
Gravitational-wave source dynamics and population inference

by

Matthew Mould



UNIVERSITY OF
BIRMINGHAM

A thesis submitted to the University of Birmingham for the degree of
DOCTOR OF PHILOSOPHY

Institute for Gravitational Wave Astronomy
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham

May 2023

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

Gravitational-wave astronomy has yielded an unprecedented treasure trove of scientific discoveries in recent years, from the first detection of a black-hole merger and strong-field tests of gravity, to the observation of colliding neutron stars and corresponding global multi-messenger followup. The future yet holds great promise. As the observational catalogue continues to grow in the coming years with improvements in sensitivity and the introduction of new detectors, our available inferential power in the gravitational landscape of the universe will grow in tandem. In this Thesis, we leverage this increasingly informative dataset to investigate the evolution of binary black holes from formation to merger, deriving new results on the influence of astrophysics and relativity on binary inspirals and, conversely, revealing constraints from gravitational-wave observables on the past lives of observed sources. We place a particular focus on Bayesian methods in analysing gravitational-wave catalogues and use targeted models, simulation-based inference, and deep learning to measure the properties of the astrophysical population of merging binary black holes. Our work highlights the crucial interplay between compact-binary formation, relativistic dynamics, and statistical inference. The tools we develop will enable deeper astrophysical insights as gravitational-wave astronomy enters into the big-data era.

*“Man is not alone on this planet.
He is part of a community,
upon which he depends absolutely.”*

Daniel Quinn, *Ishmael*

Acknowledgments

Thanks

I owe many thanks to everyone I interacted with during the course of my PhD. Thanks to all of my officemates and neighbours in the ASR and SASP/SSE groups for their unrivalled company during the good/bad/lockdown/code-debugging/paper-writing/lockdown-again/late-night times. Thanks to all the members, past and present, of the gravitational-wave group at the University of Birmingham for the interesting science and beyond. I am lucky to have collaborated with and been mentored by many excellent researchers during the course of my projects and publications.

Thank you to my mum, my brother, my sister, my family — there's no better feeling than coming home. Thanks to four years of housemates who each managed to put up with me in my home away from home. To Katherine: thank you for being you, thank you for the trips to the kitty cafe, thank you for everything and more.

I am most grateful to my supervisor, Davide Gerosa, for his insights, patience, and dedication. I feel very honoured to be the first member of his wonderful and rapidly-growing research group (now in the double digits!), and am amazed that he still makes time for each of us — but that's just the kind of mentor he is. I always look forward to our next project, so here's to future whiteboard scribble sessions, or the next hike, or my next parkrun victory!

Funding

My PhD was supported by: a studentship in the School of Physics and Astronomy at the University of Birmingham; the Science and Technology Facilities Council [grant number 2322486]; the Engineering and Physical Sciences Research Council [grant number 2287305]; European Union's H2020 ERC Starting Grant No. 945155-GWmining; Cariplo Foundation Grant No. 2021-0555; the ICSC National Center on High-Performance Computing, Big Data, and Quantum Computing funded by NextGenerationEU; Royal Society Grant No. RGS-R2-202004.

Computational work was performed using: University of Birmingham's BlueBEAR HPC service; the Baskerville Tier 2 HPC service funded by the EPSRC and UKRI through the World Class Labs scheme (EP/T022221/1) and the Digital Research Infrastructure programme (EP/W032244/1) and operated by Advanced Research Computing at the University of Birmingham; the Athena cluster at HPC Midlands+ funded by EPSRC Grant No. EP/P020232/1.

I received travel support from: the COST Action GWverse CA16104; the COST Action G2net CA17137; Project HPC-EUROPA3 (INFRAIA-2016-1-730897), with the support of the EC Research Innovation Action under the H2020 Programme, with acknowledgements to the University of Milan-Bicocca and the computer resources and technical support provided by CINECA; the Simons foundation; the LSSTC Enabling Science program; the Royal Astronomical Society.

Contents

1	Introduction	1
	Summary and contributions	1
1.1	Gravitational-wave astronomy	2
1.2	Gravity and compact binaries	5
1.2.1	General relativity	5
1.2.2	Perturbed spacetime	12
1.2.3	Gravitational-wave sources	14
1.3	Multi-timescale binary dynamics	18
1.3.1	Post-Newtonian primer	18
1.3.2	Timescale hierarchy	20
1.3.3	Spin precession	22
1.3.4	Inspirals	27
1.4	Detection and inference	30
1.4.1	Searching for signals	32
1.4.2	Bayesian parameter estimation	36
1.4.3	Hierarchical population models	39
1.5	Main findings	49
2	Endpoint of the up–down spin-precession instability	53
	Abstract	54

Summary and contributions	54
2.1 Introduction	54
2.2 Instability threshold	58
2.2.1 Binary black-hole spins as harmonic oscillators	58
2.2.2 Stable or unstable?	60
2.2.3 Numerical verification of the instability	65
2.3 Resonant configurations	69
2.3.1 Locating resonances	69
2.3.2 Number of resonances	71
2.3.3 Resonant evolution	74
2.3.4 Asymptotic resonances	74
2.4 Up–down endpoint	77
2.4.1 Instability limit	77
2.4.2 Stability-to-instability transition	81
2.4.3 A simple astrophysical population	83
2.5 Conclusions	89
3 Binary black holes time travelling back to the future	91
Abstract	92
Summary and contributions	92
3.1 Introduction	92
3.2 Time travel for population inference	94
3.2.1 Modelling and inference	94
3.2.2 Spin propagation	96
3.3 Gravitational-wave observations at past time infinity	99

3.4	Conclusions	104
4	Which black hole formed first?	107
	Abstract	108
	Summary and contributions	108
4.1	Introduction	108
4.2	Modelling mass-ratio reversal	110
	4.2.1 Catalogue and statistical tools	110
	4.2.2 Astrophysically-motivated spin models	110
4.3	Measuring mass-ratio reversal	112
	4.3.1 Nonidentical spins	112
	4.3.2 Zero-spin peaks	119
4.4	Gravitational-wave observations against astrophysical predictions	125
	4.4.1 To reverse or not to reverse	125
	4.4.2 Modelling and astrophysical uncertainties	128
4.5	Conclusions	130
5	Simulation-based deep learning of gravitational-wave populations	133
	Abstract	134
	Summary and contributions	134
5.1	Introduction	134
5.2	Hierarchical-merger simulations	138
	5.2.1 Simulation design	141
	5.2.2 First-generation black holes	141
	5.2.3 Repeated mergers	142
	5.2.4 Cosmic placement	144

5.2.5	Resulting populations	145
5.3	Enhancing population inference with deep learning	150
5.3.1	Hierarchical Bayesian inference	150
5.3.2	Population model	153
5.3.3	Selection function	160
5.3.4	Merger-generation fractions	166
5.4	Validation with mock catalogues	171
5.5	Hierarchical-merger inference with gravitational-wave observations . . .	176
5.5.1	Host escape speeds	178
5.5.2	Mass distribution	180
5.5.3	Spin distribution	185
5.5.4	Merger generations	191
5.6	Conclusions	194
6	Conclusions	197
	Bibliography	200

List of Figures

1.1	Non-inertial reference frame to study binary black-hole spins	24
2.1	The four aligned-spin configurations for binary black holes	55
2.2	Spin evolution of a single unstable up–down binary	57
2.3	Oscillation frequencies of perturbed aligned-spin binary black holes . . .	63
2.4	Up–down instability thresholds in the post-Newtonian regime	64
2.5	Evolutions of four aligned-spin binary black holes	66
2.6	Response of up–down sources to increasing perturbations	68
2.7	Spin-precession phase space bounded by resonances	73
2.8	Evolution of binary black holes in resonant spin configurations	78
2.9	Representative distributions of the analytic up–down endpoint	80
2.10	Classification of unstable up–down sources over a population	82
2.11	Effective-spin distribution for aligned-spin binaries	85
2.12	Evolution of tilt angles starting from nearly aligned spins	87
2.13	Distribution of stable and unstable up–down populations	88
3.1	Back-propagating gravitational-wave data to past time infinity	98
3.2	Gravitational-wave population parameters inferred at past time infinity	100
3.3	Gravitational-wave spin populations inferred at past time infinity . . .	101
3.4	Posterior fraction of detectable gravitational-wave sources	103

4.1	Population parameters of nonidentical black-hole spins	113
4.2	Population distributions of nonidentical black-hole spins	114
4.3	Probability that primary black holes spin more rapidly	116
4.4	Probabilities that black holes have negligible spins	118
4.5	Population parameters of nonidentical and nonspinning black holes	120
4.6	Population distributions of nonidentical and nonspinning black holes	121
4.7	Fraction of mass-ratio reversal binaries	124
5.1	Simulation-based deep-learning framework for population inference	139
5.2	Fractions of heirarchical merger generations	147
5.3	Example simulated hierarchical black-hole merger populations	148
5.4	Accuracy of the population-model neural network	157
5.5	Example predictions of the population-model neural network	159
5.6	Accuracy of the selection-function neural network	163
5.7	Example predictions of the selection-function neural network	165
5.8	Accuracy of the branching-fraction neural network	168
5.9	Example predictions of the branching-fraction neural network	170
5.10	Population inference on mock gravitational-wave catalogues	173
5.11	Population inference on mock gravitational-wave catalogues	174
5.12	Posterior of the population-model parameters	177
5.13	Posterior populations of escape speeds and gravitational kicks	179
5.14	Posterior populations of binary black-hole masses	183
5.15	Posterior populations of first-generation binary black-hole masses	184
5.16	Posterior populations of binary black-hole spins	187
5.17	Posterior populations of mass-dependent precessing spins	190
5.18	Posterior distributions of merger-generation branching fractions	192

List of Tables

1.1	Classification of gravitational-wave triggers	35
5.1	Population and source parameters of the hierarchical-merger model . . .	140
5.2	Architecture of the population-model neural network	155
5.3	Architecture of the selection-function neural network	162
5.4	Architecture of the branching-fraction neural network	167

List of abbreviations

1g	First generation
2g	Second generation
AGN	Active galactic nucleus
BEAR	Birmingham Environment for Academic Research
BH	Black hole
CE	Cosmic Explorer
DNN	Deep neural network
EHT	Event Horizon Telescope
ET	Einstein Telescope
FAR	False-alarm rate
GC	Globular cluster
GPR	Gaussian process regression
GR	General relativity
GW	Gravitational wave
IMF	Initial mass function
ISJ	Improved Sheather–Jones
KAGRA	Kamioka Gravitational Wave Detector
KDE	Kernel density estimate
LIGO	Laser Interferometer Gravitational-Wave Observatory
LISA	Laser Interferometer Space Antenna

LVC	LIGO and Virgo Collaboration
MAE	Mean absolute error
MRR	Mass-ratio reversal
MSE	Mean squared error
NANOGrav	North American Nanohertz Observatory for Gravitational Waves
NR	Numerical relativity
NS	Neutron star
NSC	Nuclear star cluster
O1	First LVC observing run
O2	Second LVC observing run
O3	Third LVC observing run
O3a	First half of the third LVC observing run
PE	Parameter estimation
PISN	Pair-instability supernova
PN	Post Newtonian
PPD	Posterior population distribution
PTA	Pulsar timing array
ReLU	Rectified linear unit
RReLU	Randomized leaky rectified linear unit
SNR	Signal-to-noise ratio
SR	Special relativity

Introduction

Summary and contributions

In this Chapter I introduce the background of and give context for the later Chapters that present my results. In Sec. 1.1 I summarize the status of gravitational-wave (GW) astronomy as an observational field and highlight what may lie ahead in the future. In Sec. 1.2 I state key theoretical results and evidence for general relativity (GR) as an accurate theory of gravity (1.2.1), in particular the prediction of gravitational radiation (1.2.2), and discuss the formation and observation of GW sources (1.2.3). Methods for studying the relativistic dynamics of binary black holes (BHs) are presented in Sec. 1.3, using approximations of GR (1.3.1) and multi-timescale techniques (1.3.2) to solve for the evolution of their spins (1.3.3) and inspirals (1.3.4). Section 1.4 describes the use of Bayesian methods in GW astronomy to find signals in noisy data (1.4.1), characterize the properties of GW events (1.4.2), and infer their astrophysical populations (1.4.3). Finally, Sec. 1.5 summarizes the work that constitutes the remaining Chapters of this Thesis, as well as highlighting some work that is not included.

This Chapter contains no original material.

1.1 Gravitational-wave astronomy

GW astronomy is one of the most exciting burgeoning fields in the modern era of physics and astronomy, allowing us to observe the previously unseen gravitational landscape of the universe. The theoretical base was founded following early suggestive works [1, 2] and, of course, the development of GR in the early 20th century [3, 4, 5, 6]. What followed is a century-long storied history [7, 8, 9, 10] rich in controversy and uncertainty due to a lack of understanding in the physical nature of GW propagation [11, 12, 13] and spurious claims of detection [14, 15, 16, 17] that were not independently reproducible [18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28] (see Refs. [29, 30, 31] for interesting historical perspectives). At the time of their conception even Einstein himself doubted their existence [7, 9, 12]. This mindset is perhaps best summarized by the statement of Eddington [7], that

“gravitational waves propagate at the speed of thought.”

It was not until the middle part of the century that the physical influence of GWs began to become more widely accepted [32, 33, 34], then leading to the pioneering development of GW detection methods [35, 36]. At this time there was optimism in the field [29]:

“In this century ‘Astronomy’ has become radio astronomy, microwave astronomy, infrared astronomy, ultraviolet, x-ray, and gamma-ray astronomy, (and optical astronomy!), and cosmic-ray astronomy in its own right... With any luck we will have a gravitational-radiation astronomy too — the prospect is dazzling.”

Over 40 years later and a century after the founding of GR this prospect was first

realized and it has indeed been dazzling. The observational capabilities of the Laser Interferometer Gravitational-Wave Observatory (LIGO) [37] and the Virgo detector [38] reached fruition, resulting in an avalanche of ground-breaking discoveries over the past five to ten years:

- The unprecedented first observation of GWs from a binary BH merger, GW150914 [39, 40, 41, 42], and hence the first direct detection of BHs and strong-field tests of GR [43].
- The first GW observation of a binary neutron star (NS) merger, GW170817 [44, 45, 46], made jointly with electromagnetic counterparts [47, 48, 49, 50, 51], beckoning a new era of multi-messenger astronomy [47], nuclear astrophysics [52, 53], and cosmology [54].
- The first observations of GWs from NS–BH binaries [55] — the last remaining class expected of merging compact objects.
- The detection of binary mergers whose component masses challenge the current understanding of compact-binary formation [56, 57, 58, 59].

Such observations constitute growing catalogues of GW events from the LIGO and Virgo Collaboration (LVC) [60, 61, 62, 63, 64], as well as from independent research groups [65, 66, 67, 68, 69, 70, 71] — crucially enabled by open data [72, 73, 74, 75]. These large sets of GW signals provide the statistical power to make constraints beyond single-event detection and characterization, including

- tests of GR in the strong-field regime [76, 77, 78],
- measurements of cosmological parameters [79, 80],

- searches for gravitational lensing of GWs [81, 82], and
- population studies of merging BHs and NSs [83, 84, 85] — the central topic of this Thesis.

There still remain treasure troves to be unearthed as the field continues to mature in the coming years and decades. The mergers so far observed of compact-binary coalescences produce transient signals. Another (as-yet undetected) source of GW bursts are from gravitational collapse, such as in core-collapse supernovae [86, 87, 88], which could yield counterpart signals and a direct observation of a natal BH or NS [89]. On the other hand, continuous GWs present a distinct signal morphology and are persistently present in the detector data stream. They can be produced by spinning NSs in X-ray binaries [90, 91, 92, 93, 94, 95], isolated pulsars [96, 97, 98], supernova remnants [99, 100], or the galactic center [101] (see also [102]). The cumulative background signal of individually unresolved sources also promises astrophysical and cosmological insights [103, 104].

While current detectors are sensitive to GW frequencies 10^1 – 10^3 Hz [37, 38] and open the window onto NS and stellar-mass BH mergers, the future spaced-based Laser Interferometer Space Antenna (LISA) mission [105, 106, 107] will be sensitive in the lower frequency range 10^{-5} – 10^0 Hz and may observe supermassive BH mergers [108, 109, 110, 111, 112, 113], extreme mass-ratio inspirals [114, 115], galactic compact binaries [116, 117, 118, 119, 120], and multi-band stellar-mass inspirals [121, 122, 123, 124]. Proposed observatories would fill in the deci-Hertz band gap between LISA and the current ground-based network, providing increased coverage for the very heaviest sources currently observable [125, 126, 127, 128].

An alternative technique to the laser [129, 130] and time-delay [131] interferometry

methods employed by these observatories is the use of pulsar timing arrays (PTAs) [132, 133, 134, 135, 136, 137, 138]; correlated deviations in the precise timing measurements of millisecond pulsars are used to infer the presence of GWs. PTAs are sensitive to a low-frequency astrophysical background of GWs produced by massive BH mergers across the universe [139]. The first positive result was recently reported by the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) collaboration [136] and further evidence will likely mount in the near future [140].

Upgrades to the current ground-based GW detector network — including the introduction of the Kamioka Gravitational Wave Detector (KAGRA) [141, 142, 143, 144] — will increase its sensitive horizon [145]. The reach of GW observatories will be pushed even further into the high-redshift universe by the so-called third-generation of detectors [146], namely the Einstein Telescope (ET) [147, 148, 149] and Cosmic Explorer (CE) [150, 151]. Such improvements will lead to orders of magnitudes increase in the detection rate of GW signals [145, 152].

GW astronomy is, as it stands, hugely fruitful, and the future ahead is promising.

1.2 Gravity and compact binaries

To provide a theoretical grounding in GR and its astrophysical implications I will summarize some of the salient points that are relevant to the study of GWs and binary BHs.

1.2.1 General relativity

GR is among the most successful theories in modern physics. Part of Einstein's effectiveness, besides his intuitive ability to use thought experiments and build up

a new geometric description of gravity to overthrow the centuries-long held Newtonian theory, was the production of concrete testable predictions of his theory [153]. Early tests of GR explained the anomalous precession of Mercury's orbit [154] and predicted the solar deflection of light [154, 155, 156, 157, 158]. Furthermore, it was demonstrated that Einstein's equations allowed non-stationary solutions on cosmological scales [159, 160, 161, 162, 163, 164, 165, 166], which agreed with the universal expansion identified by Lemaître and Hubble [162, 167]. Time reversal of these solutions imply a past singularity in finite time [168], forming the basis of the well-known 'Big Bang' cosmological model [169].

Later tests of the last century include the Shapiro time delay [170, 171] and the indirect measurement of gravitational radiation in the Hulse-Taylor binary pulsar [172, 173]. More modern tests of GR phenomena include: the measurement of relativistic parameters by lunar laser ranging [174, 175] and of frame dragging by the Gravity Probe B satellite [176]; the Event Horizon Telescope (EHT) radio observations of supermassive BH 'shadows' in the core of the galaxy M87 [177, 178, 179, 180, 181, 182] and in the Milky Way galactic centre [183, 184, 185, 186, 187]; and, of course, the direct detection of GWs by the LIGO and Virgo collaborations [39, 64].

It is expected for several reasons, however, that GR is not the final theory of gravity [188]. Most importantly, the universe is found to obey quantized physical laws at the smallest scales but GR is a classical theory yet to be reconciled with quantum field theory [5]. GR requires that energy and momentum content are exactly known at any location in spacetime, whereas the uncertainty principle dictates that these quantities cannot be simultaneously determined. If they were to be united, the result would be a complete quantum theory of gravity [189, 190, 191, 192, 193].

Nevertheless, GR has passed all experimental tests so far and Einstein has not yet

been proven wrong.

But before delving deeper into Einstein's theory, I first very briefly visit Newton's. A massive source of gravity with mass density ρ is described in terms of a scalar field ϕ that satisfies the second-order partial differential equation

$$\nabla^2\phi = 4\pi G\rho, \quad (1.1)$$

where G is the gravitational constant that fundamentally couples the strength of gravity to mass. From ϕ one can compute all measurable quantities of interest, such as forces and motions via the change in the position of \mathbf{x} of a particle over absolute time t :

$$\frac{d^2\mathbf{x}}{dt^2} = -\nabla\phi. \quad (1.2)$$

Newtonian gravity, however, is not consistent with the framework of special relativity (SR) and its many important results, including the principle of relativity in inertial reference frames, the constancy of the speed of light, and the equivalency of energy and matter [194, 195, 196]. The three-dimensional equation of motion above does not fit into the four-dimensional spacetime of SR. The Newtonian field equation (1.1) implies that the gravitational influence of massive bodies propagates instantaneously, in clear violation of relativistic principles.

Einstein realized a generalized relativistic theory was needed to unite gravity with the axiomatic basis of relativity [156]. The guiding principles in Einstein's search were:

1. The matter density source should be replaced by the energy-momentum content of a continuous matter distribution;
2. Following the equivalence principle [155] — the complete physical equivalence of

inertial and gravitational phenomena — the new gravitational potential should describe the geometric properties of spacetime;

3. The governing equations should be second-order in space as in the Newtonian theory, but also in time;
4. The relativistic generalization should reduce to the Newtonian description in the weak-gravity limit; and
5. The theory must be consistent with all other known natural laws (it is this last assertion that is contested by modern-day quantum theory).

The result is a geometric theory of four-dimensional spacetime manifolds equipped with a Lorentzian metric tensor. In the following, I will use Greek indices to denote any of the four spacetime coordinates $\{0, 1, 2, 3\}$ of vector and tensor quantities, while Latin indices will restrict to the spatial sectors $\{1, 2, 3\}$. I will use the $(-, +, +, +)$ metric signature.

Einstein's endeavours led to the famous field equations of GR [3, 4]:

$$R_{\mu\nu} - \frac{1}{2}R g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}, \quad (1.3)$$

where terms are written in component form and c is the speed of light that fundamentally limits the propagation of any interaction. The central object of study is a Lorentzian metric tensor g and it is found by solving the field equations (1.3). The metric determines lengths and all other geometric quantities on the spacetime. The infinitesimal line element (or spacetime separation) on the manifold with respect to the coordinates x^μ is given by $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$ and is often itself referred to as the 'metric'.

The left-hand side of the Einstein field equations (1.3) describes the geometrical curvature of spacetime with the Ricci tensor $R_{\mu\nu} = g^{\alpha\beta} R_{\alpha\mu\beta\nu}$ and the Ricci scalar $R = g^{\alpha\beta} R_{\alpha\beta}$, where $R_{\alpha\mu\beta\nu} = g_{\alpha\gamma} R_{\mu\beta\nu}^{\gamma}$ is the Riemann curvature tensor, given by

$$R_{\mu\beta\nu}^{\gamma} = \partial_{\beta}\Gamma_{\nu\mu}^{\gamma} - \partial_{\nu}\Gamma_{\beta\mu}^{\gamma} + \Gamma_{\beta\delta}^{\gamma}\Gamma_{\nu\mu}^{\delta} - \Gamma_{\nu\delta}^{\gamma}\Gamma_{\beta\mu}^{\delta}. \quad (1.4)$$

Through the Christoffel symbols

$$\Gamma_{\mu\nu}^{\alpha} = \frac{1}{2}g^{\alpha\beta}(\partial_{\mu}g_{\nu\beta} + \partial_{\nu}g_{\mu\beta} - \partial_{\beta}g_{\mu\nu}), \quad (1.5)$$

all these quantities may be written entirely in terms of the metric tensor $g_{\mu\nu}$ and its first and second derivatives (there are many GR textbooks where these results can be found [197, 198, 199, 200, 201, 202, 203, 204, 205, 206], while other textbooks will give a more mathematical grounding in the required differential geometry [207, 208, 209]). The right-hand side describes with $T_{\mu\nu}$ the energy and momentum present in the spacetime that source the gravitational field.

With the metric in hand, one can compute (again in coordinates) the path x^{μ} of any freely falling particle with the geodesic equation

$$\frac{d^2x^{\mu}}{ds^2} = -\Gamma_{\alpha\beta}^{\mu} \frac{dx^{\alpha}}{ds} \frac{dx^{\beta}}{ds}, \quad (1.6)$$

where $\Gamma_{\alpha\beta}^{\mu}$ are the Christoffel symbols and s is a scalar parametrization (often the proper time) of the motion.

Comparing the results of Newton and Einstein, one can interpret Eqs. (1.3, 1.6) as the relativistic analogues of the Newtonian equations of motion (1.1, 1.2). The gravitational potential ϕ is replaced with the metric g and the mass density ρ is replaced with

the stress–energy T . In Newtonian gravity, Eq. (1.2) implies that particles of mass m move due to the influence of a gravitational force $\mathbf{F} = -m\nabla\phi$. In Einsteinian gravity, Eq. (1.6) implies that motions are determined entirely by the geometric properties of the spacetime manifold according to the metric g . This is the significant conclusion that gravity reveals itself as curved spacetime [210]:

“The gravitational field influences and even determines the metrical laws of the spacetime continuum. If the laws of configuration of ideal rigid bodies are to be expressed geometrically, then in the presence of a gravitational field the geometry is not Euclidean.”

Altogether, these expressions detail the influence of the metric and mass–energy content on each other; according to a phrase attributed to Wheeler:

“Spacetime tells matter how to move; matter tells spacetime how to curve.”

An immediate consequence of the Einstein field equations is the non-linear nature of gravity; gravitational energy can self gravitate [211]. This is the significant feature of GR that makes finding solutions to the Einstein field equations so difficult. Eq. (1.3) may be rewritten in the trace-reversed form $R_{\mu\nu} = 8\pi Gc^{-4}(T_{\mu\nu} - Tg_{\mu\nu}/2)$. With this form of the Einstein equations one may note that vacuum solutions $T_{\mu\nu} = 0$ satisfy $R_{\mu\nu} = 0$. Important known vacuum solutions include:

- The Minkowski metric [212],

$$ds^2 = \eta_{\mu\nu}dx^\mu dx^\nu = -c^2dt^2 + dx^2 + dy^2 + dz^2, \quad (1.7)$$

where (t, x, y, z) are the usual time coordinate plus Cartesian coordinates. This is the flat metric on which SR is constructed and it is the spacetime extension of

three-dimensional Euclidean space.

- The Schwarzschild metric [213, 214], which in Schwarzschild coordinates (t, r, θ, φ) reads

$$ds^2 = - \left(1 - \frac{2GM}{c^2 r}\right) c^2 dt^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 (d\theta^2 + \sin^2\theta d\varphi^2), \quad (1.8)$$

describes the static gravitational field in the region outside a spherical mass M (though can be extended to the interior [215]) with no angular momentum or electric charge [216, 217, 218]. This applies to the spacetime region surrounding a non-rotating BH, which has a physical singularity at $r = 0$ and a coordinate singularity at the Schwarzschild radius $r = 2GM/c^2$ [215, 219, 220, 221, 222] — the event horizon; below this radius no energy can escape to the exterior region of the spacetime, which would require faster-than-light propagation.

- The Kerr metric [223] is the rotating extension of the Schwarzschild BH. Non-zero angular momentum induces new relativistic effects, including frame-dragging — in which the spacetime surrounding the central body is rotated along with it — and the ergoregion — the spacetime region surrounding the singularity inside which all energy cannot help but co-rotate with it, as to do otherwise would require a super-luminal rotational velocity. There is a maximum possible rotation rate above which the horizon ceases to exist, leaving a naked singularity and thus thought to be an impossibility due to cosmic censorship [224]. The Kerr solution describes astrophysical BHs because matter in the universe has on the whole neutral charge but non-zero angular momentum.
- The Reissner–Nordström metric [225, 226, 227, 228] is the electrically charged

extension of the Schwarzschild solution, while the Kerr–Newman metric [229, 230] is the charged extension of the Kerr solution.

It was far from clear that the BH singularities found in these solutions early after the development of GR were real and could be found in nature. However, these predictions came to be another great success of Einstein’s theory. Oppenheimer and Snyder showed that a black hole could be born from the death of a star undergoing gravitational collapse [231]. The singularity theorems of Penrose and Hawking [222, 232, 233, 234, 235] firmly confirmed the existence of physical singularities in solutions of the Einstein equations. Before recent results from GW astronomy cast the question of BH existence beyond a doubt, necessarily indirect observations (one cannot directly observe a BH electromagnetically due to its event horizon) had already all but confirmed the existence of astrophysical BHs [236, 237]. The evidence for supermassive BHs included observations of extremely energetic phenomena in active galactic nuclei (AGN) and quasars [238, 239, 240, 241, 242, 243], and the tight orbits of stars around the galactic centre Sagittarius A* [244, 245, 246]. The evidence for stellar-mass BHs came primarily from compact X-ray binaries, in which a visible star is bound to an unseen compact companion inferred to be a BH [247, 248].

1.2.2 Perturbed spacetime

One final great success of Einstein’s GR — and the one most important to this work (in addition to the prediction of BHs) — is the prediction of GWs [5]. While the above textbooks on GR typically contain chapters dedicated to GWs, for pedagogical sources dedicated to the subject see, e.g., Refs. [249, 250, 251, 252, 253].

From the flat Minkowski solution, a natural curiosity is spacetimes containing matter but with small curvatures, which can be considered as perturbations of the original

flat metric. One thus declare a form of these solutions

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.9)$$

where h is a metric perturbation satisfying $|h| \ll 1$. By linearization it may be shown that, in the Lorenz gauge, the Ricci tensor reduces to $R_{\mu\nu} = -\frac{1}{2}\square\bar{h}_{\mu\nu}$, where $\square = \eta^{\mu\nu}\partial_\mu\partial_\nu = -\frac{1}{c^2}\partial_t^2 + \partial_i^2$ is the d'Alembert operator or, suggestively, the wave operator and the trace-reversed metric perturbation is $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\eta^{\alpha\beta}h_{\alpha\beta}$. One can show that the solution satisfies

$$\square\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu}. \quad (1.10)$$

That is, the metric perturbation $\bar{h}_{\mu\nu}$ obeys a sourced wave equation. For compact sources, the linearized Einstein equations permit wave-like solutions

$$\bar{h}_{\mu\nu}(\mathbf{x}, t) = \frac{16\pi G}{c^4} \int_{\mathbb{R}^3} \frac{T_{\mu\nu}(\mathbf{x}', t - |\mathbf{x} - \mathbf{x}'|/c)}{4\pi|\mathbf{x} - \mathbf{x}'|} d^3\mathbf{x}', \quad (1.11)$$

where $ct = x^0$ is the temporal coordinate and $\mathbf{x} = (x^1, x^2, x^3)$ is the spatial vector. Furthermore, considering distances far from the source one arrives at the famous mass-quadrupole formula

$$\bar{h}_{ij}(\mathbf{x}, t) = \frac{2G}{c^4|\mathbf{x}|} \partial_t^2 \int_{\text{source}} x'^i x'^j \rho \left(\mathbf{x}', t - \frac{|\mathbf{x}|}{c} \right) d^3\mathbf{x}', \quad (1.12)$$

where ρ is the mass density of the source over space and time.

Put simply, a flat background spacetime produces GWs that radiate to infinity if an accelerating (non-zero second-order time derivative) asymmetric mass distribution

(non-zero integral of density over the source) is present.

1.2.3 Gravitational-wave sources

Binary stars are astrophysical systems that naturally satisfy these requirements; the accelerating asymmetry is provided by the orbit of the binary components. For a binary of mass M , orbital separation r , and period T at a distance D from Earth, the mass quadrupole is $\sim Mr^2$ and $\partial_t^2 \sim 1/T^2$. Its components orbit with velocity $v = \sqrt{GM/r}$, producing a GW strain that scales with (i) the ratio of the Schwarzschild radius to the distance and (ii) the square of the orbital velocity [42]:

$$h = |\bar{h}| \sim \frac{GM/c^2}{D} \left(\frac{v}{c}\right)^2. \quad (1.13)$$

For a given mass, the strain is therefore maximized for high orbital velocities. Fortunately, compact-object binaries — such as those containing BHs — can reach $v \sim 0.1c$ near the end of their lives and are thus promising gravitational radiators.

As an example, the first GW detection — GW150914 — had a total mass $\approx 66M_\odot$ and was at a distance ≈ 410 Mpc, producing a GW strain $h \sim 10^{-21}$ [39, 42]. The strain produces a fractional change $\delta L \sim hL$ in distances over a region of length L . The ≈ 4 km arm lengths of the LIGO interferometers [37] therefore required sensitivities $\delta L \sim 10^{-18}$ m to make this detection; to set this astounding feat in context, a single proton has size $\sim 10^{-15}$ m!

Since gravitational radiation carries energy away from a binary system, it loses orbital energy and hence its components must inspiral toward each other. The approximate GW flux for a quasi-circular BH binary with total mass M and mass ratio

$0 < q \leq 1$ causes the orbital radius r to shrink over time t at a rate [254, 255]

$$\frac{dr}{dt} = -\frac{64G^3q}{5c^5(1+q)^2} \frac{M^3}{r^3}. \quad (1.14)$$

At small separations, the emission of GWs becomes more and more effective at driving the inspiral but, conversely, it is much weaker at larger separations. After integrating and approximating that the two BHs merge at zero separation (in reality their horizons will meet before this), one finds that in order for an equal-mass binary BH to merge within the age of the universe — a Hubble time t_H — its initial separation r_0 must satisfy

$$r_0 \leq \left(\frac{64G^3t_H}{5c^5} M^3 \right)^{1/4} \sim 10^9 \left(\frac{M}{M_\odot} \right)^{3/4} \text{ m}. \quad (1.15)$$

Stellar-mass ($M \sim 10M_\odot$) and supermassive ($M \sim 10^8M_\odot$) binary mergers driven entirely by gravitational radiation therefore require initial separations $\sim 10R_\odot$ and ~ 0.01 pc, respectively. However, these are both much smaller than the separations at which such binaries naturally form. GW emission alone cannot merge astrophysical BHs. As stellar-mass BH mergers have been observed [64] and there is some tentative evidence for supermassive BH mergers [256, 257, 258, 259, 260, 261, 262, 263, 264, 265, 266, 267, 268], additional astrophysical mechanisms must exist that initially dissipate orbital energy before GW emission becomes the dominant source of radiation.

Several proposed binary-formation scenarios offer solutions to this problem on the stellar-mass scale [269, 270, 271]:

- Binary stars born through isolated stellar evolution in galactic fields may undergo a ‘common-envelope’ phase, during which one component has already collapsed

to a BH and orbits the companion star within a shared gaseous envelope [272, 273, 274, 275, 276]. The envelope is expelled as it transfers its binding energy to the binary orbit, thus shrinking its separation by orders of magnitude.

- Stellar mergers in short-period binary systems can be prevented by chemically homogeneous evolution which keeps the stars compact [277, 278, 279].
- In dense stellar environments [280, 281, 282] — such as globular clusters [283, 284, 285, 286, 287, 288, 289], nuclear star clusters [290, 291, 292, 293, 294], and young star clusters [295, 296, 297, 298] — multi-body encounters lead to the dynamical hardening of binaries.
- A particular dynamical effect is the Kozai-Lidov mechanism, in which a tertiary perturber pumps the eccentricity of a companion binary system [299, 300]. Since eccentricity enhances the GW flux [254, 255], this accelerates the merger of the inner binary [291, 301, 302, 303].
- In galactic cores, orbiting compact objects may be captured into an AGN disk and form binaries in dense migration traps [304, 305, 306, 307]. The viscous drag of the gaseous disk and dynamical interactions with other disk components reduces the orbital separation of these embedded binaries before gravitational radiation takes over [308, 309, 310], leading to a population of merging binary BHs [311, 312, 313].

For supermassive BH mergers, dynamical friction [314, 315] and viscous drag [316, 317] may solve the so-called ‘final-parsec problem’ [318, 319, 320, 321].

Astrophysical BHs are either the endpoints of the lives of massive stars or the remnants of previous compact-binary coalescences. Binary BHs, therefore, represent

the final product of the co-evolution of two stars, or the assembly of one or two merger remnants into a new binary system. It is unlikely that just a single one of the above channels is responsible for producing all merging binary BHs [269, 270, 271]. The astrophysical processes that assemble such binaries determine their properties, such as masses, which in turn influence the GW signal emitted during coalescence. From GW data it may therefore be possible to decode the formation histories of binary BHs.

For example, the two canonical formation pathways are isolated binary evolution and dynamical evolution in dense stellar clusters (see above). Stellar theory predicts the existence of a mass range in which supernovae do not result in a BH — the pair instability mass gap [322, 323, 324, 325, 326, 327]. Therefore, in the former scenario, BH mergers from the isolated channel should not have masses inside this gap. The two component stars are born from the same gaseous medium and inherit its angular momentum. There is theoretical motivation to suggest that efficient angular momentum transport through the stellar interiors to their envelopes results in slowly rotating BHs [328, 329]. While some processes can act to tilt the direction of BH spins with respect to the binary orbit [330], it is expected that isolated binary evolution leads to spin alignment [331]. On the other hand, in the latter scenario of dynamical environments the purported mass gap can be populated by BHs that merge multiple times [332, 333]. These repeated mergers also result in BHs with larger spins [334, 335], while the spin directions are randomized by dynamical encounters. Consequently, these two formation pathways predict two distinct populations of merging binary BHs: a mass-limited, low- and aligned-spin population from isolated binary evolution and a larger-mass, larger- and tilted-spin population from dense clusters. Using GW observations it may be possible to identify such subpopulations and therefore place constraints on the astrophysics of compact-binary formation [336, 337, 338].

This highlights the essential interplay between astrophysics — which predicts how stars evolve and form compact binaries — and GW physics — which predicts the ultimate fate of those systems in compact-object mergers, and hence the need for a detailed understanding of the spin dynamics of binary BHs.

1.3 Multi-timescale binary dynamics

Astrophysical formation is only one half of the story of binary BH mergers, however. After forming and once the orbital separation has sufficiently shrunk, the effects of GR become prominent and influence the binary dynamics. The post-formation inspiral that maps astrophysical properties to GW-observable properties is clean, however, in the sense that it is only mathematical in nature, free from astrophysical uncertainties.

1.3.1 Post-Newtonian primer

A significant complication is the non-linear nature of GR — no analytic solution to the two-body problem is known. Approximate techniques are necessary to find solutions. The most accurate method is numerical relativity (NR), in which the full Einstein equations are solved with advanced computational methods [339, 340, 341, 342, 343]. However, such simulations typically require at least days of supercomputer time to complete ~ 10 orbits before the merger.

On the other hand, post-Newtonian (PN) techniques are part of a long-standing [154] semi-analytic framework used to derive weak-field approximants of GR. While this approach is less accurate than NR in the strong-field regime and cannot capture the full non-linearities inherent to GR, progress can be made in closed form and numerical implementations are much more efficient.

In the following I will use geometric units $G = c = 1$ (unless otherwise stated). In these units, distance and time have dimensions of mass while angular momenta have dimensions of mass squared. As such, quantities such as r/M and L/M^2 that are dimensionless will be encountered, where r , L , and M represent a distance, an angular momentum, and a mass, respectively. I denote vectors, e.g., $\mathbf{v} \in \mathbb{R}^3$, in boldface, their magnitude with $v = |\mathbf{v}|$, and normalized vectors $\hat{\mathbf{v}} = \mathbf{v}/v$ with a hat. Considering quasi-circular binary BHs in vacuum, the heavier primary (lighter secondary) BH has constant mass m_1 (m_2). The BH spin angular momenta are \mathbf{S}_i and the dimensionless spin magnitudes (Kerr parameters) are given by $\chi_i = S_i/m_i^2$ ($i \in \{1, 2\}$). The total BH spin is $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$. The binary has total mass $M = m_1 + m_2$, mass ratio $q = m_2/m_1$, reduced mass $\mu = m_1 m_2 / M$, and symmetric mass ratio $\eta = m_1 m_2 / M^2$. The separation vector pointing from the primary BH to the secondary is \mathbf{r} . The Newtonian orbital angular momentum is $\mathbf{L} = \mu \mathbf{r} \times d\mathbf{r}/dt$ and its magnitude is $L = \eta \sqrt{M^3 r}$. The total angular momentum is $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$.

