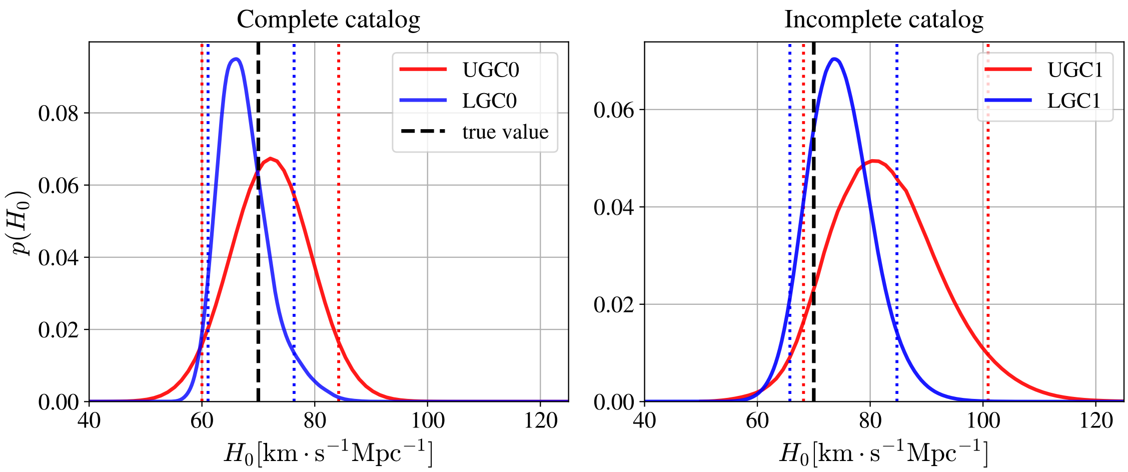

Figure 9.7: Posteriors of H_0 for the four analysis scenarios. The left panels use complete galaxy catalogues (UGC0 and LGC0), while the right panels incorporate catalogue incompleteness (UGC1 and LGC1).

In cases involving incomplete catalogues (UGC1 and LGC1), the H_0 posteriors are broader than those from the corresponding complete-catalogue cases, as the uncertainty in identifying the true host galaxy increases. In particular, in the LGC1 scenario, delensing accuracy is further reduced due to incomplete information on the lensing systems. This limitation directly increases uncertainty in the luminosity distance.

Overall, our results demonstrate that analysing even a small number of strongly lensed dark sirens with a dedicated galaxy–galaxy lensing catalogue can improve the precision of H_0 by approximately 50 % compared to a standard analysis using several hundred unlensed dark sirens. Detailed results are provided in Table 9.4. To achieve a few-percent fractional uncertainty ($\delta_{H_0} < 0.1$), comparable to the constraints provided by Type Ia supernova measurements, our framework indicates that a few tens of lensed GW systems would be required, assuming a complete catalogue. For instance, in a scenario using a randomly selected sample of 32 doubly lensed and 8 quadruply lensed events, we obtain a relative uncertainty of $\delta_{H_0} = 0.09$.

Table 9.4: Median values and 90 % C.I.s of recovered H_0 posteriors and their relative C.I. values (δ_{H_0}) for each catalogue case.

catalogue	Completeness	$H_0[\mathrm{km} \cdot \mathrm{s}^{-1} \cdot \mathrm{Mpc}^{-1}]$	δ_{H_0}
UGC0	100	$72.12^{+12.12}_{-12.12}$	0.34
UGC1	50	$82.73^{+18.18}_{-14.55}$	0.40
LGC0	100	$66.85^{+9.49}_{-5.77}$	0.23
LGC1	34	$74.29^{+10.45}_{-8.53}$	0.26

9.4 Challenges in Measurement due to Lensing Bias

In this section, we investigate potential systematic biases in H_0 inference arising from two scenarios: (i) when lensing effects are either entirely ignored by adopting a conventional unlensed framework or not fully captured by a given lens model, and (ii) when a specific number of signals are misclassified, such that lensed events are identified as unlensed and vice versa.

9.4.1 Systematic Bias in Lens Reconstruction

Accurate measurements of H_0 using lensed GW signals depend fundamentally on reliable lens reconstruction. If lensing effects are ignored or if the reconstructed lens model fails to fully capture the foreground galaxy's lensing potential, e.g., due to unresolved mass substructure, significant biases can be introduced into the inferred parameters. These biases primarily affect the recovered luminosity distance to the GW source and, by extension, the derived value of H_0. Consequently, minimising these systematic biases is essential for robust cosmological inference.

To quantify the impact of incorrect lens reconstruction and imperfect delensing, we analyse multiple simulated GW signals sharing the same underlying lensing configuration. Specifically, we adopt the galaxy-galaxy lensing configuration of the Double 1 system shown in Table 9.2. We keep the lens parameters, source distance, and all other source parameters fixed, while varying only the BBH component masses and noise realisations, thereby generating different signals that probe the same lensing geometry. This controlled setup allows us to study the impact of lens-modelling errors while accounting for variations in the source waveform and detector noise.

We simulate 30 such realisations and analyse them under three distinct modelling assumptions: (1) an unlensed configuration (no delensing applied); (2) an SIS profile, representing an incorrect lens model; and (3) an SIE profile, corresponding to the correct lens model. Figure 9.8 illustrates the H_0 posteriors inferred from individual lensed BBH events

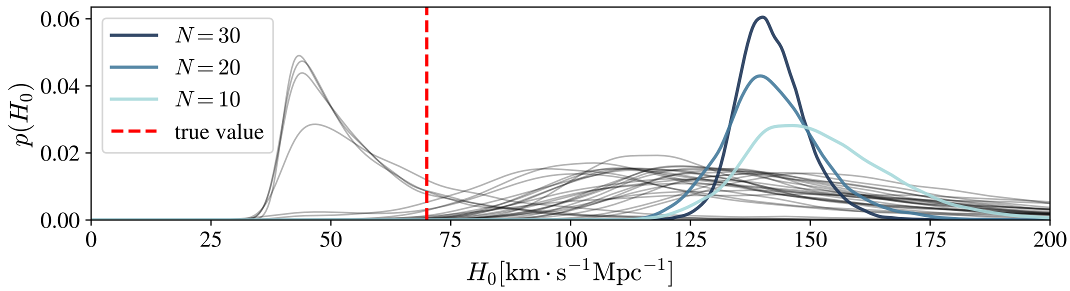

(a) No delensing applied

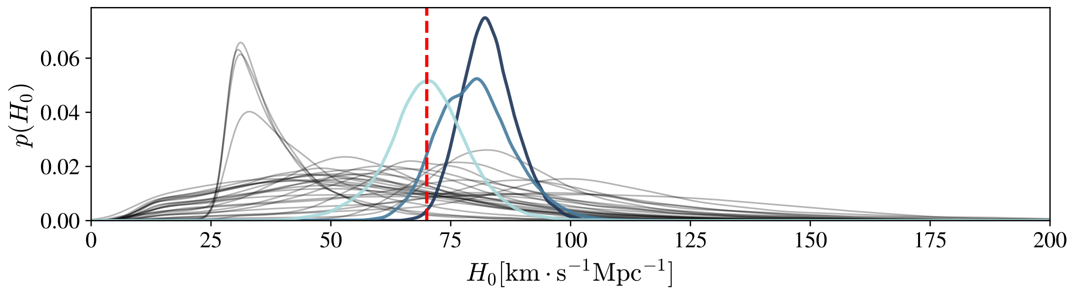

(b) Delensing with SIS model

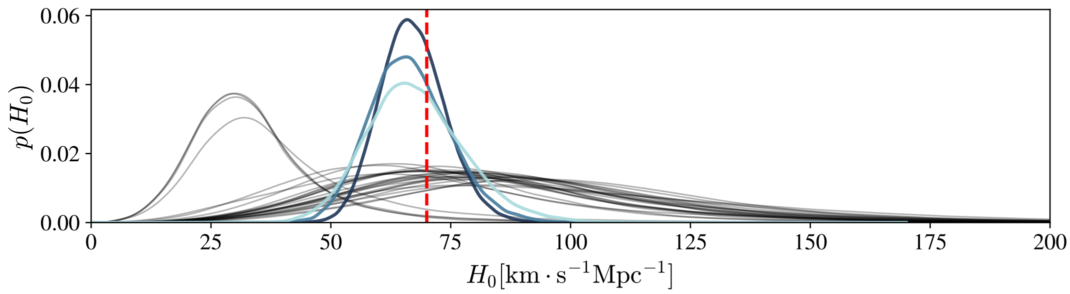

(c) Delensing with SIE model

Figure 9.8: H_0 posteriors obtained from multiple simulated GW signals sharing the same doubly lensed lensing configuration with a true SIE density profile. The BBH component masses and noise realisations are varied while keeping the lens parameters and the other source parameters fixed. The dashed red lines indicate the true value of H_0. The top, middle, and bottom panels correspond to the cases of no delensing, SIS delensing, and delensing with the correct SIE model, respectively.

(gray curves), together with the progressively combined posteriors for each case (sky-blue for $N = 10$, blue for $N = 20$, and navy for $N = 30$). While systematic biases may be difficult to diagnose from the posterior of a single realisation, combining multiple events reveals clear trends.

In the no-delensing case (Figure 9.8a), the inferred H_0 is significantly biased for all numbers of combined realisations, as uncorrected magnification leads to a systematic underestimation of the luminosity distance. For the SIE case (Figure 9.8c), the posteriors converge consistently toward the true value, with the C.I. narrowing as more realisations are accumulated. In contrast, for the SIS-delensing case (Figure 9.8b), the bias remains subtle when only $N = 10$ realisations are combined but becomes increasingly evident, as more realisations are accumulated. These results demonstrate that lens-reconstruction errors, while potentially masked by statistical noise in individual dark-siren measurements, can accumulate into significant systematic biases in joint analyses, particularly within the `gwcosmo` framework.

9.4.2 Selection Bias in Lensing Identification

In the lensed dark-siren case, selection effects are more intricate than in the unlensed case. For unlensed events, detectability is determined solely by whether a signal exceeds the threshold $\rho_{\rm thr}$, and only the population of dark sirens needs to be considered. In contrast, selection effects for lensed events must account not only for the detectability of each lensed signal but also for the requirement that multiple signals are correctly associated with the same underlying GW source. This identification process is typically quantified by the FAR of the lensing hypothesis.

Consequently, misidentification can occur in two ways: first, independent unlensed signals may be falsely identified as lensed signals from a single source (false positives); and second, a genuine lensed signal—whose counterparts are missed, delayed beyond the observation window, or too faint to be detected—may be misclassified as unlensed (false negatives). Such misidentification leads to incorrect selection-bias corrections and failures in lens reconstruction, resulting in the biased recovery of the source luminosity distance. As the number of observed events increases, the lensing misidentification can introduce substantial systematic biases into the inferred H_0, potentially dominating the statistical uncertainty. Figure 9.9 illustrates the bias in the H_0 posteriors arising from misidentification in the two scenarios described above. To isolate the effects of misidentification from uncertainties related to catalogue incompleteness, we use the complete UGC0 and LGC0 catalogues.

In the first scenario, we replace 4, 8, or 12 unlensed signals from the 250-event baseline with two, four, or six pairs of double-image signals ($N_{\rm pair}^{\rm DBL} = 2, 4$, or 6). These pairs are intentionally misclassified as unlensed events to mimic pipeline identification failures. We then estimate H_0 using the UGC0 and the selection-effect corrections appropriate for unlensed signals. Figure 9.9a shows the H_0 posteriors for these three cases, together with the reference unlensed H_0 posterior from Figure 9.7. Because these lensed signals remain undelensed, their inferred luminosity distances are systematically underestimated, as the magnification factor is typically greater than unity ($\mu > 1$). This bias shifts the inferred H_0 to higher values, as shown in Figure 9.8a. However, the impact remains modest: for every additional two $N_{\rm pair}^{\rm DBL}$, the median H_0 increases by approximately one unit (from the red solid line to the blue dotted line), as the fraction of contaminated signals remains small compared to the overall population.

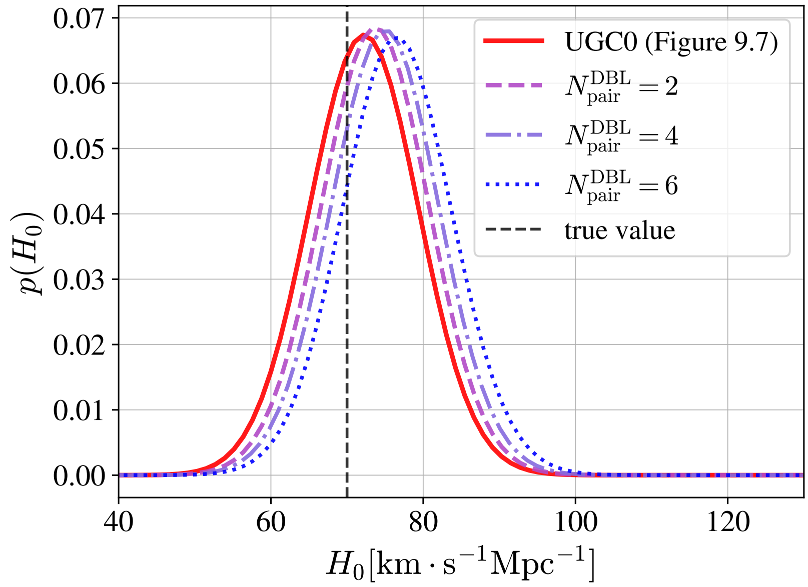

(a) H_0 inferred from an unlensed population containing hidden lensed signals

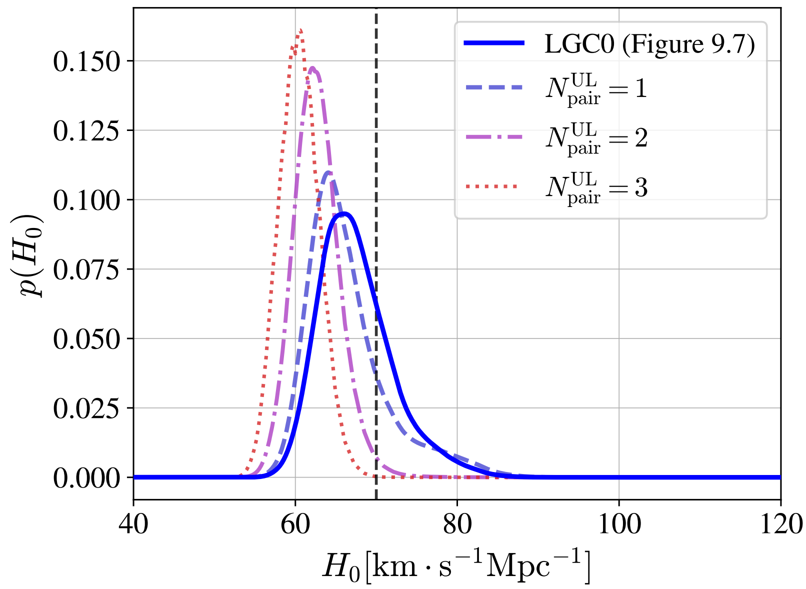

(b) H_0 inferred from a lensed population containing hidden unlensed signals

Figure 9.9: Biased H_0 posteriors obtained when a subset of signals is misclassified in a population where the majority are correctly identified. The top and bottom panels show biases arising from classifying unlensed signals as lensed and lensed signals as unlensed, respectively. The dashed black lines denote the true value of H_0, while the correctly inferred posterior from Figure 9.7 is shown for reference. As the number of misclassified events increases ($N_{\mathrm{pair}}^{\mathrm{DBL}}$ in the left panel and $N_{\mathrm{pair}}^{\mathrm{UL}}$ in the right panel), the inferred posteriors progressively shift away from the true value.

In the second scenario, we replace one, two, or three of the double-image systems out of the six doubles and two quads with an equivalent number of independent unlensed signals ($N_{\mathrm{pair}}^{\mathrm{UL}} = 1, 2,$ or 3) that are coincidentally and incorrectly identified as lensed pairs by the search pipeline. H_0 inference is then performed using the LGC0 and the selection-effect corrections for lensed signals. Figure 9.9b shows the results for these three cases, with the reference lensed H_0 posterior from Figure 9.7 included for comparison. Since the unlensed signals are erroneously delensed, their inferred distances are systematically larger than their true values. Consequently, as the number of false positive unlensed pairs increases, the H_0 posterior shifts progressively toward lower values (from the blue solid line to the red dotted line). Compared to the first scenario, the fraction of misclassified signals within this smaller population is significantly higher, resulting in a more pronounced systematic bias in the H_0 inference.

9.5 Conclusion

We have demonstrated that strong lensing significantly enhances the cosmological utility of dark sirens. When paired with a dedicated galaxy-galaxy lensing catalogue, the improved sky localisation and tighter luminosity-distance constraints provided by lensing enable significantly more informative redshift inference than in the unlensed case. Our results show that even a single strongly lensed dark siren, particularly a quadruply lensed event, can yield competitive constraints on H_0, approaching the precision of bright sirens. Furthermore, combining just a few lensed dark sirens can lead to a substantial improvement in the Hubble constant estimate, achieving constraints comparable to, or even tighter than, those obtained from several hundred unlensed events.

Robust cosmological inference from strongly lensed dark sirens fundamentally requires accurate lens reconstruction to recover the intrinsic luminosity distance to the GW source. While lens reconstruction is generally reliable with both EM and GW observations when the galactic mass distribution follows a simple profile, such as the SIE model adopted in this work, real galaxies exhibit structural complexities beyond this idealisation. Substructures and small-scale mass components embedded within galaxies can introduce additional lensing effects that may induce wave-optics signatures in GW signals. Such effects can modify the observed waveform morphology [312, 313, 314, 315, 316, 317, 318, 319]. Consequently, accurate reconstruction in these regimes requires incorporating wave-optics corrections associated with such small-scale structures.

Beyond these small-scale complexities, lens reconstruction for cosmological inference is subject to a more fundamental large-scale limitation: the mass-sheet degeneracy (MSD) [414]. The MSD arises because observed image positions and lensing observables remain invariant under a transformation that both rescales the projected mass density and introduces a uniform mass sheet. As a result, lensing data alone cannot uniquely determine the absolute mass normalisation or the time-delay distance, leading to a direct degeneracy with the inferred value of H_0.

The MSD can be broken by multiple sources at different redshifts lensed by the same galaxy, although such configurations are exceptionally rare, or by introducing independent physical constraints. A powerful tool is stellar kinematics, specifically the measurement of the velocity dispersion of stars bound to the lens galaxy [415]. These measurements probe the lens's gravitational potential independently of the lensing geometry. By jointly modelling the lensing configuration and stellar kinematics, the mass normalisation inferred from lensing (e.g., the projected mass within the Einstein radius) can be directly compared with that inferred from the dynamics, thereby constraining the mass-sheet transformation parameter. Furthermore, in the wave-optics regime, distortions induced by substructure can provide additional constraints to help break the degeneracy [416].

As discussed, strongly lensed GW systems can provide tight constraints on H_0, even when their sample size is significantly smaller than that of unlensed events. Increasing the number of lensed detections to several tens of systems further tightens the precision to the few-percent level. Next-generation ground-based detectors are expected to detect a large number of SLGWs, potentially $\mathcal{O}(100)$ events, assuming a relative lensed-to-unlensed detection rate of 10^{-3}, as the lensing optical depth increases with the detector horizon. These facilities may even detect demagnified lensed signals that currently lie below the detection threshold. With improved sky localisation and higher mean SNRs, the identification and characterisation of lensed GW sources will be significantly enhanced compared to current-generation detectors [417].

Furthermore, space-based detectors such as LISA [418] will achieve sub-degree-level localisation even for unlensed events [419]. Strongly lensed supermassive black hole mergers and extreme mass-ratio inspirals will thus provide exceptionally powerful probes for constraining H_0 [420]. However, galaxy catalogue incompleteness is expected to worsen as detector horizons extend to higher redshifts. The impact of this incompleteness will become increasingly significant, as cosmological inference will be progressively dominated by prior assumptions in regions where the catalogue is sparse.

In addition, since our work focuses exclusively on H_0, we do not consider hyperparameters describing the population properties of BBHs and lens galaxies, nor additional cosmological parameters such as the matter density, Ω_m, or the dark energy equation-of-state, w_0 introduced in Section 3.2.1. Strong lensing, however, has the potential to provide complementary constraints on these parameters [417, 421]. A fully hierarchical Bayesian analysis that incorporates population-level parameters for both sources and lenses may yield a more reliable inference of H_0 alongside Ω_m and w_0. In particular, realistic population models for lens galaxies, including the halo mass function, redshift distribution, and structural and luminosity distributions, should be incorporated to improve the framework.

In summary, we have developed distinct frameworks for inferring H_0 from single and multiple lensed dark sirens, respectively, incorporating lensed galaxy catalogues and accounting for lensing selection effects, including detection efficiency and catalogue incompleteness. We systematically quantify how the SNR of lensed GW signals, intrinsic source redshift, and delensing uncertainties affect the precision of recovered luminosity distances and the resulting H_0 posterior. Our analysis further demonstrates that even a small fraction of misclassified events can induce measurable systematic biases in H_0. These results highlight the importance of accurate lens modelling, robust lens identification, and a consistent treatment of selection effects in future GW cosmology involving strong lensing.

Chapter 10

Gravitational-Wave Microlensing as a Probe of Compact Object Populations

Disclaimer

The following chapter is based on the paper *Gravitational-Wave Microlensing as a Probe of Compact Object Populations*. The author is the first author and contributed to all stages of this work, including data generation, development of the analysis framework, data analysis, interpretation, scientific discussion of the results, manuscript preparation, and generation of all figures. This paper is currently in preparation and will be submitted to the *Astrophysical Journal Letters*.

Unlike galaxy surveys or star-cluster searches in the local Universe, detecting small and faint compact objects at cosmological distances remains exceptionally challenging. This difficulty is particularly severe for non-luminous objects that emit no EM radiation, leaving significant uncertainties in the compact-object population. GW lensing provides a unique observational probe of such otherwise inaccessible compact objects. In particular, compact lenses can produce characteristic wave-optics signatures in GW signals, imprinting distinctive frequency-dependent diffraction patterns on the waveform, as introduced in Section 2.4.2. As discussed in Section 5.1.1, these signatures can additionally provide direct constraints on the redshifted lens mass, offering information that is difficult to obtain through traditional EM microlensing observations, which primarily probe magnification variability [422, 423].

Constraining the abundance and mass function of these compact objects would provide important insight into their formation channels and their contribution to the overall matter content of the Universe [424, 425]. Such measurements could further help determine whether the observed compact-object population can be explained solely by conventional stellar remnants or requires additional exotic components, such as IMBHs or PBHs [426, 427]. However, individual GW lensing events generally provide only limited information. In most cases, the waveform primarily constrains the redshifted lens mass, leaving a degeneracy between the intrinsic lens mass and lens redshift. As a result, robust population inference requires a statistical framework that combines information from many detections.

Hierarchical Bayesian inference introduced in Section 4.4 provides a natural approach to this problem. By jointly modelling the distribution of lens masses and redshifts across multiple events, one can infer the underlying population properties and constrain the hyperparameters governing the compact-object population. In particular, incorporating lens-redshift information can help alleviate the mass–redshift degeneracy and enable direct constraints on the intrinsic lens-mass distribution and its cosmological evolution.

As mentioned in Chapter 6, no compelling evidence for GW microlensing has yet been identified in current observations. Current analyses, therefore, primarily place constraints through the absence of microlensing detection [428]. Future GW observatories, however, are expected to substantially increase the number of detected GW events lensed by compact objects whose masses fall within the range that can produce observable wave-optics effects in their frequency band [429]. This motivates a quantitative investigation of how accurately the underlying compact-object population can be reconstructed from a large sample of microlensed GW events.

In this Chapter, we investigate the detectability of microlensed GW signals and assess how populations of compact lenses can be constrained using hierarchical inference from multiple microlensing detections. Throughout this chapter, we adopt the PLANCK15 cosmology [174] as our fiducial baseline model.