In the PN approximation, the binary BH dynamics are determined by a set of equations given to a particular PN order, corresponding to a series expansion in terms of a small parameter such as v/c , where $v = \sqrt{M/r}$ is the Newtonian orbital velocity. In other words, an approximation to GR is built perturbatively upon a Newtonian foundation. For large orbital separations $r \gg M$ this is a suitable approximation since $v \ll c$, but it breaks down at late times in the binary inspiral near merger where $v \sim 0.1c$; a typical limit for the domain of validity is taken to be $r \gtrsim 10M$. An expansion including terms up to order $\mathcal{O}(v^{2n})$ beyond the Newtonian order is referred to as being of n PN order.

Without delving into the full mathematical details and derivations here (comprehensive reviews may be found in, e.g., Refs. [344, 345]), the 2PN equations determining

the evolution of the binary angular momenta are [346]

$$\frac{d\mathbf{L}}{dt} = \frac{1}{2r^3} \left[(\mathbf{S} + 3\mathbf{S}_0) \times \mathbf{L} + 6\frac{\mu}{M}(\hat{\mathbf{r}} \cdot \mathbf{S}_0)\hat{\mathbf{r}} \times \mathbf{S}_0 \right], \quad (1.16)$$

$$\frac{d\mathbf{S}_1}{dt} = \frac{1}{2r^3} \left[(4 + 3q)\mathbf{L} - 2\mathbf{S}_2 + 6\frac{\mu}{m_1}(\hat{\mathbf{r}} \cdot \mathbf{S}_0)\hat{\mathbf{r}} \right] \times \mathbf{S}_1, \quad (1.17)$$

$$\frac{d\mathbf{S}_2}{dt} = \frac{1}{2r^3} \left[\left(4 + \frac{3}{q}\right)\mathbf{L} - 2\mathbf{S}_1 + 6\frac{\mu}{m_2}(\hat{\mathbf{r}} \cdot \mathbf{S}_0)\hat{\mathbf{r}} \right] \times \mathbf{S}_2, \quad (1.18)$$

where $\mathbf{S}_0 = (1 + q)\mathbf{S}_1 + (1 + 1/q)\mathbf{S}_2$ [347]. These equations are the basis for the multi-timescale approach to PN binary BH inspirals that will follow in this Section.

One should note in particular that all three angular momenta influence each other through spin–orbit and spin–spin couplings; the 2PN order is the lowest in which all of these effects are captured. This results in the GR effect of angular-momentum precession [348, 349, 350] — the relative orientations of \mathbf{L} , \mathbf{S}_1 , and \mathbf{S}_2 change over time, causing a tilting and wobbling of the orbital plane. This precession impacts the GW signal emitted by inspiralling binary BHs and, therefore, proper characterization of the properties of GW observations must account for its effect for when modelling source dynamics.

1.3.2 Timescale hierarchy

Altogether there are three identifiable constituents to the binary dynamics: orbits, precession, and inspiral. Entrenched within these motions is a timescale hierarchy that can be leveraged to simplify the PN equations of motion. The timescales of these constituents can be estimated as follows:

- Kepler's third law informs us that the orbital timescale is

$$t_{\text{orb}} \sim r^{3/2}. \quad (1.19)$$

- Taking the dominant term $\propto \mathbf{L}$ in Eqs. (1.17, 1.18), the precessional timescale is

$$t_{\text{pre}} \sim \frac{S_i}{|dS_i/dt|} \sim r^{5/2}. \quad (1.20)$$

- From Eq. (1.14) one infers that the radiation-reaction timescale over which the orbit shrinks is

$$t_{\text{RR}} \sim \frac{r}{|dr/dt|} \sim r^4. \quad (1.21)$$

In the PN approximation the early inspiral is studied, where $r \gg M$, implying that the three timescales are disjoint:

$$t_{\text{orb}} \ll t_{\text{pre}} \ll t_{\text{RR}}. \quad (1.22)$$

In words, a binary completes many orbital cycles before completing a precession cycle, and completes many precession cycles before the orbital separation decreases appreciably. Therefore, the constituent dynamics may be treated with a multi-timescale approach, in which motion on a short timescale may be viewed in an averaged manner and secular motion on a long timescale may be viewed as quasi-adiabatic. I now treat each constituent of the binary dynamics in turn.

1.3.3 Spin precession

Numerically solving the equations of motion (1.16–1.18) over long timescales — e.g., over a binary inspiral — becomes prohibitive as many orbital cycles must be tracked. To study long inspirals — in particular the evolution of the BH spins over many orbits — one can instead average over the orbital motion. For a binary with orbital period T , the orbit-averaged value of some scalar X is given by

$$\langle X \rangle_{\text{orb}} = \frac{1}{T} \int_0^T X dt. \quad (1.23)$$

For the sake of clarity I forgo the notation $\langle \cdot \rangle_{\text{orb}}$ and implicitly assume all quantities that vary on the orbital timescale are averaged.

Performing the orbit-averaging procedure for Eqs. (1.16–1.18) yields the 2PN orbit-averaged spin-precession equations [346, 351]

$$\frac{d\mathbf{L}}{dt} = \boldsymbol{\Omega}_L \times \mathbf{L} + \frac{dL}{dt} \hat{\mathbf{L}}, \quad (1.24)$$

$$\frac{d\mathbf{S}_1}{dt} = \boldsymbol{\Omega}_1 \times \mathbf{S}_1, \quad (1.25)$$

$$\frac{d\mathbf{S}_2}{dt} = \boldsymbol{\Omega}_2 \times \mathbf{S}_2, \quad (1.26)$$

where the precession vectors are

$$\boldsymbol{\Omega}_L = \frac{1}{2r^3} \left\{ \left[4 + 3q - 3(q\mathbf{S}_1 + \mathbf{S}_2) \cdot \frac{\hat{\mathbf{L}}}{L} \right] \mathbf{S}_1 + \left[4 + \frac{3}{q} - 3 \left(\mathbf{S}_1 + \frac{\mathbf{S}_2}{q} \right) \cdot \frac{\hat{\mathbf{L}}}{L} \right] \mathbf{S}_2 \right\}, \quad (1.27)$$

$$\boldsymbol{\Omega}_1 = \frac{1}{2r^3} \left\{ (4 + 3q)\mathbf{L} - \left[3(q\mathbf{S}_1 + \mathbf{S}_2) \cdot \hat{\mathbf{L}} \right] \hat{\mathbf{L}} + \mathbf{S}_2 \right\}, \quad (1.28)$$

$$\boldsymbol{\Omega}_2 = \frac{1}{2r^3} \left\{ \left(4 + \frac{3}{q} \right) \mathbf{L} - \left[3 \left(\mathbf{S}_1 + \frac{\mathbf{S}_2}{q} \right) \cdot \hat{\mathbf{L}} \right] \hat{\mathbf{L}} + \mathbf{S}_1 \right\}. \quad (1.29)$$

The conservative part of the dynamics is supplemented with a radiative part dL/dt corresponding to GW emission; see Eq. (1.14) for the leading-order term [this was neglected in Eq. (1.16) since its PN order is higher than that of the 2PN equations, but is included here for completeness].

Binaries in which the component BH spins are aligned with the orbital angular momentum are equilibrium solutions of Eqs. (1.24–1.26) — sources that form in such configurations remain in them over time. In general, however, the mutual orientations of the three angular momenta vary over time, resulting in the aforementioned spin precession.


The 2PN spin-precession equations present a nine-dimensional problem corresponding to the components of the orbital angular momentum \mathbf{L} and BH spins \mathbf{S}_1 and \mathbf{S}_2 . I begin simplifying the problem by noting that $dS_i/dt = \hat{\mathbf{S}}_i \cdot (\boldsymbol{\Omega}_i \times \mathbf{S}_i) = 0$, implying the spin magnitudes S_i ($i \in \{1, 2\}$), and thus the Kerr parameters χ_i , are conserved. This reduces the dimensionality to seven.

Next, define a non-inertial frame, seen in Fig. 1.1, in which

- the orbital angular momentum lies along the z axis, so that $\mathbf{L} = (0, 0, L)$,
- the primary spin vector is taken to lie within the x - z plane, so that $\mathbf{S}_1 = (S_{1x}, 0, S_{1z})$, and
- the y axis completes the right-handed orthonormal frame.

This frame, which follows the motion of \mathbf{L} , imposes three restrictions — $L_x = L_y = S_{1y} = 0$ — and hence reduces the dimensionality of the 2PN spin precession problem

Figure 1.1: Non-inertial frame aligned with the orbital angular momentum \mathbf{L} . The angle between each spin vector \mathbf{S}_i and \mathbf{L} is denoted by θ_i ($i = 1, 2$), while the angle between the projections of the two spins onto the orbital plane is denoted by $\Delta\Phi$.



to four. The only time-varying quantities are the BH spin directions relative to the non-inertial frame and the magnitude of the orbital angular momentum or, equivalently, the orbital separation. The mutual orientations of the three angular momenta can be specified by the polar angles θ_1 and θ_2 between the BH spins and the orbital angular momentum (z axis) and the azimuthal angle $\Delta\Phi$ between the projection of the secondary spin onto the orbital plane and that of the primary spin (x axis):

$$\theta_1 = \arccos(\hat{\mathbf{S}}_1 \cdot \hat{\mathbf{L}}), \quad (1.30)$$

$$\theta_2 = \arccos(\hat{\mathbf{S}}_2 \cdot \hat{\mathbf{L}}), \quad (1.31)$$

$$\Delta\Phi = \arccos\left(\frac{\hat{\mathbf{S}}_1 \times \hat{\mathbf{L}}}{|\hat{\mathbf{S}}_1 \times \hat{\mathbf{L}}|} \cdot \frac{\hat{\mathbf{S}}_2 \times \hat{\mathbf{L}}}{|\hat{\mathbf{S}}_2 \times \hat{\mathbf{L}}|}\right). \quad (1.32)$$

Simplifying further still, one can make use of the separation $t_{\text{pre}} \ll t_{\text{RR}}$ between the precessional and inspiral timescales. Quantities which undergo only secular variation over the longer inspiral may be treated quasi-adiabatically, such that they are constant over the shorter precessional cycles. In particular, the magnitude L of the orbital

angular momentum is conserved over t_{pre} , meaning $dL/dt = 0$ in Eq. (1.24), as is that of the total angular momentum $J = |\mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2|$ (they do vary on the longer inspiral timescale t_{RR} , however). This reduces the number of parameters to two.

One final simplification is made by introducing the effective aligned spin [346]

$$\chi_{\text{eff}} = \frac{\mathbf{S}_0}{M^2} \cdot \hat{\mathbf{L}} = \left(\frac{\mathbf{S}_1}{m_1} + \frac{\mathbf{S}_2}{m_2} \right) \cdot \frac{\hat{\mathbf{L}}}{M}, \quad (1.33)$$

which measures the mass-weighted component of the BH spins along the direction of the orbital angular momentum. This is the best-measured spin parameter in current GW detections [61, 62, 63, 64]. One can use Eqs. (1.24–1.29) to show that χ_{eff} is conserved at 2PN.

The original nine-dimensional spin-precession problem has been reduced to solving for a single parameter that varies on the precession timescale [352, 353]. A convenient choice for this parameter is $S = |\mathbf{S}_1 + \mathbf{S}_2|$ (though recent works suggest an alternative parametrization [354, 355]). Conversions [352, 353] back to the more geometrically intuitive angles are given by

$$\cos \theta_1 = \frac{1}{2(1-q)S_1} \left(\frac{J^2 - L^2 - S^2}{L} - \frac{2M^2 q \chi_{\text{eff}}}{1+q} \right), \quad (1.34)$$

$$\cos \theta_2 = \frac{-q}{2(1-q)S_2} \left(\frac{J^2 - L^2 - S^2}{L} - \frac{2M^2 \chi_{\text{eff}}}{1+q} \right), \quad (1.35)$$

$$\cos \Delta\Phi = \frac{1}{\sin \theta_1 \sin \theta_2} \left(\frac{S^2 - S_1^2 - S_2^2}{2S_1 S_2} - \cos \theta_1 \cos \theta_2 \right), \quad (1.36)$$

while the inverses are

$$S^2 = S_1^2 + S_2^2 + 2S_1 S_2 (\cos \theta_1 \cos \theta_2 + \sin \theta_1 \sin \theta_2 \cos \Delta\Phi), \quad (1.37)$$

$$J^2 = L^2 + 2L(S_1 \cos \theta_1 + S_2 \cos \theta_2) + S^2, \quad (1.38)$$

$$\chi_{\text{eff}} = \frac{1+q}{M^2 q} (qS_1 \cos \theta_1 + S_2 \cos \theta_2) . \quad (1.39)$$

Using Eqs. (1.24–1.29), one finds that the temporal evolution of the chosen parameter S is determined by [352, 353]

$$\frac{dS}{dt} = -\frac{3(1-q^2)}{2q} \left(\frac{M^3 \chi_{\text{eff}}^2}{L^2} \right)^3 \left(1 - \frac{M^2 \eta \chi_{\text{eff}}}{L} \right) \frac{LS_1 S_2}{S} \sin \theta_1 \sin \theta_2 \sin \Delta \Phi . \quad (1.40)$$

Using Eqs. (1.34–1.36), Eq. (1.40) can be rewritten in terms of only S and quantities that do not vary on the precession timescale [356, 357]:

$$\left(\frac{dS^2}{dt} \right)^2 = -A^2 (S^6 + BS^4 + CS^2 + D) = A^2 (S_+^2 - S^2) (S^2 - S_-^2) (S^2 - S_3^2) , \quad (1.41)$$

where

$$A^2 = -\frac{6(1-q)^2}{q\sqrt{\eta}} \left(\frac{M^3 \eta^2}{L^2} \right)^3 \left(1 - \frac{M^2 \eta \chi_{\text{eff}}}{L} \right)^2 L^2 S_1^2 S_2^2 , \quad (1.42)$$

$$B = -2J^2 + \frac{1+q^2}{q} L^2 + 2M^2 \chi_{\text{eff}} L - \frac{1-q}{q} (qS_1^2 - S_2^2) , \quad (1.43)$$

$$C = \frac{2(1-q)}{q} [J^2 (qS_1^2 - S_2^2) - L^2 (S_1^2 - qS_2^2)] + (J^2 - L^2)^2 - 2M^2 \chi_{\text{eff}} L \left[(J^2 - L^2) + \frac{1-q}{1+q} (S_1^2 - S_2^2) \right] + \frac{4M^4 q \chi_{\text{eff}}^2 L^2}{(1+q)^2} , \quad (1.44)$$

$$D = \frac{1-q}{q(1+q)} [L^2 (1-q^2) (S_1^2 - S_2^2)^2 - (1+q) (qS_1^2 - S_2^2) (J^2 - L^2) + 2M^2 q \chi_{\text{eff}} L (S_1^2 - S_2^2) (J^2 - L^2)] , \quad (1.45)$$

and $S_+^2 \geq S_-^2 \geq S_3^2$ are the roots of the third-degree polynomial in S^2 from Eq. (1.41).

Since the roots represent local extrema of the motion, for spin precession to be

physical a bounded interval corresponding to two real roots is required, in which case all three roots must be real. This physical interval is $S^2 \in [S_-^2, S_+^2]$, while S_3^2 is a spurious root introduced when squaring Eq. (1.40).

In fact, Eq. (1.41) can be solved in closed form [354, 355, 356, 357] (at least or in terms of well known integrals [358]). The solution is an oscillatory function for which the two halves of the precession cycle are symmetric. Ultimately, with this solution the entire intrinsic binary BH spin dynamics on the precessional timescale is known semi-analytically (at 2PN order).

1.3.4 Inspirals

Just as orbits on the timescale t_{orb} were averaged over to study spin precession over t_{pre} , so too can precession cycles be averaged to study the secular binary inspiral over t_{RR} . Owing to the symmetry of a precession cycle, the double average of a scalar function X over orbits and, in turn, precession is given by [352, 353]

$$\langle X \rangle_{\text{pre}} = \frac{2}{\tau} \int_0^\tau \langle X \rangle_{\text{orb}} dt = \frac{2}{\tau} \int_{S_-}^{S_+} \frac{\langle X \rangle_{\text{orb}}}{|dS/dt|} dS, \quad (1.46)$$

where the precession period is $\tau = 2 \int_{S_-}^{S_+} |dS/dt|^{-1} dS$.

The effective aligned spin χ_{eff} is conserved also at lowest PN order in radiation reaction [346], implying that the only quantities varying over the quasi-adiabatic inspiral are J , L , and S . The orbit-averaged angular momentum flux at the highest spin-independent PN order has the form $d\mathbf{J}/dt = f(L)\hat{\mathbf{L}}$ [350], where f is a function of the magnitude of the orbital angular momentum L . Since L is monotonic (in particular, decreasing as the binary inspirals) it can be used itself to parametrize the inspiral, and one can show that $dJ/dL = \hat{\mathbf{J}} \cdot \hat{\mathbf{L}}$ [352, 353]. As shown above, the solution

for S over the precession timescale t_{pre} depends on quantities that then vary over the longer inspiral timescale t_{RR} — in particular on J and L . To evolve binary BHs on the inspiral timescale as a function of L , or equivalently the orbital separation r , it remains to solve only for J .

Since $S^2 = J^2 + L^2 - 2JL\hat{\mathbf{J}} \cdot \hat{\mathbf{L}}$, the precession-averaged evolution is determined by [352, 353]

$$\left\langle \frac{dJ}{dL} \right\rangle_{\text{pre}} = \frac{J^2 + L^2 - \langle S^2 \rangle_{\text{pre}}}{2JL}. \quad (1.47)$$

While this can be integrated analytically by neglecting higher-order PN terms [354, 356, 357], the precession average $\langle S^2 \rangle_{\text{pre}}$ can again be computed with standard (semi-analytic) integrals [354, 355]. This means numerical solutions of Eq. (1.47) can be computed very efficiently from arbitrarily large initial orbital separations, orders of magnitude faster than orbit-averaged solutions [353, 355]. The price one pays is losing track of the orbit-averaged precessional phase S , much in the same way one loses knowledge of the orbital phase in the orbit-averaged precession equations.

An interesting feature of the precession-averaged evolution of J is that it can be regularized for infinite orbital separations, meaning the binary dynamics can be solved all the way from or to past time infinity [353]. This limit serves as a useful approximation of the large separations at which astrophysical formation of binary BHs occurs [359].

Consider the change of variables

$$\kappa = \frac{J^2 - L^2}{2L}, \quad (1.48)$$

$$u = \frac{1}{2L}, \quad (1.49)$$

such that the infinite orbital separation limit $r \rightarrow \infty$ corresponds to $u = 0$. Eq. (1.47) simplifies to

$$\frac{d\kappa}{du} = \langle S^2 \rangle_{\text{pre}}. \quad (1.50)$$

On the other hand, $\kappa = \mathbf{S} \cdot \hat{\mathbf{L}} + S^2/(2L)$ and $d\kappa/dL = -\langle S^2 \rangle_{\text{pre}}/(2L^2) \rightarrow 0$ as $r \rightarrow \infty$, implying that κ converges to a finite value

$$\kappa_\infty = \lim_{r \rightarrow \infty} \mathbf{S} \cdot \hat{\mathbf{L}} \quad (1.51)$$

at large separations and thus the solution is regular at $u = 0$.

Eq. (1.50) is a convenient ordinary differential equation to numerically solve for precession-averaged inspirals. By providing initial conditions at $u = 0$, it can be used to evolve binary BHs from infinite orbital separations to late times at which the PN approximation breaks down. Conversely, by setting initial conditions at some finite orbital separation, a binary can be evolved backwards in time, all the way to infinite orbital separations. The former scenario is relevant when modelling astrophysical formation and propagating forward to detectable GW predictions, while the latter can be used to map GW detections backwards to binary formation.

Given the constant binary mass M , mass ratio q , and spins S_1 and S_2 , precession-averaged inspirals can be initialized with values of the effective aligned spin χ_{eff} and the new angular momentum parameter κ . Conversions between these and the spin angles are given by Eqs. (1.34, 1.35, 1.38, 1.39, 1.48). In particular, at infinite orbital separations,

$$\theta_{1\infty} = \lim_{r \rightarrow \infty} \theta_1 = \arccos \frac{(1+q)\kappa_\infty - M^2 q \chi_{\text{eff}}}{(1-q^2)S_1}, \quad (1.52)$$

$$\theta_{2\infty} = \lim_{r \rightarrow \infty} \theta_2 = \arccos \frac{M^2 q \chi_{\text{eff}} - q(1+q)\kappa_\infty}{(1-q^2)S_2}, \quad (1.53)$$

$$\kappa_\infty = S_1 \cos \theta_{1\infty} + S_2 \cos \theta_{2\infty}, \quad (1.54)$$

$$\chi_{\text{eff}} = \frac{1+q}{M^2 q} (qS_1 \cos \theta_{1\infty} + S_2 \cos \theta_{2\infty}). \quad (1.55)$$

Specifying the angles as initial conditions is often more convenient as formation models typically make predictions for those quantities (see, e.g., Sec. 1.2.3). From there, the entire precession-averaged binary BH inspiral is specified.

Since the precessional phase S was averaged over, the orbit-averaged spin dynamics is no longer deterministic. Reconstructing the spins at finite orbital separations is done probabilistically by sampling from the probability density function (PDF; see below) $\propto |dS/dt|^{-1}$, representing that a binary is more likely to have lower phase velocity dS/dt [352, 353]. Eqs. (1.34–1.36) can then be combined with the precession-averaged solution $J = J(L)$ to compute the BH spin orientations.

1.4 Detection and inference

Now one knows how to propagate astrophysical predictions of binary BH formation to the small orbital separations at which binaries merge and can be detected by GW observatories (and vice versa). Characterization the source properties of a GW event given the observed data is still required before astrophysical conclusions can be drawn. This is the role of Bayesian inference [249, 360, 361, 362, 363].

The goal is to make measurements from noisy observations, where measurement uncertainty is represented by PDFs. In words, while exact measurements are unavailable, one can attempt to infer the probability that a range of values are correct. I will denote the probability of an event, say X , as $p(X)$. If X and Y are mutually exclusive, then

the total probability must be $p(X) + p(Y) = 1$; this normalization condition readily extends to multiple events, i.e., $\sum_i p(X_i) = 1$.

Often, it is of interest to measure a continuous random variable, say x . In this limit, the probability that the value of x lies within an infinitesimal interval $[x, x + dx]$ of length dx is $p(x)dx$. Then, the normalization condition for the PDF becomes $\int p(x)dx = 1$. According to the total law of probability, the joint distribution of two random variables, x and y , may by symmetry be written as

$$p(x, y) = p(x|y)p(y) = p(y|x)p(x). \quad (1.56)$$

The conditional distribution, e.g., $p(x|y)$ (read as the probability density of x given y), represents the dependence of the random variable x on the conditioning variable y , and is normalized over the dependent variable, i.e., $\int p(x|y)dx = 1 \forall y$. Given a value of y , the uncertainty in x may be altered, but if x does not depend on y , then $p(x|y) = p(x)$.

As an example, consider a fair coin that was flipped and landed heads up, corresponding to a discrete random variable Y . When flipped again, landing on heads or tails (X) does not depend on the previous result — the flips and their probabilities are independent, i.e., $p(X, Y) = p(X)p(Y)$. On the other hand, imagine rolling a fair six-sided die such that the probability of it landing on any given face is $1/6$. If instead one is told that it landed on an odd number, the conditional probability for any one of the odd faces is $1/3$, while that for the even faces is 0. Such conditioning can also apply to continuous random variables. Using Eq. (1.56), one finds the result for conditional distributions known as Bayes' theorem [364]:

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)}. \quad (1.57)$$

1.4.1 Searching for signals

The current ground-based observatories are L-shaped Michelson interferometers with two perpendicular arms [365, 366]. Each arm contains a vacuum tube with a light beam produced from a common laser source, after being passed through a beam splitter. For the two LIGO observatories, each arm measures 4 km [37] (though the effective path length is increased 300-fold by reflecting the light within each arm [39]). The instruments are tuned such that the light beams returning from the two arms are out of phase and thus interfere destructively when recombined at the beam splitter. In the absence of distortions to this experimental configuration, no light will ultimately arrive at a receiving photodiode. On the other hand, passing GWs distort the two light paths and cause an interference pattern, leading to an electric readout from the detector from which the presence of a signal can be inferred and its morphology reconstructed.

The GW interferometers collect this time-series data $d(t)$ (the time t is the coordinate time of distant observers from the source). But of course, the detectors are not perfect instruments and as such the data contain noise; the central problem of detecting a real GW signal is digging it out of this noise [367, 368].

Let \mathcal{H} denote the hypothesis that the data contain a signal $h(t)$, while the null hypothesis \mathcal{N} is that the data contain only noise $n(t)$. Using Bayes' theorem, the posterior probability — i.e., having taken data — in favour of the signal hypothesis is

$$p(\mathcal{H}|d) = \frac{p(d|\mathcal{H})p(\mathcal{H})}{p(d)}, \quad (1.58)$$

where $p(\mathcal{H})$ is the prior belief in the signal hypothesis and $p(d|\mathcal{H})$ is the likelihood that the data was produced in the presence of the signal h . Since the two hypotheses are

mutually exclusive, the marginal likelihood is

$$p(d) = p(d|\mathcal{H})p(\mathcal{H}) + p(d|\mathcal{N})p(\mathcal{N}), \quad (1.59)$$

where $p(\mathcal{N})$ is the prior belief in the null hypothesis and $p(d|\mathcal{N})$ is the likelihood that the data is just noise. The signal hypothesis is favoured over its complement by a factor determined with the posterior odds, $p(\mathcal{H}|d)/p(\mathcal{N}|d) = [p(d|\mathcal{H})/p(d|\mathcal{N})][p(\mathcal{H})/p(\mathcal{N})]$, where $p(\mathcal{N}|d) = 1 - p(\mathcal{H}|d)$ is the posterior for the null hypothesis. Using Eqs. (1.58, 1.59), this posterior can also be rewritten in terms of the likelihood ratio (also known as the Bayes factor) $\Lambda_{\mathcal{H}/\mathcal{N}}(d) = p(d|\mathcal{H})/p(d|\mathcal{N})$ between \mathcal{H} and \mathcal{N} :

$$p(\mathcal{H}|d) = \frac{\Lambda_{\mathcal{H}/\mathcal{N}}(d)}{\Lambda_{\mathcal{H}/\mathcal{N}}(d) + p(\mathcal{N})/p(\mathcal{H})}. \quad (1.60)$$

If $p(\mathcal{H}|d) > 1/2$, then \mathcal{H} is preferred over \mathcal{N} a posteriori, but to confidently claim the detection of a signal in the data requires a stronger threshold. Note from Eq. (1.60) that $p(\mathcal{H}|d)$ is an increasing function of the likelihood ratio $\Lambda_{\mathcal{H}/\mathcal{N}}(d)$ (as is the posterior odds). Therefore, the signal that maximizes the odds in favour of \mathcal{H} is also that which maximizes $\Lambda_{\mathcal{H}/\mathcal{N}}(d)$, meaning one can equivalently threshold the likelihood ratio.

The central limit theorem implies that the many stochastic noise processes present in the GW detectors [369, 370, 371] result in an overall noise distribution that is approximately Gaussian. Assuming also that the noise is stationary (i.e., that the noise distribution does not change over time) with zero mean $\langle n(t) \rangle = 0$, the likelihood of producing data containing only noise is given by [367]

$$p(d|\mathcal{N}) \propto \exp\left(-\frac{\|d\|^2}{2}\right), \quad (1.61)$$

where, since the detector data is a real function of time, one can define an inner product

$$\langle x, y \rangle = 2 \int \frac{\tilde{x}(f)\tilde{y}^*(f)}{S_n(f)} df, \quad (1.62)$$

$\tilde{x}(f) = \int x(t)e^{-2\pi ift} dt$ denotes the Fourier transform of a function $x(t)$ and $\tilde{x}^*(f)$ its complex conjugate, $S_n(f)$ is the one-sided power spectral density of the noise, and $\|\cdot\| = \sqrt{\langle \cdot, \cdot \rangle}$ is the norm induced by the inner product. If this noise model is a good description of the detector, the residuals $d - h$ in the signal hypothesis \mathcal{H} — which should just be noise — will also follow this distribution:

$$p(d|\mathcal{H}) \propto \exp\left(-\frac{1}{2}\|d - h\|^2\right). \quad (1.63)$$

By linearity of the inner product, one can now write the likelihood ratio as

$$\log \Lambda_{\mathcal{H}/\mathcal{N}} = -\frac{1}{2} (\langle d - h, d - h \rangle - \langle d, d \rangle) = \langle d, h \rangle - \frac{\|h\|^2}{2}. \quad (1.64)$$

Only the term on the left of Eq. (1.64) depends on the data and it is monotonic with the likelihood ratio. One can therefore use $\langle d, h \rangle$ — the matched-filter signal-to-noise ratio (SNR) — together with a selected threshold value as the detection statistic — the matched filter. The other term, $\langle h, h \rangle$, is the optimal SNR, which is the variance of the matched filter [372]. Though not explicitly included above, it can be seen that the matched-filter SNR is a function of time. Let t represent a reference time for the signal h . Shifting in time with $h_t(t') = h(t' - t)$, the Fourier transform is $\tilde{h}_t(f) = e^{-2\pi ift}\tilde{h}(f)$, and therefore the time-series matched-filter SNR is

$$\langle d, h_t \rangle = 2 \int e^{2\pi ift} \frac{\tilde{d}(f)\tilde{h}^*(f)}{S_n(f)} df. \quad (1.65)$$

	GW signal	Noise transient
Trigger	True positive	False positive (false alarm)
No trigger	False negative	True negative

Table 1.1: Classification of triggers in GW data according to whether the data truly contain a signal of astrophysical origin or not.

A trigger corresponds to a peak in this time series. By maximizing over time, one can search for above-threshold triggers and find the trigger time.

As the morphology of the signal h depends on the source properties, a broad range of templates covering the parameter space must be compared to the data [373]. Matched filtering over large banks of ($\sim 10^5$ – 10^6) precomputed signal templates [374, 375, 376] can be performed to rapidly identify triggers and their approximate source properties, while unmodelled searches are agnostic to the particular signal model [377, 378].

The noise model chosen in Eq. (1.63) reveals that \mathcal{H} actually represents the hypothesis that the data contains non-Gaussian features. In reality, such features may not result from astrophysical sources (i.e., false positive detections, or false alarms); see Table 1.1 for the different trigger possibilities. Significant sources of non-Gaussianities in the data are glitches — instrumental artefacts that can mimic the chirping waveform of a compact-binary coalescence. These signals must be identified (by, e.g., employing environmental and instrumental sensors at vetoes) and removed from the data to prevent contamination of the GW event candidates [379]. The confidence that a candidate is of astrophysical origin is increased if multiple detectors independently identify a trigger that is coincident in time. These considerations can be folded into the detection statistic, meaning the presentation here is representative but simplified compared to full search pipelines [64, 363].

1.4.2 Bayesian parameter estimation

Upon identifying a trigger and validating it as an astrophysical candidate event, the next step is to fully characterize the source. The signal h is determined by source properties θ . For quasi-circular binary BHs, θ has 15 parameters: intrinsic to the source are its masses, m_1 and m_2 , and spins, \mathbf{S}_1 and \mathbf{S}_2 (eight parameters); extrinsic to the source (i.e., depending on the observer) are its luminosity distance, right ascension, and declination that determine the location, its orientation according to the orbital inclination and GW polarization, and the reference time and phase of coalescence (seven parameters). Eccentric sources introduce two additional parameters (the magnitude and argument of periapsis), but GW emission rapidly circularizes eccentric binaries [255] and so they are not considered here.

The intrinsic binary properties determine the waveform morphology in the frame of the source while the extrinsic properties determine how the signal is projected onto the observer as measured in the detector frame. As examples, consider the effect of distance and inclination. GWs from compact-binary inspirals are primarily beamed perpendicular to the orbital plane, such that an observer whose line of sight is aligned with this direction will measure a stronger signal than one who views the source edge on to the orbital plane. Similarly, sources closer to the observer appear ‘louder’. GWs are also subject to the influence of cosmological redshift. Detectors on Earth will measure BH masses larger than the BHs do themselves by a factor $1 + z$, where $z = z(D)$ is the redshift of a source with luminosity distance D . This means that GW observers do not have direct access to the source-frame masses and redshift, and instead must assume a cosmological model (whose parameters, e.g., the Hubble constant, are in practice determined by independent observations [380, 381, 382]) to convert from the

luminosity distance and redshifted masses observed in the detector frame.

Just as approximate solutions to the relativistic binary dynamics can be constructed (as discussed in Sec. 1.3), so too can models of the GW emission (see, e.g., Refs. [383, 384, 385]). These waveform models, which are functions of the source properties θ , can be used to measure those properties by comparing the signal predicted by the model to the observed data, much as in Sec. 1.4.1.

Consider the joint posterior of the hypothesis \mathcal{H} and the source properties θ , $p(\theta, \mathcal{H}|d) = p(d|\theta, \mathcal{H})p(\theta|\mathcal{H})p(\mathcal{H})/p(d)$. \mathcal{H} now represents not just the proposition that the data d contain an astrophysical signal but also folds in all the other implicit assumptions that go into producing the Bayesian posterior, such as the waveform and prior model used. The likelihood

$$p(d|\theta, \mathcal{H}) \propto \exp\left(-\frac{1}{2}\|d - h(\theta)\|^2\right) \quad (1.66)$$

is in fact the same as Eq. (1.63), except that the dependence of the GW signal h on the source properties θ is made explicit. One can see that the posterior in Eq. (1.58) is the marginalization of $p(\theta, \mathcal{H}|d)$, using

$$p(d|\mathcal{H}) = \int p(d|\theta, \mathcal{H})p(\theta|\mathcal{H})d\theta = \langle p(d|\theta, \mathcal{H}) \rangle_{\theta \sim p(\theta|\mathcal{H})}, \quad (1.67)$$

where $x \sim p(x)$ denotes sampling a random variable x from the distribution $p(x)$ and $\langle f(x) \rangle_{x \sim p(x)}$ is the expectation value of a function $f(x)$ over this distribution. In Sec. 1.4.1 a maximum-likelihood approximation via point estimates of this integral was used to facilitate rapid searches, i.e., Eq. (1.63) is really $p(d|\mathcal{H}) = p(d|\theta = \theta_0, \mathcal{H})$ from Eq. (1.66), where θ_0 are the parameters used in the signal template. The crucial difference between searches and parameter estimation (PE) that the latter requires

inference of full posterior for θ .

The joint posterior may also be written as $p(\theta, \mathcal{H}|d) = p(\mathcal{H}|d)p(\theta|d, \mathcal{H})$; in words, one first identifies an event candidate with the techniques described in Sec. 1.4.1 — $p(\mathcal{H}|d)$ — before inferring its source properties assuming that the signal is of astrophysical origin — $p(\theta|d, \mathcal{H})$. This posterior is

$$p(\theta|d, \mathcal{H}) = \frac{p(d|\theta, \mathcal{H})p(\theta|\mathcal{H})}{p(d|\mathcal{H})}. \quad (1.68)$$

The marginal likelihood $p(d|\mathcal{H})$ is the evidence for the data containing an astrophysical signal under the assumptions about the waveform, prior, and any other analysis methods. $p(\theta|\mathcal{H})$ is the chosen prior for the source properties — in GW data analysis this is usually set to be an ‘uninformative’ prior, being taken as uniform in a suitable set of parameters (which may not be the parameters enumerated above but transformations of them). The notion of an uninformative prior is perhaps misleading, however, since the posterior is always influenced by the choice of prior.

The typical result of a GW PE run is a discrete set of samples — representing the uncertain measurement of the source properties — drawn from the Bayesian posterior of Eq. (1.68) [72, 73, 74, 75] using stochastic samplers [386, 387, 388, 389]. These may or may not also estimate the evidence in Eq. (1.67), a quantity which is not important for inference — where it acts only to properly normalize the posterior distribution — but is crucial for Bayesian model comparison — e.g., if one were to use the posterior odds to infer which of two different waveform models or prior distributions is preferred in describing some GW data. This sampling problem is difficult with traditional methods [390, 391, 392] due to the high dimensionality, and each run can take weeks to reach sufficient convergence. Therefore, recent attempts have sought efficient single-source

GW inference using hardware acceleration and machine learning [393, 394, 395, 396, 397, 398, 399, 400].

1.4.3 Hierarchical population models

At this point, I have covered how to (i) search for signals and identify triggers in GW data (Sec. 1.4.1) and (ii) make a Bayesian measurement for the source properties of event candidates (Sec. 1.4.2). I have also discussed several pathways for the formation of merging binary BHs detectable with GW observatories (Sec. 1.2.3) and that the properties of GW progenitors are connected to GW observables (Sec. 1.3). A natural question to ask, then, is:

“Can one constrain the formation pathways of merging binary BHs with GW data?”

The statistical answer is provided by a hierarchical Bayesian analysis, in which one aims to infer not just the properties of a single event based on observed data, but those of an astrophysical population from which a catalogue of multiple events has been detected [362, 401, 402, 403]. Since the population of merging binary BHs is determined by astrophysical formation, measuring the properties of this population through GW data may reveal identifiers of the underlying astrophysical processes that one does not have direct knowledge of.

Bayesian hierarchical models

The parameters θ of GW sources — such as binary component masses — are drawn from an unknown global distribution — such as a power law — determined by population-level parameters λ — such as the power-law slope — of an astrophys-

ical model \mathcal{A} : $p(\theta|\lambda, \mathcal{A}) \propto \theta^{-\lambda}$. These models can be informative about any prior astrophysical knowledge, but often in practice are taken to be compositions of simple functional forms (power laws, Gaussians, etc.; see, e.g., Refs. [83, 84, 85]). Much of the work in the later Chapters is focused on introducing astrophysically-motivated priors.

The purpose of population inference is to measure both the single-source parameters θ and the global population-level parameters λ . Given a catalogue of GW data $\{d\} = \{d_i\}_{i=1}^n$, where n is the number of observations in the catalogue that have source parameters $\{\theta_i\}_{i=1}^n$, the posterior measurement required is $p(\{\theta\}, \lambda|\{d\}, \mathcal{A}, \mathcal{H})$.

In the following, I will forgo the explicit inclusion of the astrophysical model conditional \mathcal{A} to prevent unnecessarily lengthy expressions; its presence is anyway implied by the population parameters λ . I will do the same for the signal hypothesis \mathcal{H} as it is implicitly assumed that all data under consideration contain an astrophysical GW signal.

The differential rate of GW events is

$$\frac{dN}{d\theta} = \bar{N}p(\theta|\lambda), \quad (1.69)$$

where \bar{N} is the expected number of events over the observational duration T under consideration and N is the actual number; while $p(\theta|\lambda)$ determines the shape of the population, \bar{N} determines the overall scale. A useful quantity often quoted in the literature is the volumetric merger rate — i.e., the number of mergers per unit comoving volume V_c and per unit time t_s in the frame of the source. This can be computed from the total integrated number as

$$R = \frac{d^2N}{dV_c dt_s} = \frac{dt}{dt_s} \frac{dz}{dV_c} \frac{d^2N}{dz dt} = \frac{1+z}{T} \left(\frac{dV_c}{dz} \right)^{-1} \frac{dN}{dz}, \quad (1.70)$$

where a uniform rate in the source frame is redshifted to the observer and the comoving volume has been converted to redshift.

Full hierarchical posterior

The arrival of all signals — not just those which are successfully detected — is modelled as an inhomogeneous Poisson process [402, 404]. The total number of signals, N , can be split into n data–parameter pairs $\{d_i, \theta_i\}_{i=1}^n$ that are actually observed and \varkappa pairs $\{\mathcal{d}_j, \theta_j\}_{j=1}^{\varkappa}$ that contained a signal but were interpreted as just noise (false-negative detections), i.e., $N = n + \varkappa$. The strikeout notation here indicates non-detection as opposed to successful detection. The joint likelihood of all data and parameters given the population is

$$p(\{d, \theta\}, \{\mathcal{d}, \theta\}, N | \lambda, \bar{N}) = \frac{\bar{N}^N e^{-\bar{N}}}{N!} p(\{d\}, \{\mathcal{d}\} | \{\theta\}, \{\theta\}) p(\{\theta\}, \{\theta\} | \lambda), \quad (1.71)$$

where the likelihood of the data depends only on the source parameters, as in Sec. 1.4.2. Assuming each signal is statistically independent and does not overlap with another data segment, this separates as

$$p(\{d, \theta\}, n, \{\mathcal{d}, \theta\}, \varkappa | \lambda, \bar{N}) = \frac{e^{-\bar{N}}}{\varkappa!} \left[\prod_{i=1}^n \bar{N} p(d_i | \theta_i) p(\theta_i | \lambda) \right] \left[\prod_{j=1}^{\varkappa} \bar{N} p(\mathcal{d}_j | \theta_j) p(\theta_j | \lambda) \right]. \quad (1.72)$$

While the detected signals can be labelled (e.g., by time of arrival), the non-detected signals are not distinguishable and are thus penalized by the over-counting factor $\varkappa!$.

To condition only on the detected astrophysical signals, one must marginalize over the \varkappa non-detections, the actual number of which is inherently unknown. The expected number of signals can also be split as $\bar{N} = \bar{n} + \bar{\varkappa}$, where \bar{n} ($\bar{\varkappa}$) is the expected number

of (non-) detections. Marginalizing over non-detected signals gives

$$p(\{\theta\}, \{d\}, n | \lambda, \bar{N}) = e^{-\bar{N}} \left[\prod_{i=1}^n \bar{N} p(d_i | \theta_i) p(\theta_i | \lambda) \right] \sum_{\bar{n}=0}^{\infty} \frac{\bar{n}^n}{\bar{n}!} \quad (1.73)$$

$$= e^{-\bar{N} p(\text{det} | \lambda)} \prod_{i=1}^n \bar{N} p(d_i | \theta_i) p(\theta_i | \lambda), \quad (1.74)$$

where $\int p(d | \theta) p(\theta | \lambda) d d d \theta = \bar{n} / \bar{N}$, $e^{-\bar{n}} \sum_{\bar{n}=0}^{\infty} \bar{n}^n / \bar{n}! = \sum_{\bar{n}=0}^{\infty} p(\bar{n} | \bar{n}) = 1$ (by definition), and $\bar{n} = \bar{N} p(\text{det} | \lambda)$ is given by the fraction of detectable sources in the population

$$p(\text{det} | \lambda) = \int p(\text{det} | \theta) p(\theta | \lambda) d \theta. \quad (1.75)$$

This term, also referred to as the selection function $\sigma(\lambda) = p(\text{det} | \lambda)$, encodes the observational biases of the detectors. Without its inclusion, one would only be able to infer the detectable population of sources available to the GW observatories, as opposed to the intrinsic population resulting from the astrophysical formation of binary BHs. The probability of detecting a single source θ is found by integrating the likelihood

$$p(\text{det} | \theta) = \int p(\text{det} | d) p(d | \theta) d d, \quad (1.76)$$

where detectability $p(\text{det} | d)$ is usually taken as a threshold classification (see Sec. 1.4.1).

Summing up, the goal is to make measurements — i.e., compute the Bayesian posterior — for $\{\theta\}$, λ , and \bar{N} from the data $\{d\}$. Using the law of total probability, one can rewrite Eq. (1.74) as

$$p(\{\theta\}, \lambda, \bar{N} | \{d\}, n) = \frac{p(\lambda, \bar{N})}{p(\{d\}, n)} p(\{\theta\}, \{d\}, n | \lambda, \bar{N}), \quad (1.77)$$

where $p(\lambda, \bar{N})$ is the prior on the shape and scale parameters of the population, which

are usually taken to be independent of each other, i.e., $p(\lambda, \bar{N}) = p(\lambda)p(\bar{N})$. The explicit conditional dependence on the number of observations n is often excluded as it is already implied by the number of events in the catalogue, i.e., by the fact that there are n lots of data $\{d\} = \{d_i\}_{i=1}^n$. With this in mind, one arrives at the hierarchical GW population posterior often presented in the literature [362, 402, 403, 85]:

$$p(\{\theta\}, \lambda, \bar{N}|\{d\}) = \frac{p(\lambda)p(\bar{N})}{p(\{d\})} e^{-\bar{N}p(\text{det}|\lambda)} \prod_{i=1}^n \bar{N} p(d_i|\theta_i) p(\theta_i|\lambda). \quad (1.78)$$

In the derivation presented here I have neglected the inclusion of false alarms — the false-positive detections in which a noise transient tricks us into believing there is an astrophysical GW signal present in the data (see Table 1.1). The justification for this is that the effect of noise transients can be mitigated by data cleaning and subtraction, and by imposing a false-alarm threshold on events that are included in the population catalogue [85]. I instead derived the Bayesian posterior inferred from the true-positive observations while marginalizing over the false-negative ones (the fourth case of true negatives is a trivial marginalization). Ref. [405] shows how to account for the sub-threshold false-negative events and, on the other hand, Ref. [406] shows how to include the false-positive events (false alarms).

Population-only posterior

Computing or sampling from the posterior of Eq. (1.78) is, in general, very difficult. Since each set of parameters θ has 15 dimensions for quasi-circular binary BH mergers, there are at least $15n$ parameters in total, plus the scale parameter \bar{N} and those corresponding to the population shape λ . Current catalogues contain ~ 100 GW events [64, 85, 68], meaning this posterior contains $> 10^3$ parameters.

Rather than tackling the problem of an entire catalogue all at once, if only the population parameters are of interest then the assumed independence of events can be used to marginalize over individual source properties:

$$p(\lambda, \bar{N}|\{d\}) = \int p(\{\theta\}, \lambda, \bar{N}|\{d\}) d\{\theta\} \quad (1.79)$$

$$= \frac{p(\lambda)p(\bar{N})}{p(\{d\})} e^{-\bar{N}p(\det|\lambda)} \bar{N}^n \prod_{i=1}^n \int p(d_i|\theta_i)p(\theta_i|\lambda) d\theta_i. \quad (1.80)$$

The strategy will now be to estimate the parameters of each event individually, using the posterior of Eq. (1.68) to write

$$p(\lambda, \bar{N}|\{d\}) = \frac{p(\lambda)p(\bar{N})}{p(\{d\})} e^{-\bar{N}p(\det|\lambda)} \bar{N}^n \prod_{i=1}^n \int p(d_i|\mathcal{H})p(\theta_i|d_i, \mathcal{H}) \frac{p(\theta_i|\lambda)}{p(\theta_i|\mathcal{H})} d\theta_i \quad (1.81)$$

$$\propto p(\lambda)p(\bar{N}) e^{-\bar{N}p(\det|\lambda)} \bar{N}^n \prod_{i=1}^n \left\langle \frac{p(\theta_i|\lambda)}{p(\theta_i|\mathcal{H})} \right\rangle_{\theta_i \sim p(\theta_i|d_i, \mathcal{H})}, \quad (1.82)$$

where I have reintroduced the notation \mathcal{H} to make explicit the different prior dependence on the PE analysis as distinct from the population analysis.

Each integral can now be computed as the expectation value of the population-to-PE prior ratio over the event's posterior. This is far more efficient than computing the expectation value $\langle p(d_i|\theta_i) \rangle_{\theta_i \sim p(\theta_i|\lambda)}$ — evaluating the likelihood in Eq. (1.66) requires computationally costly Fourier transforms and further integrals. Recall that the outcome of each single-event PE run is a discrete set of samples drawn from the posterior. Therefore, the posterior samples can be reused to approximate $p(d_i|\lambda)$ with tractable Monte-Carlo averages.

One may be tempted to write $\prod_{i=1}^n p(d_i|\mathcal{H}) = p(\{d\}|\mathcal{A})$ since the data are assumed to be independent. However, one must again remember that the total evidence in the denominator of Eq. (1.81) is conditional on the population whilst the event evi-

dences in the numerator are not. This distinction is usually unimportant in practice as these terms are constant with respect to the posterior parameters — indicated by the proportionality in Eq. (1.82).

Shape-only posterior

As a final simplification, one might be interested in inferring only the shape of the population and not the rate. To do this one can marginalize over the scale parameter \bar{N} ,

$$p(\lambda|\{d\}) = \left[\frac{p(\lambda)}{p(\{d\})} \prod_{i=1}^n \int p(d_i|\theta_i)p(\theta_i|\lambda) d\theta_i \right] \int p(\bar{N})e^{-\bar{N}p(\det|\lambda)}\bar{N}^n d\bar{N}, \quad (1.83)$$

which requires the choice of a suitable prior (note that here the marginalization is an integral rather than a summation because \bar{N} is the expectation value for the number of mergers, rather than the actual realization). The integral can be computed in closed-form in the simple case of the reciprocal prior $p(\bar{N}) \propto 1/\bar{N}$ [407, 402, 403]:

$$\int e^{-\bar{N}p(\det|\lambda)}\bar{N}^{n-1}d\bar{N} = p(\det|\lambda)^{-n} \int e^{-\bar{n}}\bar{n}^{n-1}d\bar{n} = \frac{\Gamma(n-1)}{p(\det|\lambda)^n}, \quad (1.84)$$

where Γ is the gamma function. Substituting this result into Eq. (1.83) yields

$$p(\lambda|\{d\}) \propto p(\lambda) \prod_{i=1}^n \frac{\int p(d_i|\theta_i)p(\theta_i|\lambda)d\theta_i}{\int p(\det|\theta)p(\theta|\lambda)d\theta} = p(\lambda) \prod_{i=1}^n \frac{p(d_i|\lambda)}{p(\det|\lambda)}. \quad (1.85)$$

With the hierarchical posterior in hand, one can go back to the chosen population prior and view the constraints on the population distribution of binary BH parameters. A set of posterior samples $\{\lambda\} \sim p(\lambda|\{d\})$ implies a posterior of distributions $\{p(\theta|\lambda)\}$, one for each λ . Uncertainty in the population constraint can be visualized with, e.g.,

credible intervals over $\{p(\theta|\lambda)\}$. Maximum-likelihood or posterior-median values for λ can be used to give likely representations for the underlying population. Another common distribution used is the posterior mean, also called the posterior population distribution (PPD), given by

$$p(\theta|\{d\}) = \int p(\theta|\lambda)p(\lambda|\{d\})d\lambda = \langle p(\theta|\lambda) \rangle_{\lambda \sim p(\lambda|\{d\})}. \quad (1.86)$$

I will present examples of these distributions — measured using real GW data — in Chapters 3, 4, and 5.

Parameter transformations

There may be situations in which only a subset of all the binary parameters θ are modelled at the level of the population. For example, one may be interested in constraining the mass function of binary BH mergers but uninterested in their distribution across the sky. Suppose a population model describes a subset φ of $\theta = (\varphi, \psi)$, with the remaining ψ being nuisance parameters. This is tantamount to assuming φ and ψ are independent in the population, where the prior on ψ matches that of the original PE prior; i.e.,

$$\frac{p(\theta|\lambda)}{p(\theta|\mathcal{H})} = \frac{p(\varphi|\lambda)}{p(\varphi|\psi, \mathcal{H})}. \quad (1.87)$$

Any dependency between φ and ψ in the PE prior must still be accounted for; if they are independent then $p(\theta|\lambda)/p(\theta|\mathcal{H}) = p(\varphi|\lambda)/p(\varphi|\mathcal{H})$, so one may replace the full set θ with the modelled parameters φ .

Similarly, instead of defining a population model in terms of the source properties inferred from the PE runs, it might be defined in terms of others parameters that

are more astrophysically relevant. For example, while the detectors are sensitive to the redshifted masses, from an astrophysical perspective the source-frame masses are more relevant. Under such a parameter transformation $\theta' = \theta'(\theta)$, a PDF transforms as $p(\theta') = p(\theta)|d\theta/d\theta'|$, where (abusing notation) the scaling is given by the Jacobian determinant of the transformation.

Computing selection effects

Just as evaluating the single-event likelihoods $p(d_i|\theta_i)$ makes averaging $\langle p(d_i|\theta_i) \rangle_{\theta_i \sim p(\theta_i|\lambda)}$ an expensive operation, so too does $p(\text{det}|\theta)$ for $p(\text{det}|\lambda) = \langle p(\text{det}|\theta) \rangle_{\theta \sim p(\theta|\lambda)}$. Computing the detection probability $p(\text{det}|\theta)$ for a single source with parameters θ requires injecting it into the entire search pipeline to test whether it is detectable. One would have to undertake this task for many sources drawn from the population $p(\theta|\lambda)$ and for many populations λ drawn from $p(\lambda)$; this is completely infeasible for stochastic sampling of the posterior. Instead, one can write

$$p(\text{det}|\lambda) = \int p(\text{det}|\theta)p(\theta|\lambda)d\theta = \left\langle p(\text{det}|\theta) \frac{p(\theta|\lambda)}{p(\theta|\lambda_0)} \right\rangle_{\theta \sim p(\theta|\lambda_0)}, \quad (1.88)$$

where I have multiplied and divided through by a reference population with parameters λ_0 . With this approach, pipeline injections can be performed just once ahead of time to compute $p(\text{det}|\theta)$ for sources drawn from $p(\theta|\lambda_0)$ and then apply the appropriate weighting factors from Eq. (1.88) to approximate $p(\text{det}|\lambda)$ during inference [408, 409].

There do exist approximations to speed up computation of selection effects. Recent approaches have utilized machine learning [410, 411, 412, 413, 414]. Another example that already introduced is the SNR threshold; rather than injecting a source into the full detection pipeline, one can compute its SNR, $\rho = \langle h, h \rangle$, and classify it as detectable

if it exceeds the threshold ρ_{det} , typically taken equal to 8 for a single detector and 12 for a network [415, 145].

However, the SNR depends on all of the source properties — both intrinsic (say, φ) and extrinsic (say, ψ) — but one might only be interested in modelling the population of the former. The detection probability can be marginalized to partially alleviate this dependency in the specific case of non-spinning and aligned-spin binaries in an idealized network of identical detectors [416]. In this simplified case, the GW strain can be written $h(t) = \omega A(t) \cos[\phi(t) - \phi_0]$. The GW amplitude $A(t)$ and phase $\phi(t)$ are predicted by the chosen waveform model, while $0 \leq \omega \leq 1$ and ϕ_0 depend only on extrinsic parameters. The maximum SNR ρ_{max} for given intrinsic parameters and distance is well defined and occurs for face-on sources directly overhead the detector. The SNR can therefore be scaled down from the maximum, i.e., $\rho = \omega \rho_{\text{max}}$, such that the detection probability averaged over an extrinsic population independent of the intrinsic parameters is

$$p(\text{det}|\varphi, \lambda) = \int p(\text{det}, \psi|\varphi, \lambda) d\psi = \int p(\text{det}|\theta) p(\psi|\lambda) d\psi = \int_{\frac{\rho_{\text{det}}}{\rho_{\text{max}}}}^1 p(\omega|\lambda) d\omega, \quad (1.89)$$

where $p(\omega|\lambda)$ is the distribution of ω implied by the choice of $p(\psi|\lambda) = p(\psi|\varphi, \lambda)$, often taken to be isotropic over the sky. The advantage of this approach is that the averaged detectability $p(\text{det}|\varphi, \lambda)$ requires only a single SNR evaluation, while the one-dimensional cumulative distribution on the right of Eq. (1.89) can be easily interpolated [417, 418].

One final subtlety to consider is that detector sensitivities change over time. Though the assumption of a constant sensitivity is reasonable within a given observing run, detector upgrades lead to improvements between runs [145] and thus higher detection

probabilities. Assuming that the rate of arrival of signals does not change over time, one can write the population detection fraction — accounting explicitly for the time dependence of detectability — as

$$p(\text{det}|\lambda) = \int p(\text{det}|\theta, t)p(\theta|\lambda)p(t)dt d\theta = \sum_r \frac{T_r}{T} \int p(\text{det}|\theta, r)p(\theta|\lambda)d\theta, \quad (1.90)$$

where r denotes the disjoint observing runs, T_r their lengths with $\sum_r T_r = T$, and $p(\text{det}|\theta, r)$ signifies that the detectability of the same source θ can change depending on which observing run it was detected in.

While these last sections covered some of the more technical details of implementing a GW population analysis, they are important to get right to ensure accurate inference.

1.5 Main findings

With all this in hand, the complete process of modelling the formation of binary BHs, following the relativistic dynamics of their inspirals toward merger, searching for their signals with GW detectors, characterizing their source properties with Bayesian inference, and finally closing the loop by constraining the astrophysical population of GW sources is available. The remaining Chapters of this thesis are made up of my first-author publications to date and make significant use of the theoretical results presented in Sections 1.3 and 1.4 of this Chapter.

In Chapter 2 I present results from Mould and Gerosa (2020) [419], where the influence of relativistic binary BH dynamics on GW observables are studied. In particular, even if a binary forms with aligned BH spins, due to a dynamical instability it can still undergo spin precession. Furthermore, this instability directs sources to a specific endpoint in the parameter space at small orbital separations, which are derived in closed

form. This prediction provides a smoking-gun signature of a peculiar consequence of GR on two-body dynamics that can be tested with GW signals. A summary of a series of works on this topic was produced by Mould (2021) [420].

In the reverse direction, the results of Mould and Gerosa (2022) [421] presented in Chapter 3 show how GW observables can be propagated from detection to astrophysical references. The frameworks from Sections 1.3 and 1.4 are connected to constrain the pre-inspiral binary BH population directly from observations of GW mergers, thus bridging the two extrema of binary evolution. Additionally, it is shown that BH spins have a significant influence on population inference through selection effects; indeed, LVC population studies now account for spin precession [85] where they previously neglected it [84]. They also provide posterior measurements propagated to infinite orbital separations [63, 64, 74, 75], as suggested in my work.

The effect of astrophysical formation on GW populations is further investigated in Chapter 4. There, the work of Mould *et al.* (2022) [422] focuses on the influence of mass-transfer processes in the isolated binary BH formation channel. Simple arguments are used to show that sufficient mass transfer during the co-evolution of two massive stars leaves an identifiable imprint on the resulting binary BH merger. Targeted population models that reveal there is no strong evidence for this process — which, if efficient, should result in a large sub-population of such sources — using current GW catalogues. Our results also show significant support in the population for BHs born with non-zero spins.

As laid out in Chapter 5, Mould *et al.* (2022) [413] develop a novel approach to GW population analysis using simulation-based inference. Simple parametric population models are replaced with astrophysics-informed models constructed with deep-learning emulators of population-synthesis simulations. Rather than inferring posteriors of ar-

bitrary shape parameters, key astrophysical parameters can now be measured directly from GW data. Unlike simplified models, simulation-based priors account for correlations between the binary properties (such as masses and spins) — a crucial ingredient to capture for realistic populations. Our application shows that features in the intrinsic distribution of binary BHs, such as surprisingly large masses, can be explained by hierarchical mergers where (at least) one of the BHs as observed in a GW signal is the remnant of a previous merger (or multiple previous mergers).

Finally, Chapter 6 wraps up this Thesis with some concluding remarks.

Not included in this Thesis are several publications on which I appear as second author or lower, some of which I made significant contributions to. [Reali *et al.* \(2020\) \[423\]](#) study the correlations between asymptotic-binary and merger-remnant source properties. The existence of the precessional instability discussed in Chapter 2 was verified in the strong-field regime of full NR simulations by [Varma *et al.* \(2021\) \[424\]](#) and the detectability of its endpoint with GW observations was assessed by [De Renzi *et al.* \(2023\) \[425\]](#). Detection of higher-order spin effects with a metric — developed by [Gerosa *et al.* \(2021\) \[426\]](#) — that consistently includes all spin degrees of freedom was investigated by [De Renzi *et al.* \(2022\) \[427\]](#). An updated parametrization for binary BH spin precession was developed by [Gerosa *et al.* \(2023\) \[355\]](#). [Killestein *et al.* \(2023\) \[95\]](#) placed precise observational constraints on the orbital properties of Sco X-1, a promising candidate of continuous GW emission [94]. My publication list also includes Refs. [428, 429].

As part of these works I developed or contributed to several public codes [430, 431, 432, 433]. I presented my research at a joint total of over 25 conferences or invited seminars. I (co-) supervised three undergraduate students on summer research projects that have now either started graduate studies or will embark on further re-

search projects within GW astronomy before graduating. Beyond research, I took part in the ‘Astronomy in the City’ outreach activities within the astrophysics group, acted as a representative of the Birmingham Environment for Academic Research (BEAR), and organized a local BEAR conference for postgraduate research students.

Endpoint of the up–down spin-precession instability

Abstract

We investigate the stability of equilibrium solutions of the binary BH spin precession equations. We show that small perturbations evolve under simple-harmonic motion and confirm the existence of a dynamical instability in up–down sources, where the spin of the heavier (lighter) BH is initially (anti-) aligned to the orbital angular momentum. Numerical evolutions surprisingly reveal that unstable binaries do not disperse in the available parameter space, but rather converge to a well-defined endpoint. We derive this endpoint in closed form as a special case of solutions locked into spin–orbit resonances; specifically, unstable up–down binaries approach merger with two equally misaligned, co-linear, and co-precessing BH spins. Applying our results to a simple astrophysically-motivated population of spinning binary BHs, we find that a significant fraction of sources become unstable prior to entering the sensitive band of current ground-based GW detectors. Detection of the up–down endpoint with GWs is therefore a smoking-gun signature of this unique feature of relativistic two-body dynamics.

Summary and contributions

This Chapter is a reformatted version of our paper Ref. [419]. I completed it in collaboration with Davide Gerosa who proposed investigating the endpoint of the instability, the existence of which he originally derived [434]. The orbit-averaged spin precession formalism is based on some of his previous work [352, 353] and numerical implementation [430, 435]. The new derivation of the instability in terms of simple harmonic oscillators was performed by me. The derivation of the instability endpoint was a joint work. Davide constructed the proofs for the number of resonances and the fact the resonances remain resonant. I derived the polynomial expressions for the resonant locations, the asymptotic limit, and the up–down endpoint, and these calculations were also checked by Davide. I made the figures and performed the computational runs.

In Sec. 2.1 we give the background context for this Chapter. In Sec. 2.2 we prove that binary black holes in the up–down spin configuration are unstable. In Sec. 2.3 we perform new analytic post-Newtonian calculations to derive the locations of resonant binary BH spin configurations. In particular, we use this formalism to find the formal endpoint of the up–down spin precession instability, which we give in Sec. 2.4. We also study the effect of the up–down instability on astrophysical populations of binary BHs. We conclude in Sec. 2.5 and point to later NR and PE verifications of our predictions.

2.1 Introduction

There are four distinct configurations in which the BH spins are aligned to the orbital angular momentum (see Fig. 2.1). We label each of these cases “up–up”, “down–down”, “down–up” and “up–down”, where “up” (“down”) refers to co- (counter-)

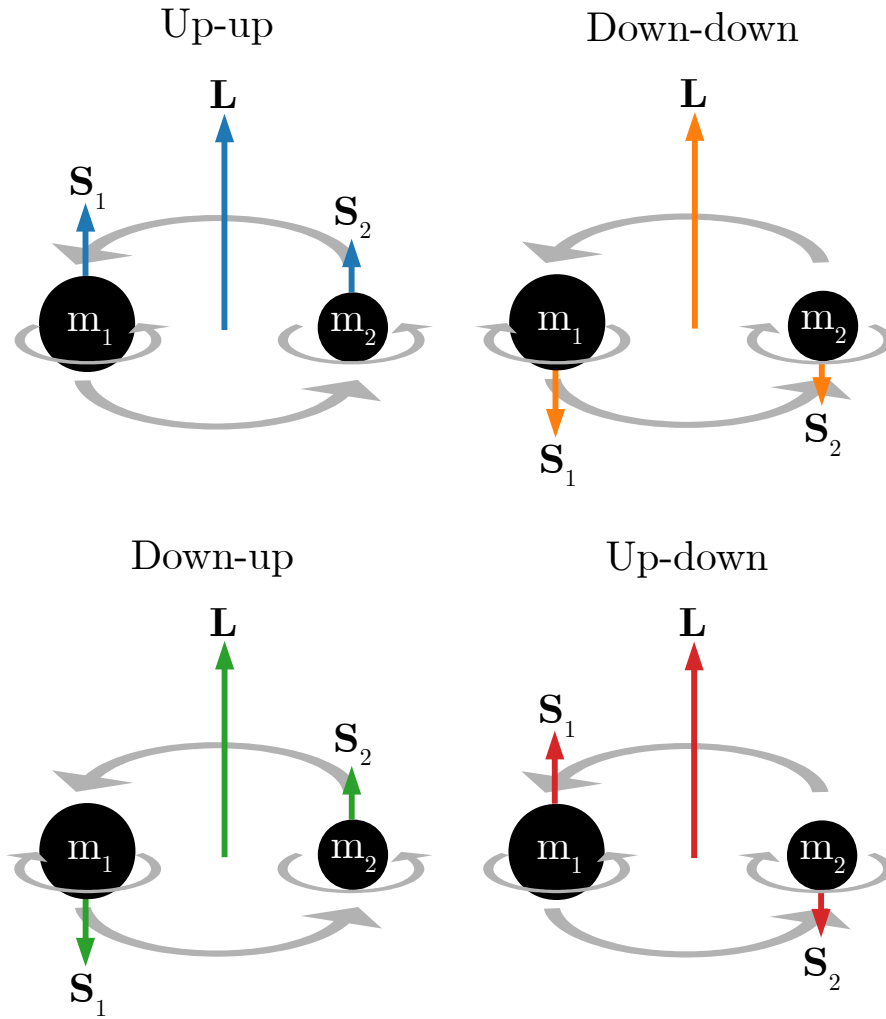


Figure 2.1: The four binary BH configurations with aligned spins. The BH with higher (lower) mass is indexed by the number 1 (2). We refer to the orientation of a BH whose spin vector \mathbf{S}_i is parallel (antiparallel) to the orbital angular momentum vector \mathbf{L} as “up” (“down”). The four distinct binary configurations are then labelled with the orientation of the primary (secondary) BH appearing before (after) the hyphen.

alignment with the orbital angular momentum and the label before (after) the hyphen refers to the spin alignment of the primary (secondary) BH. It is straightforward to show that all four of these configurations are equilibrium, non-precessing solutions of the relativistic spin-precession equations (1.24–1.26) [350] — a BH binary initialized in exactly one of these configurations remains so over its inspiral. Here, we tackle their stability: if an arbitrarily small misalignment is present, how do such configurations behave?

Employing the parametrization of generic spin precession in terms of an effective potential at 2PN order [352, 353], Gerosa *et al.* [434] investigated the robustness of aligned spin binary BH configurations (see also Ref. [436] for a subsequent study). They found that the up–up, down–down and down–up configurations are stable, remaining approximately aligned under a small perturbation of the spin directions. This is not the case for up–down binaries, i.e., those where the heavier BH is aligned with the orbital angular momentum while the lighter BH is anti-aligned. They report the presence of a critical orbital separation,

$$r_{\text{UD}+} = \frac{(\sqrt{\chi_1} + \sqrt{q\chi_2})^4}{(1-q)^2} M, \quad (2.1)$$

that defines the onset of the instability. An up–down binary that is formed at large orbital separations $r > r_{\text{UD}+}$ will at first inspiral much as the other stable aligned binaries do, with the spins remaining arbitrarily close to the aligned configuration. However, upon reaching the instability onset at $r = r_{\text{UD}+}$, the binary becomes unstable to spin precession, leading to large misalignments of the spins.

Figure 2.2 shows the orbit-averaged PN evolution of the spins for a binary BH in the up–down configuration. The binary is evolved from an orbital separation of




Figure 2.2: Numerical evolution of the normalized spins $\hat{\mathbf{S}}_i$ of a BH binary with mass ratio $q = 0.5$ and dimensionless spins $\chi_1 = \chi_2 = 1$. The blue (red) curve traces the path of the primary (secondary) BH spin over time. The integration is performed from a binary separation $r = 1000M$ to $10M$; the colours of the curves darken with decreasing separation. The binary is initialized with misalignments of 1° from the up-down configuration. The vertical z axis is initially aligned to the total angular momentum, the x axis is constructed such that the initial orbital angular momentum lies in the x - z plane, and the y axis completes the orthogonal frame. The black dots show the location of the spins for $r > r_{\text{UD}^+} \approx 34M$, before the onset of instability. The arrows show the orientation of the spins at the final separation $r = 10M$.

$r = 1000M > r_{\text{UD}+}$ to $r = 10M$. At the initial separation, the spin directions are perturbed such that there is a misalignment of 1° in the spins from the exact up–down configuration. The response to this perturbation is initially tight polar oscillations (black dots in Fig. 2.2) of the BH spins around the aligned configuration. After the onset of instability, precession induces large spin misalignments (coloured tracks).

A key question so far unanswered is the following: after becoming unstable, to what configuration do up–down binaries evolve? In other words, what is the endpoint of the up–down instability?

We obtain a surprisingly simple result: unstable up–down binaries tend to the very particular configuration where the two BH spins are co-aligned with each other and equally misaligned with the orbital angular momentum. Given explicitly, the endpoint of the up–down instability is a precessing configuration with

$$\hat{\mathbf{S}}_1 = \hat{\mathbf{S}}_2 \quad \text{and} \quad \hat{\mathbf{S}}_1 \cdot \hat{\mathbf{L}} = \hat{\mathbf{S}}_2 \cdot \hat{\mathbf{L}} = \frac{\chi_1 - q\chi_2}{\chi_1 + q\chi_2}. \quad (2.2)$$

In this Chapter, we refer to the effective aligned spin previously introduced with χ_{eff} in Eq. (1.33) as ξ , consistent with the notations of Refs. [353, 419, 434].

2.2 Instability threshold

2.2.1 Binary black-hole spins as harmonic oscillators

The stability of solutions is determined by their response to small perturbations. Ref. [434] indicated that the up–down instability develops on the short precessional timescale t_{pre} . In this regime, all variables are constant but S , as we saw in Sec. 1.3.3. Its evolution is determined by Eq. (1.40). Using Eqs. (1.41–1.45) we find the second

derivative $d^2 S^2/dt^2 = -A^2(3S^4 + 2BS^2 + C)/2$. Rearranging in terms of a perturbation $S^2 - S_*^2$ to some solution S_* of Eq. (1.40) gives

$$\frac{d^2}{dt^2}(S^2 - S_*^2) = -\frac{A^2}{2} [3(S^2 - S_*^2)^2 + 2(3S_*^2 + B)(S^2 - S_*^2) + 3S_*^4 + 2BS_*^2 + C] . \quad (2.3)$$

For binary configurations with the BH spins aligned with the orbital angular momentum, we may write the magnitude of the total spin as $S_* = |\alpha_1 S_1 + \alpha_2 S_2|$ where $\alpha_i = \cos \theta_{i*} = \pm 1$ discriminates between parallel ($\alpha_i = +1$) and antiparallel ($\alpha_i = -1$) alignment of \mathbf{S}_{i*} with \mathbf{L} ; for instance, up–down corresponds to $\alpha_1 = -\alpha_2 = 1$. As J and ξ are constant on t_{pre} , we have

$$J \approx J_* = |L + \alpha_1 S_1 + \alpha_2 S_2| , \quad (2.4)$$

$$\xi \approx \xi_* = \frac{1}{M^2} \left[(1+q)\alpha_1 S_1 + \left(1 + \frac{1}{q}\right)\alpha_2 S_2 \right] , \quad (2.5)$$

which implies that $3S_*^4 + 2BS_*^2 + C = 0$. Therefore, to leading order $\mathcal{O}(S^2 - S_*^2)$ in the perturbation (i.e., assuming small misalignments between the BH spins and the orbital angular momentum), the total spin magnitude S of binary BHs with nearly aligned spins satisfies

$$\frac{d^2}{dt^2}(S^2 - S_*^2) + \omega^2(S^2 - S_*^2) = 0 . \quad (2.6)$$

Equation (2.6) has the form of a simple harmonic oscillator equation, where we identify the oscillation frequency $\omega = A\sqrt{3S_*^2 + B}$.

2.2.2 Stable or unstable?

The stability of the aligned-spin configurations is determined by the sign of ω^2 :

- When $\omega^2 > 0$, Eq. (2.6) describes simple harmonic oscillations in S^2 around S_*^2 . The configuration is stable; small perturbations will cause precessional motion about the alignment.
- When $\omega^2 < 0$, the oscillation frequency becomes complex, corresponding to an instability in the precessional motion leading to large misalignments of \mathbf{S}_1 and \mathbf{S}_2 with \mathbf{L} .

The points during the evolution of the binary BH at which the precession motion transitions from stable to unstable, or vice-versa, correspond to the solutions of $\omega^2 = 0$. Since L (or equivalently r) is a monotonically decreasing function of time on the radiation-reaction timescale t_{RR} , such a point is a stable-to-unstable transition if $d\omega^2/dL > 0$ ($d\omega^2/dt < 0$) and an unstable-to-stable transition if $d\omega^2/dL < 0$ ($d\omega^2/dt > 0$).

The square of the oscillation frequency depends on L according to

$$\omega^2(L) = \left[\frac{3M^9 q^5 (1-q)}{2(1+q)^{11} L^7} \right]^2 \left(L - \frac{q\alpha_1 S_1 + \alpha_2 S_2}{1+q} \right)^2 \times \left[L^2 - 2 \frac{q\alpha_1 S_1 - \alpha_2 S_2}{1-q} L + \left(\frac{q\alpha_1 S_1 + \alpha_2 S_2}{1-q} \right)^2 \right]. \quad (2.7)$$

It is clear from Eq. (2.7) that ω^2 always has four roots, two being the repeated root

$$L_0 = \frac{q\alpha_1 S_1 + \alpha_2 S_2}{1+q}. \quad (2.8)$$

The corresponding value of the binary separation $r_0 = M^{-3} \eta^{-2} L_0^2$ always satisfies

$r_0 \leq M$ and is thus unphysical. The other two roots are

$$L_{\pm} = \frac{q\alpha_1 S_1 - \alpha_2 S_2 \pm 2\sqrt{-q\alpha_1\alpha_2 S_1 S_2}}{1 - q}. \quad (2.9)$$

For L_{\pm} to be real, we require that $\alpha_1\alpha_2 = -1$, leaving only the up–down and down–up cases. If $\alpha_1 = -\alpha_2 = -1$ (down–up), then $L_{\pm} = -(\sqrt{qS_1} \pm \sqrt{S_2})^2/(1 - q)$ which is always non-positive and can be discarded as unphysical.

The only combination of α_1 and α_2 which makes L_{\pm} both real and non-negative, thus indicating a physical precession instability, is $\alpha_1 = -\alpha_2 = 1$, which corresponds to the up–down configuration. Therefore, the up–up, down–down and down–up binary BH configurations are stable, whereas the up–down configuration can become unstable at separations where $\omega^2 < 0$. Any small misalignment of the BH spins with the orbital angular momentum leads to small oscillations of the spin vectors around the aligned configuration in the former three cases, but might cause large misalignments in the latter case.

In terms of only the parameters M , q , χ_1 and χ_2 of the BH binary, the expressions for the binary separations corresponding to the roots L_{\pm} in the case of up–down spin alignment are

$$r_{\text{UD}\pm} = \frac{(\sqrt{\chi_1} \pm \sqrt{q\chi_2})^4}{(1 - q)^2} M, \quad (2.10)$$

which are precisely those derived in Ref. [434] by other means.

The oscillation frequency of the up–down configuration is given in terms of r by

$$M^2 \omega_{\text{UD}}^2(r) = \frac{9}{4} \left(\frac{1 - q}{1 + q} \right)^2 \left(\frac{M}{r} \right)^5 \left(1 - \sqrt{\frac{r_{\text{UD}0}}{r}} \right)^2 \left(1 - \sqrt{\frac{r_{\text{UD}+}}{r}} \right) \left(1 - \sqrt{\frac{r_{\text{UD}-}}{r}} \right), \quad (2.11)$$

where $r_{\text{UD}0} = M(\chi_1 - q\chi_2)^2/(1+q)^2$ is the repeated root identified previously. One has

$$\lim_{r/M \rightarrow \infty} M^2 \omega_{\text{UD}}^2(r) = \frac{9}{4} \left(\frac{1-q}{1+q} \right)^2 \left(\frac{M}{r} \right)^5 > 0, \quad (2.12)$$

such that the up–down configuration tends to stability at large orbital separations (past time infinity). Since $r_{\text{UD}+} > r_{\text{UD}-}$, the point $r = r_{\text{UD}+}$ is a stable-to-unstable transition and $r = r_{\text{UD}-}$ is an unstable-to-stable transition. In other words, $d\omega_{\text{UD}}^2/dt|_{r_{\text{UD}+}} < 0$ and $d\omega_{\text{UD}}^2/dt|_{r_{\text{UD}-}} > 0$. The up–down configuration is unstable for orbital separations $r_{\text{UD}+} > r > r_{\text{UD}-}$. An example of the behaviour of ω^2 is given in Fig. 2.3. There and in the following we plot quantities as a function of decreasing orbital separation r , corresponding to evolutions forward in time.

In the equal-mass limit $q \rightarrow 1$, the precessional motion of up–down binaries tends to stability, since the time derivative of the total spin magnitude S vanishes [437]. In the test-particle limit $q \rightarrow 0$, the behaviour also tends to stability because $S \rightarrow S_1$ becomes constant.

For an up–down binary to undergo the precessional instability, its parameters q , χ_1 , and χ_2 must be such that the resulting instability onset satisfies $r_{\text{UD}+} > 10M$, as this threshold represents the breakdown of the PN approximation [438, 439, 440]. Figure 2.4 shows contours in the χ_1 – χ_2 plane for values of q where $r_{\text{UD}+} = 10M$. For mass ratios close to unity, binaries with smaller dimensionless spins still result in a physical ($r_{\text{UD}+} > 10M$) onset of instability. As the mass ratio becomes more extreme ($q \rightarrow 0$), only binaries with $\chi_i \approx 1$ are affected by the instability, though much later in the inspiral.




Figure 2.3: Oscillation frequencies for the four aligned configurations as a function of the binary separation r . The squared frequency ω^2 is scaled by r^5 for clarity; see Eq. (2.11). The up-down configuration (red line) shows qualitatively different behaviour to the other aligned configurations, with its oscillation frequency ω_{UD} becoming complex (i.e., $\omega_{\text{UD}}^2 < 0$) between $r_{\text{UD}\pm}$ (dashed lines). In this example, the mass ratio is $q = 0.9$, the dimensionless spins are $\chi_1 = 1.0$ and $\chi_2 = 0.1$, and the region of instability is given by $r_{\text{UD}+} \approx 290M$ and $r_{\text{UD}-} \approx 24M$.




Figure 2.4: Contours of constant mass ratio q for values of the dimensionless spins χ_i that result in an instability threshold $r_{\text{UD}+} = 10M$. Above each curve is the region of parameter space in which an up–down binary will experience the precessional instability at an orbital separation $r > 10M$. Below the curves, the instability takes place later in the inspiral where our PN approach is not valid.