10.1 Methodology

10.1.1 Mock Microlensed GW Signals

To simulate microlensed GW signals, we consider only BBHs, whose detection horizon with current ground-based detectors extends to the largest distances among CBC sources. Source parameters are drawn from the population models described in Section 4.2.2 with hyperparameters inferred from GWTC-3 analysis [293] to generate a synthetic BBH population. GW signals are then generated using the XPHM-ST waveform model [288, 289].

We draw microlensing parameters based on a PM population described by two hyperparameters: n_0, the comoving number density of microlenses, and $\alpha_{\rm ML}$, the slope of the lens-mass distribution. Following [427], which demonstrated that IMBH number densities are well constrained around $n_0 = 10^6$ Mpc^{-3}, this fiducial value is adopted. We assume a power-law distribution for the intrinsic lens mass (M_l), motivated by stellar lenses following the initial mass function and compact objects observed by the LVK network [237, 430]:

$$p(M_l \mid \alpha_{\rm ML}) \propto M_l^{-\alpha_{\rm ML}}, \tag{10.1}$$

where we set a limited lens-mass range $M_l \in [1, 10^5]M_\odot$. Above this upper bound, lensing effects shift from the wave-optics regime into the geometrical-optics regime within the LVK frequency band. We adopt $\alpha_{\rm ML} = [1, 2]$ for the M_l distribution to explore two representative population scenarios. The choice $\alpha_{\rm ML} = 1$ corresponds to a log-uniform mass distribution, providing a relatively broad and agnostic baseline model. In contrast, $\alpha_{\rm ML} = 2$ represents a steeper distribution in which low-mass lenses dominate the population more strongly. Comparing these two cases allows us to assess how constraints on the underlying lens population depend on the steepness of the mass function of that population.

Since only a fraction of GW signals are expected to undergo detectable microlensing, we first estimate the probability that a BBH signal generated at source redshift z_s is microlensed by a compact object located at $z_l < z_s$.

The probability that a GW signal generated at z_s encounters a microlens at $z_l < z_s$ is related to the optical depth, given by Eq. (2.28). For a given comoving number density n_0, the optical depth can be written as:

$$\tau(z_s) = \int dM_l \int_0^{z_s} dz_l n_0 \pi y_{\rm max}^2 R_E^2 \frac{dV_c}{dz_l}, \tag{10.2}$$

where V_c and R_E are the comoving volume and Einstein radius, respectively. The factor $y_{\rm max}$ is introduced because, in the wave-optics regime, diffraction effects allow detectable microlensing signatures to extend beyond R_E commonly associated with EM microlensing ($y_{\rm max} = 1$). Consequently, the effective lensing cross-section expands to $\sigma = \pi y_{\rm max}^2 R_E^2$ [147]. Using Eq. (2.31) together with the simulated BBH population and Eq. (10.2), we calculate the population-averaged optical depth and determine the expected number of microlensed GW events in our simulations. As will be discussed later (see Figure 10.1), we adopt $y_{\rm max} = 2.75$ and $y_{\rm max} = 2.07$ for the $\alpha_{\rm ML} = 1$ the $\alpha_{\rm ML} = 2$ populations, respectively.

Since GW microlensing measurements constrain the detector-frame lens mass, i.e., the redshifted lens mass M_l^z, it is convenient to express the microlensing probability directly in terms of M_l^z. Combining the optical depth with the assumed M_l distribution $p(M_l|\alpha_{\rm ML})$, the probability that a BBH at redshift z_s is microlensed by a lens with detector-frame mass M_l^z is given by:

$$\begin{aligned} p(M_l^z \mid \alpha_{\rm ML}, z_s) &= \tau(z_s) p(M_l|\alpha_{\rm ML}) \\ &\propto \int_0^{z_s} dz_l \, (M_l^z)^{1-\alpha_{\rm ML}} \, (1+z_l)^{\alpha_{\rm ML}-2} \frac{D_l D_{ls}}{D_s} \frac{dV_c}{dz_l}. \end{aligned} \tag{10.3}$$

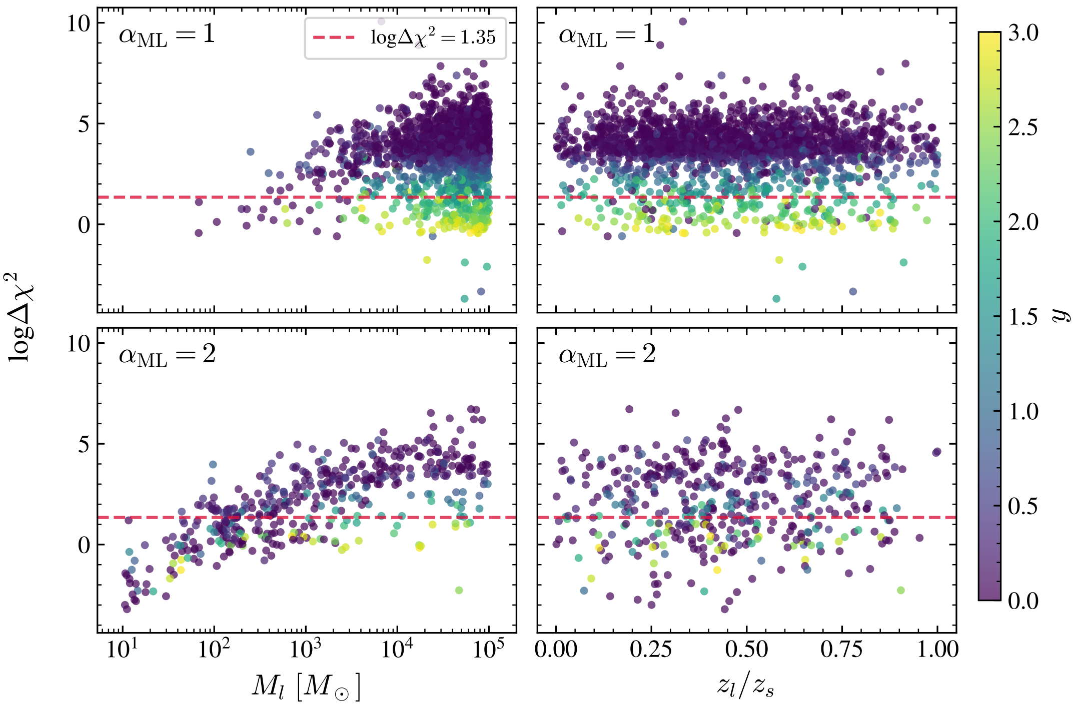

Figure 10.1: Distributions of $\log\Delta\chi^2$ between microlensed and corresponding unlensed waveforms for signals satisfying the network detection threshold $\rho_{\rm net} > 8$, shown as a function of M_l and the redshift ratio z_l/z_s. The colour scale represents y. The red line indicates the adopted distinguishability criterion.

Using Eq. (10.3), we sample M_l^z within the range $[1, 10^5]M_\odot$. Next, y values are sampled from a geometric prior distribution, $p(y) \propto y$, within the range $[0.01, 3.0]$. This standard distribution naturally arises from the differential lensing area under the assumption of a spatially uniform source population in the source plane. Extremely small values of y are statistically disfavoured because they correspond to a small geometric lensing cross-section, whereas observable microlensing distortions become increasingly suppressed at large y.

Although the optical depth determines the expected number of microlensed GW events, only a subset of these events produces detectable wave-optics signatures in GW observations. The detectability of GW microlensing depends not only on the network SNR, $\rho_{\rm net}$, but also on the waveform mismatch ϵ between microlensed and unlensed signals, which quantifies the degree of waveform distortion induced by wave-optics effects (see Eqs. (6.2) and (7.9)).

Following [431], we adopt a fiducial distinguishability criterion based on the statistic $\Delta\chi^2$, which is equivalent to Λ in Eq. (1.43) but is interpreted here as a measure of how distinguishable a microlensed waveform is from its unlensed counterpart. In the limit of small waveform mismatches, $\Delta\chi^2$ can be approximated as:

$$\Delta\chi^2 \simeq 2\rho_{\rm net}^2\epsilon_{\rm net}, \tag{10.4}$$

where $\epsilon_{\rm net}$ denotes the network-averaged mismatch computed from the individual detector signals. We assume the conventional GW detection threshold $\rho_{\rm net} > 8$ together with a fiducial mismatch threshold $\epsilon_{\rm net} > 0.03$, corresponding to the approximate level at which waveform distortions become distinguishable in GW observations. These values yield $\Delta\chi^2 = 3.84$ ($\log\Delta\chi^2 \simeq 1.35$).

Figure 10.1 shows the resulting $\Delta\chi^2$ values between simulated microlensed waveforms and their corresponding unlensed counterparts for signals satisfying the SNR detection threshold ($\rho_{\rm net} > 8$), as a function of M_l and z_l/z_s for both $\alpha_{\rm ML} = 1$ and $\alpha_{\rm ML} = 2$ populations. In both cases, $\Delta\chi^2$ shows a decreasing trend with increasing y, reflecting the progressive suppression of wave-optics distortions at larger y. We estimate $y_{\rm max}$ used in Eq. (10.2) based on the adopted $\Delta\chi^2$ threshold.

In addition, a positive correlation between M_l and $\Delta\chi^2$ is evident, as expected from the stronger and more distinguishable waveform distortions produced by more massive lenses. Since the $\alpha_{\rm ML} = 1$ population contains a larger fraction of high M_l lenses than the steeper $\alpha_{\rm ML} = 2$ distribution, a correspondingly larger fraction of simulated events satisfy the final distinguishability criterion in the $\alpha_{\rm ML} = 1$ case. The minimum lens mass capable of producing detectable microlensing signatures is approximately $\mathcal{O}(2)M_\odot$ for $\alpha_{\rm ML} = 1$, whereas detectable signals in the $\alpha_{\rm ML} = 2$ case can extend down to approximately $\mathcal{O}(1)M_\odot$. Moreover, no clear correlation is observed between $\Delta\chi^2$ and z_l/z_s. While the lensing probability is maximised for lenses located roughly midway between the source and observer due to geometric effects, the strength of observable wave-optics distortions depends more strongly on $M_l^z$and y.

It is important to note that some simulated microlensed signals may satisfy the distinguishability criterion while still failing the network SNR requirement ($\rho_{\rm net} > 8$). Such events would not be initially detectable. We therefore retain only events satisfying both detectability criteria and classify them as detectable microlensed GW signals. From this detectable population, we randomly select 200 microlensed events and inject them into simulated aLIGO–Livingston, aLIGO–Hanford, and AdV detector data, assuming their projected O4a sensitivity and Gaussian-noise realisations [365].

10.1.2 Hierarchical Bayesian Inference of Hyperparameters

To estimate the parameters of microlensed waveforms, we adopt the `Gravelamps` [298]. As shown in Chapter 6, the pipeline yields posterior distributions for the PM lens parameters, M_l^z and y, together with the Bayesian evidence for the microlensing hypothesis, $\mathcal{B}_{\rm ML}$. Performing parameter estimation for the same signal under the unlensed hypothesis additionally yields the corresponding evidence $\mathcal{B}_{\rm U}$, allowing the Bayes factor between microlensing and unlensed hypotheses, $\mathcal{B}^{\rm ML}_{\rm U}$, to be evaluated. A microlensed signal is expected to favour the lensing hypothesis when $\mathcal{B}^{\rm ML}_{\rm U} > 1$. In realistic searches, microlensing detections would additionally require a robust FPP assessment based on a large population of unlensed background events, as discussed in Section 5.1. In this work, however, we adopt a pragmatic criterion and regard microlensed events with $\mathcal{B}^{\rm ML}_{\rm U} > 1$ and well-constrained posteriors for M_l^z and y as confidently identified microlensing events, yielding a total of 132 events.

From the identified microlensing events, the posterior distribution of the hyperparameters (λ) can be inferred through a hierarchical Bayesian framework that combines information from all detected events ($\boldsymbol{d}$) (see Eqs. (4.27) and (4.29)). The corresponding posterior is given by:

$$\begin{aligned} p(\lambda \mid \boldsymbol{d}) \propto\ & p(\lambda) p(N_{\rm ML,det} \mid \lambda) \\ & \times \frac{\prod_i^{N_{\rm ML,det}} p(d_i | \boldsymbol{\theta_{s,i}}, M^z_{l,i}, y_i) p_{\rm pop}(\boldsymbol{\theta_{s,i}}, M^z_{l,i}, y_i \mid \lambda)}{\beta(\lambda)^{N_{\rm ML,det}}}, \end{aligned} \tag{10.5}$$

where $N_{\rm ML,det}$ is the number of detected microlensed events and $\beta(\lambda)$ indicates the detection efficiency for a given hyperparameter set, corresponding to a selection-bias correction [291, 292]. The hyperparameter set consists of both source-population and lens-population parameters, i.e. $\lambda = (\lambda_s, \lambda_l)$. Here λ_s describes the BBH population e.g., introduced in Eqs. (4.35)–(4.39), while $\lambda_l = (\alpha_{\rm ML}, n_0)$ in this work.

The first term in Eq. (10.5) is the prior on the hyperparameters. In this work, we fix λ_s parameters to their injected values to focus exclusively on the inference of λ_l. We adopt broad priors for λ_l, namely $\alpha_{\rm ML} \in$ U$[-4, 12]$, consistent with the power-law index priors for the primary BH of a BBH commonly used in LVK population studies [237, 293], and $\log_{10} n_0 \in$ U$[0, 10]$, following the range adopted in [427].

The second term, $p(N_{\rm ML,det} \mid \lambda)$, describes the probability of detecting $N_{\rm ML,det}$ microlensing events, which is modelled using a Poisson-distributed variable with an expected number given by Eq. (4.28):

$$\langle N_{\rm ML,det} \rangle = N_{\rm UL} \bar{\tau}(\lambda) \int p_{\rm pop}(\boldsymbol{\theta_s}, M_l^z, y \mid \lambda) P_{\rm det}(\boldsymbol{\theta_s}, M_l^z, y)\, d\boldsymbol{\theta_s} dM_l^z dy \tag{10.6}$$

where $N_{\rm UL}$ is the mean of the Poisson process for unlensed events (i.e., the expected total number of unlensed events, including both detected and undetected ones) and $\bar{\tau}(\lambda)$ is the population-averaged optical depth defined in Eq. (2.31), which depends on the given hyperparameters. $P_{\rm det}$ denotes the indicator that a microlensed signal with the given source and lens parameters satisfies the adopted detection and distinguishability criteria.

The third term in Eq. (10.5) represents the joint likelihood of all detected microlensed events, constructed as a product of individual event contributions, assuming that each event is statistically independent given λ. This assumption allows the joint likelihood to be factorised into a product of event-level terms. In the numerator, $p(d_i \mid \theta_{s,i}, M^z_{l,i}, y_i)$ denotes the event-level likelihood of each set of observed microlensed GW data, evaluated using posterior samples obtained from parameter estimation. The term $p_{\text{pop}}(\theta_s, M_l^z, y \mid \lambda)$ encodes the underlying astrophysical population model for both the source and lens parameters. In this work, the lens population is described by M_l^z distribution given in Eq. (10.3), together with the geometric prior on y.

In the denominator, $\beta(\lambda)$ represents the probability that a source drawn from the underlying population is successfully identified as a detectable microlensing event under the adopted detector sensitivity and distinguishability criteria. It can be interpreted as an overall detection efficiency that accounts for selection effects induced by both the SNR threshold and waveform distinguishability requirements. Based on Eq. (4.30), the detection efficiency for microlensed signals can be rewritten as:

$$\begin{aligned}\beta(\lambda) &\equiv \frac{\langle N_{\text{ML,det}}(\lambda)\rangle}{N_{\text{UL}}\bar{\tau}(\lambda)} \\ &= \int p_{\text{pop}}(\boldsymbol{\theta_s}, M_l^z, y \mid \lambda) P_{\text{det}}(\boldsymbol{\theta_s}, M_l^z, y)\, d\boldsymbol{\theta_s} dM_l^z dy.\end{aligned} \tag{10.7}$$

As shown in Eq. (4.32), $\beta(\lambda)$ is estimated using numerous microlensed injections, as follows:

$$\beta(\lambda) \simeq \frac{1}{N_{\text{inj}}} \sum_i^{N_{\text{inj,det}}} \frac{p_{\text{pop}}(\boldsymbol{\theta_{s,i}}, M_{l,i}^z, y_i \mid \lambda)}{\pi(\boldsymbol{\theta_{s,i}}, M_{l,i}^z, y_i)}, \tag{10.8}$$

where N_{inj} and $N_{\text{inj,det}}$ denote the total number of injected and successfully detected microlensed signals, respectively, and π is the prior distribution used in the injection campaign.

10.2 Results

10.2.1 Inference of the Compact-Lens Distribution

From Eq. (10.5), we infer the joint posterior distribution for α_{ML} and n_0. To evaluate the minimum sample size required to obtain meaningful constraints on these population parameters, we analysed three scenarios with varying numbers of detected microlensed events: $N_{\text{ML,det}} = 50, 100$, and the full sample of 132. As shown in Figure 10.2, the widths of posteriors for both hyperparameters become progressively narrower as the number of detections increases, while the posterior peaks converge toward their respective true values. Notably, n_0 is less well constrained than α_{ML}. This difference arises from its strong degeneracy with y, which more directly governs the microlensing probability and magnification than the lens mass M_l^z.

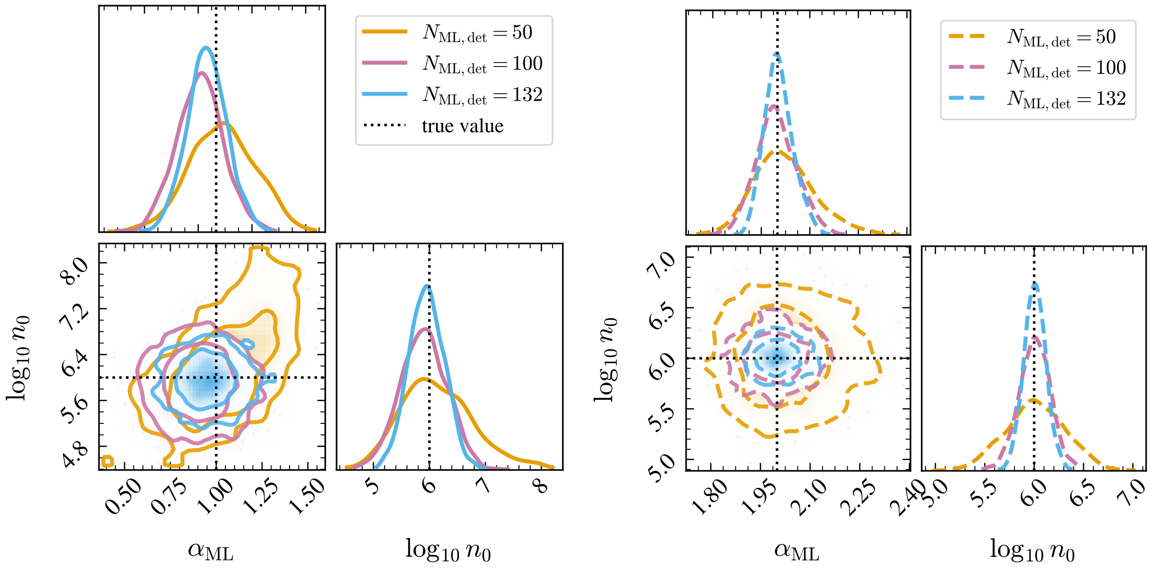

Figure 10.2: Corner plots showing the posteriors of α_{ML} and $\log_{10} n_0$ inferred from the $N_{\mathrm{ML,det}} = 50$, 100, and 132 cases for the two fiducial scenarios $\alpha_{\mathrm{ML}} = 1$ (left panel) and $\alpha_{\mathrm{ML}} = 2$ (right panel). The black dotted lines denote the corresponding true values.

Furthermore, in the $\alpha_{\mathrm{ML}} = 1$ scenario, the $N_{\mathrm{ML,det}} = 50$ case exhibits a positive correlation between α_{ML} and n_0. When the number of detected events is limited, uncertainties in α_{ML} can be partially compensated by shifts in n_0. Physically, larger values of α_{ML} correspond to lens populations dominated by lower-mass compact objects, which generally produce weaker microlensing signatures. Consequently, a higher lens number density is required to maintain a comparable detection rate, giving rise to the observed α_{ML}–n_0 degeneracy. In contrast, for larger catalogues ($N_{\mathrm{ML,det}} = 100$ and 132), the shape of the lens-mass distribution becomes much more tightly constrained by the data itself. The inferred value of α_{ML} is then effectively determined, substantially reducing the allowed compensation through n_0. As a result, the α_{ML}–n_0 degeneracy is progressively broken, yielding nearly circular 2D posteriors.

In the $\alpha_{\mathrm{ML}} = 2$ scenario, the intrinsic lens-mass distribution is strongly concentrated toward M_{min}. Even the $N_{\mathrm{obs}} = 50$ case is therefore dominated by low-mass lenses, and the absence of high-mass lenses provides strong constraints on α_{ML}, thereby also constraining n_0.

10.2.2 Constraining the Microlens Fraction in Dark Matter

According to [432], the comoving number density n_0 appearing in Eq. (10.2) can be written as

$$\begin{aligned} n_0 &= \frac{\rho_l}{\langle M_l \rangle} \\ &= \frac{\rho_c \Omega_l}{\langle M_l \rangle} \\ &= \frac{3H_0^2 \Omega_l}{8\pi G \langle M_l \rangle}, \end{aligned} \tag{10.9}$$

where ρ_l and ρ_c denote the compact-lens density and the critical density of the Universe, respectively, and $\Omega_l \equiv \rho_l/\rho_c$ is the corresponding density parameter of compact lenses. $\langle M_l \rangle$ denotes the expectation value of the intrinsic lens mass, analytically derived from the assumed mass distribution $p(M_l|\alpha_{\rm ML})$ within the mass range $[M_{\rm min}, M_{\rm max}]$. For a given $\alpha_{\rm ML}$, the theoretical $\langle M_l \rangle$ is given by:

$$\langle M_l \rangle = \begin{cases} \left(\dfrac{1-\alpha_{\rm ML}}{2-\alpha_{\rm ML}}\right) \dfrac{M_{\rm max}^{2-\alpha_{\rm ML}} - M_{\rm min}^{2-\alpha_{\rm ML}}}{M_{\rm max}^{1-\alpha_{\rm ML}} - M_{\rm min}^{1-\alpha_{\rm ML}}}, & \alpha_{\rm ML} \neq 1,\ 2, \\ \dfrac{M_{\rm max} - M_{\rm min}}{\ln(M_{\rm max}/M_{\rm min})}, & \alpha_{\rm ML} = 1, \\ \left(\dfrac{M_{\rm min} M_{\rm max}}{M_{\rm max} - M_{\rm min}}\right) \ln\left(\dfrac{M_{\rm max}}{M_{\rm min}}\right), & \alpha_{\rm ML} = 2. \end{cases} \tag{10.10}$$

If the compact lenses are interpreted as dark compact objects, such as PBHs, the lens density parameter can be expressed as $\Omega_l \equiv f_{\rm DM} \Omega_{\rm DM}$, where $f_{\rm DM}$ denotes the integrated fraction of total DM contained in compact lenses across the prescribed mass range. Consequently, constraints on n_0 and $\alpha_{\rm ML}$ obtained from the hierarchical Bayesian analysis can be translated into constraints on $f_{\rm DM}$ assuming a fiducial DM density parameter $\Omega_{\rm DM} \simeq 0.26$. Under these assumptions, the DM fraction associated with compact lenses can be written as:

$$f_{\rm DM} = \frac{8\pi G n_0 \langle M_l \rangle}{3H_0^2 \Omega_{\rm DM}}. \tag{10.11}$$

Figure 10.3 shows the posteriors of $f_{\rm DM}$, calculated from the posterior samples of n_0 and $\alpha_{\rm ML}$ for all detected microlensed signals. Here, $f_{\rm DM}$ denotes the fraction of DM in the compact objects within the mass range $[M_{\rm min}, 10^5]\, M_\odot$, where $M_{\rm min}$ represents the minimum lens mass capable of producing detectable microlensing signatures for a given $\alpha_{\rm ML}$ scenario (see Figure 10.1). Note that $f_{\rm DM} = 1$ implies that all DM is comprised of compact objects within this mass range.