2.2.3 Numerical verification of the instability

The analysis of Sec. 2.2.1 is valid up to the onset of the precessional instability at the value of the binary separation $r = r_{\text{UD}+}$, at which point spin precession invalidates the approximation of small misalignments between the BH spins and the orbital angular momentum. We therefore verify the existence of the instability with evolutions of binary BH spins performed via direct numerical integrations of the orbit-averaged spin precession equations (1.24–1.26). The integrations are performed using the PRECESSION Python package [430, 435].

The binaries are evolved from an initial separation $r = 1000M$ down to a final separation $r = 10M$. The integrations are initialized by setting θ_1 , θ_2 , and $\Delta\Phi$ (or equivalently ξ , J and S) at the initial separation. The initial value of $\Delta\Phi$ is irrelevant (for these evolutions it was set to $\pi/2$). We introduce an initial perturbation to each configuration by setting the initial values of θ_1 and θ_2 to be 5° from the aligned configuration. A number of binary BHs with varied mass ratios and dimensionless spins were evolved in this way to verify the existence of the instability. As an example, the evolution of four binaries, one in each of the aligned-spin configurations, with $q = 0.8$, $\chi_1 = 1$ and $\chi_2 = 0.5$ is displayed in Fig. 2.5.

In the exactly-aligned configurations, each of θ_1 , θ_2 , and S is constant, since such configurations are equilibrium solutions of Eqs. (1.24–1.26). In the absence of the precessional instability, a small perturbation to θ_1 or θ_2 causes small amplitude oscillations around the equilibrium solutions. For a perturbation in the angles as small as 5° , a binary acts essentially as it would in the equilibrium configurations, as seen in the first three panels of Fig. 2.5; the angles $\theta_{1,2}$ remain approximately fixed at their initial values. For the configurations in which the two BH spin vectors have the same alignment




Figure 2.5: Numerical evolutions of the total spin magnitude S and misalignment angles $\cos \theta_{1,2}$ of four binary BHs with parameters $q = 0.8$, $\chi_1 = 1.0$ and $\chi_2 = 0.5$, starting from an initial separation $r = 1000M$ and ending at $r = 10M$. Each panel shows a binary initially in a configuration with aligned spins up to a small perturbation of 5° in the angles $\theta_{1,2}$. The vertical dashed line in the right-most panel (up–down) shows the location of the instability onset $r_{\text{UD}+} \approx 180$. The horizontal dashed lines mark the formal endpoint of the up–down instability obtained in the $r \rightarrow 0$ limit (Sec. 2.4.1).

as each other with respect to \mathbf{L} (up–up and down–down), the total spin magnitude remains at the initial value $S \approx S_1 + S_2$. In the down–up configuration, the total spin magnitude remains at the initial value $S \approx |S_1 - S_2|$. However, as is clear in the rightmost panel of Fig. 2.5, in the up–down configuration the values of S and $\theta_{1,2}$ are not constant. Though initially $S \approx |S_1 - S_2|$ and $\cos \theta_1 = -\cos \theta_2 \approx 1$, after reaching the onset of the instability at $r = r_{\text{UD}+} \approx 180M$ the precessional motion moves the binary away from the initial up–down configuration.

In Fig. 2.6 we test the response of the up–down instability to the amplitude of the initial perturbation. We evolve samples of binaries from $r = 1000M$ to $r = 10M$ and show their values of S at both the initial and the final separations. Binaries are initialized by extracting the misalignments from half-Gaussian distributions in $\cos \theta_{1,2}$ with widths $1 - \cos \delta\theta$ centred on the exact up–down configuration, where $\delta\theta = 1^\circ, 5^\circ, 10^\circ, 20^\circ$. The initial value of $\Delta\Phi$ is irrelevant for the presence of the instability and is here sampled uniformly within $[-\pi, \pi]$. In this example we fix $q = 0.8$ and $\chi_1 = \chi_2 = 0.9$.

Our numerical evolutions show a somewhat surprising result: binaries do not tend to disperse in parameter space as one would expect from an instability, but present a well-defined endpoint. This effect is sharper for binaries very close to up–down. Increasing the initial misalignment $\delta\theta$ dilutes both the initial and the final spin distributions, although the same trend remains present up to $\delta\theta \lesssim 20^\circ$. Binaries that undergo the up–down instability at some large separation are likely to be found in a different, but very specific region of the parameter space at the end of the inspiral. We now aim to find this location analytically.




Figure 2.6: The response of the up–down instability to different initial perturbations. Each panel shows a set of 1000 orbit-averaged evolutions. Binaries are initialized at $r = 1000M$ with misalignments extracted from truncated Gaussians centred on the up–down configuration with widths $\delta\theta = 1^\circ, 5^\circ, 10^\circ, 20^\circ$, increasing progressively from the left to the right panel. Blue (orange) histograms show the corresponding values of the total spin S at $r = 1000M$ ($r = 10M$). In this example we fix $q = 0.8$ and $\chi_1 = \chi_2 = 0.9$. Vertical dashed lines at $S = |S_1 \pm S_2|$ mark the asymptotic locations of up–down binaries before and after the instability.

2.3 Resonant configurations

Spin-orbit resonances [351] are special configurations where the three vectors \mathbf{L} , \mathbf{S}_1 , and \mathbf{S}_2 are coplanar and jointly precess about \mathbf{J} . There are two families of resonant solutions, defined by $\Delta\Phi = 0$ and $\Delta\Phi = \pi$. The previous analysis of Ref. [434] indicated that the up-down configuration at separations $r > r_{\text{UD}+}$ ($r < r_{\text{UD}-}$) is a $\Delta\Phi = 0$ ($\Delta\Phi = \pi$) resonance. The endpoint of the up-down instability is thus deeply connected to the evolution of these special solutions. As a building block to analyse the up-down configuration, in this section we present new advances toward understanding spin-orbit resonances in a semi-analytic fashion.

2.3.1 Locating resonances

For fixed values of q , χ_1 , χ_2 , J , and L , geometrical constraints restrict the allowed values of S and ξ to [353]

$$S_{\min} \leq S \leq S_{\max}, \quad (2.13)$$

$$\xi_-(S) \leq \xi \leq \xi_+(S), \quad (2.14)$$

where

$$S_{\min} = \max(|J - L|, |S_1 - S_2|), \quad (2.15)$$

$$S_{\max} = \min(J + L, S_1 + S_2), \quad (2.16)$$

$$\begin{aligned} \xi_{\pm} = & \left\{ \frac{J^2 - L^2 - S^2}{4qM^2S^2L} [(1+q)^2S^2 - (1-q^2)(S_1^2 - S_2^2)] \right. \\ & \left. \pm (1-q^2) [J^2 - (L-S)^2]^{1/2} [(L+S)^2 - J^2]^{1/2} \right\} \end{aligned}$$

$$\times [S^2 - (S_1 - S_2)^2]^{1/2} [(S_1 + S_2)^2 - S^2]^{1/2} \}. \quad (2.17)$$

Together, the functions $\xi_{\pm}(S)$ form a closed convex loop in the S – ξ plane, which implies that the inequalities in Eqs. (2.13–2.14) can be rewritten as $S_- \leq S \leq S_+$, where S_{\pm} are the solutions of $\xi = \xi_{\pm}(S)$. One can see using Eq. (1.40) that the conditions $\xi = \xi_{\pm}(S)$ are equivalent to either alignment ($\sin \theta_i = 0$) or co-planarity ($\sin \Delta\Phi = 0$). Generic spin precession can be described as a quasi-periodic motion of S between the two solutions S_{\pm} . Spin–orbit resonances correspond to the specific case where $S_- = S_+$, i.e., $d\xi_{\pm}/dS = 0$. In this case, S is constant; the three momenta are not just co-planar, but stay co-planar on the precession timescale t_{pre} . As we will see later in Sec. 2.3.3, co-planarity is also preserved on the longer radiation-reaction timescale t_{RR} .

The conditions $\xi = \xi_{\pm}(S)$ can be squared and cast into the convenient form

$$\Sigma(S^2) = \sigma_6 S^6 + \sigma_4 S^4 + \sigma_2 S^2 + \sigma_0 = 0, \quad (2.18)$$

where the coefficients σ_i are real multiples of the coefficients in Eqs. (1.42–1.45) and are given explicitly in Appendix B of Ref. [419]. Eq. 2.18 can equivalently be found using the condition $dS^2/dt = 0$. The existence of physical solutions can be characterized using the cubic-polynomial discriminant,

$$\Delta = \sigma_4^2 \sigma_2^2 - 4\sigma_6 \sigma_2^3 - 4\sigma_4^3 \sigma_0 - 27\sigma_6^2 \sigma_0^2 + 18\sigma_6 \sigma_4 \sigma_2 \sigma_0. \quad (2.19)$$

In particular:

- If $\Delta > 0$, then $\Sigma(S^2)$ has three distinct real roots. These are the physical solutions S_- and S_+ identified in Ref. [353], plus a spurious root that does not satisfy

Eqs. (2.13–2.14).

- If $\Delta = 0$, the two solutions S_- and S_+ coincide and correspond to a spin–orbit resonance.
- If $\Delta < 0$, the polynomial $\Sigma(S^2)$ only admits one real root, thus implying that the geometrical constraints in Eqs. (2.13–2.14) cannot be satisfied for the assumed set of parameters $(q, \chi_1, \chi_2, J, \xi, L)$.

Therefore, physical spin precession takes place whenever $\Delta \geq 0$. The limiting case of the spin–orbit resonances can be located by solving $\Delta = 0$. The discriminant reported in Eq. (2.19) may be recast as a fifth-degree polynomial in J^2 ,

$$\Delta(J^2) = \delta_{10}J^{10} + \delta_8J^8 + \delta_6J^6 + \delta_4J^4 + \delta_2J^2 + \delta_0, \quad (2.20)$$

where the coefficients δ_i are lengthy (but real and algebraic) expressions containing q , S_1 , S_2 , ξ , and L ; see Appendix B of Ref. [419]. In particular,

$$\delta_{10} = -4q^3(1-q)^2(1+q)^8L^2 \leq 0. \quad (2.21)$$

2.3.2 Number of resonances

Any fifth-degree polynomial has at most two bound intervals and one unbound interval in which it is positive. The two bounded intervals are the only possible locations in which spin precession can occur. We now prove that only one of these can be physical.

To this end, it is useful to look at the asymptotic limit $r \rightarrow \infty$. While J diverges in this limit, recall from Sec. 1.3.4 that κ_∞ is constant [353]. The constraints $|\cos \theta_1| \leq 1$

and $|\cos \theta_2| \leq 1$ can be translated into

$$\kappa_\infty \geq \max \left\{ \frac{M^2 q \xi - (1 - q^2) S_1}{1 + q}, \frac{M^2 q \xi - (1 - q^2) S_2}{q(1 + q)} \right\}, \quad (2.22)$$

$$\kappa_\infty \leq \min \left\{ \frac{M^2 q \xi + (1 - q^2) S_1}{1 + q}, \frac{M^2 q \xi + (1 - q^2) S_2}{q(1 + q)} \right\}. \quad (2.23)$$

Therefore, the support of $(J^2 - L^2)/2L$ (and thus J) is a single bounded interval at large separations — only one range of J is allowed and it is bounded by two resonances. Proving by contradiction, let us now assume that the support of J does not remain a single interval. A bifurcation would be present at some finite separation where the number of valid ranges goes from one to two. At this bifurcation point, two different values of dJ/dr must coexist for the same values of $q, \chi_1, \chi_2, \xi, J$. This is only possible if the two configurations have different values of S . However, at the bifurcation point one necessarily has $\Sigma(S^2) = 0$ and thus only one value of S is allowed (i.e., a repeated root).

Our proof is consistent with the extensive numerical exploration presented in Refs. [353, 352] — there are always two spin–orbit resonances for any values of q, χ_1, χ_2, ξ , and r . The two resonances are characterized by $\Delta\Phi = 0$ and $\Delta\Phi = \pi$. In particular, the $\Delta\Phi = 0$ ($\Delta\Phi = \pi$) resonance corresponds to the maximum (minimum) value of J , i.e.,

$$J^{(\Delta\Phi=\pi)} \leq J \leq J^{(\Delta\Phi=0)}. \quad (2.24)$$

An example is shown in Fig. 2.7. The region of J^2 where physical spin precession takes place is characterized by the polynomial discriminant from Eq. (2.20), $\Delta(J^2) > 0$. The spin–orbit resonances correspond to two of the roots of $\Delta(J^2) = 0$.




Figure 2.7: The discriminant Δ of the third-degree polynomial Σ as a function of J^2 for a binary BH at a separation $r = 10M$ with mass ratio $q = 0.8$, dimensionless spins $\chi_1 = \chi_2 = 1$ and effective spin $\xi = -0.2$. The discriminant has three real roots (vertical dashed lines). The shaded area between the two roots in which $\Delta \geq 0$ corresponds to the region in which physical spin precession takes place. This region is bounded on the right (left) by the $\Delta\Phi = 0$ ($\Delta\Phi = \pi$) resonance.

2.3.3 Resonant evolution

Next, we prove that a binary in a resonant configuration remains resonant under radiation reaction.

Let us label two binaries A and B . The binaries share the same values of the radiation-reaction constants of motions q , χ_1 , χ_2 , and ξ . Suppose binary A is a $\Delta\Phi = 0$ resonance at separation r and binary B is a $\Delta\Phi = 0$ resonance at $r + \delta r$. Again by contradiction, let us now assume that A and B do not coincide. From Eq. (2.24) one has $J_A(r) > J_B(r)$ and $J_A(r + \delta r) < J_B(r + \delta r)$. At some location $r < \tilde{r} < r + \delta r$ one must have $J_A(\tilde{r}) = J_B(\tilde{r})$, but $dJ_A/dr|_{\tilde{r}} \neq dJ_B/dr|_{\tilde{r}}$. In other terms, the inspiral trajectory of the two binaries must cross in the J – r plane. This is possible only if the two binaries have different values of S at \tilde{r} , i.e., $S_A(\tilde{r}) \neq S_B(\tilde{r})$. Taking the limit $\delta r \rightarrow 0$, the location of the crossing point can be made arbitrarily close to the initial separation r . At this location, $J_A = J_B$ identifies a resonance, where only one value of S is allowed. It follows that the two binaries A and B must coincide. An analogous proof can be carried out for $\Delta\Phi = \pi$.

2.3.4 Asymptotic resonances

Further progress can be made by studying the dynamics of resonant configurations at infinitesimal separations $r \rightarrow 0$ (or equivalently $L \rightarrow 0$). Although unphysical, this limit provides the asymptotic conditions of our PN evolutions.

As $r \rightarrow 0$, one has that $J \rightarrow S$ and Δ is increasingly dominated by the term with the least power of L . In particular, one gets

$$\lim_{L \rightarrow 0} \frac{\Delta}{\delta_{10}} = \prod_{j=1}^5 (S^2 - \lambda_j). \quad (2.25)$$

The roots λ_i of this expression are given by

$$\lambda_1 = \lambda_2 = \frac{(1-q)(qS_1^2 - S_2^2)}{q}, \quad (2.26)$$

$$\lambda_3 = \frac{(1-q)(qS_1^2 - S_2^2)}{q} + \frac{M^4 q \xi^2}{(1+q)^2}, \quad (2.27)$$

$$\lambda_4 = (S_1 - S_2)^2, \quad (2.28)$$

$$\lambda_5 = (S_1 + S_2)^2. \quad (2.29)$$

Let us denote the effective spin of the up–up and up–down configuration with respectively

$$\xi_{\text{UU}} = \frac{1}{M^2} \left[(1+q) S_1 + \left(1 + \frac{1}{q}\right) S_2 \right], \quad (2.30)$$

$$\xi_{\text{UD}} = \frac{1}{M^2} \left[(1+q) S_1 - \left(1 + \frac{1}{q}\right) S_2 \right]. \quad (2.31)$$

The constraint $|\xi| \leq \xi_{\text{UU}}$ implies the following series of inequalities:

$$\lambda_1 = \lambda_2 \leq \min(\lambda_3, \lambda_4) \leq \max(\lambda_3, \lambda_4) \leq \lambda_5, \quad (2.32)$$

with

$$\max(\lambda_3, \lambda_4) = \begin{cases} \lambda_4 & \text{if } |\xi| \leq |\xi_{\text{UD}}|, \\ \lambda_3 & \text{if } |\xi| > |\xi_{\text{UD}}|. \end{cases} \quad (2.33)$$

Since $\Delta \leq 0$ as $J \rightarrow \infty$ [cf. Eq. (2.21)], the two bounded intervals of J^2 in which $\Delta \geq 0$ are $[\lambda_1, \min(\lambda_3, \lambda_4)]$ and $[\max(\lambda_3, \lambda_4), \lambda_5]$. Furthermore, in this limit Eq. (2.13) reduces to $\lambda_4 \leq S^2 \leq \lambda_5$, which implies that the single physical interval in which spin

precession takes place is given by $S^2 \in [\max(\lambda_3, \lambda_4), \lambda_5]$.

The boundaries $S = \max\{\sqrt{\lambda_3}, \sqrt{\lambda_4}\}$ and $S = \sqrt{\lambda_5}$ of this region identify the asymptotic locations of the $\Delta\Phi = \pi$ and $\Delta\Phi = 0$ resonances, respectively. Thus, the value $S^{(\Delta\Phi=0)}$ of S in the $\Delta\Phi = 0$ spin-orbit resonance asymptotes to

$$\lim_{r \rightarrow 0} S^{(\Delta\Phi=0)} = S_1 + S_2, \quad (2.34)$$

and the value $S^{(\Delta\Phi=\pi)}$ of S in the $\Delta\Phi = \pi$ resonance asymptotes to

$$\lim_{r \rightarrow 0} S^{(\Delta\Phi=\pi)} = \begin{cases} |S_1 - S_2| & \text{if } |\xi| \leq |\xi_{\text{UD}}|, \\ \sqrt{\frac{(1-q)(qS_1^2 - S_2^2)}{q} + \frac{M^4 q \xi^2}{(1+q)^2}} & \text{if } |\xi| > |\xi_{\text{UD}}|. \end{cases} \quad (2.35)$$

The corresponding values of the misalignment angles $\theta_{1,2}$ are found by imposing the co-planarity condition $\sin(\Delta\Phi) = 0$ that characterizes the resonances. Using Eq. 1.36, this yields

$$\sin \theta_1 \sin \theta_2 + \cos \theta_1 \cos \theta_2 = \frac{S^2 - S_1^2 - S_2^2}{2S_1 S_2}, \quad (2.36)$$

which can be solved together with Eq. (1.39) to find $\cos \theta_1$ and $\cos \theta_2$. For the $\Delta\Phi = 0$ resonance one gets

$$\lim_{r \rightarrow 0} \cos \theta_1^{(\Delta\Phi=0)} = \lim_{r \rightarrow 0} \cos \theta_2^{(\Delta\Phi=0)} = \frac{\xi}{\xi_{\text{UU}}}. \quad (2.37)$$

In words, the two spins tend to be equally misaligned with \mathbf{L} but co-aligned with each other. Hints of this trend had been reported in Refs. [351, 441, 442]. For $\Delta\Phi = \pi$, the

angles asymptote to

$$\lim_{r \rightarrow 0} \cos \theta_1^{(\Delta\Phi=\pi)} = \begin{cases} \frac{\xi}{\xi_{\text{UD}}} & \text{if } |\xi| \leq |\xi_{\text{UD}}|, \\ \frac{\xi^2 + \xi_{\text{UU}}\xi_{\text{UD}}}{2(1+q)S_1\xi} M^2 & \text{if } |\xi| > |\xi_{\text{UD}}|, \end{cases} \quad (2.38)$$

$$(2.39)$$

$$\lim_{r \rightarrow 0} \cos \theta_2^{(\Delta\Phi=\pi)} = \begin{cases} -\frac{\xi}{\xi_{\text{UD}}} & \text{if } |\xi| \leq |\xi_{\text{UD}}|, \\ \frac{q(\xi^2 - \xi_{\text{UU}}\xi_{\text{UD}})}{2(1+q)S_2\xi} M^2 & \text{if } |\xi| > |\xi_{\text{UD}}|. \end{cases} \quad (2.40)$$

Figure 2.8 shows the evolution of four resonant configurations for $\Delta\Phi = 0, \pi$, and the two cases $|\xi| \leq |\xi_{\text{UD}}|$ and $|\xi| > |\xi_{\text{UD}}|$. At each separation we locate the roots of $\Sigma(S^2) = 0$ numerically using the algorithm implemented in the PRECESSION code [430, 435]. As resonant binaries remain resonant during the inspiral (Sec. 2.3.3), those curves also correspond to individual evolutions. As $r \rightarrow 0$, binaries asymptote to the limits predicted above.

2.4 Up–down endpoint

2.4.1 Instability limit

The analysis of Sec. 2.3 allows us to find the asymptotic endpoint of the up–down configuration. As first shown in Ref. [434], the up–down configuration is a $\Delta\Phi = 0$ resonance for $r > r_{\text{UD}+}$. This can be immediately seen using the expressions in




Figure 2.8: Evolutions of resonant configurations to small separations. Top, middle, and bottom panels show S/M^2 , $\cos \theta_1$, and $\cos \theta_2$, respectively. Blue (orange) curves correspond to resonances with $\Delta\Phi = 0$ ($\Delta\Phi = \pi$), while binaries in the left (right) panel satisfy $|\xi| \leq |\xi_{UD}|$ ($|\xi| > |\xi_{UD}|$). We fix $q = 0.9$, $\chi_1 = 0.8$, $\chi_2 = 0.4$, $\xi = 0.1$ (left), 0.4 (right). For this set of parameters one has $\xi_{UD} \approx 0.23$ and $\xi_{UU} \approx 0.61$. Dashed grey lines indicate the $r \rightarrow 0$ limits predicted in Sec. 2.3.4. The region $r \lesssim 10M$ should be considered unphysical but is included to test our analytical calculations.

Sec. 2.3.2. As $r \rightarrow \infty$, the up–down configuration corresponds to $\kappa_\infty = S_1 - S_2$ which maximizes the allowed range of κ_∞ given in Eqs. (2.22–2.23), and hence that of J . The largest value of J for a given ξ corresponds to the $\Delta\Phi = 0$ resonance [cf. Eq. (2.24)].

A binary which is arbitrarily close to up–down before the instability onset, therefore, will be arbitrarily close to a $\Delta\Phi = 0$ spin–orbit resonance. As shown in Sec. 2.3.3, resonant binaries remain resonant during the entire inspiral. The formal $r \rightarrow 0$ limit of the up–down instability is that of a $\Delta\Phi = 0$ resonance with the correct value of the effective spin. This can be obtained directly from Eqs. (2.34) and Eq. (2.37) by setting $\xi = \xi_{\text{UD}}$.

The key result of this work is that the endpoint of the up–down instability consists of a binary configuration with


$$\cos\theta_1 = \cos\theta_2 = \frac{\chi_1 - q\chi_2}{\chi_1 + q\chi_2} \quad \text{and} \quad \Delta\Phi = 0, \quad (2.41)$$

which is equivalent to Eq. (2.2). Up–down binaries start their inspiral with $S = |S_1 - S_2|$ and asymptote to $S = S_1 + S_2$ as given by Eq. (2.34), thus spanning the entire range of available values of S [cf. Eq. (2.13)].

An example is reported in Fig. 2.5. Despite being obtained for $r \rightarrow 0$, the spin configuration in Eq. (2.41) describes the inspiral endpoint well. Similarly, Fig. 2.6 shows that binaries initially close to the up–down configuration all evolve to this precise location in parameter space.

Figure 2.9 illustrates the formal $r \rightarrow 0$ distribution for two simple BH populations. In particular, we distribute mass ratios q either uniformly or according to the astrophysical population inferred from the first GW events, $p(q) \propto q^{6.7}$ (cf. Model B in Ref. [83], and see also Ref. [443]). In both cases, we take $q \in [0.1, 1]$ and assume

Figure 2.9: Analytic up-down endpoint distribution in the $r \rightarrow 0$ limit. Top and bottom panels show the endpoint distributions of S and $\cos \theta_1 = \cos \theta_2$. Blue histograms are obtained distributing q uniformly; orange histograms assume $p(q) \propto q^{6.7}$ as observed by LIGO/Virgo [83]. In both cases, we distribute χ_1 and χ_2 uniformly and assume $q, \chi_1, \chi_2 \in [0.1, 1]$.



spin magnitudes χ_i are distributed uniformly in $[0.1, 1]$. The LVC-motivated population strongly favours equal mass events. For $q \approx 1$ the instability endpoint is given by $S/M^2 \approx (\chi_1 + \chi_2)/4$ and $\cos \theta_{1,2} \approx (\chi_1 - \chi_2)/(\chi_1 + \chi_2)$, which implies that the corresponding distributions are peaked at $S/M^2 \approx (0.1 + 1)/4 = 0.275$ and $\cos \theta_{1,2} \approx 0$. If q differs from unity, the endpoint values of both S and $\cos \theta_i$ are, on average, larger. For the case where mass ratios are drawn uniformly, unequal-mass binaries populate the region of Fig. 2.9 with $S/M^2 \gtrsim 0.5$ and $\cos \theta_{1,2} \gtrsim 0.7$.

2.4.2 Stability-to-instability transition


During the inspiral, unstable up–down binaries evolve from $S = |S_1 - S_2|$ to $S = S_1 + S_2$. The transition between the two values can only start after binaries enter the instability regime ($r < r_{\text{UD}+}$) and is halted by the merger (or, to be more conservative, by the PN breakdown). To quantify the transition properties, it is useful to define the parameter

$$\delta S = \frac{S - |S_1 - S_2|}{(S_1 + S_2) - |S_1 - S_2|}, \quad (2.42)$$

such that $\delta S = 0$ corresponds to stability and $\delta S = 1$ corresponds to the formal $r \rightarrow 0$ endpoint.

Figure 2.10 shows the distribution of δS and $r_{\text{UD}+}$ resulting from numerical integrations of up–down binaries. We distribute q , χ_1 , and χ_2 uniformly in $[0.1, 1]$ and evolve from $r_i = 1000M$ to $r_f = 10M$. Binaries with $r_{\text{UD}+} < r_i$ are initialized as up–down and might become unstable during the integration. Binaries with $r_{\text{UD}+} > r_i$, on the other hand, are already unstable at the start of our integrations. We therefore initialized them as $\Delta\Phi = 0$ resonances at $r = r_i$. In both cases, we introduce a misalignment

Figure 2.10: Distribution of $\max_r \delta S$ and δr as function of the instability onset $r_{\text{UD}+}$ for a statistical sample of 1000 up-down binaries. We distribute q , χ_1 , and χ_2 uniformly in $[0.1, 1]$, evolve from $r_i = 1000M$ to $r_f = 10M$, and initialize the spin misalignments from Gaussian distributions with widths $\delta\theta = 10^\circ$ centred on the up-down configuration. The blue (orange) sub-population indicates sources that do (not) reach $\delta S = 0.5$ by the end of the inspiral. By definition, δr can only be computed for the sub-population with $\max_r \delta S > 0.5$, with a minimum value of $\delta r \leq r_{\text{UD}+} - r_f$ (dashed line).



perturbation $\delta\theta = 10^\circ$ following the same procedure of Sec 2.2.3.

We consider the largest value of δS reached between $r_i = 1000M$ and $r_f = 10M$; in practice, this is very similar to its value at the end of the evolution, i.e., $\max_r \delta S(r) \approx \delta S(r_f)$. If $r_{\text{UD}+} \lesssim 10M$, up–down binaries are still stable at the end of our evolutions and thus $\max_r \delta S \simeq 0$. If the instability onset occurs earlier, binaries start transitioning toward larger values of δS . We find that the vast majority of sources with $r_{\text{UD}+} \gtrsim 50M$ are able to reach the predicted endpoint ($\max_r \delta S \gtrsim 0.95$) before the PN breakdown. As long as the instability has enough time to develop, the formal $r \rightarrow 0$ limit appears to provide a faithful description of dynamics. In the intermediate cases with $10M \lesssim r_{\text{UD}+} \lesssim 50M$, the instability takes places shortly before the PN breakdown and, consequently, δS does not have enough time to reach unity.

The transition between the two regimes appears to be rather sharp, taking place over a short interval in r . To better quantify this observation, we define the instability growth “time” as the difference between the instability onset $r_{\text{UD}+}$ and the separation where $\delta S = 0.5$, i.e., $\delta r = r_{\text{UD}+} - r_{\delta S=0.5}$. The bottom panels of Fig. 2.10 illustrate the behaviour of δr for the same population of BHs. The quantity δr can only be computed for binaries that reach $\delta S = 0.5$ before the end of the evolution, thus setting the constrain $\delta r \leq r_{\text{UD}+} - r_f$. The fraction of unstable binaries (those that reach $\delta S \geq 0.5$) in this population is 35%. We find that the typical transition intervals are $\delta r \lesssim 100M$, with a peak at $\delta r \simeq 25M$, so the instability develops over a short period and unstable binaries quickly reach values of S close to the endpoint.

2.4.3 A simple astrophysical population

We now study the effect of the instability on an astrophysically-motivated population of binary BHs. We model a formation channel that leads to the alignment of the

BH spins with the orbital angular momentum, but where co-alignment and counter-alignment are equally probable. This might be the case, for instance, for stellar-mass BHs brought together by viscous interactions in AGN disks [305, 306, 308, 309, 311, 444, 445, 446]. Unlike BH binaries formed from binary stars (where the initial cloud imparts its angular momentum to both objects favouring co-alignment), or systems formed in highly interacting environments like globular clusters (where frequent interactions tend to randomize the spin directions), an accretion disk defines an axisymmetric environment without a preference for co- or counter-alignment. Ref. [446] specifically modelled this scenario by assuming that 1/4 of the population is found in each of the up–up, down–down, down–up, and up–down configurations; naively, one could expect that $\sim 25\%$ of the stellar-mass BH binaries formed in AGN disks are subject to the up–down instability.

As before, we distribute mass ratios using the astrophysical population inferred from the O1+O2 GW events [83], $p(q) \propto q^{6.7}$ with $q \in [0.1, 1]$, and sample the dimensionless spins χ_i uniformly in $[0.1, 1]$. We simulate 10^4 binaries in each of the four aligned configurations, and integrate the precession equations numerically from an initial orbital separation $r_i = 1000M$ to a final separation $r_f = 10M$. Binaries are initialized by sampling $\cos\theta_{1,2}$ from truncated Gaussians with $\delta\theta = 20^\circ$. If the corresponding parameters q , χ_1 and χ_2 are such that $r_{\text{UD}+} > r_i$ (i.e., if the source went unstable before the beginning of our integrations), the initial configuration is set to be that of a $\Delta\Phi = 0$ resonance, again with a $\delta\theta = 20^\circ$ perturbation.

The resulting distribution of ξ is shown in Fig. 2.11. The effective spin ξ is a constant of motion; these curves are independent of the orbital separation. Up–up (down–down) binaries tend to pile up at positive (negative) large values of the effective spins, while both up–down and down-up sources contribute to a peak at $\xi \approx 0$.




Figure 2.11: Distribution of effective spin ξ for the four populations of aligned-spin binaries. Mass ratios q are sampled according to the power-law distribution $p(q) \propto q^{6.7}$ [83] and dimensionless spins χ_i are sampled uniformly. We set $q, \chi_1, \chi_2 \in [0.1, 1]$ and introduce an initial misalignment $\delta\theta = 20^\circ$. The dotted empty histogram indicates the full population. Colour shaded histograms differentiate between the four aligned cases. For the up-down population (red), we further separate the contributions of two sub-populations: sources that remain stable during the entire PN inspiral ($\max_r \delta S < 0.5$, dashed grey) and sources that undergo the instability ($\max_r \delta S > 0.5$, dashed black).

Figure 2.12 shows the joint distributions of $\cos \theta_1$ and $\cos \theta_2$ at the initial (left) and final (right) separations for each of the four populations. Up–up, down–down, and down–up binaries largely retain their initial, aligned orientation. Up–down binaries segment into two clear sub-populations: those which remain stable (lower-right corner in Fig. 2.12) and those which become unstable (centre of Fig. 2.12). The dispersion of the stable up–down binaries increases compared to the initial distribution owing to a proportion of these binaries that reach the onset of the instability but do not reach the formal endpoint by the end of the evolution.

The sub-population that becomes unstable presents a clear trend in the misalignment distribution: binaries pile up along the $\cos \theta_1 = \cos \theta_2$ diagonal as predicted by Eq. (2.41). As before, we characterize the two populations using δS . At $r = 1000M$, only $\approx 34\%$ binaries are in the unstable sub-population ($\delta S > 0.5$); for the vast majority, these are the cases with $r_{\text{UD}+} > r_i$. By the time binaries reach $r = 10M$, the unstable fraction goes up to $\approx 91\%$ (cf. $\approx 35\%$ for the population with mass ratios instead distributed uniformly in $[0.1, 1]$ presented in Fig. 2.10). Compared to the distribution of analytic endpoints of Fig. 2.9, the numerical population is skewed toward the initial configuration $\cos \theta_1 = -\cos \theta_2 = 1$, again due to a proportion of binaries that do not fully reach $\delta S \approx 1$.

Figure 2.13 shows the distribution of $r_{\text{UD}+}$, q , χ_1 , and χ_2 for the two up–down sub-populations. Only binaries with either $q \lesssim 0.6$ or $\chi_{1,2} \lesssim 0.2$ are still stable at the end of the evolution. These values correspond to $r_{\text{UD}+} \lesssim 50M$. All other sources belong to the unstable subpopulation and approach merger near their predicted endpoints ($\max_r \delta S \gtrsim 0.5$). An orbital separation of $50M$ corresponds to a GW frequency $f = \sqrt{M/\pi^2 r^3} \approx 20$ Hz for a typical LIGO source ($M = 10M_\odot$) and $\sim 10^{-4}$ Hz for a supermassive BH binary ($M = 10^6 M_\odot$) detectable by LISA.




Figure 2.12: Joint distribution of $\cos \theta_1$ and $\cos \theta_2$ for binary BHs with initially aligned spins. Sources are evolved numerically from a separation $r = 1000M$ (top panel) to $r = 10M$ (bottom panel). Mass ratios q are sampled according to the power law distribution $p(q) \propto q^{6.7}$ [83]; dimensionless spins χ_i are sampled uniformly. We set $q, \chi_1, \chi_2 \in [0.1, 1]$. The populations, each containing 10^4 binaries, of up-up (blue), down-down (orange) and down-up (green) binaries remain in their initial distributions whereas the up-down (red) population does not, thus highlighting the precessional instability. By the end of the evolutions the up-down binaries split into two sub-populations: those which remain stable (bottom right corner) and those which do not (central region). The trend observed in the unstable sub-population matches the prediction $\cos \theta_1 = \cos \theta_2$ of Sec. 2.4.1.

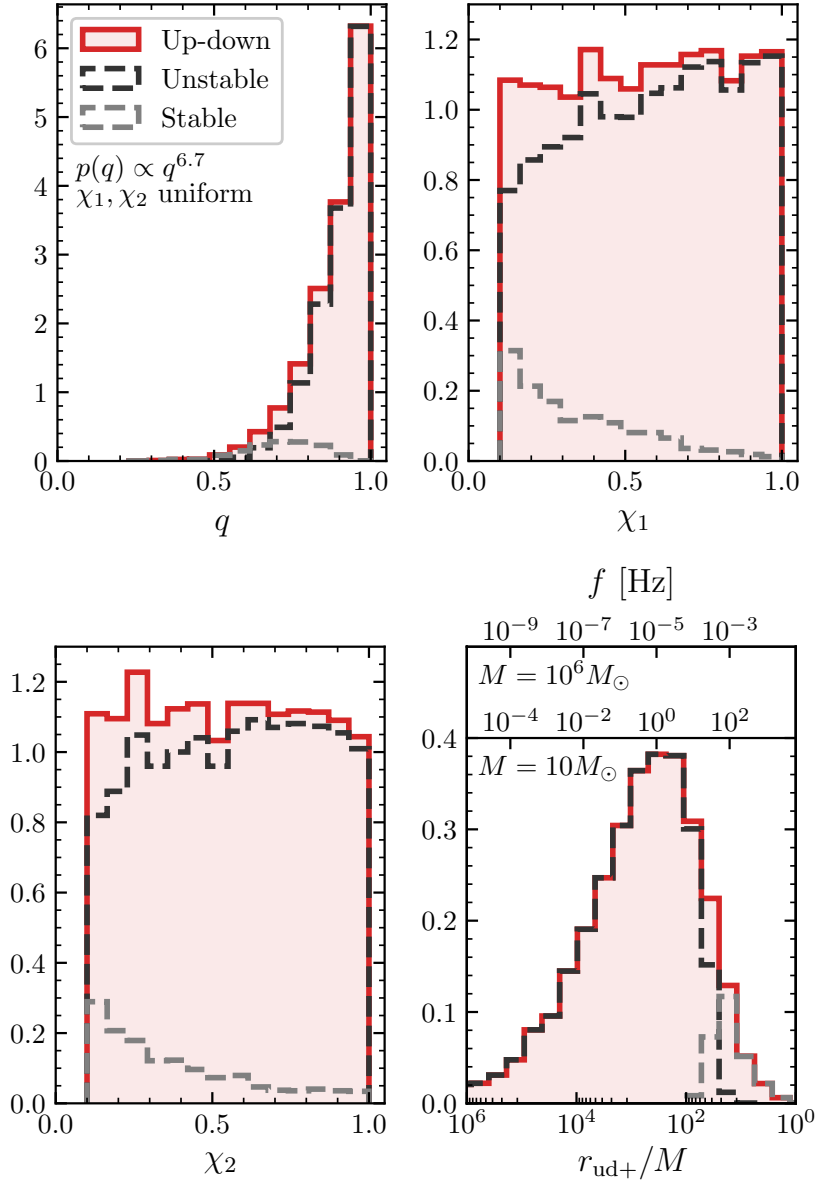


Figure 2.13: Distribution of, from left to right, mass ratio q and primary spin magnitudes χ_1 (top row), and secondary spin magnitude χ_2 and instability onset $r_{\text{UD+}}$ (bottom row), for a set of up–down binaries. Mass ratios are sampled according to a power law distribution that strongly favours equal masses [83], while spins are sampled uniformly. We set $q, \chi_1, \chi_2 \in [0.1, 1]$. Red histograms show the full up–down population, while dashed black (grey) histograms indicate systems that do (not) become unstable. Stability is defined as $\max_r \delta_S < 0.5$, integrating from $r_i = 1000M$ to $r_f = 10M$. In the bottom-right panel, the top axes show the value of the GW frequency $f = \sqrt{M/\pi^2 r^3}$ for systems with total mass $M = 10M_\odot$ and $M = 10^6 M_\odot$.

A tantalizing possibility would be the development of the precessional instability while a binary is being observed. For LIGO, we expect that such a situation is possible for only a small number of sources. To estimate this fraction, we produce a distribution of the total mass again according to Ref. [83] with the distributions of q , χ_1 , and χ_2 as in Fig. 2.13. The sub-population for which the instability develops in band is then determined by the conditions $f_{\text{UD}+} > f_{\text{LIGO}}$ (lower LIGO frequency cutoff), $r_{\text{UD}+} > 10M$ (validity of the PN approximation, which also provides the upper frequency cutoff), and $\max_r \delta S > 0.5$ (appreciable development of the instability). Unsurprisingly, this fraction depends strongly on the lower frequency cut-off: for $f_{\text{LIGO}} = 20$ Hz (10 Hz), only 0.7% (2.8%) of the total population develops the precessional instability while in the LIGO band. LISA might provide better prospects, as some supermassive BH binaries will remain visible for several precession cycles [447].

The condition $q \rightarrow 0$ and $\chi_{1,2} \rightarrow 0$ identifies the single-spin limit. In practice, we expect that the vast majority of up-down sources where two-spin effects are prominent will become unstable before entering the sensitivity window of our detectors. Proper modelling of two-spin effects appears to be crucial. The ξ distributions of the two sub-populations does not present evident systematic trends (Fig. 2.11) and largely reflects that of the full up-down sample. This suggests that it will be challenging to distinguish stable and unstable binaries by measuring only one effective spin.