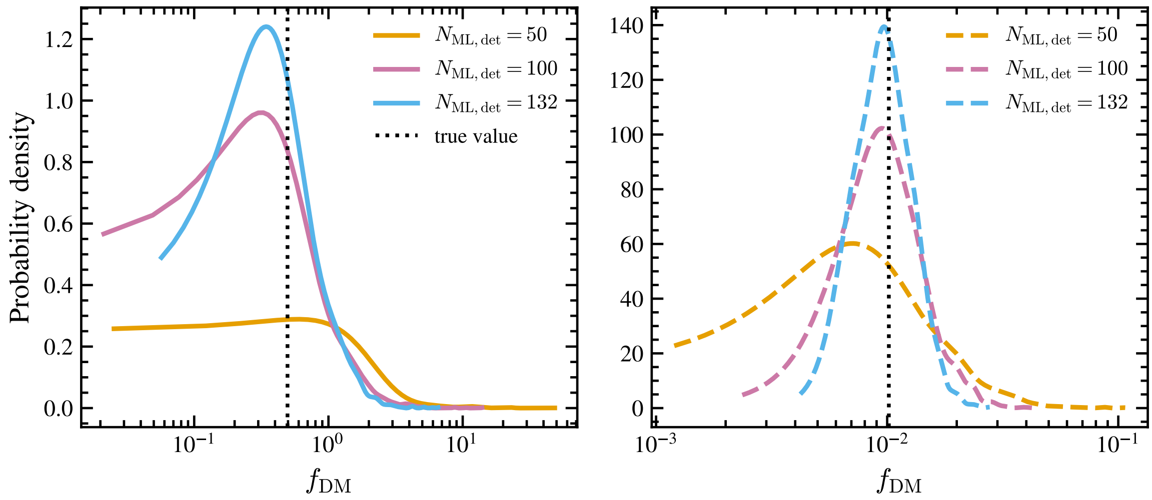

Figure 10.3: Posteriors of f_{DM} inferred for the $\alpha_{\mathrm{ML}} = 1$ (left panel) and $\alpha_{\mathrm{ML}} = 2$ (right panel) scenarios using catalogues with $N_{\mathrm{ML,det}} = 50$, 100, and 132 detected microlensing events.The black dotted lines denote the corresponding true values.

In the $\alpha_{\mathrm{ML}} = 1$ scenario, the average lens mass is relatively high, corresponding to a true DM fraction of $f_{\mathrm{DM}} \simeq 0.5$. The posterior obtained from the $N_{\mathrm{ML,det}} = 50$ catalogue remains relatively flat over a broad range of f_{DM}. As discussed in Figure 10.2, this behaviour arises from the degeneracy between α_{ML} and n_0. Nevertheless, the lower bound on f_{DM} becomes progressively tighter as $N_{\mathrm{ML,det}}$ increases.

By contrast, for the $\alpha_{\mathrm{ML}} = 2$ scenario, the average lens mass is substantially lower, yielding a correspondingly lower true DM fraction of $f_{\mathrm{DM}} \simeq 1.02 \times 10^{-2}$. The posterior distributions peak near the true value even for the $N_{\mathrm{ML,det}} = 50$ catalogue, reflecting the reduced α_{ML}–n_0 degeneracy discussed previously.

It is worth noting that while the physical boundary strictly requires $f_{\mathrm{DM}} \leq 1$, our posterior samples for the $\alpha_{\mathrm{ML}} = 1$ scenario extend into the unphysical regime ($f_{\mathrm{DM}} > 1$). This issue arises because the analysis was performed using broad hyperparameter priors that do not explicitly impose the physical constraint on f_{DM}. In our simplified population model, the physically allowed region in the $(n_0, \alpha_{\mathrm{ML}})$ parameter space can be estimated relatively straightforwardly. In more realistic compact-object population models, however, the mapping between hyperparameters and the implied value of f_{DM} is expected to become substantially more complex. For this reason, adopting broad priors at the hyperparameter level and subsequently excluding posterior samples with $f_{\mathrm{DM}} > 1$ provides a more flexible and model-independent approach than imposing hard physical boundaries a priori.

We can also estimate the differential DM fraction as a function of lens mass, enabling us to quantify how much DM is contributed by lenses of different masses. The differential DM fraction per $\log_{10} M_l$ is given by

$$\frac{df_{\rm DM}}{d\log_{10} M_l} = \frac{8\pi G n_0}{3H_0^2\Omega_{\rm DM}} M_l^{2-\alpha_{\rm ml}} \ln(10). \tag{10.12}$$

Figure 10.4 illustrates the resulting differential spectra. In the $\alpha_{\rm ML} = 1$ scenario, the spectrum is tightly constrained at the low-mass end but exhibits progressively broader C.I.s toward the high-mass regime. This behaviour directly reflects the error propagation in Eq. (10.12): since the mass dependence remains close to $\sim M_l^1$, even small statistical fluctuations in the inferred slope $\alpha_{\rm ML}$ become exponentially amplified at large M_l. A qualitatively similar trend is also observed in the $\alpha_{\rm ML} = 2$ scenario, although the differential spectrum remains substantially flatter. In this case, the mass dependence in Eq. (10.12) becomes negligible, thereby suppressing the propagation of errors toward larger M_l.

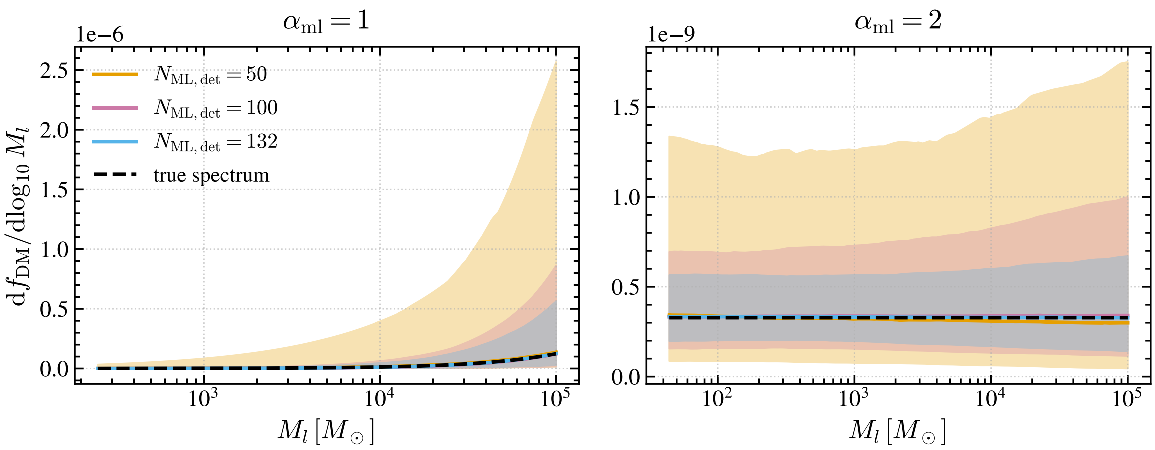

Figure 10.4: The differential fraction of DM, $\frac{\mathrm{d}f_{\rm DM}}{\mathrm{d}\log_{10} M_l}$, for the $\alpha_{\rm ML} = 1$ (left panel) and $\alpha_{\rm ML} = 2$ (right panel) scenarios using catalogues with $N_{\rm ML,det} = 50$, 100, and 132 detected microlensing events. The black dashed lines denote the corresponding true spectra. The solid coloured lines denote the median recovered spectra.

10.3 Conclusion

In this chapter, we have investigated the feasibility of constraining compact-object populations through hierarchical Bayesian inference applied to a catalogue of confidently identified GW microlensing events, while consistently accounting for selection effects. Using simulated microlensed BBH signals drawn from two representative power-law lens-mass distributions, with $\alpha_{\rm ML} = 1$ and $\alpha_{\rm ML} = 2$, we demonstrated that the accumulation of microlensed detections can place meaningful constraints on both the mass-function slope $\alpha_{\rm ML}$ and the comoving lens number density n_0, while simultaneously breaking the degeneracy between these parameters.

We have further translated these constraints into the compact-DM fraction, $f_{\rm DM}$, and its differential mass spectrum. The hierarchical framework allows the hyperparameter space to be explored without imposing explicit hard priors derived from the physical bounds on $f_{\rm DM}$. Such priors become increasingly difficult to construct as the mapping between the population hyperparameters and the physically allowed DM fraction grows more complicated for realistic population models. Instead, the physical and unphysical regions can be identified directly at the level of the inferred $f_{\rm DM}$ posterior, making the framework more flexible and model-agnostic.

We found that steeper lens-mass distributions yield tighter constraints on the inferred $f_{\rm DM}$ for the same number of detections. This behaviour arises because the parameter degeneracy between $\alpha_{\rm ML}$ and n_0 is more efficiently broken for steeper mass functions, while the smaller characteristic lens masses lead to weaker amplification of the statistical uncertainties in n_0 when propagating them into $f_{\rm DM}$. In addition, steeper mass functions suppress the contribution from the high-mass regime, where statistical uncertainties are more strongly amplified in the inferred differential DM spectrum, thereby producing more stable constraints across the parameter space. Consequently, compact-object populations dominated by relatively low-mass lenses can, in principle, yield more robust constraints on the compact-DM fraction.

Although our framework provides a robust foundation for population inference, several simplifying assumptions adopted in this work motivate important directions for future refinement. First, we employed a simplified population configuration in which the source-population hyperparameters are fixed, thereby focusing exclusively on the lens population. Incorporating uncertainties in the source population would weaken the constraints on the population properties. However, such an extension would also enable a direct investigation of potential degeneracies between the source and lens populations, particularly in scenarios where both originate from dark compact objects. This possibility becomes especially relevant if a fraction of the BBH population detected by LVK is itself primordial in origin [433, 434], since the source and lens populations could then share common formation channels and population characteristics.

In addition, we adopted a simple power-law model for the intrinsic lens-mass distribution. If the compact lenses are interpreted as PBHs, more physically motivated mass functions should be considered[1]. We also assumed a constant n_0, despite the possibility that the compact-lens population evolves over cosmic time. A more realistic treatment could incorporate redshift evolution through a parameterisation such as $n_l(z) = n_0(1+z)^k$, analogous to Eq. (4.38), where k characterises the evolution of the lens abundance. Such models generally require additional hyperparameters and would introduce further correlations between the source and lens population properties.

1. For a recent overview of PBH mass-function models and population studies, we refer readers to [435].

Nevertheless, once the underlying hyperparameters are sufficiently constrained, it may become possible to infer the lens redshift distribution statistically. In particular, the recovered intrinsic lens-mass distribution may also help partially break the degeneracy between the detector-frame lens mass M_l^z and the lens redshift z_l through population-level marginalisation. Such a partial separation of z_l from M_l^z would allow the detected compact-lens population to be grouped into redshift bins, enabling the evolution of the compact-DM fraction, $f_{\mathrm{DM}}(z_l)$, to be investigated as a function of cosmic time.

Constraining the redshift evolution of the compact-lens abundance may provide an important clue to their astrophysical origin. If the inferred compact-object density increases significantly following the peak of the cosmic star-formation history ($z \sim 2$) [295], such behaviour would strongly favour a conventional stellar-remnant origin. In contrast, if compact lenses continue to be detected at higher redshifts while the inferred f_{DM} remains approximately constant over cosmic time, this would be difficult to explain through stellar evolution alone. Instead, they would provide compelling evidence for an additional exotic compact-object population, such as PBHs formed in the early Universe before the onset of star formation.

From the perspective of detectability, the microlensing selection criteria should be refined using more realistic FPP estimates derived from dedicated background injections, rather than relying solely on Bayes factors, as discussed in Chapters 5 and 6. Our injection-and-recovery framework also relied exclusively on isolated point-mass lens templates. However, as discussed in Chapter 7, strongly lensed GW signals are expected to undergo additional microlensing perturbations from stellar fields embedded within a galaxy potential, producing waveform distortions substantially more complex than those generated by isolated point-mass lenses alone. Addressing these limitations will require both more robust statistical detection frameworks and the development of more sophisticated waveform models in future analyses.

Part IV

Conclusions

Chapter 11

Executive Summary and Future Prospects

General Relativity, formulated by Albert Einstein, stands as a monumental triumph of twentieth-century physics [9, 10], predicting two cornerstone astrophysical phenomena: gravitational waves [7, 8] and gravitational lensing [88, 89]. The intricate interplay between these two predictions offers a unique avenue for probing the distribution of matter across a wide range of scales, as well as the evolutionary history of the Universe. This thesis aims to translate these profound theoretical frameworks into observational methodologies. In particular, this work addresses two central questions: *"how gravitationally lensed GW signals can be reliably identified in detector data."* (Part II) and *"how such signals can be used to extract meaningful astrophysical and cosmological information once detected."* (Part III)

The direct detection of GWs has marked a paradigm shift in modern astrophysics and multi-messenger astronomy, opening an entirely new window onto the distant Universe [14, 16, 17]. Unlike EM observations, GWs enable direct access to the most extreme and highly dynamical gravitational environments, such as CBCs, many of which are inaccessible to conventional EM observations due either to the intrinsic absence of EM emission or to observational limitations at high redshift. ***Chapter 1*** established the theoretical and observational foundations of GW physics. The mathematical framework of GW generation was introduced together with the astrophysical properties of compact binaries and other potential GW sources [21, 24, 30]. We also reviewed the operating principles of ground-based interferometric detectors and the methodologies used to identify GW signals within noisy detector data [24, 61, 66, 72].

Gravitational lensing further extends the scientific reach of GW observations. Intervening matter distributions can act as cosmic telescopes that magnify distant sources, thereby providing a unique probe of the distribution of both baryonic and dark matter along the line of sight [103]. At the same time, lensing introduces characteristic signatures into GW signals, such as relative magnifications, time delays, and wave-optics interference patterns, which encode information about both the lensing structures and the underlying source population. Motivated by these possibilities, ***Chapter 2*** developed the theoretical framework of gravitational lensing for GWs, beginning with general lensing theory and extending to lens models, lensing probabilities, and the physical distinctions between EM and GW lensing [103, 114, 131, 147]. Within this framework, the key observables of lensed astrophysical signals and their associated implications were systematically explored.

11.1 Toward Realistic Searches for GW Lensing

To extract physical source information from GW observations and identify possible lensing signatures, we first reviewed the Bayesian inference framework for GW parameter estimation in ***Chapter 4*** [248, 253]. Bayesian inference plays a central role in GW astronomy because the underlying source properties cannot be directly observed, but must be statistically inferred from noisy detector data through waveform modelling [236, 252]. In contrast to many EM observations, GW analyses rely on coherent matched-filtering and likelihood exploration, making Bayesian techniques indispensable for robust source characterisation and uncertainty quantification. Within this framework, we further examined the physics of CBC signals and the high-dimensional parameter space describing them. Since CBC signals are characterised by a large number of correlated parameters, efficient stochastic samplers become essential for accurately obtaining posterior distributions [236, 240, 243, 244, 245, 246, 247].

Building upon this theoretical and computational foundation, we addressed the first central question of this thesis. ***Chapter 5*** investigated practical searches for gravitational lensing signatures in GW observations using advanced Bayesian inference techniques and modern computational pipelines [238, 327]. In particular, we explored two complementary classes of lensing searches: (1) searches for single distorted signals and (2) searches for multiple lensed images originating from the same source. Through these analyses, we highlighted the conceptual and methodological differences between waveform-distortion-based searches and correlation-based multiple-image searches.

Despite these methodological advances, our isolated PM microlensing analysis of BBH events from recent observing runs, as presented in ***Chapter 6***, has not yet yielded decisive and unambiguous evidence for GW microlensing. GW231123 [329], however, remains a particularly intriguing candidate. When interpreted under the microlensing hypothesis, this event exhibits the highest Bayes factors obtained to date ($\log_{10} \mathcal{B}^{\mathrm{Mod}}_{\mathrm{U}} \simeq 3.8$), improved agreement with model-agnostic reconstructed waveforms, and source parameters that are more consistent with the currently inferred astrophysical BBH population [238]. Despite these suggestive features, more rigorous validation, including extensive robustness tests against waveform systematics [331] and detector artifacts [332], together with broader population-level analyses across future GW catalogues, will be required before a confident microlensing detection can be established.

A major limitation of many existing lensing searches is that they rely on idealised lensing scenarios, such as isolated PM lenses or lens-model-independent multiple-image searches. Consequently, current search pipelines may not fully capture the complexity of realistic astrophysical lensing environments [314, 317, 318, 320, 321]. As discussed in ***Chapter 7***, a more physically realistic scenario involves the simultaneous presence of strong lensing and microlensing. To address this complexity, we introduced a residual-based framework designed to identify hidden microlensing signatures embedded within strongly lensed GW signals. Rather than performing computationally expensive, fully model-dependent parameter estimation over complex lens configurations, this method exploits characteristic structures remaining in the waveform residuals through our statistics $\Lambda_{\mathcal{H}}$ (see Eq. (7.7)). This approach enables a low-latency search strategy while retaining sensitivity to subtle wave-optics features generated by realistic lensing systems.

Nevertheless, achieving robust and quantitatively reliable identification of such complexly lensed signals ultimately requires fully lens-model-dependent Bayesian inference employing realistic astrophysical lens models. Consequently, alongside the continued development of low-latency non-parameter-estimation-based search pipelines, future efforts should prioritise the efficient and accurate computation of the oscillatory amplification factor, $F(f)$, in the wave-optics regime (see Eq. (2.71)).

11.2 Strong- and Micro- Lensing: Inferences on H_0 and $f_{\rm DM}$

The successful identification and interpretation of lensed GW signals provide a unique opportunity to connect GW astronomy directly with precision cosmology [202]. GWs themselves serve as standard sirens, providing a completely independent probe of cosmology. Since the GW waveform directly encodes the luminosity distance to the source without relying on the traditional cosmic distance ladder, GW observations offer a powerful and largely unbiased method for measuring the expansion history of the Universe, as discussed in ***Chapter 3***. This cosmological potential is further amplified when these standard sirens undergo gravitational lensing. Lensing magnification increases the effective sensitivity of GW detectors to high-z sources. In addition, both magnification and the joint analysis of multiple lensed signals from a single event can improve source localisation and luminosity-distance ($d_{\rm L}$) estimation. Ultimately, combining strong lensing with standard sirens offers a promising pathway to constrain H_0 and address the ongoing Hubble tension.

To fully exploit the potential of these lensed standard sirens, ***Chapter 8*** established the necessary methodologies for analysing individual strong lensing systems. While multiple lensed signals can yield exceptional precision in $d_{\rm L}$ and sky localisation as mentioned above, translating these measurements into astrophysical and cosmological insight requires reliable host-galaxy identification and accurate reconstruction of the lens mass distribution.

We demonstrated that combining EM observations of the host–lens system with precise GW measurements enables tight constraints on the lens-model parameters. Such accurate lens reconstruction is an essential prerequisite for the delensing process, which aims to recover the intrinsic luminosity distance required for unbiased H_0 inference.

Building upon these enhanced parameter constraints, ***Chapter 9*** explored the cosmological applications of lensed dark sirens by integrating galaxy catalogues with the hierarchical Bayesian framework introduced in ***Chapter 4***. We first demonstrated how the improved sky localisation and luminosity-distance measurements discussed in the previous chapter enable robust H_0 inference even from a single strongly lensed dark siren via the Hubble–Lemaître law. In particular, we showed that a single quadruply lensed dark siren can provide an H_0 constraint comparable to that obtained from a bright siren event.

Next, extending the analysis from individual systems to population-level cosmological inference requires a statistical framework capable of self-consistently incorporating source populations, selection effects, and lensing probabilities. To address this, we developed a comprehensive hierarchical Bayesian framework for lensed dark siren populations. Our analysis demonstrated that a small number of lensed GW events can achieve an H_0 precision comparable to that from hundreds of unlensed signals, and that future populations of multiply lensed GW events can constrain H_0 at a level competitive with existing EM observations. Finally, we investigated potential systematic biases arising from missed or incorrect lensing classifications, highlighting the importance of robust lens-identification pipelines to prevent cosmological misinterpretations.

Shifting the focus from the expansion history of the Universe to compact-object populations, ***Chapter 10*** explored the astrophysical implications of GW microlensing. In particular, we investigated the information that could be extracted from a future catalogue of confidently identified microlensed GW events. By statistically combining multiple detections within a hierarchical Bayesian framework, while consistently accounting for microlensing selection effects, we demonstrated that the mass and spatial distributions of the compact lenses can be tightly constrained.

Furthermore, if we assume that the microlensing events originate from intrinsically dark compact objects, their inferred population properties provide a direct probe of the DM fraction, $f_{\rm DM}$, composed of compact objects spanning the stellar- to IMBH-mass regime, in contrast to previous studies that have primarily constrained compact-object populations through the non-detection of microlensing signatures (e.g., [428]). Our framework illustrates how future GW microlensing observations can evolve beyond individual-event studies into a population-level probe of otherwise inaccessible compact-object populations, opening a new observational window onto both astrophysics and fundamental cosmology.

11.3 Challenges and Future Directions in GW Lensing

The current absence of definitive lensed GW detections should not be interpreted as evidence against the existence of such systems, but rather as a consequence of the intrinsic observational limitations of the present detector era. The primary limitation arises from the low intrinsic lensing probability for events detectable by existing second-generation detector networks. The optical depth for strong lensing increases rapidly with source redshift, approximately scaling as $\tau(z_s) \propto z_s^3$ at low redshift, where the Universe is approximately Euclidean. However, current GW detectors are primarily sensitive to BBH mergers at $z \lesssim 1$–1.5, where the expected strong-lensing probability remains only $\tau \sim 10^{-4}$–10^{-3} [238]. Given that present GW catalogues contain only several hundred detected BBH mergers, the absence of an unambiguous lensed detection in current observing runs is therefore statistically unsurprising.

The optical depth for GW microlensing remains substantially more uncertain because the underlying population of compact lenses is not yet well constrained. Nevertheless, compact objects capable of producing microlensing are expected to be more abundant than galaxy-scale strong lenses, potentially implying a higher occurrence rate of microlensing signatures. Identifying such events in observational data, however, remains highly challenging. Microlensing signatures typically manifest as subtle frequency-dependent modulations, such as interference fringes or beating patterns, whose detectability is fundamentally limited by the distinguishability between lensed and unlensed waveforms in the presence of detector noise. Weak microlensing configurations, such as those involving low lens masses ($M_l \lesssim 100 M_\odot$) or large impact parameters ($y > 1$), are expected to occur more frequently than configurations with high lens masses or small impact parameters. In such cases, the induced waveform distortions typically produce only very small mismatches ($\mathrm{MM} \ll 1$) between the lensed and unlensed templates (see Eq. (6.2)). In the presence of realistic detector noise, these weak distortions can be easily obscured by statistical fluctuations, preventing Bayesian model selection from decisively favouring the microlensing hypothesis. Consequently, a substantial region of the physically relevant microlensing parameter space remains effectively inaccessible to current detector sensitivities.

These challenges, however, are not expected to remain fundamental limitations in the era of third-generation detectors. The Einstein Telescope [393] and Cosmic Explorer [436] will dramatically extend the observable volume of the Universe, substantially increasing the lensing optical depth and enabling the routine detection of high-z binary mergers. Simultaneously, the exceptionally high SNRs ($\rho_{\text{net}} > 100$) and long in-spiral durations accessible to next-generation detectors will allow subtle wave-optics modulations to be resolved with far greater precision.

Nevertheless, the transition toward precision GW lensing science will require advances beyond detector sensitivity alone. As discussed in ***Chapter 6***, GW lensing signatures can exhibit significant phenomenological degeneracies with intrinsic complexities of CBC signals (e.g., [331, 340, 369]). If the frequency-dependent amplitude and phase modulations induced by lensing significantly overlap with waveform features produced by non-standard binary configurations, such as orbital eccentricity, spin precession, or overlapping GW signals, standard Bayesian inference may misidentify these effects as evidence for lensing, potentially leading to biased parameter estimation or false-positive lensing detections. Therefore, as detector sensitivities improve and statistical uncertainties decrease in the high-SNR regime, increasingly accurate waveform models incorporating both realistic source physics and lensing effects will become essential. In parallel, low-latency, computationally efficient inference pipelines capable of exploring realistic lens models will be critical for fully exploiting the scientific potential of future GW observations.