2.5 Conclusions

We reinvestigated the precessional instability in BH binaries first reported in Ref. [434]. Perturbations to equilibrium solutions of the spin precession equations lead to small-amplitude oscillations governed by simple harmonic motion. However, for

up–down sources, we found a critical threshold in the inspiral at which they become unstable and can evolve with large spin tilts — in precise agreement with Ref. [434]. We verified our analytic calculations with numerical integrations that also showed the unstable binaries do not disperse in the available parameter space. Phrased in terms of spin–orbit resonances [351], we derived in closed-form expressions to locate these extremal bounds of spin precession. The endpoint of the up–down instability is found as a particular limiting case, in which we showed the two BH spins are aligned with each other and equally misaligned with the orbital angular momentum.

For orbital separations above the instability threshold, the up–down configuration is a spin–orbit resonance. But this is no longer true when the separation decreases due to GW emission. This is precisely the cause of the instability: since resonant configurations remain resonant, binaries initially having up–down spins must evolve toward the new resonance location after passing the threshold.

We found that the instability develops quickly, so that the majority of sources which do become unstable reach the endpoint before merger. The observational consequence is that up–down binaries whose instabilities trigger before the sensitivity window of GW detectors will be observed if not right in the predicted endpoint then as precessing sources, rather than with aligned spins, whereas those with onsets at small separations will remain with the initial up–down spins. The presence of these two sub-populations could be a clear indicator of the up–down instability in GW observations.

The analysis presented here is limited to the approximate PN regime. We later verified that the instability exists in the strong-field regime of NR [424]. We also showed that its endpoint is distinguishable with realistic detector sensitivities in the near future [425]. A remaining question is whether or not merging up–down binary BHs form in nature. We look forward to a possible positive discovery in the future.

Binary black holes time travelling back to the future

Abstract

The assumptions underlying GW population studies are that the targeted source parameters refer to the same quantities for all events in the catalogue and are included when modelling selection effects. Both these points have so far been neglected when estimating the orientations of binary BH spins. In particular, the detector-frame GW frequency (e.g., 20 Hz) used to define frequency-dependent quantities introduces an inconsistent reference between events at the population level. We solve these issues by modelling binary BH populations and selection effects at past time infinity, corresponding to the well-defined reference frequency of 0 Hz. We show that, while current GW measurement uncertainties obfuscate the influence of reference frequency in population inference, ignoring spins when estimating selection effects leads to an over-prediction of spin alignment in the underlying astrophysical distribution of merging BHs.

Summary and contributions

This work is a reformatted version of our paper from Ref. [421] published in collaboration with Davide Gerosa and is based on the PN formalism and code that he developed [352, 353, 355, 430, 435] (see Sec. 1.3). I generated the back-propagated GW posterior samples and the injections used to compute selection effects. I performed the population inference runs and made Figs. 3.2, 3.3, and 3.4. Davide created Fig. 3.1. The paper was written jointly.

In Sec. 3.1 we explain systematic issues present in current GW population analyses. We present the solution to these issues and describe the population models and statistical methods we use in Sec. 3.2. Our results are given in Sec. 3.3. We wrap up and point out further directions for future work along the lines presented throughout this Chapter in Sec. 3.4.

3.1 Introduction

A key underlying assumption when performing a hierarchical Bayesian analysis of GW sources at the population level is that one is combining measurements of the same parameters for all events. This is trivially the case for constant quantities such as masses and spin magnitudes. The spin orientations, however, are subject to relativistic precession and vary as the binary inspirals toward merger [349], as we saw in Sec. 1.3. If present, the orbital eccentricity also changes during the inspiral [254].

State-of-the-art GW population inference is performed at a fixed detector-frame frequency [84]. For the LVC, the reference frequency is typically set to $f_{\text{ref}} = 20$ Hz. This choice is perfectly acceptable when analysing a single event but questionable at

the population level, mainly for two reasons:

- As currently defined, f_{ref} is a detector-frame quantity, but the parameters targeted in the inference (e.g., masses) are taken in the source frame.
- The binary configuration at a given reference frequency depends on the other parameters. Systems with larger (smaller) masses and/or anti-aligned (co-aligned) spins are closer (further) from merger.

These issues were made evident for GW190521, where the high mass imposed a lower value $f_{\text{ref}} = 11$ Hz, different from all other events in the catalogue [62].

The spin directions at f_{ref} thus describe different quantities for each event: one is not allowed to put them together in a population fit as if they were the same. To the best of our knowledge, this issue plagues all current GW population studies which make use of quantities that vary during the binary evolution.

A first step in the right direction was recently taken in Ref. [448], where spins are quoted at a fixed dimensionless time $t_{\text{ref}} = -100M$ from the peak of the GW strain. While this tackles the first issue highlighted above, it does not fully address the second point. One can still construct several dimensionless quantities that vary monotonically along the inspiral (e.g., time, orbital frequency, orbital separation, etc.) and select any of those when quoting the spin directions.

Current spin inference suffers from another pressing issue. When reconstructing the observable population of sources from the observed catalogue, one must account for selection effects [402, 403]. Although it is well known that sources with spins co-aligned (counter-aligned) with the binary orbital angular momentum are easier (harder) to observe [347, 449], this detection bias is often deemed unimportant and neglected. While (the aligned components of) the spins are included in the pipeline injections used

to estimate the LIGO and Virgo detectabilities, their dependence on the population parameters is then neglected when sampling the hierarchical likelihood [84].

We present the first complete solution to these conceptual issues: the spin directions needed in both the population and selection-effect models are those at past time infinity or, equivalently, $f_{\text{ref}} = 0$ Hz. As we saw in Sec. 1.3.4, this asymptotic configuration uniquely determines the entire history of a binary up to an orbital and a precessional phase [353, 423]. This reference puts the spin directions of all the events on equal footing, thus allowing for a consistent implementation of the population fit. Two points need to be tackled:

- Data from GW events are, by definition, taken when the binary is detectable and thus need to be propagated backward to past time infinity.
- Modelling selection effects instead requires propagating the tested population forward from past time infinity to detection.

We first travel back in time ($20 \mapsto 0$ Hz) when treating the event likelihoods and then “back to the future” ($0 \mapsto 20$ Hz) when handling selection effects. Our “DeLorean” is the precession-averaged PN formalism.

3.2 Time travel for population inference

3.2.1 Modelling and inference

The statistical problem we tackle is that presented in Sec. 1.4.3. We seek to constrain the shape parameters λ of our chosen population model from the observed GW data $\{d\}$ by sampling the posterior in Eq. (1.85). We use the parametrized model referred to as POWER LAW + PEAK and DEFAULT SPIN in Ref. [84], which returns

the highest Bayes factor among the options they tested. The model covers $\dim \theta = 6$ event parameters and $\dim \lambda = 12$ population parameters.

The distribution of the primary mass m_1 is a superposition of a power-law component with index α truncated outside m_{\min} and m_{\max} and a Gaussian component with mean μ_m , width σ_m , and mixing fraction λ_m . The secondary mass m_2 conditioned on m_1 follows a power-law distribution with index β_q . The distributions of $m_{1,2}$ are smoothed over a range δ_m near m_{\min} .

The spin magnitudes $\chi_{1,2}$ follow a beta distribution $p(\chi|\mu_\chi, \sigma_\chi) \propto \chi^{\alpha_\chi-1}(1-\chi)^{\beta_\chi-1}$ with mean μ_χ and variance σ_χ^2 , where $\alpha_\chi = [\mu_\chi(1-\mu_\chi) - \sigma_\chi^2]\mu_\chi/\sigma_\chi^2$ and $\beta_\chi = \alpha_\chi(1-\mu_\chi)/\mu_\chi$. The cosines $\cos \theta_{1,2}$ of the angles between the spins and the orbital angular momentum are distributed assuming a superposition of a uniform distribution and a truncated Gaussian with a peak at $\cos \theta_{1,2} = 1$, width σ_t , and mixing fraction ζ . Crucially, while we adopt the same functional form of Ref. [84], the spin tilts $\theta_{1,2}$ are here inserted at past time infinity and not at detection.

The distributions of all other parameters (distance, sky location, etc.) is assumed to be independent of λ and equal to the prior used in the underlying single-event analyses [see Eq. (1.87)].

The integrals at the numerator of Eq. (1.85) are approximated with Monte Carlo summations using posterior samples from the data release accompanying Refs. [450] (O1+O2) and [62] (O3a), which in total include 44 GW events with false-alarm rates (FARs) $< 1 \text{ yr}^{-1}$. The single-event priors are handled analytically with suitable conversion factors [451].

For the POWER LAW + PEAK and DEFAULT SPIN model, BH masses and spins are not correlated and, consequently, the population model can be written as the product of two terms, one including only masses and the other only spins. In Ref. [84], the

spin part was included only in the likelihood integral of Eq. (1.85) but not in that of Eq. (1.75). When computing $\sigma(\lambda)$, they instead used a fixed spin distribution, thus neglecting some λ dependencies and introducing a bias. This was motivated by the large computational cost of the search injections required to compute $p(\text{det}|\theta)$.

We find that a simpler prescription for detectability (as used previously in, e.g., Ref. [83]) fully reproduces the results of Ref. [84] while allowing for a consistent inclusion of spin effects. In particular, we use the semi-analytic approximation of Ref. [416] explained in Sec. 1.4.3, assuming two data-taking periods of approximately 166 days (O1+O2 [60, 61]) and 150 days (O3a [62]), and a single-detector SNR threshold of 8 [145, 415]. SNRs are computed with representative noise curves¹ and the IMPHENOMPV2 waveform model [452]. The integral in the denominator of Eq. (1.85) is approximated with a Monte Carlo sum using samples drawn from an injected population with $p(m_1) \propto m_1^{-2.35}$, $p(m_2|m_1) \propto m_2^2$ [62], uniform spin magnitudes, spin directions with equally weighted isotropic and preferentially aligned components ($\zeta = 0.5$ and $\sigma_t = 0.02$), and redshifts distributed uniformly in comoving volume and source-frame time.

The prior $p(\lambda)$ is uniform over all 12 population parameters with limits and additional cuts as in Ref. [84]. We sample from the posterior $p(\lambda|\{d\})$ using GWPOPULATION [453], BYNESTY [387], and BILBY [388].

3.2.2 Spin propagation

We propagate BH spin orientations across GW frequencies using the precession-averaged PN formalism first developed in Refs. [352, 353] and discussed in detail in

¹From dcc.ligo.org/LIGO-P1200087-v47 (“early high”, for O1+O2) and dcc.ligo.org/LIGO-T2000012 (“Livingston”, for O3a).

Sec. 1.3.4. We use the PRECESSION code [430, 435] which, leveraging new analytical advancements and numerical recipes [355] has been made significantly more efficient, thus facilitating the large-scale studies presented here (an alternative implementation of the formalism can be found in Ref. [359]).

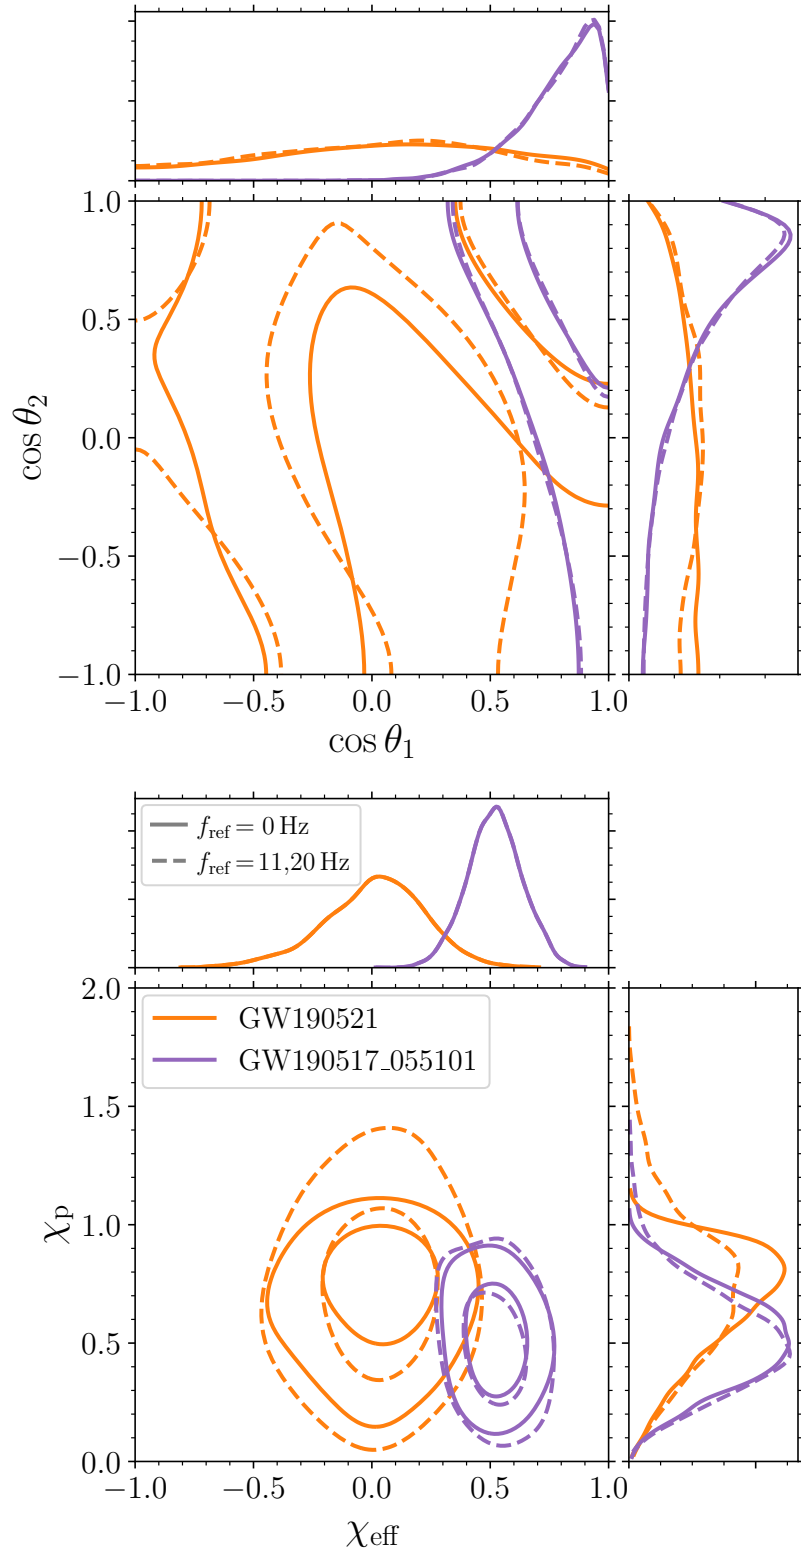
Recall that, at a finite orbital separation r , a BH binary is specified by the component masses $m_{1,2}$ and the dimensionless spin vectors $\boldsymbol{\chi}_{1,2}$. At past time infinity, a BH binary is fully specified by six quantities: $m_{1,2}$, $\chi_{1,2}$, and $\theta_{1,2}$.

First, we propagate binaries backward in time to estimate the likelihoods $\mathcal{L}(d_i|\theta_i)$. For a given posterior sample at $f_{\text{ref}} > 0$ Hz, we obtain the corresponding orbital separation r using Eq. (4.13) of Ref. [350] and the Newtonian angular momentum $L = m_1 m_2 \sqrt{r/M}$. The $f_{\text{ref}} \mapsto r$ conversion needs to be performed using the redshifted mass $M(1+z)$ because f_{ref} is a detector-frame quantity. For current LIGO events with $f_{\text{ref}} = 11, 20$ Hz, we find $6M \lesssim r \lesssim 25M$, which is roughly in the PN regime of validity. We then compute $J = |\mathbf{L} + m_1^2 \boldsymbol{\chi}_1 + m_2^2 \boldsymbol{\chi}_2|$, which serves as the initial condition for the precession-averaged evolution in Eq. (1.47) [equivalently Eq. (1.50)].

We only propagate posterior samples, and not prior samples. Current single-event priors are isotropic in the spin directions [62, 450]. Since isotropicity is preserved very accurately by PN evolutions [353, 454], the prior $p(\cos \theta_{1,2}|\mathcal{H})$ is uniform both in the detector band and at $f_{\text{ref}} = 0$ Hz.

Figure 3.1 shows back-propagated posterior distributions for the events with the largest total mass (GW190521) and the largest effective spin (GW190517_055101) in the catalogue considered. For GW190521, the values of $\theta_{1,2}$ at $f_{\text{ref}} = 0$ Hz show a slightly stronger preference for the corner of the parameter space near $(\theta_1, \theta_2) = (0, \pi)$. We also show χ_{eff} and the precession-averaged precessing spin $\chi_{\text{p}} \in [0, 2]$ defined in Ref. [426] — while χ_{eff} measures the component BH spins along the direction of \mathbf{L}

Figure 3.1: Back-propagating GW data to past time infinity. We show the event with the largest mass (GW190521, orange) and that with the largest effective spin (GW190517_0555101, purple). Top and bottom panels show the joint posterior distribution of the spin tilts (θ_1, θ_2) and the effective spins ($\chi_{\text{eff}}, \chi_p$), respectively. Dashed curves show the distributions in the LIGO/Virgo band ($f_{\text{ref}} = 11$ Hz for GW190521 and $f_{\text{ref}} = 20$ Hz for GW190517_0555101). Solid curves are computed at $f_{\text{ref}} = 0$ Hz. The central panels show two-dimensional contours at the 50% and 90% levels.



and is conserved, χ_p measures the orthogonal components and is precession averaged such that it only varies on the longer inspiral timescale t_{RR} . While χ_{eff} is a constant of motion and therefore does not depend on reference frequency, the back-propagated χ_p distributions are slightly narrower (see also [429]).

Second, we evolve binaries from the injected population forward in time to compute the detectability. A binary configuration described by $m_{1,2}$, $\chi_{1,2}$, and $\theta_{1,2}$ at $r = \infty$ is first integrated down to $f_{\text{ref}} = 20$ Hz with corresponding redshift z . At this location, we sample the precessional phase from its PDF using the 2PN result of Refs. [352, 353] (see Sec. 1.3.4) and the orbital phase uniformly in $[0, 2\pi]$. The resulting binary system, now well in the LIGO/Virgo band, is used to compute the SNR and thus $p(\text{det}|\theta)$.

3.3 Gravitational-wave observations at past time infinity

Our results are presented in Figs. 3.2 and 3.3 where we show, for the first time, the asymptotic population of BH spins inferred from current LVC data ($f_{\text{ref}} = 0$ Hz, red). We also run a control case with an identical setup but taking the spin tilts at detection ($f_{\text{ref}} = 11, 20$ Hz, blue). As in Ref. [84], in this case we use the single-event posterior samples at face value, inconsistently mixing reference frequencies. The third distribution (grey) is obtained from hierarchical posterior samples publicly released by the LVC [84] — in this case, spins are neglected when evaluating selection effects.

The corner plot of Fig. 3.2 shows that including spins in the selection function has a visible effect on the spin population parameters (compare the coloured distributions to the grey). The posterior weights are increased at small ζ — where the isotropic component of the spin-tilt distribution is favoured over spin alignment — and large

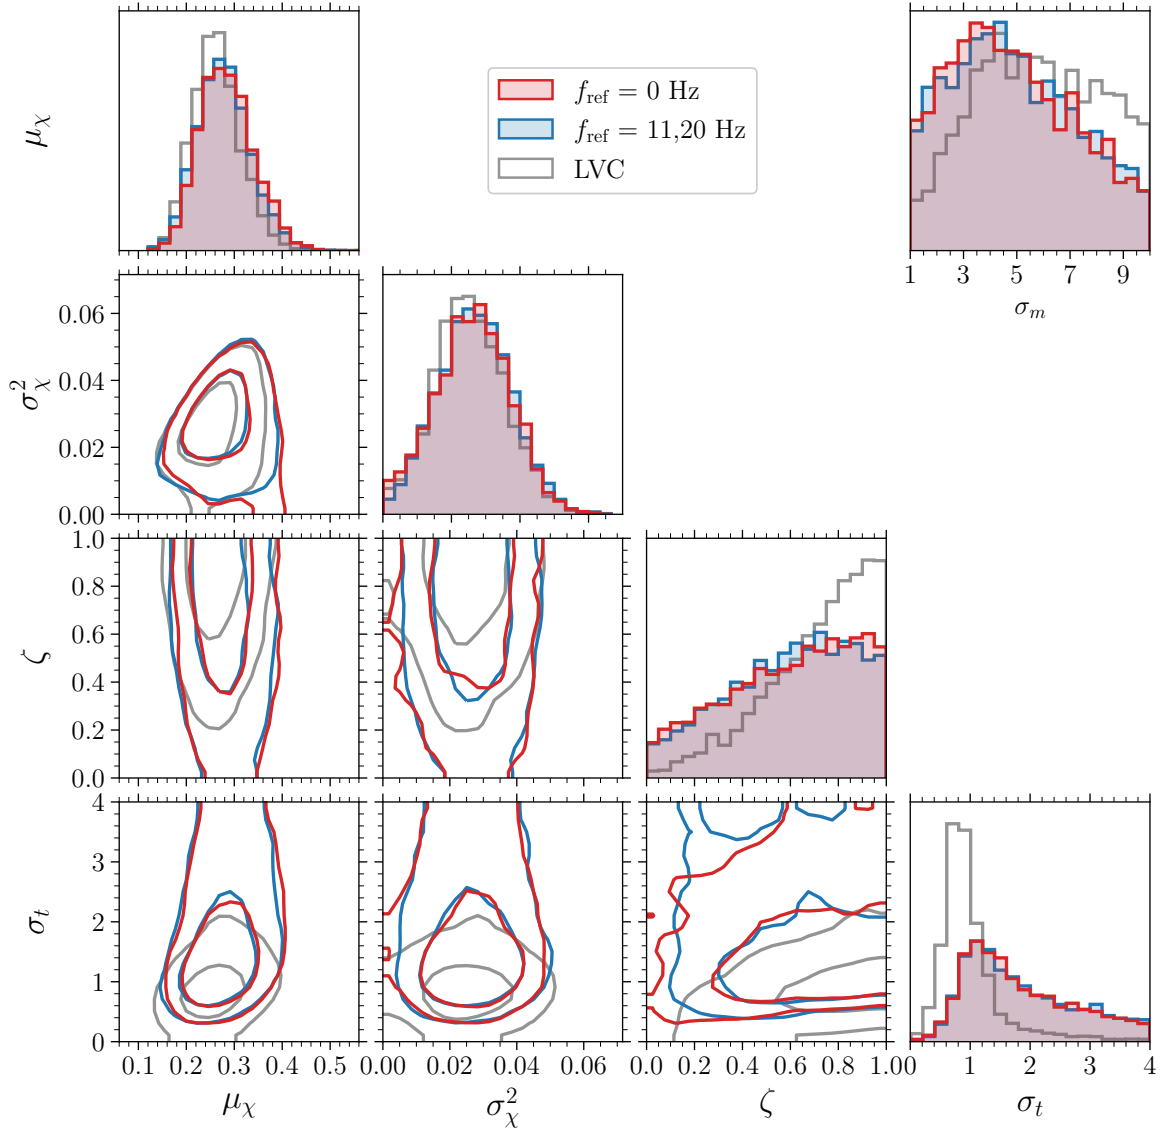


Figure 3.2: The one- and two-dimensional marginal distributions of the spin population parameters. The red curves show our results for population inference with spins defined at 0 Hz. The blue curves represent a control case in which spins are introduced to the selection function but the parameter estimation samples are quoted at the original reference frequency of 20 Hz (11 Hz for GW190521). The grey curves show results from the LVC [84] which neglect spins when estimating selection effects. The population parameters μ_χ and σ_χ^2 (ζ and σ_t) describe the distributions of spin magnitudes (spin directions). The one-dimensional distribution for the width σ_m of the Gaussian component in primary mass appears at the top-right. Contours indicate the 50% and 90% credible regions.




Figure 3.3: The PPDs (bold lines) of the dimensionless spin magnitudes χ and spin tilts θ . The shaded regions mark the symmetric 90% confidence intervals bounded by the 5% and 95% quantiles (lower and upper dashed lines, respectively). The red curves show our results for population inference with spins defined at 0 Hz. The blue curves represent a control case in which spins are introduced to the selection function but the parameter estimation samples are quoted at the original reference frequency of 20 Hz (11 Hz for GW190521). The grey curves show results from the LVC [84] which neglect spins when estimating selection effects.

σ_t — increasing the width of the aligned-spin component. In particular, we (the LVC [84]) find medians and 90% intervals of $\zeta = 0.63_{-0.52}^{+0.33}$ ($\zeta = 0.76_{-0.48}^{+0.22}$) and $\sigma_t = 1.60_{-0.91}^{+1.95}$ ($\sigma_t = 0.87_{-0.45}^{+1.29}$). Therefore, we infer somewhat larger spin misalignments compared to Ref. [84].

This is demonstrated in the PPD of Fig. 3.3, where we find $p(\theta < 45^\circ) = 0.18$ (red and blue) compared to 0.24 (grey) for Ref. [84]. Aligned-spin binaries are easier to detect, implying that their intrinsic distribution is more heavily suppressed compared to other regions of the parameter space. This is a form of Malmquist bias equivalent to the more familiar case of the BH masses, where numerous observed events at $m \approx 30M_\odot$ imply a lower intrinsic merger rate compared to $m \approx 5M_\odot$, even if we have only observed a few low-mass events.

The impact on the parameters μ_χ and σ_χ^2 governing the spin magnitude (χ) distribution, and hence the PPD for χ , is less prominent. This is not surprising, as the best measured spin parameters are the effective quantities χ_{eff} and χ_p . The mass parameters (which are not all shown in Fig. 3.2 for clarity) are largely unaffected by our analysis. We report a minor shift in the Gaussian component of the primary mass distribution as determined by σ_m in Fig. 3.2, which is narrower ($\sigma_m = 4.62_{-3.10}^{+4.32}$) compared to the LVC result ($\sigma_m = 5.69_{-3.60}^{+3.78}$).

Figure 3.4 shows our Bayesian measurement of $\sigma(\lambda)$, i.e., the distribution of Eq. (1.75) across samples of Eq. (1.85), which sets the fraction of events from the inferred population that are observable. Including spins in the detectability returns a lower fraction $\sigma(\lambda) = 0.44_{-0.13}^{+0.18}\%$ compared to $\sigma(\lambda) = 0.57_{-0.19}^{+0.25}\%$ as obtained from the samples of Ref. [84]. In other words, by neglecting spins in the inference, one is tempted to think that the intrinsic and observed population are more similar to each other than they really are.




Figure 3.4: Distributions of the selection function $\sigma(\lambda)$ across the inferred population. The grey curve displays the values provided in the data release of Ref. [84], which neglects spins when computing σ . Including spins in the selection function, in blue is the result for the control case with spins defined at the inconsistent reference frequency of 20 Hz (11 Hz for GW190521), while in the red spins are defined at 0 Hz.

Once selection effects are properly accounted for, we find that considering spins at 20 Hz or 0 Hz has a minor impact (compare blue and red distributions in Figs. 3.2, 3.3, and 3.4). The systematic error one incurs when putting together spin directions corresponding to different evolutionary stages of the binary inspiral is, at present, subdominant compared to statistical errors. While our findings validate the assumptions made so far in the GW literature, we stress that this is an important conceptual point which is addressed here for the first time. Furthermore, while statistical errors are bound to decrease with increasing catalogue size and improved detector sensitivity, systematics will be amplified.

3.4 Conclusions

We solved two conceptual issues that affected all GW population studies so far. An unbiased inference requires the BH spins to be (i) fully included when estimating selection effects and (ii) measured at past time infinity. We find that, with the same GW events and population model, the preference for spin alignment inferred in Ref. [84] is reduced. Notably, this result is the opposite to the recent findings of Refs. [455, 456], although for a very different reason. There it is claimed that the `DEFAULT SPIN` model is inappropriate to describe the current set of events and under-predicts spin alignment. Future applications of non-parametric population modelling (e.g., [457]) will inevitably require a consistent reference among all events.

In the future, systematic biases in population studies will become more influential, for two reasons: (i) single-source measurement uncertainties will decrease as detector sensitivities increase, and (ii) there will be more detections stacked together in the GW catalogue. It will be important, therefore, to project when biases due to inconsistent parameter references are expected to become significant; this is an interesting avenue for further work.

While here we focused on BH spins, the orbital eccentricity is another key property that evolves during the binary inspiral. As GW analyses extend from quasi-circular to eccentric sources [458, 459, 460, 461, 462], we will need to face the issue of building a population model for the eccentricity in a coherent fashion [463]. We anticipate this could be tackled using suitable extensions to the averaging techniques proposed here (cf. Ref. [464]).

The spin directions at past time infinity are equal, to an excellent approximation [359], to those at BH formation [330, 331, 442, 465, 466]. As GW astronomy enters

the large-statistics regime, travelling backward and forward in time when performing population analyses will allow for a direct, systematics-free comparison between GW data and astrophysical models of compact-object formation. Indeed, the importance of these results is being reflected in the work of the LVC who now account for the impact of precessing spins on selection effects in population analyses [84, 85] and have started to propagate their released posterior samples to infinite orbital separations [63, 64, 74, 75].

Which black hole formed first?

Abstract

Different compact-binary formation channels predict identifiable signatures in the astrophysical distributions of source parameters, such as masses and spins. One example within the scenario of isolated binary evolution is mass-ratio reversal: even assuming efficient core–envelope coupling and tidal spin-up in massive stars, a compact binary with a lighter, nonspinning first-born BH and a heavier, spinning second-born BH can still form through mass transfer from the initially more to less massive progenitor. Using GW observations, we measure the fraction of sources in the underlying population with this mass–spin combination and thus identify which BH in the binary was born first. We model subpopulations of nonspinning BHs and, most importantly, nonidentical component spin distributions. We do not find evidence for a subpopulation of BHs with negligible spins and thus measure the fraction of massive binary stars undergoing mass-ratio reversal to be consistent with zero and $< 32\%$ at 99% confidence. The peaks in the distributions of dimensionless spins around 0.2 appear robust, however, and are yet to be explained by progenitor formation scenarios.

Summary and contributions

This Chapter is a reformatted version of our paper Ref. [422] published in collaboration with Davide Gerosa, Floor Broekgaarden, and Nathan Steinle. The idea for this study was motivated in part by Floor’s recent work [467], who additionally provided valuable astrophysical interpretations. In addition to writing the paper and creating the figures, I designed the population models used, wrote the analysis code, and performed the production inference runs. Davide guided the methodology and paper writing. Nathan also contributed to the astrophysical interpretation and discussion of caveats in Sec. 4.4.

We introduce the astrophysical process of mass-ratio reversal (MRR) in massive binary stars and its observational consequences with GWs in Sec. 4.1. In Sec. 4.2, we summarize our modelling and inference procedure. In Sec. 4.3, we present the results of our analyses, in particular the fractions of sources with negligible spins and undergoing MRR. We discuss our findings and their caveats in Sec. 4.4.

4.1 Introduction

In some independent population studies, different approaches have led to contradictory results. Refs. [84, 85] infer that BHs have non-zero spins with support for large misalignments with respect to the binary orbital angular momentum, while Refs. [455, 456] report a large population of nonspinning BHs or a lack of evidence for misaligned spins. Such differences become important when attempting to relate the results of GW inference with predictions of compact-binary formation, in which distinct channels leave identifiable imprints on the population of detectable mergers [269, 270, 271].

In the classical picture of isolated stellar binary evolution, the initially more massive star evolves more rapidly and thus collapses first to form the heavier BH. If angular momentum is lost from the outer envelope of the massive star after efficient transport from the stellar core [468, 469, 470] the heavier BH is expected to have negligible spin [328, 329]. Before the second BH forms, its progenitor may be spun up due to tidal torques from the first-born BH [328, 470, 471, 472, 473, 474]. The less (more) massive BH in the resulting binary therefore has significant (negligible) spin.

On the other hand, if there is significant mass transfer from the initially more to less massive star, the binary mass ratio can be reversed and the first (second) collapse of the now less (more) massive star can form the lighter (heavier) BH. This mechanism has frequently been suggested in the literature to explain Algol binaries and compact object binaries with young neutron stars (e.g., [475, 476, 477, 478, 479, 480]). Crucially, this scenario of MRR leads to GW detections of binary BHs with the more (less) massive BH having non-negligible (negligible) spin [467, 481], as well as different spin misalignments [442]. While the individual BH spins are typically poorly measured [64], in principle the occurrence of MRR can be constrained using GW observations.

To this end, we present a suite of GW population analyses with models designed to identify — via spin measurements — binary BHs that have undergone MRR. The key ingredient is that we must no longer assume the two BH spins are identically distributed as is done in current state-of-the-art analyses [85]. We consider minimal extensions of widely used spin models [85, 482, 483], accounting also for the possibility of small or nonspinning components [456]. From the underlying population we measure the fraction of sources in which the heavier BH has larger spin and, by thresholding the spin magnitudes or including a peak at zero spins, the fraction in which either one or both components have negligible spin. These measurements give insight in the

fractions of binary BHs undergoing MRR.

4.2 Modelling mass-ratio reversal

4.2.1 Catalogue and statistical tools

We take the $n = 69$ detected GW events with FARs $< 1 \text{ yr}^{-1}$ from the first (O1), second (O2), and third (O3) LVC observing runs. As described in Sec. 1.4.3, we use the public LVC posterior samples [74, 75, 484] and injection campaign [485] to evaluate the population likelihood and selection effects. For events detected in O3 the threshold for recovery is set to FAR $< 1 \text{ yr}^{-1}$ in at least one detection pipeline, while for events from O1 and O2 a network SNR > 10 is required [85]. We account for uncertainty in both integrals due to finite sampling as in Ref. [85] (see also Refs. [409, 486, 487, 488]).

The source-frame primary BH masses and mass ratios are modelled using the POWER LAW + PEAK model, while the model for redshifts z corresponds to a merger rate density per unit comoving volume and source-frame time that is a power law in $1 + z$. The parameters and their prior distributions are as in [85]. Our inference on the mass and redshift parameters is unchanged with respect to any of the following spin models. The population-level posterior is sampled using GWPOPULATION [453], BILBY [388], and DYNesty [387].

4.2.2 Astrophysically-motivated spin models

Identical spins

The DEFAULT spin parametrization used in Ref. [85] models the dimensionless spin magnitudes $\chi_{1,2}$ (corresponding to BH masses $m_1 \geq m_2$) as independent and

identical beta distributions that have mean μ and variance σ^2 , with uniform priors in $[0, 1]$ and $[0.005, 0.25]$, respectively. Additional prior constraints on the standard beta distribution shape parameters $\alpha, \beta > 1$ enforce regularity at the prior boundaries (one could alternatively sample these parameters directly as in, e.g., Ref. [489]).

The two polar spin misalignments are modelled independently of the magnitudes but not of each other, being distributed in cosine as a mixture between isotropic (flat) and preferentially aligned (unit-mean Gaussian truncated on $[-1, 1]$) components, indicative of isolated and dynamical formation, respectively. The prior on the mixing fraction ζ is uniform in $[0, 1]$ while that of the aligned-component standard deviation τ is uniform in $[0.1, 4]$.

Nonidentical spins

To assess the difference between BH spin magnitudes, the first extension we consider is to relax the assumption of identical distributions. In this `NONIDENTICAL` model the spin magnitudes employ the same independent beta parametrization but are no longer identical, with means μ_i and variances σ_i^2 ($i = 1, 2$). We additionally allow nonidentical alignments τ_i ; however, since spin magnitudes and tilts are independent, this will not directly affect inferences on $\chi_{1,2}$. The mixing fraction ζ between aligned and isotropic tilts is kept identical for both BHs since they must have formed in the same environment. We impose priors as above.

Zero-spin peaks

Since predictions of isolated binary stellar evolution include BHs with small spins, as discussed in Sec. 4.1, we also consider a model that explicitly includes such subpopulations. In the spirit of Ref. [456], we modify the `NONIDENTICAL` spin model

by adding zero-mean Gaussians truncated within $0 \leq \chi_{1,2} \leq 1$ to the spin magnitude distributions. The prior of the standard deviations ω_i is uniform with $0.01 \leq \omega_i \leq 0.05$ and the mixing fractions f_i between the beta distributions and zero-spin peaks are uniform in $[0, 1]$. We keep the NONIDENTICAL distribution of spin tilts and allow for the possibility of just one (non-) spinning BH by taking the two dimensionless spin magnitude distributions to be independent of each other (such that the joint distribution is a product of two mixtures, rather than a mixture of two products). We refer to this model as NONIDENTICAL + ZEROS.

4.3 Measuring mass-ratio reversal

Here we focus on the differences between the spin magnitude distributions of the two BHs, as is important in the MRR scenario. In the following we quote numerical results with either medians and 90% symmetric intervals, or the upper 99% confidence bound in the case of one-sided posteriors.

4.3.1 Nonidentical spins

We first consider the NONIDENTICAL model, which allows us to ask the question: do the more massive BH components spin more rapidly than the less massive components?

Figures 4.1 and 4.2 presents the results for the NONIDENTICAL model. The parametric forms are the same as in the DEFAULT model, the only difference being the two BH components are no longer identical. In general this results in larger measurement errors, particular for the secondary spin as seen in Fig. 4.2, which is unsurprising since the spin of the less massive component is more difficult to measure [64].

We infer secondary BH spin magnitudes with lower means μ_2 and widths σ_2 , result-

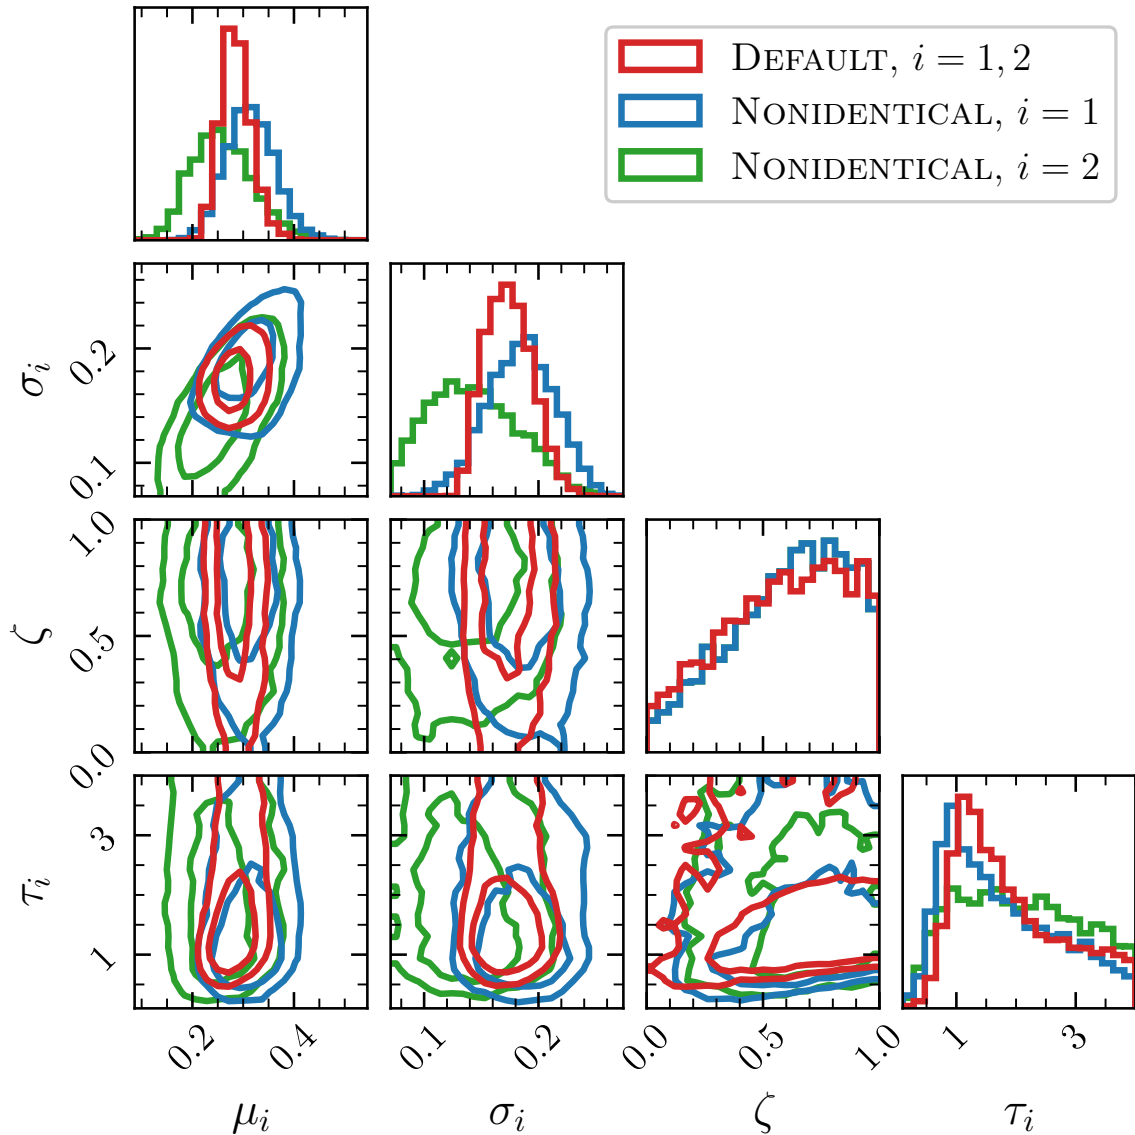


Figure 4.1: The posterior distributions for the spin hyperparameters. The primary, $i = 1$, distributions (secondary, $i = 2$) are plotted in blue (green), and the DEFAULT spin model is shown in red for comparison. The parameters of the NONIDENTICAL model are the means μ_i and standard deviations σ_i of the beta distribution in spin magnitudes, the mixing fraction ζ between aligned and isotropic spin tilt subpopulations (which is the same for primary and secondary BHs), and the widths τ_i of said aligned components. For the DEFAULT model the primary and secondary components are identical ($i = 1, 2$).

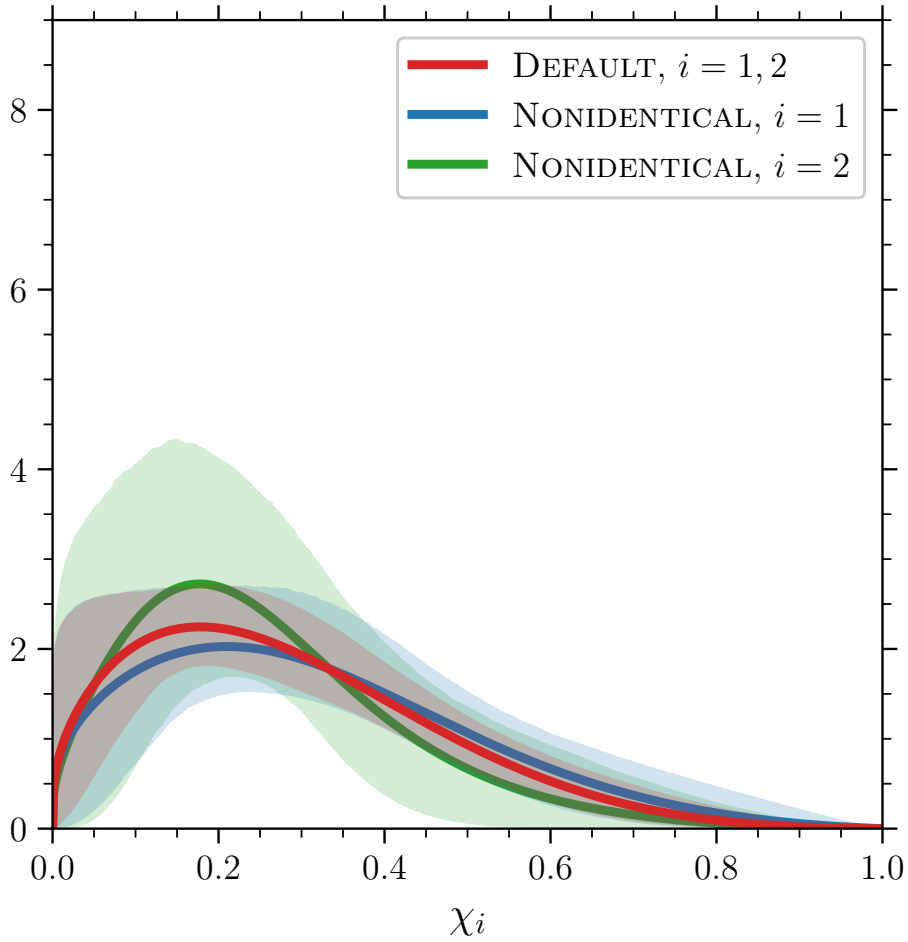


Figure 4.2: The posterior distributions of the spin magnitudes for the NONIDENTICAL model. The primary, $i = 1$, distributions (secondary, $i = 2$) are plotted in blue (green), and the DEFAULT spin model is shown in red for comparison. The solid lines represent the mean distributions, i.e., the PPDs, and the shaded regions correspond to the symmetric 90% credible regions.

ing in secondary spins which are on average lower and peak more narrowly than that of the primaries. The primary hyperparameters are consistent with the DEFAULT posteriors, albeit with larger errors. The component spin magnitudes remain consistent with being identical within the displayed 90% credible regions and feature a peak around $\chi_i \approx 0.2$. Given the beta-distribution parametrization, the probability densities are consistent with zero at $\chi_i = 0$. The mixing fraction ζ between isotropic and preferentially aligned spin tilts is identical for primary and secondary BHs by assumption — hence the overlapping of the $i = 1, 2$ curves in the corresponding one-dimensional histogram — and is unchanged with respect to the original DEFAULT analysis. The Bayes factor over the DEFAULT model is $\log_{10} \mathcal{B} = -0.18$, implying there is no strong preference for either.

Since the spin-magnitude distributions are independent of each other and not identical, i.e., $p(\chi_1, \chi_2 | \lambda) = p(\chi_1 | \lambda_1) p(\chi_2 | \lambda_2)$, the probability that the more massive BH spins more rapidly than the less massive BH is

$$P(\chi_1 > \chi_2 | \lambda) = \int_0^1 p(\chi_1 | \lambda_1) \int_0^{\chi_1} p(\chi_2 | \lambda_2) d\chi_2 d\chi_1, \quad (4.1)$$

where $\lambda_i = (\mu_i, \sigma_i)$ are the hyperparameters individually characterizing the spin distributions. Our measurement for the probability of more rapidly spinning primaries, given the Bayesian uncertainty $p(\lambda | \{d\})$ in the hyperparameters of the NONIDENTICAL model, is given by the blue distribution in Fig. 4.3. We find $P(\chi_1 > \chi_2 | \lambda) = 0.58_{-0.25}^{+0.21}$. If the two distributions were instead identical, Eq. (4.1) reduces to $P(\chi_1 > \chi_2 | \lambda) = P(\chi_1 < \chi_2 | \lambda) = 1/2$. We infer a median value above this mid-point and 71% posterior support for $P(\chi_1 > \chi_2 | \lambda) > 1/2$, indicating mild evidence that primary BHs spin more rapidly.




Figure 4.3: The probability that the heavier BH is more rapidly spinning in our NONIDENTICAL (blue) and NONIDENTICAL + ZEROS (red) models. The uncertainty is due to the uncertainty in the model hyperparameters, as measured from GW data. The coloured vertical dashed lines correspond to the medians and symmetric 90% confidence intervals. The vertical dashed black line at $P(\chi_1 > \chi_2) = 1/2$ corresponds to the result for identical spin distributions.

While this measurement compares spin components across the entire unit interval, the isolated binary evolution scenario discussed in Sec. 4.1 predicts small BH spins. To relate spin measurements more closely to these predictions we place a threshold, χ_0 , below which spin magnitudes are defined as negligible. The probabilities that a BH spin lies below or above this threshold are given by

$$P(\chi_i < \chi_0|\lambda) = 1 - P(\chi_i > \chi_0|\lambda) = \int_0^{\chi_0} p(\chi_i|\lambda)d\chi_i. \quad (4.2)$$

We find that $4.8_{-4.5}^{+7.0}\%$ ($4.0_{-3.9}^{+11}\%$) of primary (secondary) BHs have negligible spins when imposing a threshold $\chi_0 = 0.05$ [467]. These results are one-sided and peak at zero, skewing the posterior measurements away from the boundary. Instead, one can consider these as measurements of $P(\chi_i < \chi_0|\lambda) \approx 0$ with one-sided uncertainties $P(\chi_1 < \chi_0|\lambda) < 0.15$ and $P(\chi_2 < \chi_0|\lambda) < 0.20$.

Since the spin magnitudes are modelled independently, the joint probability that, e.g., the primary BH spin is non-negligible and the secondary is negligible is simply given by the product $P(\chi_1 > \chi_0, \chi_2 < \chi_0|\lambda) = P(\chi_1 > \chi_0|\lambda)P(\chi_2 < \chi_0|\lambda)$. If one assumes that the entire binary BH population formed through isolated evolution and that there is no other mechanism by which primary BHs can acquire spin, this probability can be interpreted as the fraction of MRR sources. On the other hand, sources with $\chi_1 < \chi_0$ but $\chi_2 > \chi_0$ correspond to the standard isolated binary scenario in this interpretation. Again taking $\chi_0 = 0.05$, we find $< 12\%$ ($< 10\%$), or $3.8_{-3.7}^{+11}\%$ ($4.6_{-4.3}^{+6.9}\%$), of the underlying population falls into the MRR (standard) spin scenario. The majority of the posterior support lies away from zero spins with $P(\chi_1, \chi_2 > \chi_0|\lambda) = 0.90_{-0.10}^{+0.07}$, while the fraction of sources with two negligible spins is at the sub-percent level.




Figure 4.4: The joint probabilities that the spins of the more and less massive BHs, $\chi_{1,2}$ lie above or below a given threshold, χ_0 , in the NONIDENTICAL model. The blue (orange) curve represents the case in which both components are above (below) threshold. We interpret the red and green curves as the MRR and standard isolated evolution scenarios, respectively. The solid lines give the median probabilities over the hyperparameter uncertainties as a function of χ_0 , while the shaded bands enclose the 90% credible regions. The vertical dashed black line indicates the fiducial threshold value of $\chi_0 = 0.05$.

In Fig. 4.4 we plot the measurements of these probabilities as a function of the spin threshold χ_0 ; of course, when $\chi_0 = 1$ ($\chi_0 = 0$) all sources are classified as having two (non-) negligible spins. There is no support in the population for binaries with both spins above ≈ 0.6 , while $P(\chi_1, \chi_2 < \chi_0 | \lambda)$ dominates for $\chi_0 \gtrsim 0.8$ since the component spin distributions taper off above this value (see Fig. 4.1). The fraction of sources that are placed in the MRR ($\chi_1 > \chi_0, \chi_2 < \chi_0$) and standard ($\chi_1 < \chi_0, \chi_2 > \chi_0$) scenarios peak at $\chi_0 \approx 0.3$ and $\chi_0 \approx 0.25$, respectively. The locations of these peaks are determined by the locations of the peaks in the component distributions (see Fig. 4.1) and the former is slightly higher due to the slight preference for primaries with larger spins.

4.3.2 Zero-spin peaks

The above results are inferred using the NONIDENTICAL spin model, which does not account for a subpopulation of sources with small spins. Instead of relying on the ad-hoc small-spin threshold, the NONIDENTICAL + ZEROS model explicitly allows for fractions f_i ($i = 1, 2$) occupying a peak at zero spin magnitudes [456]. In Figs. 4.5 and 4.6 we present the results for the NONIDENTICAL + ZEROS model. The posteriors on the parameters of the beta spin components are unchanged with respect to the NONIDENTICAL analysis, however, and are hence not plotted; one can simply refer to Fig. 4.1 instead. The Bayes factor over the DEFAULT model is $\log_{10} \mathcal{B} = -0.05$, again indicating that the NONIDENTICAL + ZEROS model describes the population equally well as the DEFAULT model.

The fractions f_i of both primary and secondary BHs with negligible spins are one sided and peak at zero, with one-sided 99% credible bounds $f_1 < 0.46$ and $f_2 < 0.36$ (or medians and symmetric 90% intervals of $f_1 = 0.13_{-0.12}^{+0.22}$ and $f_2 = 0.08_{-0.07}^{+0.18}$)

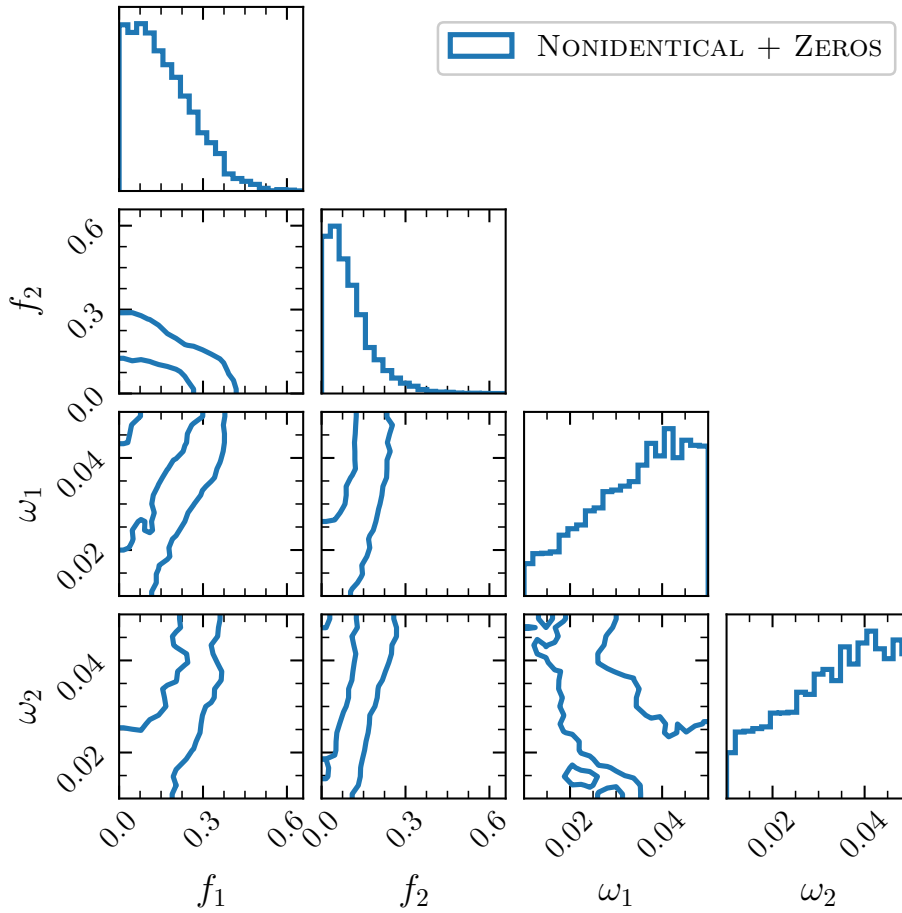


Figure 4.5: The posterior distributions for the hyperparameters of the NONIDENTICAL + ZEROS model characterizing the peaks at zero spins. The occupation fractions of these peaks are f_i and the standard deviations are ω_i ($i = 1, 2$). The posteriors of the beta distribution parameters are identical to the NONIDENTICAL case and thus not displayed.

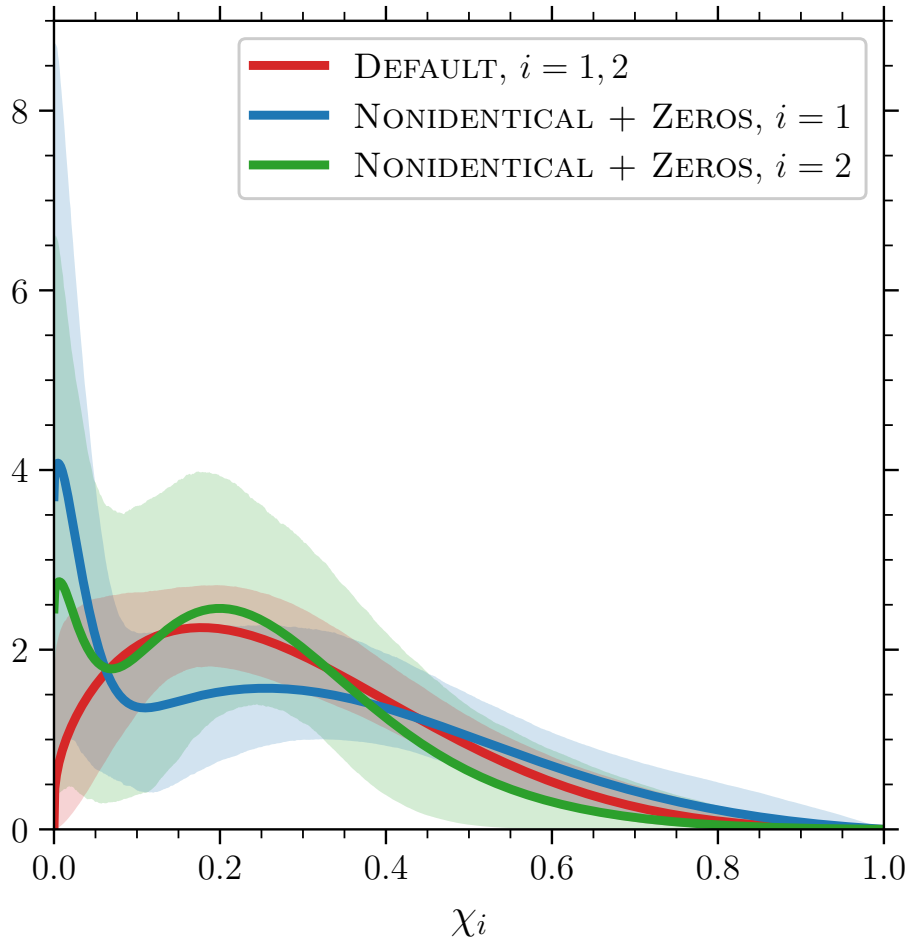


Figure 4.6: The posterior distributions for the spin magnitudes determined by the NONIDENTICAL + ZEROS model. The primary (secondary) spin is plotted in blue (green), and the DEFAULT model is displayed in red for comparison. The solid lines represent the mean distributions, i.e., the posterior population distributions, and the shaded regions correspond to the symmetric 90% credible regions.

with no posterior support above 0.6. These are consistent with the NONIDENTICAL measurements of $P(\chi_i < 0.05|\lambda)$ (recall that the prior bounds on the widths of the zero-spin peaks ω_i are uniform in $[0.01, 0.05]$). For comparison, we compute $P(\chi_1 > \chi_2|\lambda) = 0.56_{-0.21}^{+0.19}$ with $P(\chi_1 > \chi_2|\lambda) > 0.5$ at 68% confidence for the NONIDENTICAL + ZEROS model, consistent with the NONIDENTICAL model and only slightly lowered due to the additional but small contribution from the zero-spin peaks.

The posteriors of the zero-spin peak widths ω_i feature more support at the upper boundaries. We find the same result when widening the prior domain to $[0.01, 0.1]$, inevitably leading to larger inferred values of f_i but which still peak at zero. The same behaviour is observed in analyses which do not impose a prior cut on the population likelihood determined by the effective sample sizes of Monte Carlo integrals [490], as used in this work, implying the inferred posteriors are robust. In contrast to the findings of [456], these results altogether point to a lack of evidence for a subpopulation of small spins, or alternatively an inability to measure such a feature in the spin distribution. On the other hand, the peaks in χ_i between 0.2–0.3 remain apparent in Fig. 4.6. While allowing for the possibility of a subpopulation of BHs with small spins results in a visible spike in density at $\chi_i = 0$, the probability there is still low due to the small volume of occupied parameter space ($\chi_i < 0.05$). Unlike the DEFAULT and NONIDENTICAL models, the mixture parametrization of the NONIDENTICAL + ZEROS model allows the χ_i distributions to be constrained away from zero at $\chi_i = 0$ within the 90% confidence interval. Within this uncertainty, both primary and secondary BH spins are still consistent with a single continuous distribution, however.

As similarly described in Sec. 4.3.1, the independence of the dimensionless spins means that the fraction of sources with both BHs in the beta distribution component is $f_\beta = (1 - f_1)(1 - f_2)$ while that for both in the small-spin peaks is $f_0 = f_1 f_2$.

The fraction of binaries in which the primary BH has small spin and the secondary is spinning — i.e., the standard isolated binary evolution channel — is $f_S = f_1(1 - f_2)$, while the converse — the MRR scenario — is $f_R = (1 - f_1)f_2$. The spin tilt distributions separate the isolated and dynamical evolution channels with a fraction ζ of sources in a preferentially aligned component, as predicted for binaries formed in the field, while the rest are isotropic. Therefore, from the total population we can identify the fraction of sources forming in the isolated channel that are placed in each sub-channel $X \in \{\beta, 0, S, R\}$ as ζf_X (where $\sum_X f_X = 1$).

We present the measurements of these branching fractions from our NONIDENTICAL + ZEROS population analysis in Fig. 4.7, including both the total and isolated fractions, f_X and ζf_X , as well as the fraction of non-isolated sources, $1 - \zeta$. Binaries with both BHs placed in the beta spin distributions make up the majority of the population with $f_\beta = 0.77^{+0.16}_{-0.20}$. The fractions in which either or both BH spins are negligible all peak at zero: the standard and MRR fractions are $f_S < 0.43$ and $f_R < 0.32$ ($f_S = 0.12^{+0.21}_{-0.11}$ and $f_R = 0.06^{+0.17}_{-0.06}$), respectively, while $f_0 < 0.06$. The fraction of the total population within the isolated formation channel is not well measured at $\zeta = 0.62^{+0.33}_{-0.48}$ (the posterior for ζ is the same in both the NONIDENTICAL and NONIDENTICAL + ZEROS models). Since by definition $0 \leq \zeta \leq 1$, the isolated fractions are of course lowered, i.e., $\zeta f_X \leq f_X$. In short, at 99% confidence we measure $\zeta f_R < 0.23$, $\zeta f_S < 0.33$, and $\zeta f_0 < 0.04$.




Figure 4.7: Posterior fractions of binary BHs with different spin combinations from the NONIDENTICAL + ZEROS model. In blue (orange) is the fraction with two (non-) spinning BHs, f_β (f_0). In red (green) is the fraction f_R (f_S) of MRR (standard) sources. The dashed lines indicate the total fractions within the astrophysical population, f_X ($X \in \{\beta, 0, S, R\}$, $\sum_X f_X = 1$), while the solid filled histograms represent the subset which have preferentially small spin tilts, ζf_X . The solid black line gives the posterior measurement on the mixing fraction of isotropic spin tilts, $1 - \zeta$.

4.4 Gravitational-wave observations against astrophysical predictions

Altogether, our results (given the current GW catalogue) point to a lack of evidence for large subpopulations of merging stellar-mass binary BHs with negligible spins, or an inability to measure a sharp feature such as this in our models. This is in disagreement both with previous measurements from GW data [456] and predictions of isolated binary evolution [273, 328, 329, 442, 467, 481].

4.4.1 To reverse or not to reverse

Ref. [456] finds that $54_{-25}^{+21}\%$ of BHs have negligible spin,¹ whereas we find in Sec. 4.3.2 that $< 46\%$ (36%) of primary (secondary) BHs have negligible spin with fractions peaking at zero, $23_{-16}^{+20}\%$ of binaries have at least one BH with negligible spin, and $\sim 1\%$ have both BHs with negligible spins. We model negligible spins with a truncated Gaussian distribution of non-zero width centred at zero, whereas their zero-spin peak is a delta distribution at exactly $\chi_{1,2} = 0$ resulting from PE runs performed with nonspinning priors. They also analysed the binary BH mergers observed up to the end of the first half of O3, whereas the catalogue analysed here contains all binary BHs through the end of O3. To perform the equivalent O3 analysis allowing for nonidentical spin distributions would require PE runs for each event with nonspinning priors placed on either or both BHs, resulting in three additional stochastic-sampling runs per event (and per waveform); such analyses are beyond the scope of our study. However, our results are robust with respect to the choice of population model; we measure consistent

¹This value is taken from their updated [erratum](#).

fractions of BHs with negligible spins in both the NONIDENTICAL and NONIDENTICAL + ZEROS models.

Ref. [490] presented a detailed study on the intrinsic distribution of binary BH spins to address the issue of a potential subpopulation of negligible spins. They focus on the technical details of population modelling and analysis that can lead to erroneous conclusions on such features. They report findings consistent with our own and inconsistent with Ref. [456], whether modelling effective or component spin. In particular, as in this work they also find a one-sided posterior measurement peaking at zero for the fraction of BHs with small spin, with an upper limit $< 60\%$. In addition to modelling differences, Ref. [490] finds that the Bayes factor between spinning and nonspinning prior hypotheses — a crucial ingredient in the analysis of [456] — was estimated incorrectly for at least one event. We focus on the astrophysical inferences these results allow us to make, in particular constraining the occurrence of MRR, which unlike in their analysis means crucially we model BH spins as nonidentical.

Ref. [481] use rapid population synthesis to investigate the occurrence of MRR in stellar progenitors of binary BHs. They find that, depending on the efficiency of accretion from the initially more to less massive star, up to 72% of systems undergo MRR (i.e., have a heavier second-born BH). Similarly, Ref. [467] assesses both the detectable and intrinsic MRR fractions in a range of population-synthesis models. They find typical intrinsic MRR fractions $> 30\%$, lying between 11% and 82% across their models. These recent results are in general agreement with older predictions which found fractions $\sim 10\text{--}50\%$ [273, 465]. From GW data we measure the related fraction in the intrinsic population to be $< 32\%$ (99% confidence) based on our definitions and astrophysical assumptions. Consistency between these population-synthesis predictions and our measurements therefore requires lower accretion efficiencies between massive

stars in order to result in lower MRR rates. However, even assuming a conservatively low accretion efficiency of 0.25, Ref. [481] finds MRR fractions $> 18\%$, though most models have larger fractions. The low value of our fraction stems from the fact that, even when accounting for the possibility of negligible-spin subpopulations, we infer the presence of only a small number of negligible-spin BHs in the astrophysical population. This particular MRR scenario requires the secondary BH — the first-born BH that was initially the more massive star — to have small spin.

More generally, accounting for the inferred astrophysical distribution of binary BH spins solely with the isolated binary formation channel requires some mechanism to produce sufficiently large spins $\chi_i > 0.05$. For example, various scenarios including weak core-envelope coupling, tidal synchronization, and significant accretion can lead to binary BHs with larger spins [330, 481, 491, 492]. The apparent lack of negligible-spin sources in the intrinsic population may imply that, if isolated formation is the dominant contribution, the assumptions typically made in binary population synthesis (e.g., efficient angular momentum transport in massive stars) are incorrect, or other formation channels dominate the GW merger rate. Large spins can be formed in environments where binary BHs interact dynamically, such as dense stellar clusters; hierarchical mergers — in which binary components are the remnants of previous mergers (see Ref. [332] for a review) — lead to high spins with a peak in the distribution around $\chi_i \approx 0.7$ (see, e.g., Refs. [334, 335, 493, 494]).

Nevertheless, this is away from the measured peaks in the distributions of dimensionless spins at $\chi_i \approx 0.2\text{--}0.3$, which is emerging as a solid outcome from several population fits (see Refs. [84, 85, 421, 495] in addition to our own results). To the best of our knowledge this feature has so far not been identified as a preferred outcome in current stellar-binary evolution models, and therefore requires further attention from

the perspective of progenitor population modelling. We note that, though there is no evidence for large subpopulations of binary BHs with negligible spins, this does not mean that individual nonspinning sources do not exist. Indeed, in our NONIDENTICAL + ZEROS model there is posterior support for primary and secondary BHs with spins $\chi_i = 0$ at 90% confidence, a conclusion not allowed by the DEFAULT and NONIDENTICAL models (see Fig. 4.5). However, given the current catalogue and measurement uncertainties, the two BH spins remain consistent with being drawn from a single continuous distribution.

4.4.2 Modelling and astrophysical uncertainties

Inference results are subject to modelling errors; the implicit conditional in our Bayesian measurements is that of the population model. This is a crucial point, which we have explicitly indicated by reporting the conditional dependence on λ in all probabilities of Sec. 4.3. Indeed, misspecification can lead to biased results [496] and we do not pretend that the simple parametric models used here fully describe the astrophysical distribution of merging stellar-mass binary BHs.

In particular, the masses, spins, and redshifts are assumed to be independent of each other when in reality they may be correlated [497, 498, 499, 500]. While the distribution of spin tilts is taken to be a mixture of two contributions, mimicking isolated and dynamical formation, it is likely that multiple channels contribute to the merger rate with distinct features across all source parameters [501, 502]. On the other hand, observables such as spin tilts may not be such clear discriminators as previously assumed. For binary BHs formed in isolation, core-collapse natal kicks [330, 331] and mass-transfer episodes [503] may result in significant misalignments and spin precession. Stellar-binary processes such as common-envelope episodes may also

occur in dynamical environments [504, 505], leading to both aligned and misaligned spins [506]. In these cases, the fraction ζ of sources with preferentially aligned spins employed in Sec. 4.3.2 is not a robust discriminator of isolated binary formation.

We chose to model the spin magnitudes independently of the spin orientations. In reality, the spin magnitude distributions may be different even within the two channels that divide the spin orientations — e.g., the small-spin subpopulation being relevant only for binaries formed in isolation and therefore having preferentially aligned spins. We do not allow for correlations between the spin magnitudes and the mixing fraction ζ , and more complex models would include correlations between these parameters motivated by the astrophysical settings in which stellar-mass binary BH formation occurs. However, given the statistical uncertainties in the current GW catalogue and systematic uncertainties in models of binary BH formation, there is a limit to the available constraining power with a given number of population model parameters. Our NONIDENTICAL and NONIDENTICAL + ZEROS models can be considered as minimal extensions to state-of-the-art analyses that may already be demonstrating limited constraining power given the number of model parameters, as seen in the broad measurement of the mixing fraction ζ .

We assume the key MRR signature in GW observations to be spinning primary BHs that formed second from the collapse of the initially less massive stellar progenitor and nonspinning secondary BHs forming first from the initially more massive star. As discussed, several mechanisms can influence the resulting BH spin and therefore the fraction of binaries in the intrinsic population with this spin configuration may not be a clean indicator of the fraction of MRR systems. Contamination can result both from binaries that do not experience MRR but attain our assumed spin configuration and from systems that do undergo MRR but end up merging with different spins.

It is important to stress that the question we are asking the GW data here is, in a nutshell: “which BH spins the most?” The answer is then interpreted as the relative fraction of GW sources that undergo MRR by relying on a (hopefully broad, but admittedly specified) set of astrophysical expectations, thus translating the targeted question into: “which BH formed first?”

4.5 Conclusions

In isolated binary evolution, mass transfer between the massive stellar progenitors can lead to a reversal of the binary mass ratio [467, 481]. The key consequence of this is that, if ultimately forming a merging compact binary, the initially less (more) massive star becomes the (non-) spinning primary (secondary) BH observed with GWs. We constrained the occurrence of this astrophysical process using GW data and population models that allow for nonidentical distributions between the two BH spins.

In contrast to Ref. [456], we did not find evidence for large subpopulations of negligible-spin BHs. Instead, with current GW data the BH spins are consistent with being drawn from identical, continuous distributions with non-vanishing support at small spins. We measure the occupation fractions of peaks at zero spins to be $< 46\%$ and $< 36\%$ for primary and secondary BHs, respectively. Our results have now been independently confirmed [490, 507].

There is vanishing posterior support for binaries with both dimensionless spins > 0.6 . We measure the intrinsic fraction of sources in which the heavier BH has significant spin while the lighter does not to be $< 32\%$ and peak at zero, and reinterpret this as the fraction of binaries whose progenitors undergo MRR. A small proportion of MRR sources is also in tension with recent predictions from isolated binary evolution

[467, 481], implying either that such sources do not contribute to the GW merger rate or that the standard assumptions of isolated binary evolution we made do not hold.

On the other hand, the previously reported over-density in the dimensionless spin distributions around 0.2–0.3 at present appears to be a robust feature even when accounting for nonidentical spins and small-spin subpopulations. A clear explanation for this feature from common formation channels remains to be found.

Simulation-based deep learning of gravitational-wave populations

Abstract

Conventional GW population models are parametrized with simple functional forms. We propose an emulation-based approach that can compare astrophysical simulations directly against GW data. We combine state-of-the-art deep-learning techniques with hierarchical Bayesian inference to: (i) construct a flexible single-channel population model that accurately emulates multi-modal and correlated distributions in multiple dimensions, (ii) rapidly estimate population selection effects, and (iii) recover the branching ratios of subchannels within the population. Applying our approach to simplified simulations of hierarchical BH formation, using the latest LVC GW catalogue we find that: host-environment escape speeds $< 100 \text{ km s}^{-1}$ are favoured, though around 37% of first-generation merger remnants are retained in their hosts; the maximum mass of first-generation BHs is $\approx 40M_{\odot}$; there is multi-modal substructure in the astrophysical distributions of both binary BH masses and spins; and mergers with a higher-generation BH make up at least 14% of the underlying population.

Summary and contributions

This Chapter is reformatted from our paper published in Ref. [413] in collaboration with Davide Gerosa and Stephen Taylor. The genesis of this project was some of their previous work [404]. Steve provided code for an initial version of the project using Gaussian process regression (GPR) and tools for hierarchical Bayesian inference. Davide wrote the initial version of the population simulation code and first used neural networks for the interpolation. Davide and I performed the calculation to rewrite the hierarchical likelihood in terms of the observed parameter space. I completed the production code and runs, including the simulation generation, neural-network training, and population inference. The figures were made by me, as was the paper writing with invaluable guidance and inputs from Davide and Steve.

We introduce the concept of simulation-based inference and motivate its utility for GW population studies in Sec. 5.1. In Sec. 5.2 we describe our simple approach to simulating hierarchical-merger distributions. We lay out the statistical and deep-learning tools employed in Sec. 5.3. Our deep-learning-enhanced statistical pipeline is validated with mock GW catalogues in Sec. 5.4. In Sec. 5.5, we report the results of our inference on the latest catalogue of GW events, discussing the astrophysical implications and comparing to recent related works. We finish with a summary of future extensions to our work and concluding remarks in Sec. 5.6.

5.1 Introduction

Given the catalogue of GW detections, one can take two approaches to assess the underlying astrophysical population of binary BHs. In a simulation-based analysis,

sources are synthesized — accounting for as many detailed astrophysical processes as are known or are computationally feasible — to form distributions of detectable merging binaries. By varying population-level input parameters controlling binary evolution (e.g., common-envelope efficiency, strength of supernova kicks, etc.), one can assess the degree of consistency with the observed events; for reviews see, e.g., Refs. [269, 270, 271]. However, such simulations are typically computationally intensive and large uncertainties remain on key parameters; see, e.g., Refs. [508, 509].

The second approach is to first construct a model of the astrophysical distribution of source parameters — which is conditionally dependent on given population-level parameters controlling its shape (the “hyperparameters”, e.g., mass cutoffs or spectral indices) — and use the observed catalogue to perform a hierarchical Bayesian inference (Sec. 1.4.3) that accounts for observational biases (e.g., that heavier sources are easier to detect). This statistical analysis is hierarchical in the sense that one uses previous Bayesian measurements of the binary BH source parameters to then measure said hyperparameters [402, 403]. The population model used could be as in the previous approach such that the distribution is known only at discrete values of the hyperparameters, but this would allow only for single posterior evaluations for relative comparisons (e.g., via Bayes factors) and leave some of the hyperparameters unconstrained; (see, e.g., Refs. [502, 510] for examples of this approach).

On the other hand, a population model that can be continuously evaluated across the population-level parameter space can be used to make Bayesian measurements of the hyperparameters (again, see Sec. 1.4.3 for details). This requirement typically necessitates simple, quick-to-evaluate parametric forms with statistical independence between source parameters (see, e.g., the models used in Refs. [83, 84, 85] and in Chapters 3 and 4) to enable efficient sampling of the hierarchical posterior [Eqs. (1.80,

1.85)]. The disadvantage of this approach is that it is inherently phenomenological with a discretionary selection of the employed functional forms. Recent work has sought to improve parametric population models by addressing potential correlations between mass and spin parameters [85, 497, 511] and assessing the suitability of spin parametrizations [421, 455, 456] since accurate inference requires appropriate models [496]. Along other lines, the flexibility of population analyses can be improved with semi-parametric and non-parametric models [457, 512, 513, 514, 515].

Previous studies have focused on combining the simulation-based and parametric approaches: an emulator constructed with sufficient accuracy to rapidly synthesize predictive distributions over the hyperparameter space can be adopted in place of discrete simulations or parametrized models in the Bayesian inference pipeline. Such models leverage the advantages of efficient hyperposterior sampling and direct astrophysics-to-GW data comparison provided by each approach.

A first step in this direction within the context of GW population inference was taken in Ref. [404] (see also Ref. [516]). Compressed principal components of binned simulation data were emulated over low- (typically one- or two-) dimensional source- and population-level parameter spaces using GPR. However, this emulation approach was shown to be unsuitable for extension to more complex higher-dimensional modelling scenarios due to poor predictive accuracy and infeasible computational requirements [517, 518]. These issues were tackled in Ref. [519] by employing deep-learning techniques to construct simulation-informed population models; in particular, the conditional density estimator takes the form of a generative neural network known as a normalizing flow [520, 521, 522]. In general, neural networks are powerful tools that offer greater flexibility when employed as functional emulators. In their case, normalizing flows prompted population studies considering the scenarios of primordial BHs

[523] and mixture models between isolated and dynamical evolution [501].

In this work, we develop complementary deep-learning techniques that build on the advancements of Refs. [404, 519] by pushing the emulated parameter space dimensionality and introducing new neural network applications. We employ fully connected deep neural networks (DNNs; also known as multilayer perceptrons) to act as the population model and to capture the effect of GW detection biases on the population of observed binary BH events (see also Refs. [410, 411, 412, 414] for machine-learning approaches to estimating selection effects).

Motivated by evidence for large masses in the observed GW catalogue, we apply these deep-learning techniques to binary BH populations containing hierarchical mergers, in which component BHs may be the remnants of (multiple) previous mergers [332]. These so-called “higher-generation” BHs may explain the outlier properties of events such as GW190412 [56, 524, 525, 526, 527], GW190521 [58, 59, 333, 458, 528, 529, 530, 531] (though see also Refs. [532, 533], which find that these events may in fact be consistent with the population), and GW190814 [57, 534, 535, 536]. The presence of hierarchical mergers in binary BH populations is crucially dependent on the escape speeds of dynamical host environments (e.g., young star clusters, globular clusters, and nuclear star clusters [333]; see Sec. 1.2.3) and the magnitudes of gravitational recoils received due to the anisotropic emission of GWs [537, 538, 539].

Our DNN population model learns from simple simulations of cluster-like environments [431, 540], which account for the retention and ejection of merger remnants due to GW kicks. We model the joint distribution of four source parameters — two masses and two effective spins, which present identifiable features due to the influence of higher-generation BHs [429, 541] — and six population-level (hyper-) parameters. These hyperparameters control the population properties of first-generation BHs born

in stellar collapse, the binary pairing process, and host escape speeds. We also train a DNN to predict the fractional contributions of the population-dependent first-, mixed-, second-, and higher-generation binary BHs.