Ultimately, the coming decades are expected to transform GW lensing from a candidate-driven search into a precision observational framework for probing cosmology and compact structures across multiple astrophysical scales.

In the context of advanced cosmological inferences, next-generation detectors will provide substantially improved sky localisation and luminosity-distance measurements, facilitating host-galaxy identification and redshift determination for GW sources. In particular, a future catalogue of SLGW events will enable time-delay cosmography with unprecedented precision. Since GW arrival times can be measured to millisecond accuracy, the corresponding time-delay distances (see Eq. (2.25)) may provide sub-percent constraints on H_0, potentially competitive with existing EM-only probes such as the CMB and Type Ia supernovae.

Beyond constraining the local expansion rate, these lensed GW observations can also address evolving cosmological phenomena. Recent DESI results [197] have shown increasing evidence for dynamical dark-energy models in which the equation-of-state parameter evolves with redshift, rather than remaining fixed at $\omega = -1$ as assumed in the standard ΛCDM framework [437, 438] (see Eqs. (3.18) and (3.21)). In this context, strongly lensed GW observations from next-generation detector networks may provide an important and highly complementary probe of the cosmic expansion history and the possible evolution of dark energy [417].

Complementing these large-scale cosmological inferences, future GW detectors with improved sensitivity will additionally become capable of identifying lower-mismatch microlensing signals, substantially increasing the number of detectable compact-lens candidates and strengthening population-level constraints. In particular, the detection of microlensing signatures from compact DM candidates at cosmological distances would place stringent constraints on the compact-DM fraction beyond the peak of the cosmic star-formation history, where conventional stellar-remnant populations become increasingly difficult to produce. Such observations may also probe lower compact-object mass ranges that are less contaminated by ordinary stellar populations, thereby providing a cleaner window into possible primordial compact-object populations [424, 439].

Furthermore, access to lower GW frequencies in future detectors would extend the observable wave-optics regime toward more massive lenses, enabling interference signatures produced by intermediate-mass and heavier compact objects ($10^5 M_\odot < M_l^z \lesssim 10^8 M_\odot$) to be probed directly [386]. Such observations could open a new avenue for constraining the population properties of IMBHs and other massive compact objects across cosmic history.

Although the unambiguous detection of a lensed GW signal remains an outstanding milestone, the field is rapidly transitioning from a theoretical possibility to an imminent observational reality. As detector sensitivities continue to improve and the catalogue of GW events expands, the methodologies and investigations presented in this thesis serve as a useful foundation for the eventual identification and scientific exploitation of GW lensing. Looking further ahead, the emergence of multi-messenger lensing, in which GW and EM signals are jointly observed from the same lensed source, will establish a new paradigm for probing both cosmology and compact-object populations [102]. Ultimately, future GW lensing observations will transcend current methodological challenges, helping to address long-standing questions surrounding the nature of dark matter, the demographics of compact objects, and the evolution of the Universe itself.

Appendix

A Posteriors of the Hubble Constant

In this appendix, we provide a detailed description of each component in Eq. (9.10). The derivation presented here builds upon the formalism of Gray et al. [220]. Rewriting the expression, the posterior distribution of the Hubble constant, H_0, for a set of lensed dark sirens, x_{lgw}, with corresponding detection indicators, D_{lgw} (where each element is either 0 or 1), and assuming $N_{\text{lens,det}}$ detected lensed systems, can be written as:

$$
\begin{aligned}
&p(H_0|\{x_{\text{lgw}}\},\{D_{\text{lgw}}\}) \propto p(H_0)p(N_{\text{lens,det}}|H_0) \\
&\times \left[\iiiint p(D_{\text{lgw}}|\theta_s,\theta_l,z_s,z_l,H_0)p(\theta_s|H_0)\sum_j^{N_{\text{pix}}} p(z_s,z_l,\theta_l|\Omega_j,H_0,s,l)d\theta_l d\theta_s dz_l dz_s\right]^{-N_{\text{lens,det}}} \\
&\times \prod_i^{N_{\text{lens,det}}}\left[\iiiint \sum_j^{N_{\text{pix}}} p(x_{\text{lgw},i}|\Omega_j,\theta_s,\theta_l,z_s,z_l,H_0)p(\theta_s|H_0)p(z_s,z_l,\theta_l|\Omega_j,H_0,s,l)d\theta_l d\theta_s dz_l dz_s\right],
\end{aligned}
\tag{1}
$$

where θ_s, θ_l, z_s, and z_l denote the GW source parameters excluding the distance, the lens parameters of the source-lens galaxy system, the source redshift, and the lens galaxy redshift, respectively. $p(z_s,z_l,\theta_l|\Omega_j,H_0,s,l)$ is the LoS redshift distribution of source and lens galaxies obtained from the galaxy-galaxy lensing catalogue. The GW sky localisation is discretised into N_{pix} pixels, with Ω_j corresponding to the sky direction of the j^{th} pixel. $p(H_0)$ denotes the prior on the Hubble constant, which is a uniform distribution over the range of $[20, 140]$ km $\cdot$ s^{-1} $\cdot$ Mpc^{-1}. More detailed descriptions of the remaining terms are provided in the following subsections.

A.1 Selection Effect Associated with Lensed Gravitational Waves

Selection bias arises when the number of detectable signals differs from the total number of mergers due to limited detector sensitivity. The inferred H_0 becomes biased if selection effects are ignored. For example, low-redshift events are more easily detected, which would incorrectly suggest an overabundance of GW sources at low redshift. To correct for this, the likelihood of each GW event must be divided by its detection efficiency, that is, the probability that an event with given parameters would be observed.

Given the detection probability for a source with parameters $\vec{\theta} = (\theta_s, \theta_l, z_s, z_l)$, $p_{\rm det}(\vec{\theta})$ and the underlying population distribution $p_{\rm pop}(\vec{\theta})$ for a given set of hyperparameters Λ, such as Eqs. (4.35), (4.37), and (4.39), the expected number of detections of lensed GW events is given by:

$$\begin{aligned}\langle N_{\rm lens,det}(\Lambda)\rangle &= N_{\rm lens}\int p_{\rm det}(\vec{\theta})p_{\rm pop}(\vec{\theta}|\Lambda)d\vec{\theta} \\ &= N_{\rm lens}\iiiint p(D_{\rm lgw}|\theta_s,\theta_l,z_s,z_l,\Lambda)p(\theta_s,\theta_l,z_s,z_l|\Lambda)d\theta_s d\theta_l dz_l dz_s,\end{aligned} \tag{2}$$

where $N_{\rm lens}$ denotes the total number of lensed mergers [291]. The detection efficiency is the same as the ratio between the total and expected number of detection. To efficiently evaluate the detection efficiency of lensed GW events, we generated numerous injections from a fiducial distribution $\pi(\vec{\theta})$ (see [292] for more details). For a given set of detected lensed GW events, the detection efficiency, equivalent to the selection-effect term in Eq. (1), can be evaluated as follows:

$$\begin{aligned}\frac{\langle N_{\rm lens,det}(\Lambda)\rangle}{N_{\rm lens}} &= \iiiint p(D_{\rm lgw}|\theta_s,\theta_l,z_s,z_l,\Lambda)p(\theta_s,\theta_l,z_s,z_l|\Lambda)d\theta_s d\theta_l dz_l dz_s \\ &\simeq \frac{1}{N_{\rm inj}}\sum_i^{N_{\rm inj,det}}\frac{p_{\rm pop}(\vec{\theta}_i|\Lambda)}{\pi(\vec{\theta}_i)} \\ &= \frac{1}{N_{\rm inj}}\sum_i^{N_{\rm inj,det}}\frac{p_{\rm pop}(m_{1,i},m_{2,i},z_{s,i}|\Lambda)}{\pi(m_{1,i},m_{2,i},z_{s,i})}\frac{p_{\rm pop}(\theta_{l,i},z_{l,i}|z_{s,i},\Lambda)}{\pi(\theta_{l,i},z_{l,i})} \\ &= \frac{1}{N_{\rm inj}}\sum_i^{N_{\rm inj,det}}\frac{p_{\rm pop}(m_{1,i},m_{2,i},z_{s,i}|\Lambda)}{\pi(m_{1,i},m_{2,i},z_{s,i})}\frac{p_{\rm pop}(\vec{\mu}_i|\Lambda)}{\pi(\vec{\mu}_i)},\end{aligned} \tag{3}$$

where $m_{1,i}$, $m_{2,i}$, and $\vec{\mu}_i$ indicate the component masses and individual magnification factors of the $i^{\rm th}$ lensed event, respectively.

Instead of sampling the θ_l and z_l directly, we first compute $\vec{\mu}$ from the lens parameters, creating multiple images under a specific lens model. The resulting magnification factors define a distribution from which we draw samples, as these factors directly affect the SNR of the lensed signals. We adopt the population models in Section 4.4.2 for the source-frame primary and secondary masses and redshift. All other source parameters are drawn from hyperparameter-independent distributions. For these latter parameters, the population prior $p_{\rm pop}$ reduces to the prior π and is thus omitted.

Since detectability is determined in the detector frame, we draw injections in the detector frame, which modifies Eq. (3) as follows,

$$\begin{aligned}
\frac{\langle N_{\text{lens,det}}\,(\Lambda)\rangle}{N_{\text{lens}}} &\simeq \frac{1}{N_{\text{inj}}}\sum_{i}^{N_{\text{inj,det}}}\frac{p_{\text{pop}}(m_{1,i},m_{2,i},z_{s,i}|\Lambda)}{\pi(m_{1,i},m_{2,i},z_{s,i})}\frac{p_{\text{pop}}(\vec{\mu}_i|\Lambda)}{\pi(\vec{\mu}_i)} \\
&= \frac{1}{N_{\text{inj}}}\sum_{i}^{N_{\text{inj,det}}}\frac{p_{\text{pop}}(m_{1,i},m_{2,i},z_{s,i},\vec{\mu}_i|\Lambda)}{\pi(m^d_{1,i},m^d_{2,i},\vec{d}^{\,\text{eff}}_{L,i})}\left|\frac{\partial(m_{1,i},m_{2,i},z_{s,i},\vec{\mu}_i)}{\partial(m^d_{1,i},m^d_{2,i},\vec{d}^{\,\text{eff}}_{L,i})}\right| \\
&= \frac{1}{N_{\text{inj}}}\sum_{i}^{N_{\text{inj,det}}}\frac{p_{\text{pop}}(m_{1,i},m_{2,i},z_{s,i},\vec{\mu}_i|\Lambda)}{\pi(m^d_{1,i},m^d_{2,i},\vec{d}^{\,\text{eff}}_{L,i})}\prod_j\left|\frac{2\mu_{i,j}}{d^{\text{eff}}_{L,i,j}}\right|\frac{1}{(1+z_{s,i})^2}\left|\frac{\partial z_s}{\partial d_L}\right|_{z_s=z_{s,i}},
\end{aligned} \tag{4}$$

where m_1^d and m_2^d are detector-frame primary and secondary masses. Although injections are generated in the detector frame, the selection function is evaluated by reweighting them to the source-frame population model, including the full Jacobian associated with the transformation between $(m_1^d, m_2^d, \vec{d}_L^{\,\text{eff}})$ and $(m_1, m_2, z_s, \vec{\mu})$.

To determine whether a signal is detected, the unlensed case is straightforward. A single event is considered detected if its network SNR exceeds the chosen threshold, $\rho_{\text{net}} = 8$. However, for lensed events, multiple signals from the same source must not only exceed the network SNR threshold, but also be confirmed to originate from the same source and truly lensed through a lensing hypothesis test. Strong lensing can be verified using pipelines equipped with joint parameter estimation, such as [300, 301, 302, 303] employed in projects of searching for GW lensing signatures [238, 324, 325]. In this study, we assume that all lensed signals pass the hypothesis test. Hence, the detection criterion reduces to verifying that lensed signals are detected above the SNR threshold and correctly identified as lensed. Accordingly, $N_{\text{lens,det}}$ in Eq. (3) denotes the number of GW lensing systems that satisfy both the detection threshold and the lensing hypothesis test.

For doubly lensed GW systems, if one of the signals falls below the detection threshold, the lensing nature of the system cannot be inferred from the single detected signal alone. In principle, one may nevertheless perform a dedicated sub-threshold search to identify a counterpart of a detected super-threshold GW signal that is insufficiently magnified or strongly demagnified to be detected above the threshold [377, 413]. Furthermore, a lensed GW signal could be identified even from a single image by searching for waveform distortions in Type-II images induced by higher-order GW modes [157, 367] [1].

For quadruply lensed systems, only two or three of the four images may be detected in realistic scenarios. We impose the analogous requirement to the doubly-lensed case that at least two signals must exceed the detection threshold for the system to be identified as strongly lensed. Since the lensing hypothesis test determines the number of detected lensed

1. However, such effects are not directly relevant to the high-magnification regime or to improvements in sky localisation, and therefore we do not consider this possibility in this study.

signals in each system, we consider three cases: two, three, or all four images detected, when accounting for selection effects. It is important to note that partially detecting only two or three of the four images can make it initially unclear whether the system originates from a quadruply lensed source. Accounting for selection effects under the assumption that such systems are doubly or triply lensed leads to a biased inference of the Hubble constant.

Once a specific lens model is adopted, however, the possible number of lensed signals becomes fixed, and the corresponding population distributions of lensing observables, such as relative magnification factors and time delays, are determined. Under the SIE model, for example, only two or four images can be produced. Thus, if three images are detected, the event can be unambiguously identified as a quadruply-lensed system, allowing the appropriate quadruple-image selection effects to be applied. When only two images are detected, however, both the doubly-lensed and quadruply-lensed interpretations remain possible, and selection effects from both cases must be considered.

The term $p_{\text{pop}}(\vec{\theta}|\Lambda)$ in Eq. (2) can be expressed in terms of the underlying population distribution function, $dN/d\vec{\theta}$ [291, 292]. We begin by defining the number of compact binary mergers per unit observer time and unit redshift, defined as:

$$\frac{d^2N}{dtdz} = \frac{\mathcal{R}(z)}{1+z}\frac{\partial V_c}{\partial z}, \tag{5}$$

where $\mathcal{R}(z)$ is the intrinsic merger rate per comoving volume, V_c, at source redshift z, and the $(1+z)^{-1}$ factor accounts for cosmological time dilation.

To obtain the subset of these merger events that undergo strong lensing, the optical depth introduced in Eq. (2.28) can be rewritten as:

$$\tau(z_s, H_0) = \int_0^{z_s} dz_l \frac{dV_c}{d\Omega dz_l} \int d\theta_l \; n_l(z_l, \theta_l)\hat{\sigma}_{\text{SL}}(z_s, z_l, \theta_l, H_0), \tag{6}$$

where $dV_c/(d\Omega dz_l)$ denotes the differential comoving volume per steradian at the lens redshift z_l, n_l is the comoving number density of lenses with properties θ_l, and $\hat{\sigma}_{\text{SL}}$ represents the strong lensing cross-section of an individual lens for a given z_s, z_l, θ_l, and H_0.

From Eq. (5), $dN/d\vec{\theta}$ can then be related to the intrinsic merger rate, together with the corresponding optical depth in Eq. (6) as follows:

$$\begin{aligned}\langle N_{\text{lens,det}}\,(H_0)\rangle &= N_{\text{lens}} \iiiint p(D_{\text{lgw}}|\theta_s, \theta_l, z_s, z_l, H_0)\frac{1}{N_{\text{lens}}}\left(\frac{d^4N_{\text{lens}}}{d\theta_s d\theta_l dz_l dz_s}\right) d\theta_s d\theta_l dz_l dz_s \\ &= T_{\text{obs}} \iiiint p(D_{\text{lgw}}|\theta_s, \theta_l, z_s, z_l, H_0)\tau(z_s, H_0)\frac{R(z_s)}{1+z_s}\frac{\partial V_c}{\partial z_s} d\theta_s d\theta_l dz_l dz_s,\end{aligned} \tag{7}$$

where T_{obs} denotes the observation period. Eq. (7) encodes the distribution of θ_l and z_l, both of which depend on H_0 in contrast to the z_s and θ_s, which are independent of H_0. Therefore, unlike in the unlensed case, the probability of detecting $N_{\text{lens,det}}$ lensed dark siren system for a given value of H_0, $p(N_{\text{lens,det}} \mid H_0)$, cannot be neglected even when adopting a scale-invariant prior on the merger rate, $P(\mathcal{R}) \propto 1/\mathcal{R}$. For a given expected number of detections ($\langle N_{\text{lens,det}} \rangle$), $p(N_{\text{lens,det}} \mid H_0)$ in Eq. (1) can be expressed as a Poisson distribution as follows:

$$p(N_{\text{lens,det}} \mid H_0) = e^{-\langle N_{\text{lens,det}} \rangle} \left(\langle N_{\text{lens,det}} \rangle\right)^{N_{\text{lens,det}}} . \tag{8}$$

A.2 Line-of-sight Redshift Prior

The LoS redshift prior is constructed by using the measured redshift, luminosity, and apparent magnitude information of galaxies in a given catalogue to identify which galaxies lie within the GW sky localisation, to weight them by their probability of hosting a BBH merger (e.g., luminosity weighting) and being lensed, and to correct for the incompleteness of the galaxy catalogue. This framework enables the inference of the redshift distribution of the observed BBH events from GW observations.

For the purpose of identifying galaxies within a given GW localisation, a galaxy-galaxy lensing catalogue, in which both source and lens galaxy information is available for each system, can be more informative than a standard galaxy catalogue containing only unlensed galaxies. However, luminosity weighting and corrections for catalogue incompleteness require the intrinsic source-galaxy luminosity (i.e., the absolute magnitude) and the redshift. While galaxy-scale gravitational lensing is achromatic and leaves the observed redshift unchanged to within the precision of spectroscopic measurements, the absolute magnitude can be significantly altered by lensing magnification. Therefore, delensing is required to recover the intrinsic absolute magnitudes of source galaxies in order to obtain robust luminosity weights and incompleteness corrections. Accordingly, the LoS redshift prior must account for whether both the source and lens galaxies are included in the catalogue.

A.2.1 Complete catalogue

We first consider the complete catalogue case, in which there is no apparent magnitude limit, such that every lens galaxy and all lensed images of the source galaxy for each system are included in the catalogue. In this case, the LoS redshift prior in Eq. (1), $p(z_s, z_l, \theta_l | \Omega_j, H_0, s, l)$, is given by:

$$p\left(z_s, z_l, \theta_l \mid \Omega_j, H_0, s, l\right) = \iint p\left(z_s, z_l, \theta_l, M_s, m_s', M_l, m_l \mid \Omega_j, H_0, s, l\right)\, dM_s\, dm_s', dM_l, dm_l, \tag{9}$$

where M_s and M_l denote the intrinsic absolute magnitudes of source and lens galaxies, respectively. m_l and m'_s are the apparent magnitude of the lens galaxy and lensed source-galaxy images (see Eq. (2.59)), respectively. Note that M_s cannot be correctly obtained from m'_s using the standard distance modulus alone, as the magnification factor of the lensed image must be taken into account. Hereafter, we define $\vec{M} = (M_s, M_l)$ and $\vec{m} = (m'_s, m_l)$ to concisely represent the absolute and apparent magnitudes. The integrand in Eq. (9) can then be written as:

$$
\begin{aligned}
& p\left(z_s, z_l, \theta_l, \vec{M}, \vec{m} \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \\
& = \frac{1}{p\left(s, l \mid \Omega_j, H_0\right)} p\left(z_s, z_l, \theta_l, \vec{M}, \vec{m} \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) p\left(s, l \mid z_s, z_l, \theta_l, \vec{M}, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\
& = \frac{1}{p\left(s, l \mid \Omega_j, H_0\right)} \delta(M_s - M_s(z_s, z_l, \theta_l, m'_s, H_0, \mathcal{M}_{\text{SIE}})) \delta(M_l - M_l(z_l, m_l, H_0)) \\
& \times p\left(z_s, z_l, \theta_l, m'_s, m_l \mid \Omega_j, H_0\right) p\left(s, l \mid z_s, z_l, \theta_l, M_s, H_0\right),
\end{aligned}
\tag{10}
$$

where $\mathcal{M}_{\text{SIE}}$ indicates that SIE profiles are assumed to infer M_s from m'_s of their images.

Since all lensed images and the lens galaxy are observed in a complete catalogue, the intrinsic luminosity of the source galaxy can, in principle, be accurately recovered. We therefore assume that the intrinsic absolute magnitude of the source galaxy is uniquely determined once the other measured parameters are specified, such that its prior can be represented as a delta function. The absolute magnitude of the lens galaxy, M_l, is unaffected by lensing and independent of the source properties, and thus its prior can likewise be treated as a delta function[2]. Moreover, the probability that a galaxy hosts a GW source and that the detected lensed images are genuinely associated does not depend on M_l, Ω_j, or $\mathcal{M}_{\text{SIE}}$, which are omitted in the last term $p(s, l|...)$.

The LoS redshift prior for the complete catalogue case can be written as:

$$
\begin{aligned}
& p\left(z_s, z_l, \theta_l \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \\
& = \frac{1}{p\left(s, l \mid \Omega_j, H_0\right)} \frac{1}{N_{\text{Lgal}}(\Omega_j)} \\
& \times \sum_{k}^{N_{\text{Lgal}}} p\left(z_s, z_l, \theta_l \mid \hat{z}_{s,k}, \hat{z}_{l,k}, \hat{\theta}_{l,k}\right) p\left(s, l \mid z_s, z_l, \theta_l, M_s(z_s, z_l, \theta_l, \hat{m'}_{s,k}, H_0, \mathcal{M}_{\text{SIE}})\right),
\end{aligned}
\tag{11}
$$

where N_{Lgal} denotes the number of lensing systems within the sky patch Ω_j. Following the notation of [220], the hat notation denotes the measured value for the k^{th} galaxy-galaxy lensing system, and the corresponding $p(z_s, z_l, \theta_l)$ are modelled as Gaussian distributions centered on these measured values. For z_s and z_l, we adopt the spectroscopic redshift un-

2. Note that, however, in galaxy-galaxy lensing systems, the spatial overlap between foreground lenses and background lensed source-galaxy images leads to blended photometry, which can systematically bias photometric redshift estimates.

certainties introduced in Section 9.2.1, namely $\sigma_z = 0.001(1+z)$. For the lens parameters, we assume a standard deviation of 0.01, which corresponds to the grid resolution used in the delensing procedure described in [311]. As shown in Eq. (11), the magnitudes of the lens galaxies are not included, which can be neglected when deriving the LoS redshift priors, and thus we will omit M_l and m_l and not explicitly write them in the following.

Since no explicit correlation between GW source redshift and host galaxy luminosity is known, the term $p(s, l|...)$ can be factorised as follows:

$$p(s, l|z_s, z_l, \theta_l, M_s, H_0) \propto p(s|z_s)p(s|M_s)p(l|s, z_s, z_l, \theta_l, H_0). \tag{12}$$

For simplicity, one may reduce the complexity of the LoS computation by disabling the luminosity weighting, i.e., setting $p(s|M_s) \propto const.$. In this case, all source galaxies at z_s are assumed to have equal probability of hosting a BBH merger, and the resulting LoS redshift prior depends on the merger rate and lensing probability as a function of redshift and the distribution of galaxies.

A.2.2 Incomplete catalogue

In practice, galaxy catalogues are often incomplete due to observational limits. In this section, we examine the impact of an apparent magnitude threshold m_{th}, which may cause certain lens galaxies or lensed source-galaxy images to be missing from the catalogue. We consider six possibilities based on the visibility of the lens galaxy and the source-galaxy images: whether the lens is in- or out-of-catalogue, and whether all, a subset, or none of the source images are in the catalogue. Since whether a system is in the catalogue depends on whether the system is identified as lensed (GG) or not ($\overline{GG}$), these six possibilities can be grouped into the two categories (see Table A.1 for details).