We illustrate our procedure schematically in Fig. 5.1, in which each element represents a single modelling process, arrows direct the one-way flow of information between them, and rows group distinct stages of our pipeline. The first row represents simulations — controlled by population-level parameters — of binary BH mergers characterized by a complete set of source-level parameters that are condensed into a set of modelled parameters. In the second row we list the postprocessing performed on the simulated data. The outputs for each simulation are the joint probability density of modelled source-level parameters conditioned on the population-level parameters, the expected fractional number of detectable sources, and the relative contributions from each hierarchical merger generation to the total population, as listed in the third row. These discrete sets of evaluations are transformed into continuous functions using deep learning in the fourth row. These DNN functional emulators, listed in the fifth row, are employed in conjunction with data from the GW events detected to date to perform a hierarchical Bayesian inference and ultimately constrain the population of merging stellar-mass binary BHs, as illustrated in the final row. Each ingredient and the relevant symbols are defined throughout the paper.

5.2 Hierarchical-merger simulations

We model the retention and ejection of merger remnants in “clusters” — which here simply refers to a collection of BHs in an environment with constant escape speed v_{esc} . We use the setup described in Ref. [540] (see Refs. [524, 541] for additional

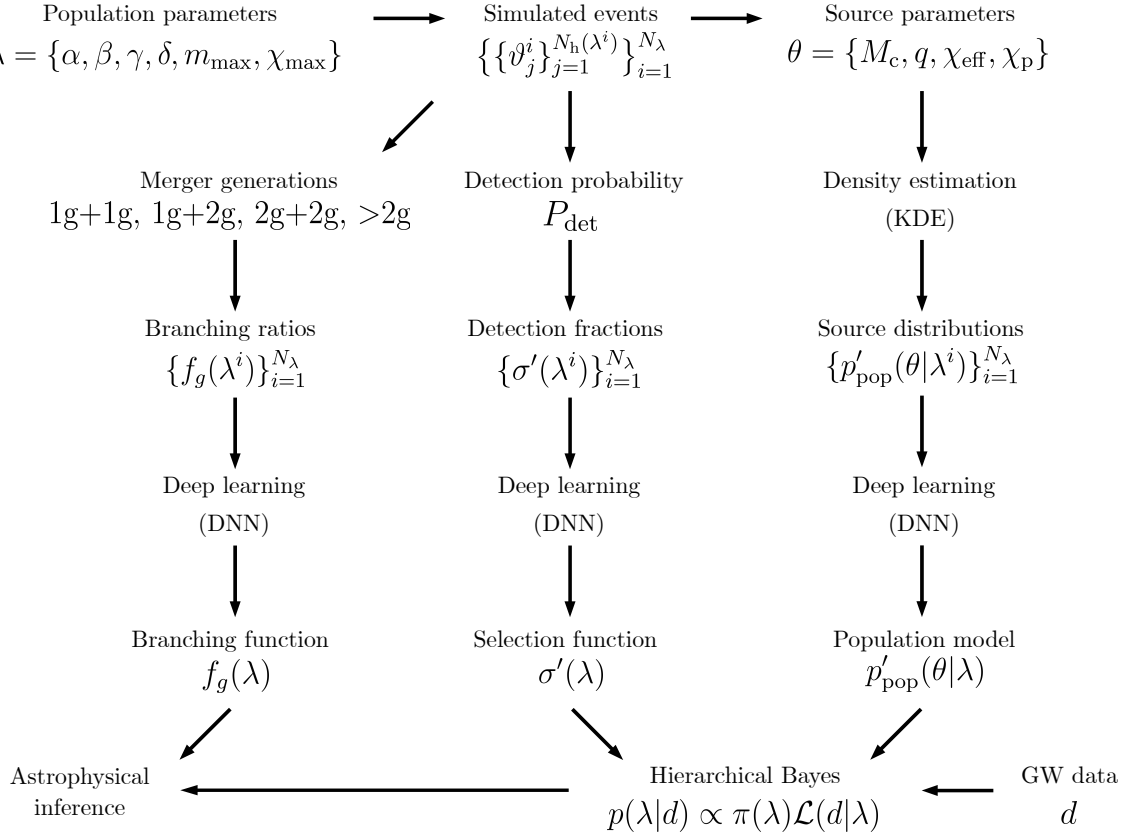


Figure 5.1: Schematic diagram of our population modelling and inference procedure. Arrows indicate information that is passed from one element to another, and elements that occur at the same stage of the pipeline are grouped into rows. The first row represents simulations of binary BH mergers, while the second lists postprocessing applied to the simulated data to produce the training data in the third row. We leverage deep learning, shown in the fourth row, by constructing DNNs to act as functional emulators for key ingredients of GW population inference, indicated within the fifth row. In the final row, the deep-learned selection function and population model are combined with data from GW catalogues to feed into a hierarchical Bayesian inference which, along with a third DNN to predict branching fractions between subpopulations, is used to make conclusions about the underlying distribution of merging stellar-mass binary BHs.

	Parameter	Symbol	Range
Population, λ	Primary pairing slope	α	$[-10, 10]$
	Secondary pairing slope	β	$[-10, 10]$
	1g mass slope	γ	$[-10, 10]$
	Escape-speed slope	δ	$[-10, 10]$
	Maximum 1g mass	m_{\max}	$[30M_{\odot}, 100M_{\odot}]$
	Maximum 1g spin	χ_{\max}	$[0, 1]$
Source, θ	Source-frame chirp mass	M_c	$[5M_{\odot}, 105M_{\odot}]$
	Mass ratio	q	$[0, 1]$
	Effective aligned spin	χ_{eff}	$[-1, 1]$
	Effective precessing spin	χ_p	$[0, 2]$

Table 5.1: Parameters in our model of hierarchical binary BH merger populations, the symbols we use to identify them, and their bounds. The population parameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$ determine the shape of the distribution of first-generation BHs and the properties of the host cluster that can lead to repeated mergers. The bounds on the power-law indices are broad such that the range of training simulations can incorporate more restrictive prior bounds. The source parameters $\theta = \{M_c, q, \chi_{\text{eff}}, \chi_p\}$ are measured by the LVC when detecting individual GW events. The bounds on chirp mass encompass the extrema of the GW catalogue posteriors and are only used when evaluating the population-level likelihood, as described in Sec. 5.3.1.

applications and Ref. [431] for the code implementation). Our model depends on six population parameters, $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$. These are reported in Table 5.1 and described below. In particular, the quantities γ , m_{\max} , and χ_{\max} parametrize the population of first-generation (1g) BHs, while the quantities α , β , and δ parametrize the pairing and merger process.

This setup is an excellent test bed for our deep-learning explorations because these simulations are not computationally intensive (thus allowing us to explore different DNN architectures) while at the same time providing a binary BH population that

ultimately is not parametric (thus making our emulation approach essential).

5.2.1 Simulation design

We generate $N_\lambda = 1000$ sets of population parameters λ using Latin hypercube sampling to efficiently cover the higher-dimensional space [542] (see, e.g., Fig. 2 of Ref. [404] for a visual representation.) With this design, the hyperparameter space (that is, the space of population-level parameters) is split into N_λ equally probable subintervals in each dimension. From the N_λ^6 possible choices, a total of N_λ unique coordinates are drawn such that, for each of the six dimensions, only one of the N_λ subintervals is selected. In general, there are multiple possible realizations of this random draw; we choose to maximize the minimum distance between samples, whose values are chosen as the centres of the intervals. Our simulation design is generated with PYDOE¹.

5.2.2 First-generation black holes

Each cluster is seeded with $N_{1g} = 5000$ BHs (this number is chosen to ensure convergence of the resulting merger distributions; see Ref. [540]). Their masses m_{1g} are drawn from a truncated power-law distribution

$$p(m_{1g}|\gamma, m_{\max}) \propto \begin{cases} m_{1g}^\gamma & \text{if } 5M_\odot < m_{1g} < m_{\max}, \\ 0 & \text{otherwise,} \end{cases} \quad (5.1)$$

with slope $\gamma \in [-10, 10]$, maximum cutoff $m_{\max} \in [30, 100]M_\odot$, and a fixed lower boundary of $5M_\odot$ (thus only describing black holes and not neutron stars). Pair-instability [543] and pulsation pair-instability supernovae (PISN) [544] prevent the

¹pythonhosted.org/pyDOE.

formation of stellar-mass BHs between about 50 and $120M_{\odot}$ [322, 325, 327]. This prediction is supported by current GW observations, which point to a decrease of the merger rate at those masses [85]. The precise details of the pair-instability mass gap are uncertain and depend on poorly constrained stellar-physics processes such as the nuclear reaction rates [325, 326, 327], rotation [327, 545], accretion [546, 547, 548, 549, 550], winds [551, 552], envelope retention [553, 554, 555] and dredge-up episodes [556]. We thus allow for a broad range of values of m_{\max} and aim to infer it from the GW data.

The BH spin directions are drawn from an isotropic distribution, as expected in dynamical environments. The dimensionless spin magnitudes are uniform in $[0, \chi_{\max}]$, where the maximum natal spin is $\chi_{\max} \in (0, 1)$. The largest spin formed from stellar collapse is uncertain and difficult to model; see Refs. [329, 470]. The spin model we use for first-generation BHs is therefore not necessarily physically well-motivated but is used for illustrative purposes.

5.2.3 Repeated mergers

At each hyperparameter coordinate, we simulate $N_{\text{cl}} = 500$ clusters with escape speeds v_{esc} drawn according to

$$p(v_{\text{esc}}|\delta) \propto \begin{cases} v_{\text{esc}}^{\delta} & \text{if } 0 \text{ km s}^{-1} < v_{\text{esc}} < 500 \text{ km s}^{-1}, \\ 0 & \text{otherwise,} \end{cases} \quad (5.2)$$

where $\delta \in [-10, 10]$. Large positive (negative) values of δ give escape-speed distributions skewed towards the maximum (minimum) value of v_{esc} . For context, the escape speed of a typical globular cluster is $10\text{--}100 \text{ km s}^{-1}$, while those of nuclear star clus-

ters are up to an order of magnitude larger [557, 292, 558]; we take an upper limit of 500 km s^{-1} to accommodate these larger escape speeds. Cases with large, negative δ essentially describe isolated stellar evolution, where repeated mergers do not take place (though we always assume isotropically distributed spins, not partial alignment as expected in isolated binary evolution [331, 465, 330]). On the other hand, $\delta = 0$ corresponds to a flat v_{esc} distribution, favouring all environments equally.

The key ingredient in our populations is the presence of so-called “higher-generation” BHs that have undergone multiple mergers due to remnant retention in the host cluster. We form circular binary systems by selectively pairing cluster members according to

$$p(m_1|\alpha) \propto m_1^\alpha, \quad p(m_2|\beta, m_1) \propto \begin{cases} m_2^\beta & \text{if } m_2 \leq m_1, \\ 0 & \text{otherwise,} \end{cases} \quad (5.3)$$

where $m_1 \geq m_2$ are the component BH masses. As for the other power-law indices, we again take $\alpha, \beta \in [-10, 10]$; this broad range is taken in each case so that the simulated populations encompass the prior bounds used later in our statistical inference of Sec. 5.3.1. One by one, BH pairs are drawn from the collection according to Eq. (5.3) and the properties of their merger remnants are estimated (assuming a uniform sampling of the orbital phase) with the implementation of Refs. [435, 430], which collects various numerical relativity fitting formulae [537, 538, 559, 560, 561, 562, 563, 564]. Upon merging, the remnant BHs receive a gravitational recoil [565, 566]. If the magnitude v_{kick} of this kick velocity exceeds the escape speed of the host cluster, i.e., $v_{\text{kick}} > v_{\text{esc}}$, the remnant BH is removed and does not merge again. Otherwise, it remains inside the cluster and can undergo subsequent mergers. The estimated rem-

nant mass and spin magnitude are retained, while the spin directions are resampled isotropically. This pairing, merger and ejection procedure is iterated until a single BH remains.

For each merger we record the source parameters $\theta = \{M_c, q, \chi_{\text{eff}}, \chi_p\}$. In particular: $M_c = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ is the chirp mass — the best-measured mass parameter for compact binaries in GWs, as it governs the leading-order evolution of the orbital frequency [see Eq. (1.14)]; $q = m_2/m_1 \leq 1$ is the mass ratio; $\chi_{\text{eff}} \in [-1, 1]$ is the effective aligned spin [346]; and $\chi_p \in [0, 2]$ is a suitable parameter encoding the dominant effect of orbital-plane precession — we use the augmented definition of Ref. [426] which consistently averages over the precessional motion including effects from both component spins. While this definition of χ_p is still a frequency-dependent quantity over the inspiral timescale, our recent work has shown that the influence of the GW reference frequency at the population level is currently subdominant compared to measurement errors [421] — see Chapter 3. In the simulated populations we measure χ_p at the reference GW frequency of 20 Hz.

Additionally, we record whether each merger is that of two first-generation BHs (1g+1g) that produces a second-generation (2g) remnant, a first- and second-generation BH (1g+2g), or two second-generation BHs (2g+2g), or whether it contains a component BH of higher generation (>2g). From these, we compute the fraction of mergers in each generation: f_{1g+1g} , f_{1g+2g} , f_{2g+2g} , and $f_{>2g} = 1 - f_{1g+1g} - f_{1g+2g} - f_{2g+2g}$.

5.2.4 Cosmic placement

The distribution of sources is assumed to be isotropic over the sky, inclination and polarization angle. We do not infer the redshift distribution of BH binaries but consider it fixed, i.e., independent of the hyperparameters λ . Each merger is placed at a redshift

z according to a distribution that is uniform in comoving volume V_c and source-frame time, i.e.,

$$p(z) \propto \frac{1}{1+z} \frac{dV_c}{dz}. \quad (5.4)$$

An immediate generalization of this work would include taking into account the longer assembly times of higher-generation binaries (e.g., Ref. [334]) via their redshift distribution. This can be implemented with an additional hyperparameter and will be tackled in future work.

Ostensibly, $z \in (0, \infty)$, but in practice, there is a detector-dependent horizon, z_{\max} , beyond which binary BH mergers are not observable. To find a conservative z_{\max} , we consider a series of binaries with aligned maximal spins, equal masses, and optimal orientation with respect to a single detector (overhead and face on). These are the loudest sources for a given total mass and redshift. We compute SNRs as described in Sec. 5.3.3 and find that the entire mass range becomes subthreshold above an upper bound $z_{\max} = 2.3$, which we thus take as the maximum of the redshift distribution (in agreement with Appendix E of Ref. [84]).

5.2.5 Resulting populations

The above prescription allows us to transform a simple parametrized description of first-generation BH populations into a complex numerical distribution containing hierarchical mergers. The combined set of hyperparameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$ are very interdependent, and changes in their values cause large variations in the distributions of source parameters $\theta = \{M_c, q, \chi_{\text{eff}}, \chi_p\}$. The total set of simulated events is $\{\{\theta_j^i\}_{j=1}^{N_h(\lambda^i)}\}_{i=1}^{N_\lambda}$, where $N_h(\lambda^i)$ is the number of mergers occurring in the simulation

with hyperparameter coordinate λ^i . The total number of mergers occurring at a given hyperparameter coordinate depends on the distribution of escape speeds, determined by δ . For the numerical setup adopted here, it ranges from $\min_{\lambda} N_{\text{h}}(\lambda) = N_{\text{cl}}N_{1\text{g}}/2 = 1.25 \times 10^6$ (when each remnant BH is ejected so only first-generation mergers occur) to $\max_{\lambda} N_{\text{h}}(\lambda) = N_{\text{cl}}N_{1\text{g}} - 1 \approx 2.5 \times 10^6$ when BHs are repeatedly paired with the same single retained remnant, i.e., a “cluster catastrophe” [541, 567]) and the upper range is populated by simulations with larger numbers of repeated mergers. This is demonstrated in Fig. 5.2, where we plot the branching fractions of different merger generations as a function of the total number of mergers.

Four representative cases among the set of $N_{\lambda} = 1000$ simulations we performed are illustrated in Fig. 5.3 and labelled based on the qualitative properties of the resulting source distributions: broad masses, narrow mass ratio, broad mass ratio, and repeated mergers.

If clusters are preferentially formed with larger escape speeds, many remnants are retained and proceed to take part in hierarchical mergers, leading to multiple modes in the mass distributions. This is the case for the red curves (repeated mergers) in Fig. 5.3, where $\delta = 5.1$. Since the sharp initial mass function (IMF) ($\gamma = 5.5$) forms first-generation BHs with masses that are all very close to the maximum $m_{\text{max}} = 70M_{\odot}$, hierarchical mergers appear as distinct peaks in the mass distributions. The first generation of mergers has $m_1 \approx m_2 \approx m_{\text{max}}$, giving $M_{\text{c}} \approx 50M_{\odot}$. Cross-generational mergers also occur. For example, there is a 1g+2g peak; the peak does not occur at $q = 0.5$ because a fraction $1 - \epsilon \approx 5\%$ of mass is lost via GWs [339] such that second-generation BHs have mass of approximately $2\epsilon m_{\text{max}}$, implying $q = 1/(2\epsilon) \approx 0.53$ and $M_{\text{c}} \approx 80M_{\odot}$. Similarly, for a 1g+3g merger, one has $q \approx 1/[\epsilon(2\epsilon + 1)] \approx 0.36$, which explains the third peak observed in the red curves of Fig. 5.3.




Figure 5.2: Fraction of mergers in our simulations from each binary generation as a function of the total number of mergers. The simulations are separated into bins equally spaced in the total number of mergers and the bin-averaged branching fraction of each binary generation — 1g+1g (blue), 1g+2g (orange), 2g+2g (green), and higher generations (red) — is plotted. At the lower (upper) end, simulations are dominated by mergers between first- (higher-) generation BHs.

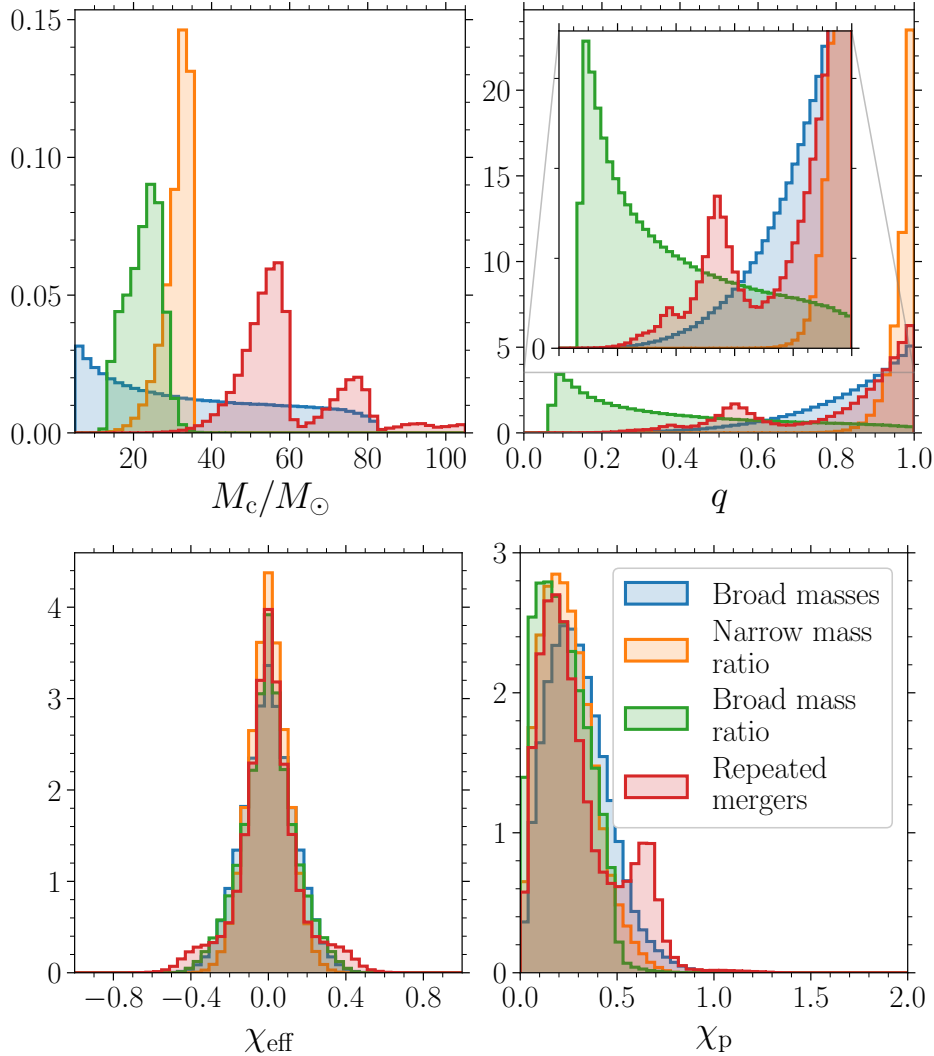


Figure 5.3: Example marginal distributions of chirp mass M_c , mass ratio q , effective aligned spin χ_{eff} , and our precession parameter χ_p for different population parameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\text{max}}, \chi_{\text{max}}\}$. We select four of our simulations to illustrate different features of the resulting binary BH distributions. In blue, we show a simulation with a population of broad masses; $\lambda = \{-1.7, 1.7, -0.5, -3.4, 96M_\odot, 0.57\}$ In orange, we show a simulation that results in a narrow range of mass ratios; $\lambda = \{-8.8, 8.3, 6.8, -4.1, 40M_\odot, 0.43\}$ Conversely, we show in green a simulation that results in a broad spread to unequal masses; $\lambda = \{9.2, -9.8, -0.5, -4.0, 74M_\odot, 0.50\}$. In red is a simulation resulting in a population of repeated mergers; $\lambda = \{4.1, 3.1, 5.5, 5.1, 70M_\odot, 0.37\}$.

When more first-generation BHs are born with large spins, set by χ_{\max} , fewer second-generation mergers occur due to the larger imparted recoils [539]. On the other hand, if natal spins are small and repeated mergers do occur, the distribution of effective spins features a sharp peak at $\chi_{\text{eff}} = 0$ from first-generation mergers as well as extended tails from high-generation mergers, as is the case for the red curve in the third column of Fig. 5.3. The χ_{eff} distributions are always symmetric about zero due to the assumption of spin isotropy. For the 1g+2g populations, the 2g BH spin is approximately 0.7 [332] and, because in this case $\chi_{\max} = 0.37$, is typically higher than the spin of the 1g BH. In this limit, one has $\chi_p \approx \sqrt{0.7^2 - 4\chi_{\text{eff}}^2} \approx 0.7$ [429], explaining the secondary peak in the χ_p distribution.

Whether higher-generation BHs pair with other BHs of equal generation or form cross-generational binaries (e.g., 1g+2g) depends on the pairing slopes α and β . If $\alpha, \beta, \gamma \approx 0$, then the first-generation mass distribution is broad and binary components are selected with uniform probabilities leading to an extended range of mass ratios, as seen in the blue “broad masses” curves of Fig. 5.3. If $\alpha, \beta \gg 0$ ($\alpha, \beta \ll 0$), the heaviest (lightest) BHs are preferentially selected for both binary components, leading to a heavier (lighter) first generation of approximately equal-mass binaries. If $\alpha \ll 0$ and $\beta \gg 0$ ($\alpha \gg 0$ and $\beta \ll 0$), then the lightest (heaviest) primaries and heaviest (lightest) secondaries are paired, leading to mass-ratio distributions that are sharply peaked at unity (broad and peaked at lower values), as seen in the orange “narrow mass ratio” (green “broad mass ratio”) curve of Fig. 5.3. In the case of narrow mass ratios, given the maximum first-generation mass $m_{\max} \approx 40M_{\odot}$ and since $q \approx 1$, the chirp mass peak is located at $M_c \approx 35M_{\odot}$.

5.3 Enhancing population inference with deep learning

Our simplified population model is indicative of realistic applications where GW events are modelled using, e.g., population-synthesis codes. As we will show, deep learning is the ideal tool for such a scenario. First, we review the key ingredients that enter our hierarchical Bayesian inference of GW data (Sec. 5.3.1). We then present our method to model the population prior (Sec. 5.3.2) and selection effects (Sec. 5.3.3) using deep learning. We use similar techniques to model the branching fractions between different merger generations (Sec. 5.3.4).

5.3.1 Hierarchical Bayesian inference

Preliminaries

We now denote with $\vartheta \supset \theta$ the full superset of the modelled source parameters $\theta = \{M_c, q, \chi_{\text{eff}}, \chi_p\}$ that additionally contains the unmodelled parameters $\bar{\theta}$, e.g., redshift, sky location, inclination, etc. (cf. the notation θ , φ , and ψ , respectively, from Sec. 1.4.3).

We approximate the sensitivity as constant within each observation period: the combined first and second run (O1+O2) and the third run (O3). The corresponding two-detector observing periods that enter the selection function of Eq. (1.90) are $T_{\text{O1+O2}} \approx 166$ days [60, 61] and $T_{\text{O3}} \approx 275$ days [62, 64], respectively.

The priors on the parameters m_{max} and χ_{max} are uniform over the ranges listed in Table 5.1. The priors of the power-law indices α , β , γ , and δ are uniform over $[-8, 8]$; these prior bounds lie within the training data range and we checked that resulting

posteriors are robust to more stringent constraints.

Factorization of the observed volume

While the integrals in Eq. (1.85) are formally defined over the entire domain of source parameters, in practice they can be safely performed within the observable volume $V_h = \{\vartheta : z < z_{\max} = 2.3\}$, beyond which the detection probability is zero, as discussed in Sec. 5.2.4. Even if $p(\vartheta|\lambda)$ models the binary BH population outside of V_h , as it appears in both the numerator and through $\sigma(\lambda)$ in the denominator of Eq. (1.85), one can safely assume that $\int_{V_h} p(\vartheta|\lambda)d\vartheta = 1$.

Since it will be useful in Sec. 5.3.2, we define the a posteriori observed volume $V_p = \{\vartheta : p(\vartheta|d_i) > 0 \forall i\} \subset V_h$ as the subset of the observable volume beyond which all single-event posteriors $p(\vartheta_i|d_i)$ vanish. For the events considered in this work we find that V_p corresponds to $M_c \in [5, 105]M_\odot$. For q , χ_{eff} , and χ_p , we maintain their natural bounded domains ($[0, 1]$, $[-1, 1]$, and $[0, 2]$, respectively). It will also be useful to define the population prior of our modelled parameters θ normalized over the observed volume,

$$p'(\theta|\lambda) = \frac{p(\theta|\lambda)}{\int_{V_p} p(\theta|\lambda)d\theta}. \quad (5.5)$$

Since $p(\theta|\lambda)$ is assumed to be normalized over V_h , we can write this extra normalization factor as

$$\int_{V_p} p(\theta|\lambda)d\theta = \frac{N_p(\lambda)}{N_h(\lambda)} \leq 1, \quad (5.6)$$

where $N_h(\lambda)$ is the number of mergers occurring within the horizon volume V_h and $N_p(\lambda)$ is the number of mergers occurring within the subset $V_p \subset V_h$, given population

parameters λ . For convenience, we refactor this term into the detection efficiency by defining the selection function

$$\sigma'(\lambda) = \frac{N_h(\lambda)}{N_p(\lambda)}\sigma(\lambda). \quad (5.7)$$

By separating the source parameters and noting that our population prior and the PE prior over the unmodelled parameters are equal, i.e., $p(\bar{\theta}|\lambda) = p(\bar{\theta}|\mathcal{H})$, the hyperposterior in Eq. (1.85) may be written as

$$p(\lambda|\{d\}) \propto p(\lambda) \prod_{i=1}^n \frac{1}{\sigma'(\lambda)} \int_{V_p} p(\theta_i, \bar{\theta}_i|d_i, \mathcal{H}) \frac{p'(\theta_i|\lambda)}{p(\theta_i|\bar{\theta}_i, \mathcal{H})} d\theta_i d\bar{\theta}_i. \quad (5.8)$$

Since the PE prior is placed on detector-frame masses, we must convert the prior on detector-frame chirp mass M_c^{det} to the source frame. In particular, we have

$$p(\theta|z, \mathcal{H}) = p(M_c^{\text{det}}, q, \chi_{\text{eff}}, \chi_p|z, \mathcal{H}) \left| \frac{\partial M_c^{\text{det}}}{\partial M_c} \right|. \quad (5.9)$$

Since the Jacobian is $|\partial M_c^{\text{det}}/\partial M_c| = 1 + z$ and the prior on detector-frame masses is independent of the prior on redshift for the PE results we use below, we have $p(\theta|z, \mathcal{H}) = p(M_c^{\text{det}}, q, \chi_{\text{eff}}, \chi_p|\mathcal{H})(1 + z)$.

Event samples

We select the confident binary BH detections made during the first (O1), second (O2) and third (O3) observing runs, employing a threshold minimum FAR $< 1 \text{ yr}^{-1}$ across all search analyses. This results in a catalogue of $n = 69$ binary BH events. For the events in O1 and O2, we use samples² from the reanalysis of Ref. [450] because the

²dcc.ligo.org/LIGO-P2000193/public.

precession parameter χ_p depends on the azimuthal spin angles whose posteriors were not released in GWTC-1 [61]. For the events in O3, we take the posterior samples combining analyses with waveforms including both precession and higher-order modes as provided by the GWTC-2³ [62], GWTC-2.1⁴ [63], and GWTC-3⁵ [64] data releases.

For each event we draw prior samples for $\{M_c^{\text{det}}, q, \chi_{\text{eff}}, \chi_p\}$ and compute the prior density $p(M_c^{\text{det}}, q, \chi_{\text{eff}}, \chi_p | \mathcal{H})$ using Gaussian kernel density estimates (KDEs), as implemented in SCIPY [568] but modified to enforce reflective boundary conditions [569, 570]. Each KDE is then evaluated on the single-event posterior samples. The hierarchical posterior is sampled using DYNESTY [387] and BILBY [388].

5.3.2 Population model

The results of our simulations are lists of binary BH mergers, characterized by source parameters $\theta = \{M_c, q, \chi_{\text{eff}}, \chi_p\}$, at each of the $N_\lambda = 1000$ population parameter coordinates, $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\text{max}}, \chi_{\text{max}}\}$. Our approach to modelling the resulting population distribution $p'(\theta|\lambda)$ employs a combination of probability-density estimation and regression algorithms.

Density estimation

At each of the hyperparameter locations $\{\lambda^i\}_{i=1}^{N_\lambda}$, we evaluate the conditional population density with a Gaussian KDE. To efficiently evaluate $p'(\theta|\lambda^i)$ with sufficient resolution in the four-dimensional space of source parameters, we use a version of the convolution-based implementation in KDEPY [571], which we modify to enforce the parameter limits (Table 5.1) with reflective boundary conditions [569]. With this method,

³gw-openscience.org/GWTC-2.

⁴[PrecessingSpinIMRHM from gw-openscience.org/GWTC-2.1](https://gw-openscience.org/GWTC-2.1).

⁵[C01:Mixed from gw-openscience.org/GWTC-3](https://gw-openscience.org/GWTC-3).

density estimations of multivariate data with millions of samples evaluated on millions of points take seconds on a standard, off-the-shelf machine, compared to hours with standard KDE routines (the evaluation points must, however, lie on a linearly spaced Cartesian grid that bounds the data extrema). Each dimension is individually scaled with bandwidths determined by the Improved Sheather–Jones (ISJ) plug-in selection rule [572, 573]. The ISJ algorithm does not make the assumption of normality on the underlying distribution and, as such, is more robust when determining optimal bandwidths for non-Gaussian multimodal distributions. We evaluate each of the N_λ KDEs on a linearly spaced Cartesian grid, including the parameter bounds, with 21 points in each axis.

Regression with a deep neural network

Elucidating the scale of the regression problem, there are $21^4 \approx 2 \times 10^5$ KDE evaluations estimating $p'(\theta|\lambda)$ over the combined ten-dimensional vectors of source and population parameters (θ, λ) at each of the $N_\lambda = 1000$ hyperparameter locations. While the KDEs approximate the N_λ functions $\{\theta \mapsto p'(\theta|\lambda^i)\}_{i=1}^{N_\lambda}$, we must also interpolate over the population parameters to find an accurate mapping $(\theta, \lambda) \mapsto p'(\theta|\lambda)$.

To achieve this result, we make use of a fully connected DNN implemented with `TENSORFLOW` [574, 575]. The network performs a regression of the KDE values of $p'(\theta|\lambda)$ over the space of (θ, λ) coordinates. As a preprocessing step, we normalize all coordinates (θ, λ) to a unit hypercube using the limits given in Table 5.1, while the values of $p'(\theta|\lambda)$ are similarly scaled between zero and their maximum over the grid and simulations. The input layer has $\dim \theta + \dim \lambda = 10$ neurons, while the output layer has one neuron with enforced non-negativity corresponding to the predicted value of the probability density. Between the input and output layers, the network architecture

Layer	Neurons	Activation	Parameters
Input	10	...	0
Dense 1	128	RReLU	1,408
Dense 2	128	RReLU	16,512
Dense 3	128	RReLU	16,512
Dense 4	128	RReLU	16,512
Dense 5	128	RReLU	16,512
Output	1	Absolute value	129
Total			67,585

Table 5.2: Architecture of the DNN that emulates the simulated populations by predicting the conditional density $p'(\theta|\lambda)$ of the source parameters θ given population-level parameters λ . Each row represents a single layer of the network and lists the number of neurons in the layer, the activation function of those neurons (RReLU for the hidden layers and absolute value for the final output), and the corresponding number of free parameters.

consists of five hidden layers, each with 128 neurons. We summarize the network architecture in Table 5.2. The number of parameters in a given layer is given by the number of weights (equal to the product of the number of neurons with that of the preceding layer) plus the number of biases (equal to the number of neurons).

We use randomized leaky rectified linear units (RReLU) [576] in each layer. This modifies the standard rectified linear unit (ReLU) activation function, that is given by $\text{ReLU}(x) := \max(0, x)$, in two ways. First, leaky ReLU activation functions are maps $x \mapsto \max(0, x) + \min(0, ax)$, where $a \in [l, u]$ is a parameter fixed to a small number; i.e., the positive region is linear with unit slope while the negative region is linear with slope a . Second, the randomized leaky variant RReLU samples a uniformly in $[l, u]$ during training and fixes $a = (l + u)/2$ when making predictions (we keep the default values of $l = 1/8$ and $u = 1/3$ [576]). Empirically, we find that, among other ReLU variants and nonlinear activations, RReLU gives the best predictive performance while

reducing overfitting to the training data.

We split the $N_\lambda = 1000$ simulations into a training data set of 900 runs and a validation set of 100 runs. The validation sample is unseen by the training process except to assess the network performance. The training data input to the network, which is randomly shuffled at each iteration, thus consists of approximately 1.75×10^8 values of the ten-dimensional vector (θ, λ) and the corresponding KDE estimates of $p'(\theta|\lambda)$. The network is trained using the Adam optimizer [577], the mean absolute error (MAE) loss function, a learning rate of 10^{-4} , and batch size equal to 0.01% of the total number of training data points. Training is performed for 10^4 epochs on an NVIDIA A100 Tensor Core GPU, taking about four days. With this setup, the number of training epochs is sufficient to ensure convergence of the MAE; the average gradient of the (smoothed) validation MAE over the penultimate 100 epochs is less than 0.1% that of the first 100.

When making predictions with the trained DNN, the values are first rescaled from the unit interval to the probability-density parameter space. While the predictions are approximately normalized, the network does not enforce unit normalization. Therefore, we estimate normalization factors $\int p'(\theta|\lambda)d\theta$ by numerically integrating the predicted distributions.

In Fig. 5.4, we summarize the training procedure and predictive performance of our DNN population model. The convergence of the MAE loss function for the training and validation samples is plotted in the top panel. The DNN fits slightly better to the training data — the validation MAE being, on average, about 1.2 times larger — but there is no significant overfitting.

In the bottom panel of Fig. 5.4, we quantify this statement by comparing the predictive accuracy of the trained population model using the Hellinger distance [578],

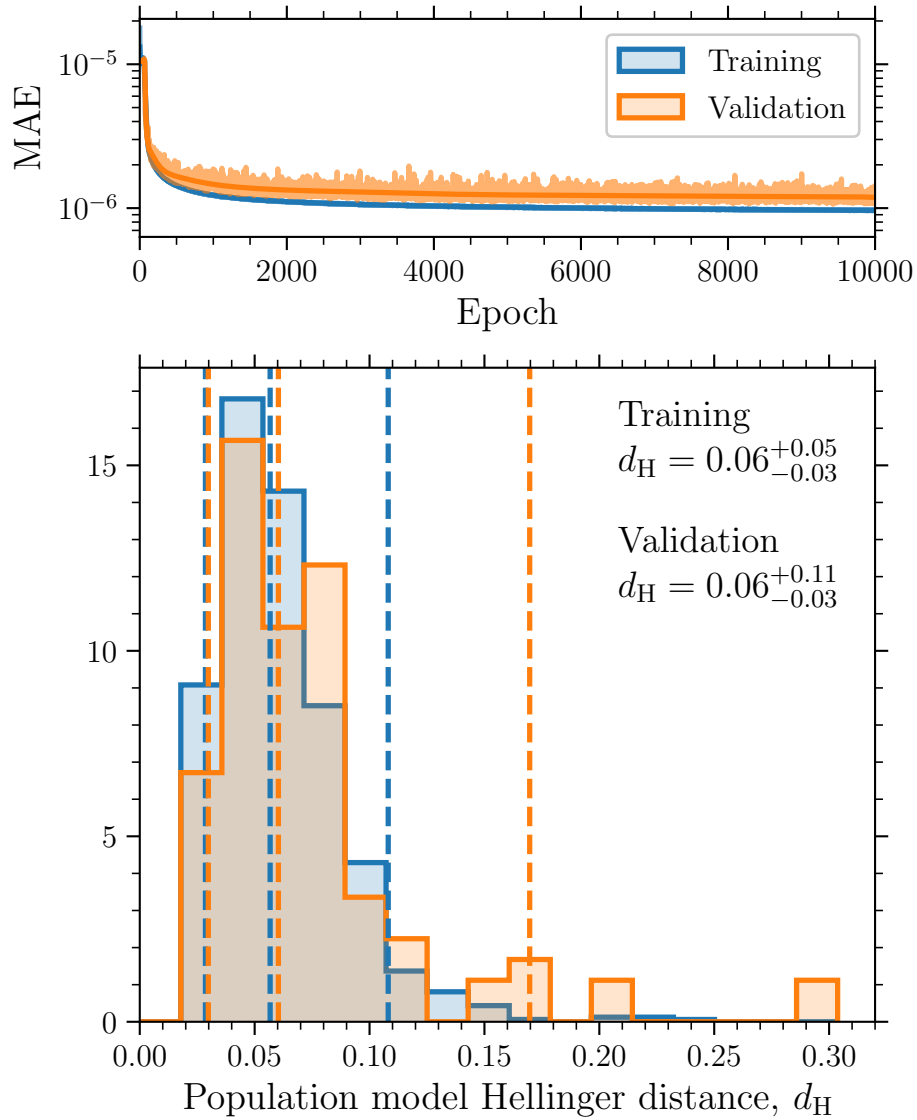


Figure 5.4: Top panel: loss functions versus epoch for the training (blue) and validation (orange) data of the population density DNN $p'(\theta|\lambda)$. Smoothed versions are overplotted in bold. Bottom panel: distribution across all simulations of the Hellinger distances d_H between the true KDE evaluations of $p'(\theta|\lambda)$ and those predicted by the DNN. The medians and 90% intervals of d_H are plotted as vertical dashed lines and listed explicitly.

a metric d_{H} over the space of probability densities that measures the “distance” between two distributions. For two probability densities p and q , it is given by

$$d_{\text{H}}(p, q)^2 = \frac{1}{2} \int \left[\sqrt{p(x)} - \sqrt{q(x)} \right]^2 dx. \quad (5.10)$$

Here, d_{H} has the desirable properties of being symmetric and bounded in $[0, 1]$, with $d_{\text{H}}(p, q) = 0$ only when $p \equiv q$ and $d_{\text{H}}(p, q) = 1$ when p and q have disjoint supports (see, e.g., Appendix C of Ref. [579] for a physics-oriented summary of the Hellinger distance). For each of our simulations we compute the distance between the KDE evaluation and the DNN prediction for the probability density. While the mild overfitting presents itself as a small number of outliers at larger values of d_{H} in the validation distribution, both the training and validation subsets have median values of approximately 0.06 and are consistent with each other.