With these definitions, Eq. (9) can be rewritten as

$$\begin{aligned} p\left(z_s, z_l, \theta_l \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) &= \iint \sum_{g=GG,\overline{GG}} p\left(z_s, z_l, \theta_l, M_s, m_s', g \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \, dM_s dm_s' \\ &= \sum_{g=GG,\overline{GG}} p\left(g \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \iint p\left(z_s, z_l, \theta_l, M_s, m_s' \mid g, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \, dM_s dm_s'. \end{aligned} \tag{13}$$

Case	catalogue	z_s, m'_s	z_l, θ_l	Delensing
(G, LG)	In	✓	✓	Very accurate (0-5 %)
(G', LG)	In	✓	✓	Accurate (5-20 %)
$(G, \overline{LG})$	In	✓	✗	Possible (20-50 %)
$(G', \overline{LG})$	Mixed	✓	✗	Poor (50 %-) or Impossible
$(\bar{G}, LG)$	Out	✗	✗	Impossible
$(\bar{G}, \overline{LG})$	Out	✗	✗	Impossible

Table A.1: Information availability for the six possibilities determined by the observational accessibility of the lens galaxy and the lensed source-galaxy images. G denotes the case in which all lensed images are included in the catalogue, $\bar{G}$ denotes the case in which none of the images are included, and G' denotes the case in which only a subset of the images is included. Similarly, LG and $\overline{LG}$ indicate whether the lens galaxy is present in or absent from the catalogue, respectively. A checkmark (✓) and a cross (✗) indicate that the corresponding quantity is observationally accessible or not, respectively. For the $(G', \overline{LG})$ case, if two or three images of a quadruple-image system are observed, the system is identified as lensed and delensing is partially possible, though limited. The accuracy of the delensing depends on the number of accessible observables. The numbers in parentheses denote the relative error ranges of the recovered magnification factors. Conversely, if only one image of a double-image system is observed, the system would not be included in the galaxy-galaxy lensing catalogue, as it may be interpreted as a weak lensing system or a special configuration of an unlensed galaxy. In this case, delensing is not possible.

The integrand of the first term in Eq. (13), corresponding to systems that are in-catalogue, has a form similar to Eq. (10) and can be expanded as

$$\begin{aligned}
&p\left(z_s, z_l, \theta_l, M_s, m'_s \mid GG, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \\
&= \frac{p\left(z_s, z_l, \theta_l, M_s, m'_s \mid GG, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right)}{p\left(s, l \mid GG, \Omega_j, H_0\right)} p\left(s, l \mid z_s, z_l, \theta_l, M_s, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\
&= \frac{p(M_s|z_s, z_l, \theta_l, m'_s, H_0, \mathcal{M}_{\text{SIE}}) p\left(z_s, z_l, \theta_l, m'_s \mid GG, \Omega_j, H_0\right)}{p\left(s, l \mid GG, \Omega_j, H_0\right)} p\left(s, l \mid z_s, z_l, \theta_l, M_s, H_0\right),
\end{aligned} \tag{14}$$

where $p(M_s|\ldots)$ denotes the conditional probability distribution based on the observed images and the SIE assumption. In the incomplete catalogue case, delensing cannot be assumed fully accurate, and the prior on the intrinsic absolute magnitude of the source galaxy can no longer be represented as a delta function, unlike in the complete catalogue case discussed in Section A.2.1. Nevertheless, the distribution of magnification factors for each source-galaxy image can be inferred through lens reconstruction, using the available information. This information may vary depending on the possibilities in Table A.1, allowing the intrinsic absolute magnitude to be recovered to some extent.

Integrating Eq. (14) over M_s and m'_s yields

$$\iint p\left(z_s, z_l, \theta_l, M_s, m'_s \mid GG, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) dM_s dm'_s = \frac{1}{p\left(s, l \mid GG, \Omega_j, H_0\right)} \frac{1}{N_{\text{Lgal}}(\Omega_j)} \times \sum_{k}^{N_{\text{Lgal}}} \left[p\left(z_s, z_l, \theta_l \mid \hat{z}_{s,k}, \hat{z}_{l,k}, \hat{\theta}_{l,k}\right) \left\{ \frac{1}{N_{\text{samp}}} \sum_{n}^{N_{\text{samp}}} p\left(s, l \mid z_s, z_l, \theta_l, M_{s,k}^{(n)}, H_0\right) \right\} \right], \tag{15}$$

where $M_{s,k}^{(n)}$ denotes the n^{th} sample among the N_{samp} samples drawn from the conditional distribution $p(M_s|..., \hat{m}'_{s,k})$, which is generally more informative than a uniform distribution over possible absolute magnitudes of source galaxies. Furthermore, for the source redshift, even if only a single lensed image is observed, z_s can still be determined. Therefore, Eq. (15) retains the same form as Eq. (11) with the conditional sampling over M_s replacing the delta function. Again, if luminosity weighting is not applied, retrieval of the true intrinsic luminosity is not required, and the LoS prior reduces to a simpler form.

Next, the integrand of the second term in Eq. (13), corresponding to systems that are out-of-catalogue, can be expanded in a manner similar to Eq. (2.11) of [220] as follows,

$$\begin{aligned}
&p\left(z_s, z_l, \theta_l, M_s, m'_s \mid \overline{GG}, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \\
&= \frac{1}{p\left(s, l \mid \overline{GG}, \Omega_j, H_0\right)} p\left(s, l \mid z_s, z_l, \theta_l, M_s, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\
&\times p\left(z_s, z_l, \theta_l, M_s, m'_s \mid \overline{GG}, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\
&= \frac{1}{p\left(s, l \mid \overline{GG}, \Omega_j, H_0\right)} p(s, l \mid z_s, z_l, \theta_l, M_s, H_0) \\
&\times \frac{p\left(\overline{GG} \mid z_s, z_l, \theta_l, M_s, m'_s, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) p\left(z_s, z_l, \theta_l, M_s, m'_s \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right)}{p\left(\overline{GG} \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right)},
\end{aligned} \tag{16}$$

where the term $p\left(\overline{GG} \mid ...\right)$ in the numerator indicates the probability that a galaxy-galaxy lensing system with the specified properties is not included in the catalogue. In the out-of-catalogue case, the observability of the lens galaxy is less relevant, and only the source-galaxy redshift is required when constructing the LoS redshift prior. Since the criterion for identifying a system as lensed is the detection of at least two lensed images of the source galaxy, the probability $p\left(\overline{GG} \mid ...\right)$ can be expressed as follows,

$$p\left(\overline{GG} \mid z_s, z_l, \theta_l, M_s, m'_s, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) = \Theta\left(\left[\sum_{k=1}^{N_{\text{img}}} \Theta\left(m'_{s,k} - m_{\text{th}}(\Omega_j)\right)\right] - 2\right). \tag{17}$$

Here, $\Theta(x)$ denotes the Heaviside step function, defined as $\Theta(x) = 1$ for $x \geq 0$ and $\Theta(x) = 0$ otherwise. For the adopted SIE lens model, the total number of images in a lensing system, N_{img}, is either 2 or 4.

Eq. (16) can then be rewritten as

$$p\left(z_s, z_l, \theta_l, M_s, m'_s \mid \overline{GG}, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) = \frac{\Theta\left(\left[\sum_{k=1}^{N_{\text{img}}} \Theta\left(m'_{s,k} - m_{\text{th}}(\Omega_j)\right)\right] - 2\right)}{p\left(s, l \mid \overline{GG}, \Omega_j, H_0\right) p\left(\overline{GG} \mid \Omega_j, H_0\right)}$$
$$\times\, p\left(z_s, z_l, \theta_l, M_s, m'_s \mid H_0, \mathcal{M}_{\text{SIE}}\right) p(s, l \mid z_s, z_l, \theta_l, M_s, H_0). \tag{18}$$

Since the systems under consideration are galaxy-galaxy lensing systems, the lens properties are explicitly included. The term $p\left(z_s, z_l, \theta_l, M_s, m'_s \mid H_0, \mathcal{M}_{\text{SIE}}\right)$ corresponds to the prior distributions of redshifts, lens parameters, and magnitudes, which are independent of the sky pixel Ω_j, as the properties of galaxies do not depend on their tangential location. We adopt minimally informative models to predict galaxy properties since not all relevant galaxy observables are directly accessible. Unlike the unlensed-catalogue case, the observed quantities correspond to the lensed source-galaxy properties and the lens properties. We assume that the redshifts of source galaxies follow a uniform distribution in comoving volume. Accordingly, the redshift distribution of lens galaxies is determined jointly by the source redshift distribution and the lensing optical depth.

For the lens parameters, we adopt the distributions described in Section 9.1.1. Note that both the lensing optical depth and the lens parameters are evaluated under the assumption of SIE profiles, and accordingly, the prior term explicitly depends on $\mathcal{M}_{\text{SIE}}$. Lastly, we adopt a Schechter function for the intrinsic absolute magnitudes of source galaxies. Given the source and lens redshifts, the source absolute magnitude, and the lens parameters, the lensed apparent magnitudes of source galaxies can then be computed.

Therefore, the prior can be written as

$$p\left(z_s, z_l, \theta_l, M_s, m'_s \mid H_0, \mathcal{M}_{\text{SIE}}\right) = \delta(m'_s - m'_s(z_s, z_l, \theta_l, M_s, H_0, \mathcal{M}_{\text{SIE}}))p(z_s, z_l, \theta_l, M_s \mid H_0). \tag{19}$$

Combining Eq. (18) with Eq. (19) and performing the integration over M_s and m'_s yields

$$\iint p\left(z_s, z_l, \theta_l, M_s, m'_s \mid \overline{GG}, \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) dM_s dm'_s$$
$$= \frac{\int_{M_s(z_s, z_l, \theta_l, m_{\text{th}}(\Omega_j), H_0, \mathcal{M}_{\text{SIE}})}^{M_{\max}(H_0)} p(z_s, z_l, \theta_l, M_s \mid H_0) p(s, l \mid z_s, z_l, \theta_l, M_s, H_0) dM_s}{p\left(s, l \mid \overline{GG}, \Omega_j, H_0\right) p\left(\overline{GG} \mid \Omega_j, H_0\right)} \tag{20}$$

The nested Heaviside step function in Eq. (17) has been converted into an integration limit over M_s, which depends on the properties of the source and lens galaxies and is determined by the apparent-magnitude threshold of pixel j.

By combining Eqs. (15) and (20), the full expression for the LoS redshift in the incomplete galaxy-galaxy lensing catalogue case (Eq. (13)) can be rewritten as

$$
\begin{aligned}
&p(z_s, z_l, \theta_l \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}) \\
&= \frac{p\left(GG \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right)}{p\left(s, l \mid GG, \Omega_j, H_0\right)} \\
&\times \frac{1}{N_{\text{Lgal}}(\Omega_j)} \sum_k^{N_{\text{Lgal}}} p\left(z_s, z_l, \theta_l \mid \hat{z}_{s,k}, \hat{z}_{l,k}, \hat{\theta}_{l,k}\right) \left[\frac{1}{N_{\text{samp}}} \sum_n^{N_{\text{samp}}} p\left(s, l \mid z_s, z_l, \theta_l, M_{s,k}^{(n)}, H_0\right)\right] \\
&\quad + \frac{p\left(\overline{GG} \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right)}{p\left(s, l \mid \overline{GG}, \Omega_j, H_0\right) p\left(\overline{GG} \mid \Omega_j, H_0\right)} \\
&\times \int_{M_s(z_s, z_l, \theta_l, m_{\text{th}}(\Omega_j), H_0, \mathcal{M}_{\text{SIE}})}^{M_{\max}(H_0)} p(z_s, z_l, \theta_l, M_s \mid H_0)\, p(s, l \mid z_s, z_l, \theta_l, M_s, H_0)\, dM_s\,.
\end{aligned}
\tag{21}
$$

By applying Bayes' theorem to $p(g|\Omega_j, H_0, s)$ terms in each in-catalogue and out-of-catalogue term, Eq. (21) can be simplified as,

$$
\begin{aligned}
&p\left(z_s, z_l, \theta_l \mid \Omega_j, H_0, s, l, \mathcal{M}_{\text{SIE}}\right) \\
&= \frac{1}{p(s, l \mid \Omega_j, H_0)} \\
&\times \left[\frac{p(GG \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}})}{N_{\text{Lgal}}(\Omega_j)} \sum_k^{N_{\text{Lgal}}} \left\{p\left(z_s, z_l, \theta_l \mid \hat{z}_{s,k}, \hat{z}_{l,k}, \hat{\theta}_{l,k}\right)\right.\right. \\
&\qquad \left. \times \left(\frac{1}{N_{\text{samp}}} \sum_n^{N_{\text{samp}}} p\left(s, l \mid z_s, z_l, \theta_l, M_{s,k}^{(n)}, H_0\right)\right)\right\} \\
&\quad \left. + \int_{M_s^{\min}}^{M_{\max}(H_0)} p(z_s, z_l, \theta_l, M_s \mid H_0)\, p(s, l \mid z_s, z_l, \theta_l, M_s, H_0)\, dM_s\right],
\end{aligned}
\tag{22}
$$

where $p(GG \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}})$ denotes the probability that a galaxy-galaxy lensing system located on the j^{th} pixel is in-catalogue. It is important to note that this expression does not depend on s, so we must consider not only the host galaxies but also all other galaxies within the given redshift cut that satisfy the magnitude threshold. The explicit definition of this probability can be written as:

$$\begin{aligned} & p\left(GG \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\ & = \int d\boldsymbol{\Xi}\, p\left(GG \mid \boldsymbol{\Xi}, \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right)\, p\left(\boldsymbol{\Xi} \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\ & = \int d\boldsymbol{\Xi}\, \Theta\left(2 - \sum_{k=1}^{N_{\text{img}}} \Theta\left(m'_{s,k} - m_{\text{th}}(\Omega_j)\right)\right) p\left(\boldsymbol{\Xi} \mid \Omega_j, H_0, \mathcal{M}_{\text{SIE}}\right) \\ & = \iiint dz_s\, dz_l\, d\theta_l\, p\left(z_s, z_l, \theta_l \mid H_0\right) \int_{M_{\min}(H_0)}^{M_s^{\min}} p(M_s \mid H_0)\, dM_s, \end{aligned} \tag{23}$$

where we denoted the full integration domain as $d\boldsymbol{\Xi} \equiv dM_s\, dm'_s\, dz_s\, dz_l\, d\theta_l$ for brevity, and the integration upper limit as $M_s^{\min} \equiv M_s(z_s, z_l, \theta_l, m_{\text{th}}(\Omega_j), H_0, \mathcal{M}_{\text{SIE}})$. The term $p\left(z_s, z_l, \theta_l \mid H_0\right)$ denotes the joint prior on the source and lens redshifts and lens parameters introduced in the out-of-catalogue case, while the term $p\left(M_s \mid H_0\right)$ is modelled as a Schechter function describing the intrinsic absolute-magnitude distribution of source galaxies.

Bibliography

[1] Scott Dodelson and Fabian Schmidt. *Modern cosmology*. Vol. 1. Academic press Amsterdam, 2003.

[2] John G Learned and Karl Mannheim. 'High-energy neutrino astrophysics'. In: *Annual Review of Nuclear and Particle Science* 50.1 (2000), pp. 679–749.

[3] Wick C Haxton, RG Hamish Robertson and Aldo M Serenelli. 'Solar neutrinos: status and prospects'. In: *Annual Review of Astronomy and Astrophysics* 51 (2013), pp. 21–61.

[4] Xun-Jie Xu, Zhe Wang and Shaomin Chen. 'Solar neutrino physics'. In: *Progress in Particle and Nuclear Physics* 131 (2023), p. 104043.

[5] IceCube Collaboration et al. 'Neutrino emission from the direction of the blazar TXS 0506+ 056 prior to the IceCube-170922A alert'. In: *Science* 361.6398 (2018), pp. 147–151.

[6] IceCube Collaboration et al. 'Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A'. In: *Science* 361.6398 (2018), eaat1378.

[7] Albert Einstein. 'Approximative Integration of the Field Equations of Gravitation'. In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften Berlin (Mathematisch-physikalische Klasse)* 1916 (1916), pp. 688–696.

[8] Albert Einstein. 'Über Gravitationswellen'. In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften Berlin (Mathematisch-physikalische Klasse)* 1918 (1918), pp. 154–167.

[9] Albert Einstein. 'Zur Allgemeinen Relativitätstheorie'. In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften Berlin (Mathematisch-physikalische Klasse)* 1915 (1915), pp. 778–786.

[10] Albert Einstein. 'The Field Equations of Gravitation'. In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften Berlin (Mathematisch-physikalische Klasse)* 1915 (1915), pp. 844–847.

[11] J. Weber. 'Detection and Generation of Gravitational Waves'. In: *Physical Review* 117 (1960), pp. 306–313.

[12] RA Hulse and JH Taylor. 'Discovery of a pulsar in a binary system'. In: *Neutron stars, black holes, and binary X-ray sources* 48 (1975), p. 433.

[13] LIGO Scientific Collaboration et al. 'Advanced LIGO'. In: *Classical and Quantum Gravity* 32.7 (2015), p. 074001.

[14] Benjamin P Abbott et al. 'Observation of gravitational waves from a binary black hole merger'. In: *Physical Review Letters* 116.6 (2016), p. 061102.

[15] Fausto Acernese et al. 'Advanced Virgo: a second-generation interferometric gravitational wave detector'. In: *Classical and Quantum Gravity* 32.2 (2015), p. 024001.

[16] Benjamin P Abbott et al. 'GW170817: observation of gravitational waves from a binary neutron star inspiral'. In: *Physical Review Letters* 119.16 (2017), p. 161101.

[17] Benjamin P Abbott et al. 'Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A'. In: *The Astrophysical Journal Letters* 848.2 (2017), p. L13.

[18] Richard Abbott et al. 'Observation of gravitational waves from two neutron star–black hole coalescences'. In: *The Astrophysical Journal Letters* 915.1 (2021), p. L5.

[19] Shreya Anand et al. 'Optical follow-up of the neutron star–black hole mergers S200105ae and S200115j'. In: *Nature Astronomy* 5.1 (2021), pp. 46–53.

[20] AG Abac et al. 'GWTC-4.0: Updating the gravitational-wave transient catalog with observations from the first part of the fourth LIGO-Virgo-KAGRA observing run'. In: *The Astrophysical Journal Letters* 1004.2 (2026), p. L22.

[21] Charles W. Misner, K. S. Thorne and J. A. Wheeler. *Gravitation*. San Francisco: W. H. Freeman, 1973.

[22] John M Lee. *Introduction to Riemannian manifolds*. Vol. 2. Springer Cham, 2018.

[23] Albert Einstein. 'Kosmologische Betrachtungen zur allgemeinen Relativitäts-theorie'. In: *Das Relativitätsprinzip: Eine Sammlung von Abhandlungen*. Springer, 1917, pp. 130–139.

[24] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press, 2007.

[25] Steven Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. New York: John Wiley and Sons, 1972.

[26] Ashley J Ruiter et al. 'The LISA gravitational wave foreground: a study of double white dwarfs'. In: *The Astrophysical Journal* 717.2 (2010), pp. 1006–1021.

[27] Seppe Staelens and Gijs Nelemans. 'Likelihood of white dwarf binaries to dominate the astrophysical gravitational wave background in the mHz band'. In: *Astronomy & Astrophysics* 683 (2024), A139.

[28] Sophie Hofman and Gijs Nelemans. 'Uncertainty of the white dwarf astrophysical gravitational wave background'. In: *Astronomy & Astrophysics* 691 (2024), A261.

[29] Thomas W Baumgarte and Stuart L Shapiro. *Numerical relativity: solving Einstein's equations on the computer*. Cambridge University Press, 2010.

[30] Luc Blanchet, Thibault Damour and Bala R Iyer. 'Gravitational waves from inspiralling compact binaries: Energy loss and waveform to second-post-Newtonian order'. In: *Physical Review D* 51.10 (1995), p. 5360.

[31] Luc Blanchet. 'Post-Newtonian theory for gravitational waves'. In: *Living Reviews in Relativity* 27.1 (2024), p. 4.

[32] William H Press and Saul A Teukolsky. 'Perturbations of a rotating black hole. II. Dynamical stability of the Kerr metric'. In: *The Astrophysical Journal* 185 (1973), pp. 649–674.

[33] Edward W Leaver. 'An analytic representation for the quasi-normal modes of Kerr black holes'. In: *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* (1985), pp. 285–298.

[34] Olaf Dreyer et al. 'Black-hole spectroscopy: testing general relativity through gravitational-wave observations'. In: *Classical and Quantum Gravity* 21.4 (2004), pp. 787–803.

[35] Robert Abbott et al. 'Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data'. In: *Physical Review D* 106.4 (2022), p. 042003.

[36] S D'Antonio et al. 'Semicoherent analysis method to search for continuous gravitational waves emitted by ultralight boson clouds around spinning black holes'. In: *Physical Review D* 98.10 (2018), p. 103017.

[37] Maximiliano Isi et al. 'Directed searches for gravitational waves from ultralight bosons'. In: *Physical Review D* 99.8 (2019), p. 084042.

[38] Robert Abbott et al. 'All-sky search for gravitational wave emission from scalar boson clouds around spinning black holes in LIGO O3 data'. In: *Physical Review D* 105.10 (2022), p. 102001.

[39] Stan Woosley and Thomas Janka. 'The physics of core-collapse supernovae'. In: *Nature Physics* 1.3 (2005), pp. 147–154.

[40] H-Th Janka et al. 'Theory of core-collapse supernovae'. In: *Physics Reports* 442.1-6 (2007), pp. 38–74.

[41] J Abadie et al. 'Search for gravitational wave bursts from six magnetars'. In: *The Astrophysical Journal Letters* 734.2 (2011), p. L35.

[42] Burkhard Zink, Paul D Lasky and Kostas D Kokkotas. 'Are gravitational waves from giant magnetar flares observable?' In: *Physical Review D* 85.2 (2012), p. 024030.

[43] Thibault Damour and Alexander Vilenkin. 'Gravitational wave bursts from cusps and kinks on cosmic strings'. In: *Physical Review D* 64.6 (2001), p. 064008.

[44] Marco Drago et al. 'Coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis'. In: *SoftwareX* 14 (2021), p. 100678.

[45] Nelson Christensen. 'Stochastic gravitational wave backgrounds'. In: *Reports on Progress in Physics* 82.1 (2019), p. 016903.

[46] Tania Regimbau. 'The astrophysical gravitational wave stochastic background'. In: *Research in Astronomy and Astrophysics* 11.4 (2011), pp. 369–390.

[47] Benjamin P Abbott et al. 'GW150914: Implications for the stochastic gravitational-wave background from binary black holes'. In: *Physical Review Letters* 116.13 (2016), p. 131102.

[48] Benjamin P Abbott et al. 'GW170817: implications for the stochastic gravitational-wave background from compact binary coalescences'. In: *Physical Review Letters* 120.9 (2018), p. 091101.

[49] Gabriella Agazie et al. 'The NANOGrav 15 yr data set: Evidence for a gravitational-wave background'. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L8.

[50] Heng Xu et al. 'Searching for the nano-hertz stochastic gravitational wave background with the Chinese pulsar timing array data release I'. In: *Research in Astronomy and Astrophysics* 23.7 (2023), p. 075024.

[51] 'The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals'. In: *Astronomy & Astrophysics* 678 (2023), A50.

[52] Daniel J Reardon et al. 'Search for an isotropic gravitational-wave background with the Parkes Pulsar Timing Array'. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L6.

[53] Michele Maggiore. 'Gravitational wave experiments and early universe cosmology'. In: *Physics Reports* 331.6 (2000), pp. 283–367.

[54] Michael S. Turner. 'Detectability of inflation-produced gravitational waves'. In: *Physical Review D* 55 (2 1997), R435(R)–R439(R).

[55] Neil Barnaby, Enrico Pajer and Marco Peloso. 'Gauge field production in axion inflation: Consequences for monodromy, non-Gaussianity in the CMB, and gravitational waves at interferometers'. In: *Physical Review D* 85 (2 2012), p. 023525.

[56] Marc Kamionkowski, Arthur Kosowsky and Michael S. Turner. 'Gravitational radiation from first-order phase transitions'. In: *Physical Review D* 49 (6 1994), pp. 2837–2851.

[57] Chiara Caprini et al. 'Science with the space-based interferometer eLISA. II: Gravitational waves from cosmological phase transitions'. In: *Journal of Cosmology and Astroparticle Physics* 2016.04 (2016), pp. 001–001.

[58] Thibault Damour and Alexander Vilenkin. 'Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows'. In: *Physical Review D* 71 (6 2005), p. 063510.

[59] Xavier Siemens, Vuk Mandic and Jolien Creighton. 'Gravitational-Wave Stochastic Background from Cosmic Strings'. In: *Physical Review Letters* 98 (11 2007), p. 111101.

[60] Wei-Tou Ni and Mark Zimmermann. 'Inertial and gravitational effects in the proper reference frame of an accelerated, rotating observer'. In: *Physical Review D* 17.6 (1978), p. 1473.

[61] Bernard F Schutz and Massimo Tinto. 'Antenna patterns of interferometric detectors of gravitational waves–I. Linearly polarized waves'. In: *Monthly Notices of the Royal Astronomical Society* 224.1 (1987), pp. 131–154.

[62] John T. Whelan. *The Geometry of Gravitational Wave Detection.* LIGO Technical Note T1300666-v3. LIGO Scientific Collaboration, 2013.

[63] Kentaro Somiya and KAGRA collaboration. 'Detector configuration of KAGRA–the Japanese cryogenic gravitational-wave detector'. In: *Classical and Quantum Gravity* 29.12 (2012), p. 124007.

[64] Yoichi Aso et al. 'Interferometer design of the KAGRA gravitational wave detector'. In: *Physical Review D* 88.4 (2013), p. 043007.

[65] Federica Santucci, Mariateresa Crosta and Mario Gilberto Lattanzi. 'Pattern functions of the Astrometric Gravitational Wave Antenna'. In: *Scientific Reports* 15.1 (2025), p. 32908.

[66] Peter R Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, 1994.

[67] Kirk McKenzie et al. 'Experimental demonstration of a squeezing-enhanced power-recycled Michelson interferometer for gravitational wave detection'. In: *Physical Review Letters* 88.23 (2002), p. 231102.

[68] Craig Cahillane and Georgia Mansell. 'Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four'. In: *Galaxies* 10.1 (2022), p. 36.

[69] TJ Balle and WH Flygare. 'Fabry–Perot cavity pulsed Fourier transform microwave spectrometer with a pulsed nozzle particle source'. In: *Review of Scientific Instruments* 52.1 (1981), pp. 33–45.

[70] BJ Meers. 'The frequency response of interferometric gravitational wave detectors'. In: *Physics Letters A* 142.8-9 (1989), pp. 465–470.

[71] Rana X Adhikari. 'Gravitational radiation detection with laser interferometry'. In: *Reviews of Modern Physics* 86.1 (2014), pp. 121–151.

[72] Benjamin P Abbott et al. 'A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals'. In: *Classical and Quantum Gravity* 37.5 (2020), p. 055002.

[73] Derek Davis et al. 'Improving the sensitivity of Advanced LIGO using noise subtraction'. In: *Classical and Quantum Gravity* 36.5 (2019), p. 055011.

[74] Benjamin J Owen and Bangalore Suryanarayana Sathyaprakash. 'Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement'. In: *Physical Review D* 60.2 (1999), p. 022002.

[75] S Pardi et al. 'Dynamic Matched Filters for Gravitational Waves Detection'. In: *American Institute of Physics Conference Proceedings* 751 (2005), pp. 236–238.

[76] Curt Cutler and Eanna E Flanagan. 'Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?' In: *Physical Review D* 49.6 (1994), p. 2658.

[77] Piotr Jaranowski and Andrzej Królak. 'Gravitational-wave data analysis. Formalism and sample applications: the Gaussian case'. In: *Living Reviews in Relativity* 15.1 (2012), p. 4.

[78] Kipp Cannon et al. 'GstLAL: A software framework for gravitational wave discovery'. In: *SoftwareX* 14 (2021), p. 100680.

[79] Virginia Klema and Alan Laub. 'The singular value decomposition: Its computation and some applications'. In: *IEEE Transactions on automatic control* 25.2 (1980), pp. 164–176.

[80] Michael E Wall, Andreas Rechtsteiner and Luis M Rocha. 'Singular value decomposition and principal component analysis'. In: *A practical approach to microarray data analysis*. Springer, 2003, pp. 91–109.

[81] Samantha A Usman et al. 'The PyCBC search for gravitational waves from compact binary coalescence'. In: *Classical and Quantum Gravity* 33.21 (2016), p. 215004.

[82] J Abadie et al. 'Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3'. In: *Physical Review D* 85.8 (2012), p. 082002.

[83] S Babak et al. 'Searching for gravitational waves from binary coalescence'. In: *Physical Review D* 87.2 (2013), p. 024033.

[84] Lee S Finn. 'Detection, measurement, and gravitational radiation'. In: *Physical Review D* 46.12 (1992), p. 5236.

[85] Benjamin P Abbott et al. 'The rate of binary black hole mergers inferred from Advanced LIGO observations surrounding GW150914'. In: *The Astrophysical Journal Letters* 833.1 (2016), p. L1.

[86] Will M Farr et al. 'Counting and confusion: Bayesian rate estimation with multiple populations'. In: *Physical Review D* 91.2 (2015), p. 023005.

[87] TA Callister et al. 'Observing gravitational waves with a single detector'. In: *Classical and Quantum Gravity* 34.15 (2017), p. 155007.

[88] Albert Einstein. 'Erklärung der Perihelbewegung des Merkur aus der allgemeinen Relativitätstheorie'. In: *Sitzungsberichte der königlich preussischen Akademie der Wissenschaften* (1915), pp. 831–839.

[89] Albert Einstein. 'Lens-like action of a star by the deviation of light in the gravitational field'. In: *Science* 84.2188 (1936), pp. 506–507.

[90] Frank Watson Dyson, Arthur Stanley Eddington and Charles Davidson. 'IX. A determination of the deflection of light by the Sun's gravitational field, from observations made at the total eclipse of May 29, 1919'. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 220.571-581 (1920), pp. 291–333.

[91] Esteban Roulet and Silvia Mollerach. 'Microlensing'. In: *Physics Reports* 279.2 (1997), pp. 67–118.

[92] Evencio Mediavilla et al. 'Microlensing-based estimate of the mass fraction in compact objects in lens galaxies'. In: *The Astrophysical Journal* 706.2 (2009), pp. 1451–1462.

[93] B Scott Gaudi. 'Microlensing surveys for exoplanets'. In: *Annual Review of Astronomy and Astrophysics* 50.1 (2012), pp. 411–453.

[94] Dennis Walsh*, Robert F Carswell and Ray J Weymann. '0957+ 561 A, B: twin quasistellar objects or gravitational lens?' In: *Nature* 279.5712 (1979), pp. 381–384.

[95] Frederic Courbin, Prasenjit Saha and Paul L Schechter. 'Quasar lensing'. In: *Gravitational Lensing: An Astrophysical Tool*. Springer, 2002, pp. 1–54.

[96] Yun Wang. ‘Supernova pencil beam survey’. In: *The Astrophysical Journal* 531.2 (2000), pp. 676–683.

[97] Daniel E Holz. ‘Seeing double: strong gravitational lensing of high-redshift supernovae’. In: *The Astrophysical Journal Letters* 556.2 (2001), pp. L71–L74.

[98] Masamune Oguri. ‘Strong gravitational lensing of explosive transients’. In: *Reports on Progress in Physics* 82.12 (2019), p. 126901.

[99] Richard Massey et al. ‘Weak lensing from space. II. Dark matter mapping’. In: *The Astronomical Journal* 127.6 (2004), pp. 3089–3101.

[100] Henk Hoekstra, Howard KC Yee and Michael D Gladders. ‘Properties of galaxy dark matter halos from weak lensing’. In: *The Astrophysical Journal* 606.1 (2004), pp. 67–77.

[101] Antony Lewis and Anthony Challinor. ‘Weak gravitational lensing of the CMB’. In: *Physics Reports* 429.1 (2006), pp. 1–65.

[102] Graham P. Smith, Tessa Baker, Simon Birrer et al. ‘Multi-messenger gravitational lensing’. In: *Philosophical Transactions of the Royal Society of London Series A* 383.2295, 20240134 (2025), p. 20240134.

[103] Peter Schneider, Christopher Kochanek and Joachim Wambsganss. *Gravitational lensing: strong, weak and micro: Saas-Fee advanced course 33*. Vol. 33. Springer Science & Business Media, 2006.

[104] Eungwang Seo et al. ‘Residual test to search for microlensing signatures in strongly lensed gravitational wave signals’. In: *The Astrophysical Journal* 988.2 (2025), p. 159.

[105] Irwin I Shapiro. ‘Fourth test of general relativity’. In: *Physical Review Letters* 13.26 (1964), p. 789.

[106] Matthias Bartelmann and Peter Schneider. ‘Weak gravitational lensing’. In: *Physics Reports* 340.4-5 (2001), pp. 291–472.

[107] Roger Blandford and Ramesh Narayan. ‘Fermat’s principle, caustics, and the classification of gravitational lens images’. In: *The Astrophysical Journal, Part 1* 310 (1986), pp. 568–582.

[108] Richard S. Ellis. ‘Gravitational lensing: a unique probe of dark matter and dark energy’. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 368.1914 (2010), pp. 967–987.

[109] Israel Kovner. ‘The quadrupole gravitational lens’. In: *The Astrophysical Journal, Part 1* 312 (1987), pp. 22–44.

[110] Tommaso Treu and Philip J. Marshall. ‘Time Delay Cosmography’. In: *The Astronomy and Astrophysics Review* 24.1 (2016), p. 11.

[111] Simon Birrer et al. ‘Time-delay cosmography: measuring the Hubble constant and other cosmological parameters with strong gravitational lensing’. In: *Space Science Reviews* 220.5 (2024), p. 48.

[112] Martin Millon et al. 'TDCOSMO-I. An exploration of systematic uncertainties in the inference of H0 from time-delay cosmography'. In: *Astronomy & Astrophysics* 639 (2020), A101.

[113] William L Burke. 'Multiple gravitational imaging by distributed masses'. In: *The Astrophysical Journal* 244 (1981), p. L1.

[114] Peter Schneider. 'Gravitational lensing statistics'. In: *Gravitational Lenses: Proceedings of a Conference Held in Hamburg, Germany 9–13 September 1991*. Springer. 2005, pp. 196–208.

[115] James Binney and Scott Tremaine. *Galactic dynamics*. Princeton university press, 2011.

[116] Julio F. Navarro, Carlos S. Frenk and Simon D. M. White. 'The Structure of Cold Dark Matter Halos'. In: *The Astrophysical Journal* 462 (1996), p. 563.

[117] Julio F Navarro, Carlos S Frenk and Simon DM White. 'A universal density profile from hierarchical clustering'. In: *The Astrophysical Journal* 490.2 (1997), p. 493.

[118] Matthias Bartelmann. 'Arcs from a universal dark-matter halo profile'. In: *arXiv preprint astro-ph/9602053* (1996).

[119] Robert Kormann, Peter Schneider and Matthias Bartelmann. 'Isothermal elliptical gravitational lens models'. In: *Astronomy and Astrophysics* 284 (1994), pp. 285–299.

[120] Rennan Barkana. 'Fast calculation of a family of elliptical gravitational lens models'. In: *The Astrophysical Journal* 502.2 (1998), p. 531.

[121] Nicolas Tessore and R Benton Metcalf. 'The elliptical power law profile lens'. In: *Astronomy & Astrophysics* 580 (2015), A79.

[122] RR Bourassa, R Kantowski and TD Norton. 'The spheroidal gravitational lens'. In: *The Astrophysical Journal* 185 (1973), pp. 747–756.

[123] RR Bourassa and R Kantowski. 'The theory of transparent gravitational lenses'. In: *The Astrophysical Journal* 195 (1975), pp. 13–21.

[124] Ghislain Golse and J-P Kneib. 'Pseudo elliptical lensing mass model: Application to the NFW mass distribution'. In: *Astronomy & Astrophysics* 390.3 (2002), pp. 821–827.

[125] YP Jing and Yasushi Suto. 'Triaxial modeling of halo density profiles with high-resolution N-body simulations'. In: *The Astrophysical Journal* 574.2 (2002), pp. 538–553.

[126] Charles R Keeton. 'A catalog of mass models for gravitational lensing'. In: *arXiv preprint astro-ph/0102341* (2001).

[127] WJG De Blok. 'The Core-Cusp Problem'. In: *Advances in Astronomy* 2010.1 (2010), p. 789293.

[128] Andrew Gould. 'Femtolensing of gamma-ray bursters'. In: *The Astrophysical Journal, Part 2-Letters* 386 (1992), pp. L5–L7.

[129] Robert J Nemiroff and Andrew Gould. 'Probing for MACHOs of Mass 10- 15 M⊙ to 10- 7 M⊙ with Gamma-Ray Burst Parallax Spacecraft'. In: *The Astrophysical Journal Letters* 452.2 (1995), pp. L111–L114.

[130] Andrey Katz et al. ‘Femtolensing by dark matter revisited’. In: *Journal of Cosmology and Astroparticle Physics* 2018.12 (2018), pp. 005–005.

[131] Peter Schneider. ‘Weak gravitational lensing’. In: *Gravitational lensing: strong, weak and micro*. Springer, 2006, pp. 269–451.

[132] Nabila Aghanim et al. ‘Planck 2018 results-viii. gravitational lensing’. In: *Astronomy & Astrophysics* 641 (2020), A8.

[133] Antony Lewis. ‘Lensed CMB simulation and parameter estimation’. In: *Physical Review D* 71.8 (2005), p. 083008.

[134] Pawan Kumar and Bing Zhang. ‘The physics of gamma-ray bursts & relativistic jets’. In: *Physics Reports* 561 (2015), pp. 1–109.

[135] Brian D Metzger. ‘Kilonovae’. In: *Living Reviews in Relativity* 23.1 (2020), p. 1.

[136] SE Woosley and JS Bloom. ‘The supernova–gamma-ray burst connection’. In: *Annu. Rev. Astron. Astrophys.* 44.1 (2006), pp. 507–556.

[137] Bohdan Paczynski. ‘Gravitational microlensing and gamma-ray bursts’. In: *The Astrophysical Journal* 317 (1987), p. L51.

[138] Abraham Loeb and Rosalba Perna. ‘Microlensing of gamma-ray burst afterglows’. In: *The Astrophysical Journal* 495.2 (1998), pp. 597–603.

[139] Sherry H Suyu et al. ‘Strong gravitational lensing and microlensing of supernovae’. In: *Space Science Reviews* 220.1 (2024), p. 13.

[140] Gregory Dobler and Charles R Keeton. ‘Microlensing of lensed supernovae’. In: *The Astrophysical Journal* 653.2 (2006), pp. 1391–1399.

[141] Daniel A Goldstein et al. ‘Precise time delays from strongly gravitationally lensed type Ia supernovae with chromatically microlensed images’. In: *The Astrophysical Journal* 855.1 (2018), p. 22.

[142] Philip F Hopkins, Gordon T Richards and Lars Hernquist. ‘An observational determination of the bolometric quasar luminosity function’. In: *The Astrophysical Journal* 654.2 (2007), pp. 731–753.

[143] M Tewes et al. ‘COSMOGRAIL: the COSmological MOnitoring of GRAvItational Lenses-XIII. Time delays and 9-yr optical monitoring of the lensed quasar RX J1131-1231’. In: *Astronomy & Astrophysics* 556 (2013), A22.

[144] M Millon et al. ‘Cosmograil-xix. time delays in 18 strongly lensed quasars from 15 years of optical monitoring’. In: *Astronomy & Astrophysics* 640 (2020), A105.

[145] Shuji Deguchi and William D Watson. ‘Diffraction in gravitational lensing for compact objects of low mass’. In: *The Astrophysical Journal, Part 1* 307 (1986), pp. 30–37.

[146] Takahiro T Nakamura. ‘Gravitational lensing of gravitational waves from inspiraling binaries by a point mass lens’. In: *Physical Review Letters* 80.6 (1998), p. 1138.

[147] Ryuichi Takahashi and Takashi Nakamura. ‘Wave effects in the gravitational lensing of gravitational waves from chirping binaries’. In: *The Astrophysical Journal* 595.2 (2003), p. 1039.

[148] Christian Baraldo, Akio Hosoya and Takahiro T Nakamura. 'Gravitationally induced interference of gravitational waves by a rotating massive object'. In: *Physical Review D* 59.8 (1999), p. 083001.

[149] Max Born, Emil Wolf et al. *Principles of optics*. Vol. 6. Pergamon press Oxford, 1964.

[150] Ernst Eduard Kummer. 'De integralibus quibusdam definitis et seriebus infinitis.' In: *Journal für die reine und angewandte Mathematik* 17 (1837), pp. 228–242.

[151] Zhoujian Cao, Li-Fang Li and Yan Wang. 'Gravitational lensing effects on parameter estimation in gravitational wave detection with advanced detectors'. In: *Physical Review D* 90 (6 2014), p. 062003.

[152] Kwun-Hang Lai et al. 'Discovering intermediate-mass black hole lenses through gravitational wave lensing'. In: *Physical Review D* 98 (8 2018), p. 083005.

[153] Pierre Christian, Salvatore Vitale and Abraham Loeb. 'Detecting stellar lensing of gravitational waves with ground-based observatories'. In: *Physical Review D* 98 (10 2018), p. 103022.

[154] Sunghoon Jung and Chang Sub Shin. 'Gravitational-Wave Fringes at LIGO: Detecting Compact Dark Matter by Gravitational Lensing'. In: *Physical Review Letters* 122 (4 2019), p. 041103.

[155] Yun Wang, Albert Stebbins and Edwin L Turner. 'Gravitational lensing of gravitational waves from merging neutron star binaries'. In: *Physical Review Letters* 77.14 (1996), p. 2875.

[156] Liang Dai and Tejaswi Venumadhav. 'On the waveforms of gravitationally lensed gravitational waves'. In: *arXiv preprint arXiv:1702.04724* (2017).

[157] Jose María Ezquiaga et al. 'Phase effects from strong gravitational lensing of gravitational waves'. In: *Physical Review D* 103.6 (2021), p. 064047.

[158] Ken KY Ng et al. 'Precise LIGO lensing rate predictions for binary black holes'. In: *Physical Review D* 97.2 (2018), p. 023012.

[159] Shun-Sheng Li et al. 'Gravitational lensing of gravitational waves: A statistical perspective'. In: *Monthly Notices of the Royal Astronomical Society* 476.2 (2018), pp. 2220–2229.

[160] Masamune Oguri. 'Effect of gravitational lensing on the distribution of gravitational waves from distant binary black hole mergers'. In: *Monthly Notices of the Royal Astronomical Society* 480.3 (2018), pp. 3842–3855.

[161] Graham P Smith et al. 'Strong-lensing of gravitational waves by galaxy clusters'. In: *Proceedings of the International Astronomical Union* 13.S338 (2017), pp. 98–102.

[162] Graham P Smith et al. 'What if LIGO's gravitational wave detections are strongly lensed by massive galaxy clusters?' In: *Monthly Notices of the Royal Astronomical Society* 475.3 (2018), pp. 3823–3828.

[163] GP Smith et al. 'Deep and rapid observations of strong-lensing galaxy clusters within the sky localization of GW170814'. In: *Monthly Notices of the Royal Astronomical Society* 485.4 (2019), pp. 5180–5191.

[164] Andrew Robertson et al. 'What does strong gravitational lensing? The mass and redshift distribution of high-magnification lenses'. In: *Monthly Notices of the Royal Astronomical Society* 495.4 (2020), pp. 3727–3739.

[165] Dan Ryczanowski et al. 'On building a cluster watchlist for identifying strongly lensed supernovae, gravitational waves and kilonovae'. In: *Monthly Notices of the Royal Astronomical Society* 495.2 (2020), pp. 1666–1671.

[166] Anna Liu et al. 'Exploring the hidden Universe: a novel phenomenological approach for recovering arbitrary gravitational-wave millilensing configurations'. In: *Monthly Notices of the Royal Astronomical Society* 525.3 (2023), pp. 4149–4160.

[167] Tessa Baker and Mark Trodden. 'Multimessenger time delays from lensed gravitational waves'. In: *Physical Review D* 95.6 (2017), p. 063512.

[168] Thomas E Collett and David Bacon. 'Testing the speed of gravitational waves over cosmological distances with strong gravitational lensing'. In: *Physical Review Letters* 118.9 (2017), p. 091101.

[169] Xi-Long Fan et al. 'Speed of gravitational waves from strongly lensed gravitational waves and electromagnetic signals'. In: *Physical Review Letters* 118.9 (2017), p. 091102.

[170] Enis Belgacem et al. 'Gravitational-wave luminosity distance in modified gravity theories'. In: *Physical review D* 97.10 (2018), p. 104066.

[171] Atsushi Nishizawa. 'Generalized framework for testing gravity with gravitational-wave propagation. I. Formulation'. In: *Physical Review D* 97.10 (2018), p. 104037.

[172] Andreas Finke et al. 'Cosmology with LIGO/Virgo dark sirens: Hubble parameter and modified gravitational wave propagation'. In: *Journal of Cosmology and Astroparticle Physics* 2021.08 (2021), p. 026.

[173] Enis Belgacem et al. 'Modified gravitational-wave propagation and standard sirens'. In: *Physical review D* 98.2 (2018), p. 023510.

[174] Peter AR Ade et al. 'Planck 2015 results-xiii. cosmological parameters'. In: *Astronomy & Astrophysics* 594 (2016), A13.

[175] N. Aghanim, Y. Akrami, M. Ashdown et al. 'Planck 2018 results. VI. Cosmological parameters'. In: *Astronomy & Astrophysics* 641, A6 (2020). [Erratum: Astronomy & Astrophysics 652, C4 (2021)].

[176] Bruce Bassett and Renée Hlozek. 'Baryon acoustic oscillations'. In: *Dark energy: observational and theoretical approaches* 246 (2010), p. 53.

[177] David H Weinberg et al. 'Observational probes of cosmic acceleration'. In: *Physics Reports* 530.2 (2013), pp. 87–255.

[178] Adam G Riess, William H Press and Robert P Kirshner. 'A precise distance indicator: Type Ia supernova multicolor light-curve shapes'. In: *The Astrophysical Journal* 473.1 (1996), pp. 88–109.

[179] Bruno Leibundgut. 'Cosmological implications from observations of type Ia supernovae'. In: *Annual Review of Astronomy and Astrophysics* 39.1 (2001), pp. 67–98.

[180] Eleonora Di Valentino et al. 'In the realm of the Hubble tension—a review of solutions'. In: *Classical and Quantum Gravity* 38.15 (2021), p. 153001.

[181] Adam G Riess et al. 'A comprehensive measurement of the local value of the Hubble constant with 1 km s^{-1} Mpc^{-1} uncertainty from the Hubble Space Telescope and the SH0ES team'. In: *The Astrophysical Journal Letters* 934.1 (2022), p. L7.

[182] Elcio Abdalla et al. 'Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies'. In: *Journal of High Energy Astrophysics* 34 (2022), pp. 49–211.

[183] Archisman Ghosh. *Charting the expansion of the Universe with gravitational waves.* Innovation News Network. June 2025. URL: `https://www.innovationnewsnetwork.com/charting-the-expansion-of-universe-with-gravitational-waves/58293/`.

[184] P. J. E. Peebles. *Principles of Physical Cosmology.* Princeton University Press, 1993.

[185] Robert J Nemiroff and Bijunath Patla. 'Adventures in Friedmann cosmology: A detailed expansion of the cosmological Friedmann equations'. In: *American Journal of Physics* 76.3 (2008), pp. 265–276.

[186] John A Peacock. *Cosmological physics.* Cambridge University Press, 1999.

[187] Steven Weinberg. *Cosmology.* OUP Oxford, 2008.

[188] David W Hogg. 'Distance measures in cosmology'. In: *arXiv preprint astro-ph/9905116* (1999).

[189] Georges Lemaître. 'Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques'. In: *Annales de la Société Scientifique de Bruxelles, A47, p. 49-59* 47 (1927), pp. 49–59.

[190] Edwin Hubble. 'A relation between distance and radial velocity among extra-galactic nebulae'. In: *Proceedings of the National Academy of Sciences* 15 (1929), pp. 168–173.

[191] DES Collaboration: TMC Abbott et al. 'The dark energy survey: Cosmology results with 1500 new high-redshift type Ia supernovae using the full 5 yr data set'. In: *The Astrophysical Journal Letters* 973.1 (2024), p. L14.

[192] Wayne Hu and Scott Dodelson. 'Cosmic microwave background anisotropies'. In: *Annual Review of Astronomy and Astrophysics* 40.1 (2002), pp. 171–216.

[193] George Efstathiou and J Richard Bond. 'Cosmic confusion: degeneracies among cosmological parameters derived from measurements of microwave background anisotropies'. In: *Monthly Notices of the Royal Astronomical Society* 304.1 (1999), pp. 75–97.

[194] Daniel J Eisenstein, Wayne Hu and Max Tegmark. 'Cosmic complementarity: H 0 and Ω m from combining cosmic microwave background experiments and redshift surveys'. In: *The Astrophysical Journal Letters* 504.2 (1998), pp. L57–L60.

[195] Chris Blake and Karl Glazebrook. 'Probing dark energy using baryonic oscillations in the galaxy power spectrum as a cosmological ruler'. In: *The Astrophysical Journal* 594.2 (2003), pp. 665–673.

[196] Hee-Jong Seo and Daniel J Eisenstein. 'Probing dark energy with baryonic acoustic oscillations from future large galaxy redshift surveys'. In: *The Astrophysical Journal* 598.2 (2003), pp. 720–740.

[197] M Abdul Karim et al. 'DESI DR2 results. II. Measurements of baryon acoustic oscillations and cosmological constraints'. In: *Physical Review D* 112.8 (2025), p. 083515.

[198] Hsin-Yu Chen et al. 'Distance measures in gravitational-wave astrophysics and cosmology'. In: *Classical and Quantum Gravity* 38.5 (2021), p. 055010.

[199] AG Abac et al. 'Open Data from LIGO, Virgo, and KAGRA through the First Part of the Fourth Observing Run'. In: *The Astrophysical Journal* 1004.2 (2026), p. 232.

[200] Jie Zheng et al. 'Investigating the dynamical models of cosmology with recent observations and upcoming gravitational-wave data'. In: *The European Physical Journal Plus* 137.4 (2022), p. 509.

[201] Shang-Jie Jin et al. 'Gravitational wave standard sirens: A brief review of cosmological parameter estimation'. In: *Science China Physics, Mechanics & Astronomy* 69.2 (2026), p. 220401.

[202] AG Abac et al. 'GWTC-4.0: Constraints on the Cosmic Expansion Rate and Modified Gravitational-wave Propagation'. In: *The Astrophysical Journal Letters* 1007.1 (2026), p. L17.

[203] Bernard F Schutz. 'Determining the Hubble constant from gravitational wave observations'. In: *Nature* 323.6086 (1986), pp. 310–311.

[204] Daniel E Holz and Scott A Hughes. 'Using gravitational-wave standard sirens'. In: *The Astrophysical Journal* 629.1 (2005), pp. 15–22.

[205] Samaya Nissanke et al. 'Exploring short gamma-ray bursts as gravitational-wave standard sirens'. In: *The Astrophysical Journal* 725.1 (2010), pp. 496–514.

[206] Samaya Nissanke et al. 'Determining the Hubble constant from gravitational wave observations of merging compact binaries'. In: *arXiv preprint arXiv:1307.2638* (2013).

[207] SJea Smartt et al. 'A kilonova as the electromagnetic counterpart to a gravitational-wave source'. In: *Nature* 551.7678 (2017), pp. 75–79.

[208] Iair Arcavi et al. 'Optical emission from a kilonova following a gravitational-wave-detected neutron-star merger'. In: *Nature* 551.7678 (2017), pp. 64–66.

[209] Salvatore Vitale and Hsin-Yu Chen. 'Measuring the Hubble constant with neutron star black hole mergers'. In: *Physical Review Letters* 121.2 (2018), p. 021303.

[210] Tanja Hinderer. 'Tidal Love numbers of neutron stars'. In: *The Astrophysical Journal* 677.2 (2008), pp. 1216–1220.

[211] C. Messenger and J. Read. 'Measuring a Cosmological Distance-Redshift Relationship Using Only Gravitational Wave Observations of Binary Neutron Star Coalescences'. In: *Physical Review Letters* 108 (9 2012), p. 091101.

[212] Kent Yagi. 'Multipole Love relations'. In: *Physical Review D* 89 (4 2014), p. 043011.

[213] B. P. Abbott et al. 'A gravitational-wave standard siren measurement of the Hubble constant'. In: *Nature* 551.7678 (2017), pp. 85–88.

[214] Gergely Dálya et al. 'GLADE+: an extended galaxy catalogue for multimessenger searches with advanced gravitational-wave detectors'. In: *Monthly Notices of the Royal Astronomical Society* 514.1 (2022), pp. 1403–1411.

[215] Walter Del Pozzo. 'Inference of cosmological parameters from gravitational waves: Applications to second generation interferometers'. In: *Physical Review D* 86.4 (2012), p. 043011.

[216] Hsin-Yu Chen, Maya Fishbach and Daniel E Holz. 'A two per cent Hubble constant measurement from standard sirens within five years'. In: *Nature* 562.7728 (2018), pp. 545–547.

[217] Maya Fishbach et al. 'A standard siren measurement of the Hubble constant from GW170817 without the electromagnetic counterpart'. In: *The Astrophysical Journal Letters* 871.1 (2019), p. L13.

[218] Marcelle Soares-Santos et al. 'First measurement of the Hubble constant from a dark standard siren using the dark energy survey galaxies and the LIGO/Virgo binary–black-hole merger GW170814'. In: *The Astrophysical Journal Letters* 876.1 (2019), p. L7.

[219] Rachel Gray et al. 'Cosmological inference using gravitational wave standard sirens: A mock data analysis'. In: *Physical Review D* 101.12 (2020), p. 122001.

[220] Rachel Gray, Freija Beirnaert, Christos Karathanasis et al. 'Joint cosmological and gravitational-wave population inference using dark sirens and galaxy catalogues'. In: *Journal of Cosmology and Astroparticle Physics* 12 (2023), p. 023.

[221] Stephen R. Taylor, Jonathan R. Gair and Ilya Mandel. 'Cosmology using advanced gravitational-wave detectors alone'. In: *Physical Review D* 85 (2 2012), p. 023535.

[222] Walter Del Pozzo, Tjonnie GF Li and Chris Messenger. 'Cosmological inference using only gravitational wave observations of binary neutron stars'. In: *Physical Review D* 95.4 (2017), p. 043502.

[223] Jose María Ezquiaga and Daniel E Holz. 'Spectral sirens: Cosmology from the full mass distribution of compact binaries'. In: *Physical Review Letters* 129.6 (2022), p. 061102.

[224] Suvodip Mukherjee et al. 'Cross-correlating dark sirens and galaxies: Constraints on H 0 from GWTC-3 of LIGO–Virgo–KAGRA'. In: *The Astrophysical Journal* 975.2 (2024), p. 189.

[225] Samsuzzaman Afroz and Suvodip Mukherjee. 'Prospect of precision cosmology and testing general relativity using binary black holes–galaxies cross-correlation'. In: *Monthly Notices of the Royal Astronomical Society* 534.2 (2024), pp. 1283–1298.

[226] Isabela Santiago de Matos et al. 'First measurement of the Hubble constant from gravitational wave-galaxy cross-correlations'. In: *arXiv preprint arXiv:2512.15380* (2025).

[227] Constantina Nicolaou et al. 'The impact of peculiar velocities on the estimation of the Hubble constant from gravitational wave standard sirens'. In: *Monthly Notices of the Royal Astronomical Society* 495.1 (2020), pp. 90–97.

[228] Chris Blake and Ryan J Turner. 'The role of peculiar velocity uncertainties in standard siren cosmology'. In: *arXiv preprint arXiv:2509.03101* (2025).

[229] Ariel J Amsellem et al. 'Probing the environment around GW170817 with DESI: insights on galaxy group peculiar velocities for standard siren measurements'. In: *The Astrophysical Journal* 1001.2 (2026), p. 157.

[230] MV Van Der Sluys et al. 'Gravitational-wave astronomy with inspiral signals of spinning compact-object binaries'. In: *The Astrophysical Journal Letters* 688.2 (2008), pp. L61–L64.

[231] V Raymond et al. 'Degeneracies in sky localization determination from a spinning coalescing binary through gravitational wave observations: a Markov-chain Monte Carlo analysis for two detectors'. In: *Classical and Quantum Gravity* 26.11 (2009), p. 114007.

[232] John Veitch and Alberto Vecchio. 'Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network'. In: *Physical Review D* 81.6 (2010), p. 062003.

[233] J Veitch et al. 'Estimating parameters of coalescing compact binaries with proposed advanced detector networks'. In: *Physical Review D* 85.10 (2012), p. 104045.

[234] Carl L Rodriguez et al. 'Basic parameter estimation of binary neutron star systems by the Advanced LIGO/Virgo network'. In: *The Astrophysical Journal* 784.2 (2014), p. 119.

[235] K Grover et al. 'Comparison of gravitational wave detector network sky localization approximations'. In: *Physical Review D* 89.4 (2014), p. 042004.

[236] John Veitch et al. 'Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library'. In: *Physical Review D* 91.4 (2015), p. 042003.

[237] A. G. Abac, I. Abouelfettouh, F. Acernese et al. 'GWTC-4.0: Population Properties of Merging Compact Binaries'. In: *The Astrophysical Journal Letters* 1005.2 (2026), p. L51.

[238] AG Abac et al. 'GWTC-4.0: Searches for Gravitational-Wave Lensing Signatures'. In: *arXiv preprint arXiv:2512.16347* (2025).

[239] R. Bellman, R.E. Bellman and Rand Corporation. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957.

[240] Gregory Ashton et al. 'BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy'. In: *The Astrophysical Journal Supplement Series* 241.2 (2019), p. 27.

[241] Michael J Williams, John Veitch and Chris Messenger. 'Nested sampling with normalizing flows for gravitational-wave inference'. In: *Physical Review D* 103.10 (2021), p. 103006.

[242] Maximilian Dax et al. 'Real-time gravitational wave science with neural posterior estimation'. In: *Physical Review Letters* 127.24 (2021), p. 241103.

[243] Nicholas Metropolis et al. 'Equation of state calculations by fast computing machines'. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092.

[244] W. K. Hastings. 'Monte Carlo sampling methods using Markov chains and their applications'. In: *Biometrika* 57.1 (1970), pp. 97–109.

[245] Andrew Gelman and Donald B Rubin. 'Inference from iterative simulation using multiple sequences'. In: *Statistical Science* 7.4 (1992), pp. 457–472.

[246] John Skilling. 'Nested sampling'. In: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering* 735 (2004), pp. 395–405.

[247] John Skilling. 'Nested sampling for general Bayesian computation'. In: *Bayesian Analysis* 1.4 (2006), pp. 833–859.

[248] Andrew Gelman et al. *Bayesian Data Analysis.* Chapman and Hall/CRC, 1995.

[249] Lee Samuel Finn. 'Issues in gravitational wave data analysis'. In: *arXiv preprint gr-qc/9709077* (1997).

[250] Devinderjit Sivia and John Skilling. *Data analysis: a Bayesian tutorial.* Oxford University Press, 2006.

[251] Eric Thrane and Colm Talbot. 'An introduction to Bayesian inference in gravitational-wave astronomy: parameter estimation, model selection, and hierarchical models'. In: *Publications of the Astronomical Society of Australia* 36, e010 (2019). [Erratum: Publications of the Astronomical Society of Australia, 37, e036 (2020)].

[252] Neil J Cornish and Tyson B Littenberg. 'Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches'. In: *Classical and Quantum Gravity* 32.13 (2015), p. 135012.

[253] Robert E Kass and Adrian E Raftery. 'Bayes factors'. In: *Journal of the American Statistical Association* 90.430 (1995), pp. 773–795.

[254] David JC MacKay. *Information theory, inference and learning algorithms.* Cambridge University Press, 2003.

[255] Kipp Cannon, Chad Hanna and Jacob Peoples. 'Likelihood-ratio ranking statistic for compact binary coalescence candidates with rate estimation'. In: *arXiv preprint arXiv:1504.04632* (2015).

[256] Phil Gregory. *Bayesian logical data analysis for the physical sciences: A comparative approach with Mathematica® support.* Cambridge University Press, 2005.

[257] Nicolas Lartillot and Hervé Philippe. 'Computing Bayes factors using thermodynamic integration'. In: *Systematic Biology* 55.2 (2006), pp. 195–207.

[258] Charles J Geyer. 'Practical markov chain monte carlo'. In: *Statistical Science* (1992), pp. 473–483.

[259] Alan Sokal. 'Monte Carlo methods in statistical mechanics: foundations and new algorithms'. In: *Functional Integration: Basics and Applications.* Springer, 1997, pp. 131–192.

[260] Stephen P Brooks and Andrew Gelman. 'General methods for monitoring convergence of iterative simulations'. In: *Journal of Computational and Graphical Statistics* 7.4 (1998), pp. 434–455.

[261] Daniel Foreman-Mackey et al. 'emcee: the MCMC hammer'. In: *Publications of the Astronomical Society of the Pacific* 125.925 (2013), pp. 306–312.

[262] Farhan Feroz et al. 'Importance nested sampling and the MultiNest algorithm'. In: *arXiv preprint arXiv:1306.2144* (2013).

[263] Joshua S Speagle. 'dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences'. In: *Monthly Notices of the Royal Astronomical Society* 493.3 (2020), pp. 3132–3158.

[264] Johannes Buchner. 'PyMultiNest: Python interface for MultiNest'. In: *Astrophysics Source Code Library* (2016), ascl–1606.

[265] Farhan Feroz, Michael P Hobson and Michael Bridges. 'MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics'. In: *Monthly Notices of the Royal Astronomical Society* 398.4 (2009), pp. 1601–1614.

[266] AG Abac et al. 'GWTC-5.0: Tests of General Relativity'. In: *arXiv preprint arXiv:2607.19293* (2026).

[267] Bangalore Suryanarayana Sathyaprakash and SV Dhurandhar. 'Choice of filters for the detection of gravitational waves from coalescing binaries'. In: *Physical Review D* 44.12 (1991), p. 3819.

[268] Thibault Damour, Piotr Jaranowski and Gerhard Schäfer. 'Dimensional regularization of the gravitational interaction of point masses'. In: *Physics Letters B* 513.1-2 (2001), pp. 147–155.

[269] Luc Blanchet et al. 'Dimensional regularization of the third post-Newtonian gravitational wave generation from two point masses'. In: *Physical Review D* 71.12 (2005), p. 124004.

[270] Walter D Goldberger and Ira Z Rothstein. 'Effective field theory of gravity for extended objects'. In: *Physical Review D* 73.10 (2006), p. 104029.

[271] Alejandro Bohé, Sylvain Marsat and Luc Blanchet. 'Next-to-next-to-leading order spin–orbit effects in the gravitational wave flux and orbital phasing of compact binaries'. In: *Classical and Quantum Gravity* 30.13 (2013), p. 135009.

[272] Alessandra Buonanno and Thibault Damour. 'Effective one-body approach to general relativistic two-body dynamics'. In: *Physical Review D* 59.8 (1999), p. 084006.

[273] Alessandra Buonanno and Thibault Damour. 'Transition from inspiral to plunge in binary black hole coalescences'. In: *Physical Review D* 62.6 (2000), p. 064015.

[274] Andrea Taracchini et al. 'Effective-one-body model for black-hole binaries with generic mass ratios and spins'. In: *Physical Review D* 89.6 (2014), p. 061502.

[275] Michael Pürrer. 'Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass ratios and spins'. In: *Physical Review D* 93.6 (2016), p. 064041.

[276] Charlie Hoy. 'Accelerating multimodel Bayesian inference, model selection, and systematic studies for gravitational wave astronomy'. In: *Physical Review D* 106.8 (2022), p. 083003.

[277] Charlie Hoy et al. 'Incorporation of model accuracy in gravitational wave Bayesian inference'. In: *Nature Astronomy* 9.8 (2025), pp. 1256–1267.

[278] Mark Hannam et al. 'Simple model of complete precessing black-hole-binary gravitational waveforms'. In: *Physical Review Letters* 113.15 (2014), p. 151101.

[279] Marcus E Lower et al. 'Measuring eccentricity in binary black hole inspirals with gravitational waves'. In: *Physical Review D* 98.8 (2018), p. 083028.

[280] Isobel M Romero-Shaw, Paul D Lasky and Eric Thrane. 'Searching for eccentricity: signatures of dynamical formation in the first gravitational-wave transient catalogue of LIGO and Virgo'. In: *Monthly Notices of the Royal Astronomical Society* 490.4 (2019), pp. 5210–5216.

[281] Larne Pekowsky et al. 'Impact of higher-order modes on the detection of binary black hole coalescences'. In: *Physical Review D* 87.8 (2013), p. 084008.

[282] Duncan A Brown, Prayush Kumar and Alexander H Nitz. 'Template banks to search for low-mass binary black holes in advanced gravitational-wave detectors'. In: *Physical Review D* 87.8 (2013), p. 082004.

[283] Collin Capano, Yi Pan and Alessandra Buonanno. 'Impact of higher harmonics in searching for gravitational waves from nonspinning binary black holes'. In: *Physical Review D* 89.10 (2014), p. 102003.

[284] Sebastian Khan et al. 'Including higher order multipoles in gravitational-wave models for precessing binary black holes'. In: *Physical Review D* 101.2 (2020), p. 024056.

[285] Juan Calderón Bustillo et al. 'Impact of gravitational radiation higher order modes on single aligned-spin gravitational wave searches for binary black holes'. In: *Physical Review D* 93.8 (2016), p. 084019.

[286] Lionel London et al. 'First higher-multipole model of gravitational waves from spinning and coalescing black-hole binaries'. In: *Physical Review Letters* 120.16 (2018), p. 161102.

[287] Samantha A Usman, Joseph C Mills and Stephen Fairhurst. 'Constraining the inclinations of binary mergers from gravitational-wave observations'. In: *The Astrophysical Journal* 877.2 (2019), p. 82.

[288] Geraint Pratten et al. 'Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes'. In: *Physical Review D* 103.10 (2021), p. 104056.

[289] Marta Colleoni et al. 'Fast frequency-domain gravitational waveforms for precessing binaries with a new twist'. In: *Physical Review D* 111 (10 2025), p. 104019.

[290] Maya Fishbach, Daniel E Holz and Will M Farr. 'Does the black hole merger rate evolve with redshift?' In: *The Astrophysical Journal Letters* 863.2 (2018), p. L41.

[291] Ilya Mandel, Will M Farr and Jonathan R Gair. 'Extracting distribution parameters from multiple uncertain observations with selection biases'. In: *Monthly Notices of the Royal Astronomical Society* 486.1 (2019), pp. 1086–1093.

[292] Will M. Farr. 'Accuracy Requirements for Empirically Measured Selection Functions'. In: *Research Notes of the AAS* 3.5 (2019), p. 66.

[293] Richard Abbott et al. 'Population of merging compact binaries inferred using gravitational waves through GWTC-3'. In: *Physical Review X* 13.1 (2023), p. 011048.

[294] Colm Talbot and Eric Thrane. 'Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization'. In: *The Astrophysical Journal* 856.2 (2018), p. 173.

[295] Piero Madau and Mark Dickinson. 'Cosmic star-formation history'. In: *Annual Review of Astronomy and Astrophysics* 52.1 (2014), pp. 415–486.

[296] Tom Callister et al. 'Shouts and murmurs: Combining individual gravitational-wave sources with the stochastic background to measure the history of binary black hole mergers'. In: *The Astrophysical Journal Letters* 896.2 (2020), p. L32.

[297] Thomas A Callister and Will M Farr. 'Parameter-free tour of the binary black hole population'. In: *Physical Review X* 14.2 (2024), p. 021005.

[298] Mick Wright and Martin Hendry. 'Gravelamps: Gravitational Wave Lensing Mass Profile Model Selection'. In: *The Astrophysical Journal* 935.2 (2022), p. 68.

[299] K Haris et al. 'Identifying strongly lensed gravitational wave signals from binary black hole mergers'. In: *arXiv preprint arXiv:1807.07062* (2018).

[300] Justin Janquart et al. 'A fast and precise methodology to search for and analyse strongly lensed gravitational-wave events'. In: *Monthly Notices of the Royal Astronomical Society* 506.4 (2021), pp. 5430–5438.

[301] Xiaoshu Liu, Ignacio Magana Hernandez and Jolien Creighton. 'Identifying strong gravitational-wave lensing during the second observing run of Advanced LIGO and Advanced Virgo'. In: *The Astrophysical Journal* 908.1 (2021), p. 97.

[302] Rico KL Lo and Ignacio Magaña Hernandez. 'Bayesian statistical framework for identifying strongly lensed gravitational-wave signals'. In: *Physical Review D* 107.12 (2023), p. 123015.

[303] Justin Janquart et al. 'The return of GOLUM: improving distributed joint parameter estimation for strongly lensed gravitational waves'. In: *Monthly Notices of the Royal Astronomical Society* 526.2 (2023), pp. 3088–3098.

[304] Ankur Barsode, Srashti Goyal and Parameswaran Ajith. 'Fast and efficient Bayesian method to search for strongly lensed gravitational waves'. In: *The Astrophysical Journal* 980.2 (2025), p. 258.

[305] Yijun Wang et al. 'Identifying type II strongly lensed gravitational-wave images in third-generation gravitational-wave detectors'. In: *Physical Review D* 103.10 (2021), p. 104055.

[306] Aditya Vijaykumar, Ajit Kumar Mehta and Apratim Ganguly. 'Detection and parameter estimation challenges of type-II lensed binary black hole signals'. In: *Physical Review D* 108.4 (2023), p. 043036.

[307] Kelsie Taylor, Derek Davis and Rico KL Lo. 'Phase consistency test to identify type II strongly lensed gravitational-wave signals using a single event'. In: *Physical Review D* 112.2 (2025), p. 024035.

[308] Larry Wasserman. *All of statistics: a concise course in statistical inference*. Vol. 26. Springer, 2004.

[309] Eilam Gross and Ofer Vitells. 'Trial factors for the look elsewhere effect in high energy physics'. In: *The European Physical Journal C* 70.1 (2010), pp. 525–530.

[310] Otto A Hannuksela et al. 'Localizing merging black holes with sub-arcsecond precision using gravitational-wave lensing'. In: *Monthly Notices of the Royal Astronomical Society* 498.3 (2020), pp. 3395–3402.

[311] Eungwang Seo, Tjonnie Guan Feng Li and Martin Anthony Hendry. 'Inferring Properties of Dark Galactic Halos Using Strongly Lensed Gravitational Waves'. In: *The Astrophysical Journal* 966.1 (2024), p. 107.

[312] José María Diego et al. 'Observational signatures of microlensing in gravitational waves at LIGO/Virgo frequencies'. In: *Astronomy & Astrophysics* 627 (2019), A130.

[313] José María Diego. 'Constraining the abundance of primordial black holes with gravitational lensing of gravitational waves at LIGO frequencies'. In: *Physical Review D* 101 (12 2020), p. 123512.

[314] Mark HY Cheung et al. 'Stellar-mass microlensing of gravitational waves'. In: *Monthly Notices of the Royal Astronomical Society* 503.3 (2021), pp. 3326–3336.

[315] Anuj Mishra et al. 'Gravitational lensing of gravitational waves: effect of microlens population in lensing galaxies'. In: *Monthly Notices of the Royal Astronomical Society* 508.4 (2021), pp. 4869–4886.

[316] Ashish Kumar Meena et al. 'Gravitational lensing of gravitational waves: Probability of microlensing in galaxy-scale lens population'. In: *Monthly Notices of the Royal Astronomical Society* 517.1 (2022), pp. 872–884.

[317] Eungwang Seo, Otto A Hannuksela and Tjonnie GF Li. 'Improving Detection of Gravitational-wave Microlensing Using Repeated Signals Induced by Strong Lensing'. In: *The Astrophysical Journal* 932.1 (2022), p. 50.

[318] Simon MC Yeung et al. 'Detectability of microlensed gravitational waves'. In: *Monthly Notices of the Royal Astronomical Society* 526.2 (2023), pp. 2230–2240.

[319] Xikai Shan et al. 'Microlensing sheds light on the detection of strong lensing gravitational waves'. In: *arXiv preprint arXiv:2301.06117* (2023).

[320] Xikai Shan et al. 'Microlensing bias on the detection of strong lensing gravitational wave'. In: *Science China Physics, Mechanics & Astronomy* 67.6 (2024), p. 269511.

[321] Anuj Mishra et al. 'Exploring the impact of microlensing on gravitational wave signals: Biases, population characteristics, and prospects for detection'. In: *Monthly Notices of the Royal Astronomical Society* 531.1 (2024), pp. 764–787.

[322] Juno CL Chan et al. 'Detectability of lensed gravitational waves in matched-filtering searches'. In: *Physical Review D* 111.8 (2025), p. 084019.

[323] Xikai Shan et al. 'Wave effect of gravitational waves intersected with a microlens field: A new algorithm and supplementary study'. In: *Science China Physics, Mechanics & Astronomy* 66.3 (2023), p. 239511.

[324] OA Hannuksela et al. 'Search for gravitational lensing signatures in LIGO-Virgo binary black hole events'. In: *The Astrophysical Journal Letters* 874.1 (2019), p. L2.

[325] Richard Abbott et al. 'Search for lensing signatures in the gravitational-wave observations from the first half of LIGO–Virgo's third observing run'. In: *The Astrophysical Journal* 923.1 (2021), p. 14.

[326] R Abbott et al. 'Search for gravitational-lensing signatures in the full third observing run of the LIGO–Virgo network'. In: *The Astrophysical Journal* 970.2 (2024), p. 191.

[327] Mick Wright et al. 'LensingFlow: An Automated Workflow for Gravitational Wave Lensing Analyses'. In: *RAS Techniques and Instruments* (2026), rzag003.

[328] AG Abac et al. 'GWTC-4.0: Methods for identifying and characterizing gravitational-wave transients'. In: *The Astrophysical Journal Letters* 1004.2 (2026), p. L21.

[329] A. G. Abac, R. Abbott, I. Abouelfettouh et al. 'GW231123: a Binary Black Hole Merger with Total Mass 190-265 $M_\odot$'. In: *The Astrophysical Journal Letters* 993.1 (2025), p. L25.

[330] Sandy L Zabell. 'The rule of succession'. In: *Erkenntnis* 31.2 (1989), pp. 283–321.

[331] Anna Liu and Kyungmin Kim. 'Can we discern millilensed gravitational-wave signals from signals produced by precessing binary black holes with ground-based detectors?' In: *Physical Review D* 110 (12 2024), p. 123008.

[332] Anarya Ray et al. 'GW231123: extreme spins or microglitches?' In: *arXiv preprint arXiv:2510.07228* (2025).

[333] Jonathan E Thompson et al. 'Phenomenological gravitational-wave model for precessing black-hole binaries with higher multipoles and asymmetries'. In: *Physical Review D* 109.6 (2024), p. 063012.

[334] Vijay Varma et al. 'Surrogate models for precessing binary black hole simulations with unequal masses'. In: *Physical Review Research* 1.3 (2019), p. 033015.

[335] Justin Janquart et al. 'Follow-up analyses to the O3 LIGO–Virgo–KAGRA lensing searches'. In: *Monthly Notices of the Royal Astronomical Society* 526.3 (2023), pp. 3832–3860.

[336] S Soni et al. 'LIGO Detector Characterization in the first half of the fourth Observing run'. In: *Classical and Quantum Gravity* 42.8 (2025), p. 085016.

[337] Xikai Shan et al. 'An interference-based method for the detection of strongly lensed gravitational waves'. In: *Nature Astronomy* 9.6 (2025), pp. 916–924.

[338] Benjamin J Owen. 'Search templates for gravitational waves from inspiraling binaries: Choice of template spacing'. In: *Physical Review D* 53.12 (1996), p. 6749.

[339] Philip Relton and Vivien Raymond. 'Parameter estimation bias from overlapping binary black hole events in second generation interferometers'. In: *Physical Review D* 104.8 (2021), p. 084039.

[340] Qian Hu et al. 'GW231123: Overlapping Gravitational Wave Signals?' In: *The Astrophysical Journal* 1003.2 (2026), p. 187.

[341] A Renske AC Wierda et al. 'Beyond the detector horizon: Forecasting gravitational-wave strong lensing'. In: *The Astrophysical Journal* 921.2 (2021), p. 154.

[342] Connor McIsaac et al. 'Search for strongly lensed counterpart images of binary black hole mergers in the first two LIGO observing runs'. In: *Physical Review D* 102.8 (2020), p. 084031.

[343] Srashti Goyal, Shasvath J Kapadia and Parameswaran Ajith. 'Rapid identification of strongly lensed gravitational-wave events with machine learning'. In: *Physical Review D* 104.12 (2021), p. 124057.

[344] Jose María Ezquiaga, Wayne Hu and Rico KL Lo. 'Identifying strongly lensed gravitational waves through their phase consistency'. In: *Physical Review D* 108.10 (2023), p. 103520.

[345] Srashti Goyal et al. 'Rapid method for preliminary identification of subthreshold strongly lensed counterparts to superthreshold gravitational-wave events'. In: *Physical Review D* 109.2 (2024), p. 023028.

[346] Sourabh Magare, Anupreeta More and Sunil Choudhary. 'SLICK: Strong Lensing Identification of Candidates Kindred in gravitational wave data'. In: *Monthly Notices of the Royal Astronomical Society* 535.1 (2024), pp. 990–999.

[347] Aniruddha Chakraborty and Suvodip Mukherjee. 'GLANCE–Gravitational Lensing Authenticator using Non-modelled Cross-correlation Exploration of Gravitational Wave Signals'. In: *Monthly Notices of the Royal Astronomical Society* 532.4 (2024), pp. 4842–4863.

[348] Sudhir Gholap et al. 'χ^2 statistic for the identification of strongly lensed gravitational waves from compact binary coalescences'. In: *Physical Review D* 112.12 (2025), p. 124074.

[349] Vishal Baibhav et al. 'Gravitational-wave detection rates for compact binaries formed in isolation: LIGO/Virgo O3 and beyond'. In: *Physical Review D* 100.6 (2019), p. 064060.

[350] Anuradha Samajdar et al. 'Biases in parameter estimation from overlapping gravitational-wave signals in the third-generation detector era'. In: *Physical Review D* 104.4 (2021), p. 044003.

[351] Francesco Iacovelli et al. 'Forecasting the detection capabilities of third-generation gravitational-wave detectors using GWFAST'. In: *The Astrophysical Journal* 941.2 (2022), p. 208.

[352] Benjamin P Abbott et al. 'Tests of general relativity with GW150914'. In: *Physical Review Letters* 116.22 (2016), p. 221101.

[353] BP Abbott et al. 'Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1'. In: *Physical Review D* 100.10 (2019), p. 104036.

[354] Robert Abbott et al. 'Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog'. In: *Physical review D* 103.12 (2021), p. 122002.

[355] R Abbott et al. 'Tests of general relativity with GWTC-3'. In: *Physical Review D* 112.8 (2025), p. 084080.

[356] Wenwen Zheng et al. 'An improved GPU-based ray-shooting code for gravitational microlensing'. In: *The Astrophysical Journal* 931.2 (2022), p. 114.

[357] Xikai Shan et al. 'Wave effect of gravitational waves intersected with a microlens field II: An adaptive hierarchical tree algorithm and population study'. In: *Science China Physics, Mechanics & Astronomy* 68.1 (2025), p. 219512.

[358] Alex Nitz et al. *gwastro/pycbc: v2.3.3 release of PyCBC*. Version v2.3.3. 2024.

[359] Fei Xu, Jose Maria Ezquiaga and Daniel E. Holz. 'Please Repeat: Strong Lensing of Gravitational Waves as a Probe of Compact Binary and Galaxy Populations'. In: *Astrophys. J.* 929.1 (2022), p. 9.

[360] B. P. Abbott, R. Abbott, T. D. Abbott et al. 'Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo'. In: *Astrophysical Journal Letters* 882.2 (2019), p. L24.

[361] Thomas E Collett. 'The population of galaxy–galaxy strong lenses in forthcoming optical imaging surveys'. In: *The Astrophysical Journal* 811.1 (2015), p. 20.

[362] Mario Spera, Michela Mapelli and Alessandro Bressan. 'The mass spectrum of compact remnants from the PARSEC stellar evolution tracks'. In: *Monthly Notices of the Royal Astronomical Society* 451.4 (2015), pp. 4086–4103.

[363] Georgios Vernardos. 'Microlensing flux ratio predictions for Euclid'. In: *Monthly Notices of the Royal Astronomical Society* 483.4 (2019), pp. 5583–5594.

[364] J. M. Diego et al. 'Microlensing and the type Ia supernova iPTF16geu'. In: *Astronomy & Astrophysics.* 662 (2022), A34.

[365] Benjamin P Abbott et al. 'Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA'. In: *Living reviews in relativity* 23.1 (2020), p. 3.

[366] Alex Nitz et al. *gwastro/pycbc: v2.0.5 release of PyCBC*. Version v2.0.5. 2022.

[367] Justin Janquart et al. 'On the identification of individual gravitational-wave image types of a lensed system using higher-order modes'. In: *The Astrophysical Journal Letters* 923.1 (2021), p. L1.

[368] Aniruddha Chakraborty and Suvodip Mukherjee. 'μ-GLANCE: A Novel Technique to Detect Chromatically and Achromatically Lensed Gravitational-wave Signals'. In: *The Astrophysical Journal* 984.2 (2025), p. 107.

[369] Anuj Mishra and Apratim Ganguly. 'Can eccentric binary black hole signals mimic gravitational-wave microlensing?' In: *Physical Review D* 113.12 (2026), p. 123036.

[370] Stephen Fairhurst et al. 'When will we observe binary black holes precessing?' In: *Physical Review D* 102 (4 2020), p. 041302.

[371] Mark Hannam et al. 'General-relativistic precession in a black-hole binary'. In: *Nature* 610.7933 (2022), pp. 652–655.

[372] AG Abac et al. 'Search for eccentric black hole coalescences during the third observing run of LIGO and Virgo'. In: *The Astrophysical Journal* (2024).

[373] Kyungmin Kim, Eungwang Seo and Chunglee Kim. 'Gravitational lensing aided luminosity distance estimation for compact binary coalescences'. In: *Physical Review D* 109.4 (2024), p. 043017.

[374] Richard Abbott et al. 'GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run'. In: *Physical Review X* 13.4 (2023), p. 041039.

[375] Isobel M Romero-Shaw et al. 'Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue'. In: *Monthly Notices of the Royal Astronomical Society* 499.3 (2020), pp. 3295–3319.

[376] Kwun-Hang Lai et al. 'Discovering intermediate-mass black hole lenses through gravitational wave lensing'. In: *Physical Review D* 98.8 (2018), p. 083005.

[377] Alvin KY Li et al. 'Targeted subthreshold search for strongly lensed gravitational-wave events'. In: *Physical Review D* 107.12 (2023), p. 123014.

[378] Simon Birrer and Adam Amara. 'lenstronomy: Multi-purpose gravitational lens modelling software package'. In: *Physics of the Dark Universe* 22 (2018), pp. 189–201.

[379] Hai Yu, Pengjie Zhang and Fa-Yin Wang. 'Strong lensing as a giant telescope to localize the host galaxy of gravitational wave event'. In: *Monthly Notices of the Royal Astronomical Society* 497.1 (2020), pp. 204–209.

[380] Ewoud Wempe et al. 'On the detection and precise localization of merging black holes events through strong gravitational lensing'. In: *Monthly Notices of the Royal Astronomical Society* 530.3 (2024), pp. 3368–3390.

[381] LSST Science Collaboration et al. 'LSST Science Book, Version 2.0'. In: *arXiv e-prints*, arXiv:0912.0201 (2009).

[382] Mesut Çalışkan et al. 'Lensing or luck? False alarm probabilities for gravitational lensing of gravitational waves'. In: *Physical Review D* 107.6 (2023), p. 063023.

[383] Justin Janquart, Anupreeta More and Chris Van Den Broeck. 'Ordering the confusion: a study of the impact of lens models on gravitational-wave strong lensing detection capabilities'. In: *Monthly Notices of the Royal Astronomical Society* 519.2 (2023), pp. 2046–2059.

[384] Giovanni Tambalo et al. 'Gravitational wave lensing as a probe of halo properties and dark matter'. In: *Physical Review D* 108.10 (2023), p. 103529.

[385] Mesut Çalıșkan et al. 'Observability of lensing of gravitational waves from massive black hole binaries with LISA'. In: *Physical Review D* 107.4 (2023), p. 043029.

[386] Mesut Çalışkan et al. 'Probing wave-optics effects and low-mass dark matter halos with lensing of gravitational waves from massive black holes'. In: *Physical Review D* 108.12 (2023), p. 123543.

[387] Eungwang Seo et al. 'The impact of strong lensing on Hubble constant measurements with gravitational-wave dark sirens'. In: *The Astrophysical Journal Supplement Series* 284.1 (2026), p. 20.

[388] Laura Elina Uronen et al. 'Finding black holes: an unconventional multi-messenger'. In: *Philosophical Transactions A* 383.2294 (2025), p. 20240152.

[389] Suvodip Mukherjee et al. 'Impact of astrophysical binary coalescence time-scales on the rate of lensed gravitational wave events'. In: *Monthly Notices of the Royal Astronomical Society* 506.3 (2021), pp. 3751–3759.

[390] Kenneth C. Wong, Sherry H. Suyu, Geoff C. F. Chen et al. 'H0LiCOW - XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes'. In: *Monthly Notices of the Royal Astronomical Society* 498.1 (2020), pp. 1420–1439.

[391] Kenneth C. Wong, Frédéric Dux, Anowar J. Shajib et al. 'TDCOSMO: XVI. Measurement of the Hubble constant from the lensed quasar WGD 2038−4008'. In: *Astronomy & Astrophysics* 689, A168 (2024), A168.

[392] M. Punturo, M. Abernathy, F. Acernese et al. 'The Einstein Telescope: a third-generation gravitational wave observatory'. In: *Classical and Quantum Gravity* 27.19, 194002 (2010), p. 194002.

[393] Adrian Abac et al. 'The science of the Einstein telescope'. In: *Journal of Cosmology and Astroparticle Physics* 2026.03 (2026), p. 081.

[394] Rachel Gray, Chris Messenger and John Veitch. 'A pixelated approach to galaxy catalogue incompleteness: improving the dark siren measurement of the Hubble constant'. In: *Monthly Notices of the Royal Astronomical Society* 512.1 (2022), pp. 1127–1140.

[395] P. Fosalba et al. 'The MICE Grand Challenge light-cone simulation – III. Galaxy lensing mocks from all-sky lensing maps'. In: *Monthly Notices of the Royal Astronomical Society* 447.2 (2014), pp. 1319–1332.

[396] J. Carretero et al. 'An algorithm to build mock galaxy catalogues using MICE simulations'. In: *Monthly Notices of the Royal Astronomical Society* 447.1 (2014), pp. 646–670.

[397] K. Hoffmann et al. 'Measuring the growth of matter fluctuations with third-order galaxy correlations'. In: *Monthly Notices of the Royal Astronomical Society* 447.2 (2014), pp. 1724–1745.

[398] P Fosalba et al. 'The MICE grand challenge lightcone simulation–I. Dark matter clustering'. In: *Monthly Notices of the Royal Astronomical Society* 448.4 (2015), pp. 2987–3000.

[399] Martin Crocce et al. 'The MICE Grand Challenge lightcone simulation–II. Halo and galaxy catalogues'. In: *Monthly Notices of the Royal Astronomical Society* 453.2 (2015), pp. 1513–1530.

[400] Michael R Blanton et al. 'The galaxy luminosity function and luminosity density at redshift z= 0.1'. In: *The Astrophysical Journal* 592.2 (2003), p. 819.

[401] J Carretero et al. 'An algorithm to build mock galaxy catalogues using MICE simulations'. In: *Monthly Notices of the Royal Astronomical Society* 447.1 (2015), pp. 646–670.

[402] Adam S Bolton et al. 'The Sloan Lens ACS Survey. I. A large spectroscopically selected sample of massive early-type lens galaxies'. In: *The Astrophysical Journal* 638.2 (2006), p. 703.

[403] Adam S Bolton et al. 'The sloan lens ACS survey. V. The full ACS strong-lens sample'. In: *The Astrophysical Journal* 682.2 (2008), p. 964.

[404] L Vujeva et al. 'lenscat: a public and community-contributed catalogue of known strong gravitational lenses'. In: *Philosophical Transactions A* 383.2294 (2025), p. 20240168.

[405] Aditya Vijaykumar et al. 'Inferring host-galaxy properties of LIGO–Virgo–KAGRA's black holes'. In: *The Astrophysical Journal* 972.2 (2024), p. 157.

[406] Zhuotao Li, Rachel Gray and Ik Siong Heng. 'Using Gravitational-wave Dark Sirens to Choose between Host Galaxy Weighting Models'. In: *The Astrophysical Journal* 1003.2 (2026), p. 176.

[407] Alexandra G Hanselman et al. 'Gravitational-wave dark siren cosmology systematics from galaxy weighting'. In: *The Astrophysical Journal* 979.1 (2025), p. 9.

[408] Gabriele Perna, Simone Mastrogiovanni and Angelo Ricciardone. 'Investigating the impact of galaxies' compact binary hosting probability for gravitational wave cosmology'. In: *Astronomy & Astrophysics* 698 (2025), A128.

[409] Maciej Bilicki et al. 'WISE$\times$ SuperCOSMOS photometric redshift catalog: 20 million galaxies over 3π steradians'. In: *The Astrophysical Journal Supplement Series* 225.1 (2016), p. 5.

[410] Mick Wright, Justin Janquart and Martin Hendry. 'Determination of lens mass density profile from strongly lensed gravitational-wave signals'. In: *The Astrophysical Journal* 959.2 (2023), p. 70.

[411] K. M. Górski et al. 'HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere'. In: *The Astrophysical Journal* 622.2 (2005), p. 759.

[412] Andrea Zonca et al. 'healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python'. In: *Journal of Open Source Software* 4.35 (2019), p. 1298.

[413] Alvin KY Li et al. 'TESLA-X: an effective method to search for subthreshold lensed gravitational waves with a targeted population model'. In: *Monthly Notices of the Royal Astronomical Society* 542.2 (2025), pp. 998–1010.

[414] EE Falco, MV Gorenstein and II Shapiro. 'On model-dependent bounds on H (0) from gravitational images Application of Q0957+ 561A, B'. In: *The Astrophysical Journal, Part 2-Letters to the Editor* 289 (1985), pp. L1–L4.

[415] Peter Schneider and Dominique Sluse. 'Mass-sheet degeneracy, power-law models and external convergence: Impact on the determination of the Hubble constant from gravitational lensing'. In: *Astronomy & Astrophysics* 559 (2013), A37.

[416] Paolo Cremonese, Jose María Ezquiaga and Vincenzo Salzano. 'Breaking the mass-sheet degeneracy with gravitational wave interference in lensed events'. In: *Physical Review D* 104.2 (2021), p. 023503.

[417] Anson Chen and Jun Zhang. 'Forecasting Constraints on Cosmology and Modified Gravitational-wave Propagation by Strongly Lensed Gravitational Waves Associating with Galaxy Surveys'. In: *arXiv preprint arXiv:2601.21820* (2026).

[418] Pau Amaro-Seoane et al. 'Laser interferometer space antenna'. In: *arXiv preprint arXiv:1702.00786* (2017).

[419] Stanislav Babak, Martin Hewitson and Antoine Petiteau. 'LISA sensitivity and SNR calculations'. In: *arXiv preprint arXiv:2108.01167* (2021).

[420] Martina Toscani et al. 'Strongly lensed extreme mass-ratio inspirals'. In: *Physical Review D* 109.6 (2024), p. 063505.

[421] Souvik Jana et al. 'Cosmography using strongly lensed gravitational waves from binary black holes'. In: *Physical Review Letters* 130.26 (2023), p. 261401.

[422] Joachim Wambsganss. 'Discovering Galactic planets by gravitational microlensing: magnification patterns and light curves'. In: *Monthly Notices of the Royal Astronomical Society* 284.1 (1997), pp. 172–188.

[423] Joachim Wambsganss. 'Gravitational microlensing'. In: *Gravitational Lensing: Strong, Weak and Micro*. Springer, 2006, pp. 453–540.

[424] Misao Sasaki et al. 'Primordial black holes—perspectives in gravitational wave astronomy'. In: *Classical and Quantum Gravity* 35.6 (2018), p. 063001.

[425] Bernard Carr and Florian Kühnel. 'Primordial black holes as dark matter: recent developments'. In: *Annual Review of Nuclear and Particle Science* 70.1 (2020), pp. 355–394.

[426] Juan Urrutia and Ville Vaskonen. 'Lensing of gravitational waves as a probe of compact dark matter'. In: *Monthly Notices of the Royal Astronomical Society* 509.1 (2022), pp. 1358–1365.

[427] Joseph Gais et al. 'Inferring the intermediate-mass black hole number density from gravitational-wave lensing statistics'. In: *The Astrophysical Journal Letters* 932.1 (2022), p. L4.

[428] Soummyadip Basak et al. 'Constraints on compact dark matter from gravitational wave microlensing'. In: *The Astrophysical Journal Letters* 926.2 (2022), p. L28.

[429] Malcolm Fairbairn, Juan Urrutia and Ville Vaskonen. 'Microlensing of gravitational waves by dark matter structures'. In: *Journal of Cosmology and Astroparticle Physics* 2023.07 (2023), p. 007.

[430] Gilles Chabrier. 'Galactic stellar and substellar initial mass function'. In: *Publications of the Astronomical Society of the Pacific* 115.809 (2003), pp. 763–795.

[431] Jose Maria Ezquiaga et al. 'Modified gravitational wave propagation with higher modes and its degeneracies with lensing'. In: *Journal of Cosmology and Astroparticle Physics* 2022.08 (2022), p. 016.

[432] Edwin L Turner, Jeremiah P Ostriker and J Richard Gott III. 'The statistics of gravitational lenses-The distributions of image angular separations and lens redshifts'. In: *The Astrophysical Journal, Part 1* 284 (1984), pp. 1–22.

[433] Zu-Cheng Chen, Chen Yuan and Qing-Guo Huang. 'Confronting the primordial black hole scenario with the gravitational-wave events detected by LIGO-Virgo'. In: *Physics Letters B* 829 (2022), p. 137040.

[434] Lang Liu et al. 'Constraining the merger history of primordial-black-hole binaries from GWTC-3'. In: *Physical Review D* 107.6 (2023), p. 063035.

[435] Eleni Bagui et al. 'Primordial black holes and their gravitational-wave signatures'. In: *Living Reviews in Relativity* 28.1 (2025), p. 1.

[436] David Reitze et al. 'Cosmic explorer: the US contribution to gravitational-wave astronomy beyond LIGO'. In: *arXiv preprint arXiv:1907.04833* (2019).

[437] Michel Chevallier and David Polarski. 'Accelerating universes with scaling dark matter'. In: *International Journal of Modern Physics D* 10.02 (2001), pp. 213–223.

[438] Eric V Linder. 'Exploring the expansion history of the universe'. In: *Physical Review Letters* 90.9 (2003), p. 091301.

[439] Suvodip Mukherjee, Matthew SP Meinema and Joseph Silk. 'Prospects of discovering subsolar primordial black holes using the stochastic gravitational wave background from third-generation detectors'. In: *Monthly Notices of the Royal Astronomical Society* 510.4 (2022), pp. 6218–6224.