In Fig. 5.5 we illustrate example predictions from our deep-learned population model. For a given set of population-level parameters λ , the DNN predicts the value of the joint four-dimensional probability density over $\theta = \{M_{\text{c}}, q, \chi_{\text{eff}}, \chi_{\text{p}}\}$. For three validation simulations, we plot the predicted values of $p'(\theta|\lambda)$ (solid lines) along with the true KDE evaluations for comparison (circle markers), numerically marginalizing to one-dimensional distributions for the purpose of visualization.

The first example (red) has good predictive accuracy, with $d_{\text{H}} = 0.10$. Here, we use the same distribution labelled “repeated mergers” in Fig. 5.3, with parameters $\alpha = 4.1$, $\beta = 3.1$, $\gamma = 5.5$, $\delta = 5.1$, $m_{\text{max}} = 70M_{\odot}$, and $\chi_{\text{max}} = 0.37$. Here, the larger escape velocities and sharp mass function and pairing probabilities lead to distinct peaks due to higher-generational mergers. Even though the Hellinger distance of this simulation is greater than the median value, the one-dimensional marginal predictions present

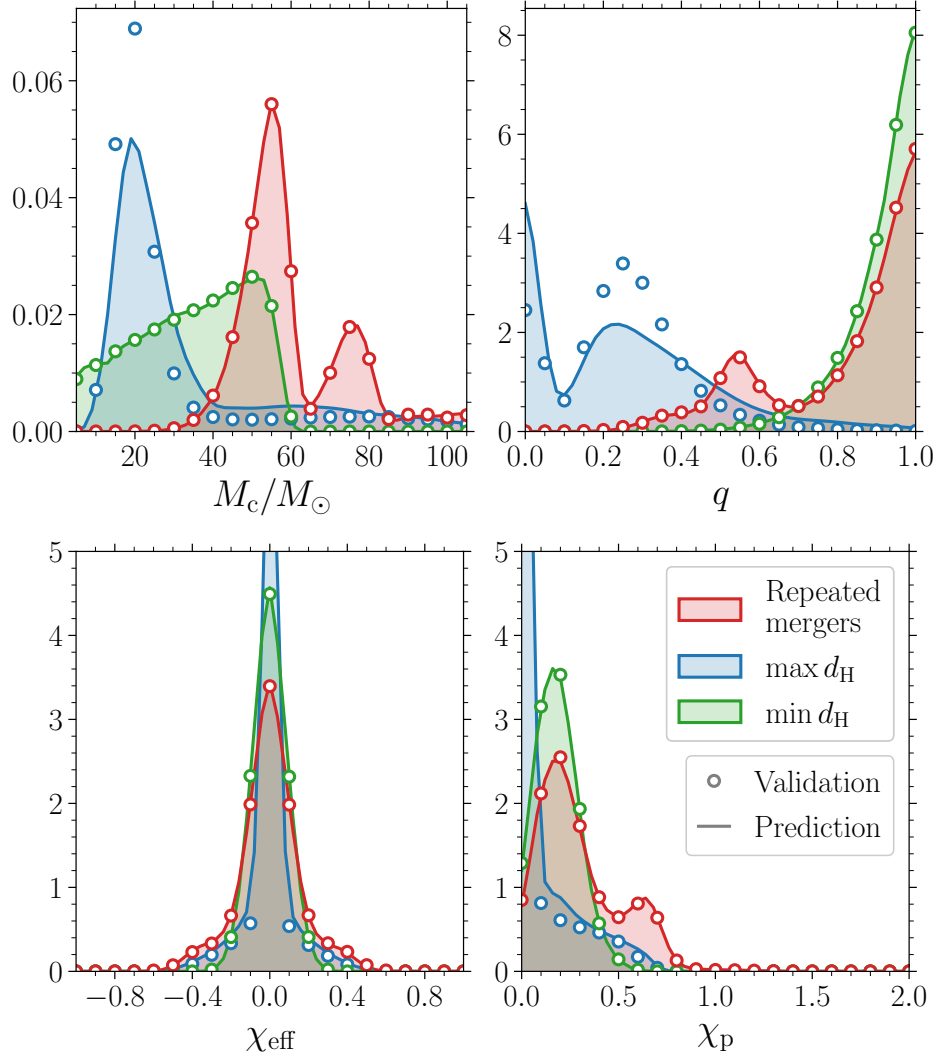


Figure 5.5: True KDE evaluations (circle markers) of the population density $p'(\theta|\lambda)$ compared against the DNN population model predictions (solid lines) for three validation simulations. The full four-dimensional distributions are marginalized to each one-dimensional event-level parameter (left to right: chirp mass M_c , mass ratio q , effective aligned spin χ_{eff} , and effective precessing spin χ_p) for the purpose of visualization. In blue, we show the validation simulation that has the worst predictive accuracy, with a Hellinger distance of $d_H = 0.30$ and population-level parameters $\alpha = 6.3$, $\beta = -7.3$, $\gamma = 1.8$, $\delta = 8.7$, $m_{\text{max}} = 46M_\odot$, and $\chi_{\text{max}} = 0.01$. In green, we show the validation simulation with the smallest Hellinger distance $d_H = 0.02$ and $\alpha = -1.9$, $\beta = 5.2$, $\gamma = 0.5$, $\delta = -9.4$, $m_{\text{max}} = 67M_\odot$, $\chi_{\text{max}} = 0.35$. In red, we show a validation simulation (as in Fig. 5.3) with $d_H = 0.10$ and whose distribution contains distinct features due to repeated mergers.

excellent matches to the true validation data, accurately capturing all sharp features.

The second case ($\max d_{\text{H}}$, in blue) is a very conservative bound on the performance of our DNN, taking the validation simulation with the largest value of the Hellinger distance $d_{\text{H}} = 0.30$ (i.e., that with the worst predictive accuracy). The population parameters are $\alpha = 6.3$, $\beta = -7.3$, $\gamma = 1.8$, $\delta = 8.7$, $m_{\text{max}} = 46M_{\odot}$, and $\chi_{\text{max}} = 0.01$. While the distributions of the spin parameters χ_{eff} and χ_{p} are still fairly well captured, the predictions in the mass distributions suffer from larger errors, though the main features have been learned. The small value of the maximum natal spin $\chi_{\text{max}} = 0.01$ leads to sharply peak effective spins $\chi_{\text{eff}}, \chi_{\text{p}} \approx 0$, while the pairing process generates smaller mass ratios. We stress that this is the worst case among the entire validation set and a rather extreme outlier (cf. Fig. 5.4). Figure 5.5 presents the marginalized distributions, while the model predicts the full four-dimensional density, meaning errors over the full source parameter are propagated to the one-dimensional marginals.

The third case ($\min d_{\text{H}}$, in green) represents the best predictive accuracy of our population model, with $d_{\text{H}} = 0.02$. In this validation simulation, the hyperparameters are $\alpha = -1.9$, $\beta = 5.2$, $\gamma = 0.5$, $\delta = -9.4$, $m_{\text{max}} = 67M_{\odot}$, and $\chi_{\text{max}} = 0.35$, which produces equal masses and a unimodal distribution in the joint four-dimensional space of source parameters. Unsurprisingly, distributions with a simple feature set like this are easier to learn by our DNN population model.

5.3.3 Selection function

Detection probability

We assume sources are distributed uniformly in sky location, inclination, and polarization angle. We estimate $p(\text{det}|\vartheta)$ with the widely used single-detector semi-analytic

approximation of Refs. [416, 580] (see Sec. 1.4.3), as implemented in the GWDET Python package [418], which relies on computing the SNR of optimally oriented sources with the same intrinsic parameters. This is estimated using PYCBC [581, 582], the IMPHENOMPV2 waveform approximant [452, 583, 584], and noise curves representative of the LIGO detector performance during O1O2⁶ and O3⁷ [145]. While the analytic marginalization of Refs. [416, 580] is, strictly speaking, only valid if one neglects spin precession and higher-order GW modes [recall that we assumed non-spinning or aligned-spin sources in Eq. (1.89)], the impact of these additional effects is subdominant [410]. Their inclusion requires further modelling which, as we previously saw, has recently been tackled using machine-learning techniques [410, 412]; we plan to include these refinements in future versions of our population-inference pipeline. We employ a SNR threshold of 8 [415] and thus set $p(\text{det}|\vartheta) = 0$ for all subthreshold binaries.

We compute the detection probabilities for all mergers in each simulation using the numerical expectation values and time averaging in Eqs. (1.88, 1.90).

Regression with a deep neural network

To evaluate the (refactored) detection efficiency at arbitrary values of the population parameters, the function $\sigma'(\lambda)$ must be emulated using the discrete evaluations at $\{\lambda^i\}_{i=1}^{N_\lambda}$. Here, we also use a DNN with TENSORFLOW [574, 575]. The network architecture consists of an input layer with $\dim \lambda = 6$ neurons and a linear output layer with one, corresponding to the predicted value of $\ln \sigma'(\lambda)$. We add three hidden layers with 128 neurons each and RReLU activation. This network architecture is summarized in Table 5.3.

We split the hyperparameter coordinates into the same 90% training and 10% vali-

⁶Early high from dcc.ligo.org/LIGO-P1200087-v47/public.

⁷LIGO Livingston from dcc.ligo.org/LIGO-T2000012/public.

Layer	Neurons	Activation	Parameters
Input	6	...	0
Dense 1	128	RReLU	896
Dense 2	128	RReLU	16,512
Dense 3	128	RReLU	16,512
Output	1	...	129
Total			34,049

Table 5.3: Architecture of the DNN that predicts the logarithmic selection function $\ln \sigma'(\lambda)$ as a function of the population-level parameters λ . Each row represents a single layer and lists its number of neurons, the activation function used, and the corresponding number of free parameters. All hidden layers employ RReLU nonlinearities.

dation simulations as in Sec. 5.3.2, though note that the training data here consist only of the hyperparameters λ rather than the joint vector (θ, λ) . As a preprocessing stage, we again normalize the input values of $\{\lambda^i\}_{i=1}^{N_\lambda}$ to a unit hypercube using the ranges in Table 5.1 and train on the output values of $\{\ln \sigma'(\lambda^i)\}_{i=1}^{N_\lambda}$, which are normalized to the unit interval according to the extrema across the simulations. Predictions are rescaled back to the relevant parameter space. We use Adam optimization [577] with a learning rate of 10^{-3} to minimize the mean squared error (MSE) loss function. At each epoch, the training data are shuffled into batches containing 1% of the training data. We train the network for 2000 epochs on a single Intel Core i5-8365U CPU, which took approximately four minutes. The training of this DNN is significantly quicker than that of $p'(\theta|\lambda)$ since it has input dimensionality $\dim \lambda = 6$, corresponding to a much smaller training sample size of 900 and a smaller network architecture.

The performance of our DNN to predict $\ln \sigma'$ is reported in Figs. 5.6 and 5.7. In the top panel of Fig. 5.6, we display the convergence of the loss function over the training epochs; the average gradient of the (smoothed) validation MSE over the final 100 epochs

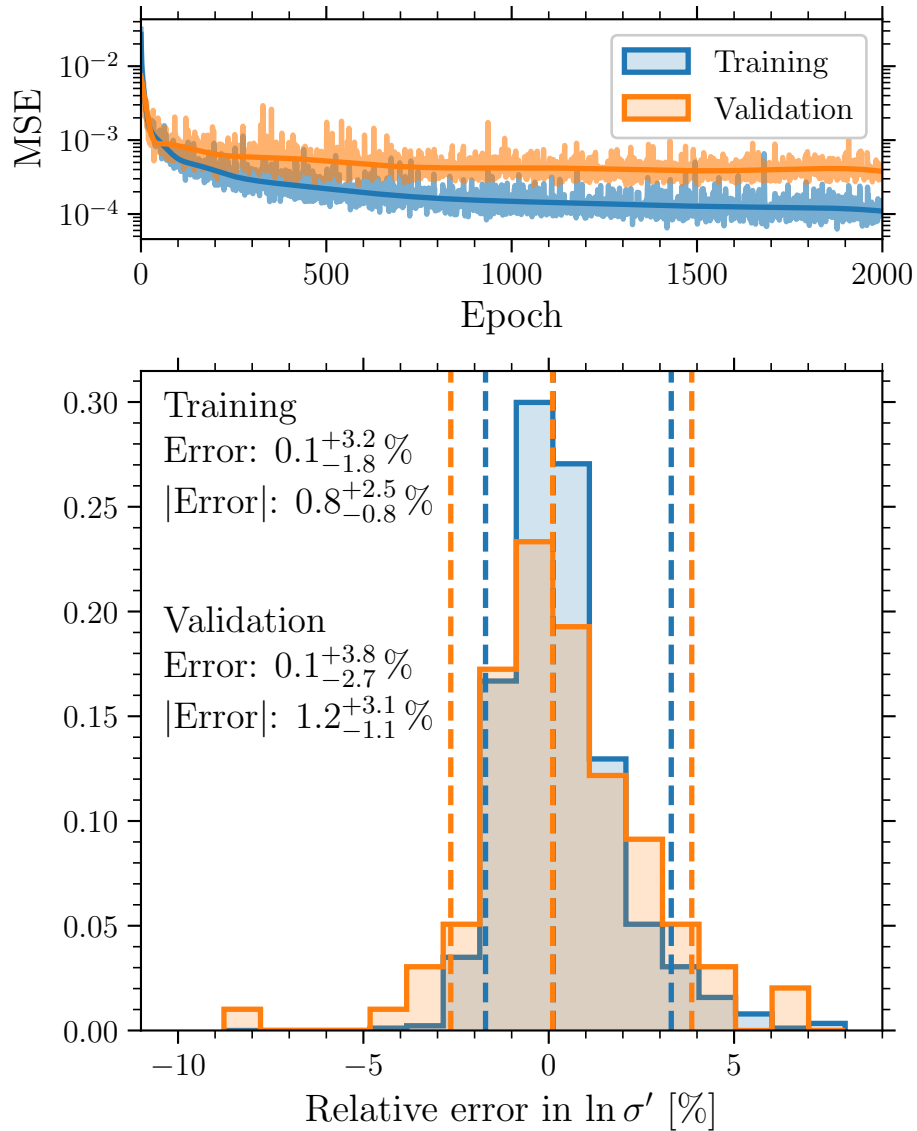


Figure 5.6: Top panel: loss function over epochs for the training (blue) and validation (orange) data of the DNN predicting the refactored detection efficiency $\sigma'(\lambda)$. Bottom panel: relative error between the true and predicted values of $\ln \sigma'$. The medians and 90% intervals of the errors are plotted as vertical dashed lines. They are also listed explicitly, as are the magnitudes of the relative errors.

is $\lesssim 0.5\%$ that over the first 100. While there is some overfitting to the training data, we verify the effect is mild, as follows. In the bottom panel, we display the relative error between the true and DNN-predicted values of $\ln \sigma'$. Since the median and 90% symmetric interval for the validation and training errors are $0.1_{-1.8}^{+3.2}\%$ and $0.1_{-2.7}^{+3.8}\%$, respectively, both are consistent with being centred on and symmetric about zero, i.e., the DNN introduces no systematic biases. The magnitudes of the relative errors for the validation and training sets are consistent with each other and typically less than or close to 5%; the medians and 90% intervals are $0.8_{-0.8}^{+2.5}\%$ and $1.2_{-1.1}^{+3.1}\%$, respectively.

In Fig. 5.7, we show the dependence of the DNN selection function on each of the hyperparameters for the same three example simulations as in Fig. 5.5. The true values of $\ln \sigma'(\lambda)$ are shown with circle markers. The predictions of the DNN are given by the solid lines, where in each panel we vary a single hyperparameter while keeping the others fixed to the values corresponding to each simulation. For all simulations, σ' is an increasing function of both γ — the power-law index of 1g BHs — and m_{\max} — the maximum 1g mass; larger γ implies a greater number of BHs born with masses closer to the maximum m_{\max} , while heavier sources emit louder signals and are thus easier to detect (though there is also a compromise with the frequency-dependent — and therefore, mass-dependent — detector sensitivity). The simulation containing repeated mergers (red) consistently features higher values — implying a larger fraction of detectable mergers in the underlying population — due to the larger average binary mass.

The mismatch for the least accurate hyperparameter coordinate of the population model ($\max d_H$, blue) is visible in the offset between truths and predictions. Here, the selection function also depends on α and β , which determine the primary and secondary pairing probabilities. For this simulation, the first-generation mass slope, $\gamma = 1.8$ is

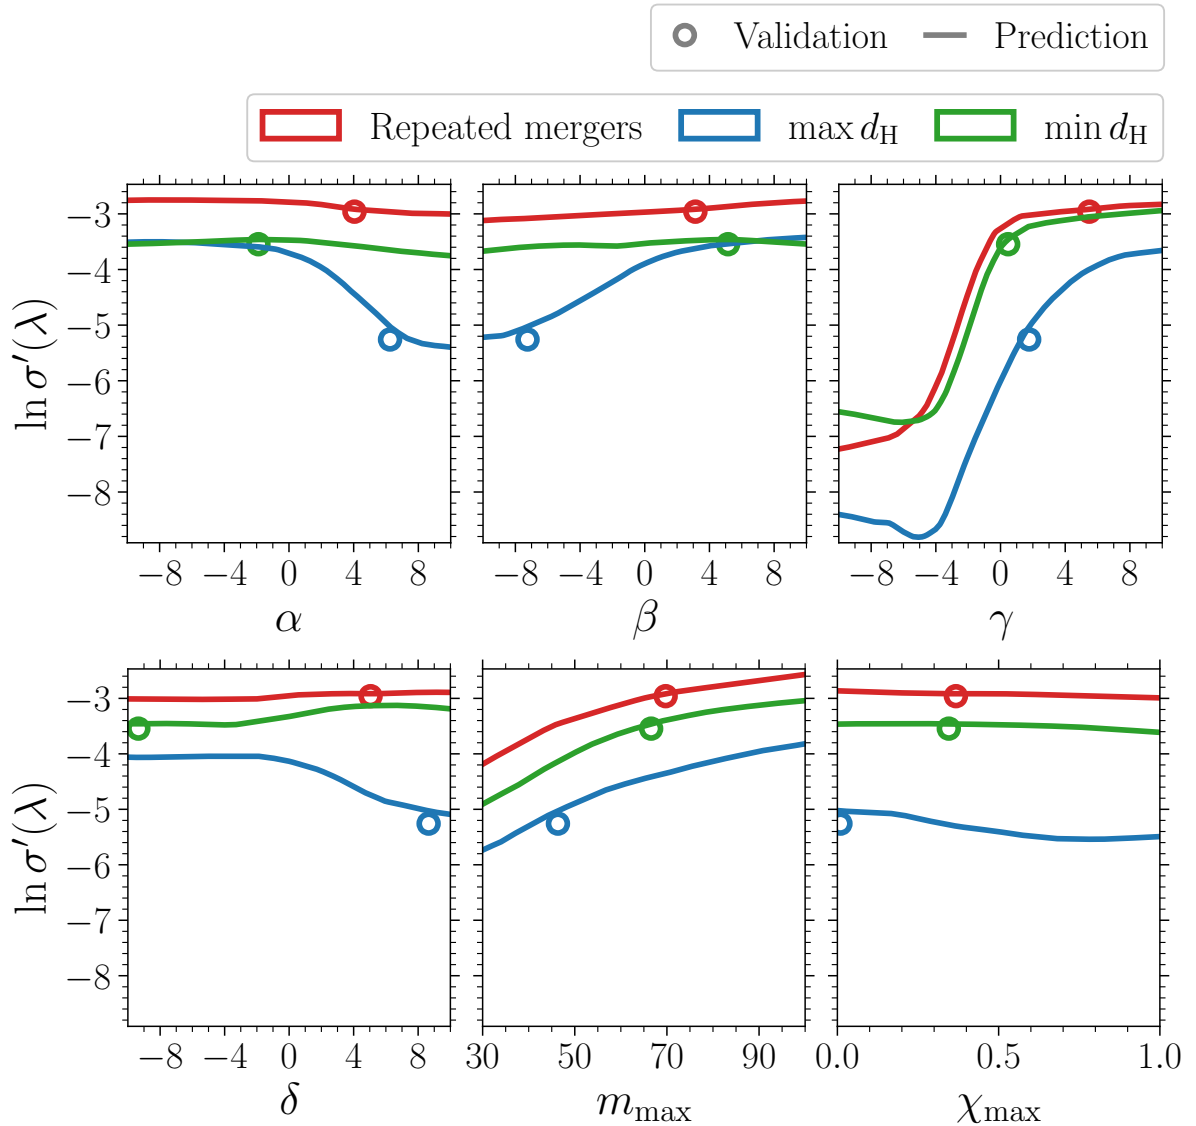


Figure 5.7: Example evaluations of the DNN selection function $\ln \sigma'(\lambda)$ at the same three hyperparameter coordinates λ displayed in Fig. 5.5: a simulation containing repeated mergers (red), and those with the least (blue) and most (green) accurate predictions for the population model DNN. The true value for each simulation is displayed as a circle marker, while predictions made by the DNN are solid lines. In each panel, we vary a single hyperparameter, while the others are fixed to values in the three simulations.

quite broad. A wider range of masses implies that a wider range of mass ratios are possible when selecting the BHs in the binary pairing procedure. Higher (lower) values of α (β) lead to higher (lower) primary (secondary) masses and more extreme mass ratios, thus decreasing the detectability. Though repeated mergers occur due to a high $\delta = 8.7$, they are preferentially of mixed generations, and therefore, larger δ also leads to lower detectability.

For the validation simulation with the highest population model accuracy (min d_H , green), first-generation masses are broad since $\gamma = 0.5$ and larger since $m_{\max} = 67M_{\odot}$ (compared to $m_{\max} = 46M_{\odot}$ for the max d_H case). Masses are paired equally since $\alpha = -1.9$ selects the lightest primaries and $\beta = 5.2$ selects the heaviest secondaries. The greater prevalence of higher-mass sources with equal mass ratios results in a selection function that is higher (corresponding to increased detectability in the binary BH population) and flatter (with respect to all hyperparameters except γ and m_{\max} , as discussed).

5.3.4 Merger-generation fractions

As a final utilization of deep-learning techniques within GW population inference, we train a DNN to infer the branching fractions f_g of the merger generations $g \in \{1g+1g, 1g+2g, 2g+2g, >2g\}$, as defined in Sec. 5.2.3. It is important to note that, unlike the case of branching ratios in mixture population models (e.g., Refs. [334, 501, 585, 586, 587]), these fractions are not hyperparameters themselves but are functions of the hyperparameters λ . In particular, $f_g(\lambda) = \int p(\theta|\lambda)\mathcal{I}_g(\theta, \lambda)d\theta$, where $\mathcal{I}_g(\theta, \lambda)$ is an indicator function that labels the merger generation, such that $\sum_g f_g(\lambda) \equiv 1$. Our application to the fraction of systems in each hierarchical generation is an example of the more generic problem of constraining formation subchannels that enter a single

Layer	Neurons	Activation	Parameters
Input	6	...	0
Dense 1	128	RReLU	896
Dense 2	128	RReLU	16,512
Dense 3	128	RReLU	16,512
Output	4	Softmax	516
Total			34,436

Table 5.4: Structure of the DNN that models the branching fractions f_{1g+1g} , f_{1g+2g} , f_{2g+2g} , and $f_{>2g}$ between the binary merger generations, where 1g (2g) denotes a first- (second-) generation component BH. The rows illustrate each layer of the network and report the number of neurons in each, their activation functions (RReLU for the hidden layers and softmax for the output layer), and the number of free parameters.

population.

We use the same training process and network architecture as in Sec. 5.3.3, with one modification. Since the four branching fractions form a discrete distribution with unit sum, the output layer here has four neurons and employs the activation function $\text{softmax}(\mathbf{x})_i := \exp(x_i) / \sum_j \exp(x_j)$, where x_i are the components of the input vector \mathbf{x} . The architecture of this DNN is summarized in Table 5.4.

In Fig. 5.8, we plot the converged MSE loss curves. We assess the accuracy of the DNN predictions against the true generation fractions on the training and validation data using the Hellinger distance, which, for discrete probability densities p and q , is

$$d_H(p, q)^2 = 1 - \sum_i \sqrt{p_i q_i}. \quad (5.11)$$

The performances on training and validation subsets are consistent with each other, representing a lack of overfitting. Both have median Hellinger distances of $d_H \approx 0.01$ with $d_H \lesssim 0.1$ for most simulations. The enforced unit summation implies the branch-




Figure 5.8: Top panel: loss functions over training epochs for the training (blue) and validation (orange) data of the DNN predicting the branching fractions f_{1g+1g} , f_{1g+2g} , f_{2g+2g} , and $f_{>2g}$. The actual loss curves are plotted with shading and smoothed versions are overplotted in bold. Bottom panel: Hellinger distances between the discrete distributions of the true and DNN-predicted merger-generation branching fractions. The medians and 90% confidence intervals are plotted as vertical dashed lines and listed explicitly.

ing fraction emulator in fact has only three independent outputs despite predicting four contributions, and in many of our simulations, one or more of the generation labels has zero contribution (e.g., no higher-generation mergers when all remnants are ejected from the host cluster). Both of these considerations produce a tendency for small values of the Hellinger distance, which explains the skew towards $d_H \lesssim 0.01$.

In Fig. 5.9, we display the dependence of the DNN to predict the branching fractions $f_g(\lambda)$ on the hyperparameters λ for the same validation simulations reported in Figs. 5.5 and 5.7. As in Fig. 5.7, the true values computed from the simulated data are given by circle markers, while predictions made by the DNN are plotted as solid lines where a single hyperparameter is varied while keeping the others fixed. We only display the variation with the power-law indices $\{\alpha, \beta, \gamma, \delta\}$ as we found each $f_g(\lambda)$ to be independent of the maximum first-generation mass m_{\max} and spin χ_{\max} for these simulations. Each branching fraction depends most strongly on the distribution of escape speeds — as determined by the power-law index δ — and the primary binary component pairing probability index α , whereas the indices of the first-generation mass distribution γ and the secondary component pairing β are less impactful.

When $\delta < 0$, the host clusters all have small escape speeds, and therefore, the branching fractions of sources with a remnant BH are close to zero, i.e., $f_{1g+1g} \approx 1$, as seen in the rightmost column of Fig. 5.9. With a fixed negative δ , as in the case of the green simulation, the branching fractions become independent of the other hyperparameters as no repeated mergers take place. On the other hand, when δ becomes positive, escape speeds are typically larger and repeated mergers can occur, so the contribution to the population from first-generation-only binaries decreases, i.e., $f_{1g+1g} < 1$.

Which binary generation then begins to dominate the population depends on the BH pairing process. When heavier (lighter) primary components form binaries due

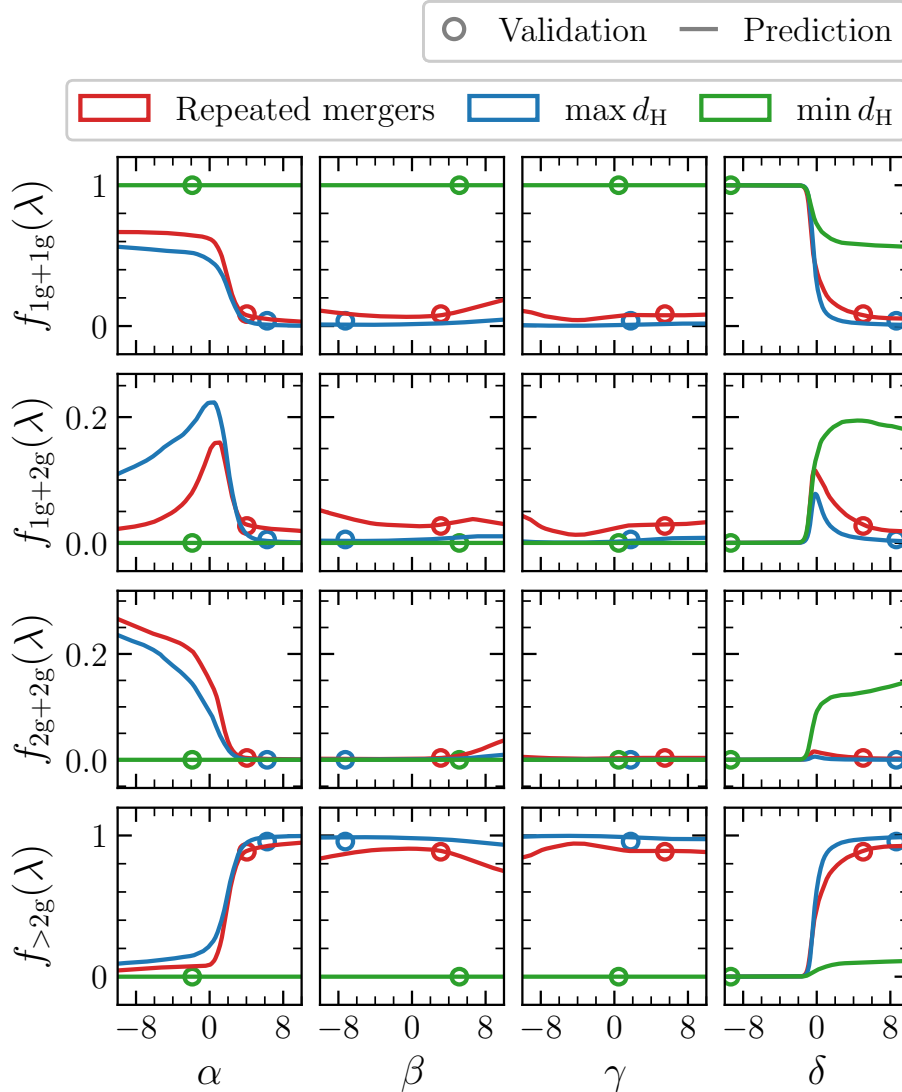


Figure 5.9: Example evaluations of the DNN predicting the binary merger-generation branching fractions $f_g(\lambda)$, $g \in \{1g + 1g, 1g + 2g, 2g + 2g, > 2g\}$ (from top to bottom rows), as a function of the hyperparameters λ . The results for hyperparameters taken from three illustrative simulations as in Fig. 5.5 — repeated mergers (red), and least and most accurate population predictions (max d_H in blue and min d_H in green, respectively) — are presented. The true values of the generation fractions are plotted as circle markers, whereas DNN predictions are given by solid lines. In each column, a single population-level parameter is varied while keeping the others fixed to those from the simulations.

to a fixed $\alpha > 0$ ($\alpha < 0$), as is the case for the red and blue (green) simulations, $>2g$ higher-generation (equal second-generation $2g+2g$) binaries preferentially populate the merger distributions. For small positive δ , the contribution from binaries of mixed first and second generations increases, but it is reduced at larger δ in favour of higher-generation mergers. For fixed positive δ (red and blue simulations), larger primary pairing indices α select, with increasingly strong preference, the heaviest remnant BHs in the population to form new binaries, thus increasing the fraction of greater-than-second-generation binaries, i.e., $f_{>2g} \approx 1$, while reducing the prevalence of other generations, as seen in the leftmost column of Fig. 5.9. The branching fractions are flat for $\alpha < 0$ if, as for the blue simulation, $\gamma < 0$ because all first-generation BHs are typically lighter; therefore, reducing the preference for low-mass primary BHs (i.e., making α less negative) has little effect. In contrast, when $\gamma > 0$ as in the red simulation, first-generation BHs are heavier, and increasing α while keeping β fixed will select heavier primaries relative to the secondaries, therefore favouring $1g+2g$ binaries.

5.4 Validation with mock catalogues

To test the inference pipeline in the absence of detection biases and single-event measurement uncertainties (equivalent to the limit of large SNRs) and without systematics due to the DNN population, we generate mock GW catalogues by drawing binary BH mergers from our DNN population model. Since for the technical reasons discussed in Sec. 5.3.1 this distribution is bounded in chirp mass, these draws are inherently taken from that range (listed in Table 5.1). This also means that the selection function constructed in Sec. 5.3.3 cannot be used in this mock inference; $\sigma'(\lambda)$ is defined over the entire range of source parameters, not just the observed range, and also accounts

for the required missing factor between $p'(\theta|\lambda)$ and $p(\theta|\lambda)$. Including selection effects would require training a different model for the detection efficiency and thus our tests would include ingredients that do not enter the actual inference of Sec. 5.5. Another technical difficulty is that we model two effective spins, χ_{eff} and χ_{p} , while the detection probability in principle depends on all six spin degrees of freedom. Creating a mock catalogue of observable GW events — i.e., taking samples from the detection-weighted population $\propto p(\text{det}|\vartheta)p(\theta|\lambda)$ — would require assuming an effective lower-dimensional dependence or resampling full spin vectors consistent with the sampled values of χ_{eff} and χ_{p} (cf. Ref. [517] for a more in-depth exploration of these issues). However, as we saw in Chapter 3, correctly including spin information in selection biases has a noticeable effect at the population level [421].

For testing purposes, we consider the high SNR limit, in which all events are detectable and their source parameters are measured exactly. This corresponds to a selection function $\sigma \equiv 1$ and single-event likelihoods $p(d_i|\theta_i) \propto \delta(\theta - \theta_i)$. From Eq. (5.8) the population-level likelihood is thus given by $p(\{d\}|\lambda) \propto \prod_{i=1}^n p'(\theta_i|\lambda)$ (where the statistical details are otherwise equivalent to Sec. 5.3.1). We draw $n = 50, 100, 200, 500$ events to create increasingly large catalogues (and in going from, e.g., 50 to 100 events, the first 50 are added to when increasing the catalogue size) with source parameters θ_i ($i = 1, \dots, n$) using rejection sampling of our neural-network population model. We repeat the analysis five times with new catalogues to assess the impact of Poisson fluctuations on the inference. To enable a conservative mock catalogue test we fix the true hyperparameters to those of the validation simulation with the lowest predictive accuracy for the DNN population model ($\max d_{\text{H}}$ in Figs. 5.5, 5.7, and 5.9): $\alpha = 6.3$, $\beta = -7.3$, $\gamma = 1.8$, $\delta = 8.7$, $m_{\text{max}} = 46M_{\odot}$, and $\chi_{\text{max}} = 0.01$.

We present the results of our mock inference runs in Figs. 5.10 and 5.11. The

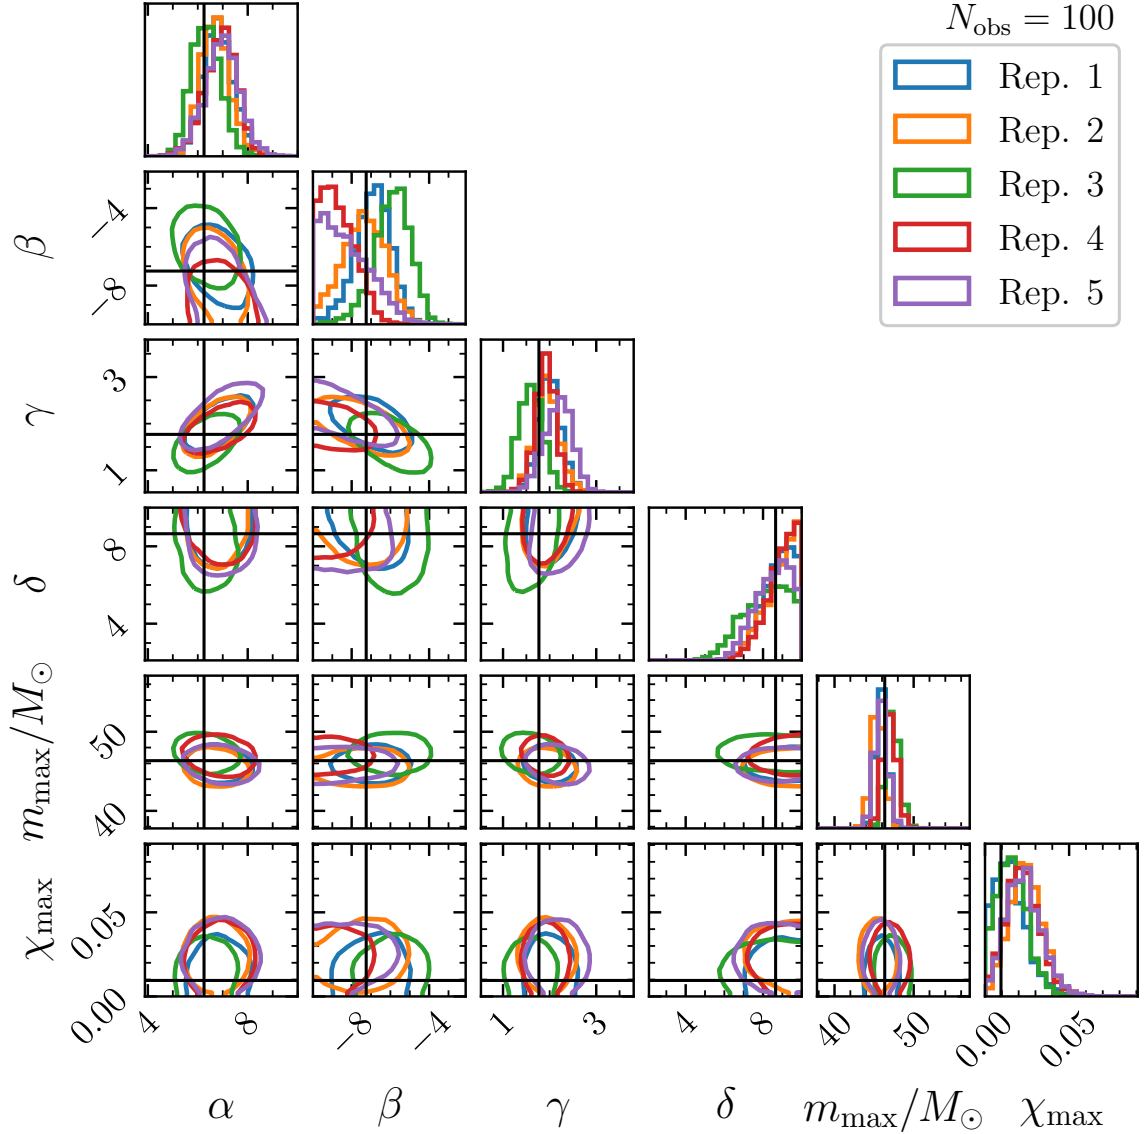


Figure 5.10: One- and two-dimensional marginalized posteriors of the population-level parameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$, corresponding to the $\max d_{\text{H}}$ simulation of Fig. 5.5, as measured from inference runs without measurement errors and selection biases (corresponding to the high SNR limit), and systematics from the DNN population model by drawing mock GW catalogues directly from $p'(\theta|\lambda)$. For the joint two-dimensional panels, each contour encloses the 90% credible region for a single analysis. Injected values are marked with black lines. The number of observations in the catalogue is fixed to $n = N_{\text{obs}} = 100$ and five independent realizations of the inference with distinct events performed, each represented with a different coloured curve (Reps. 1–5).

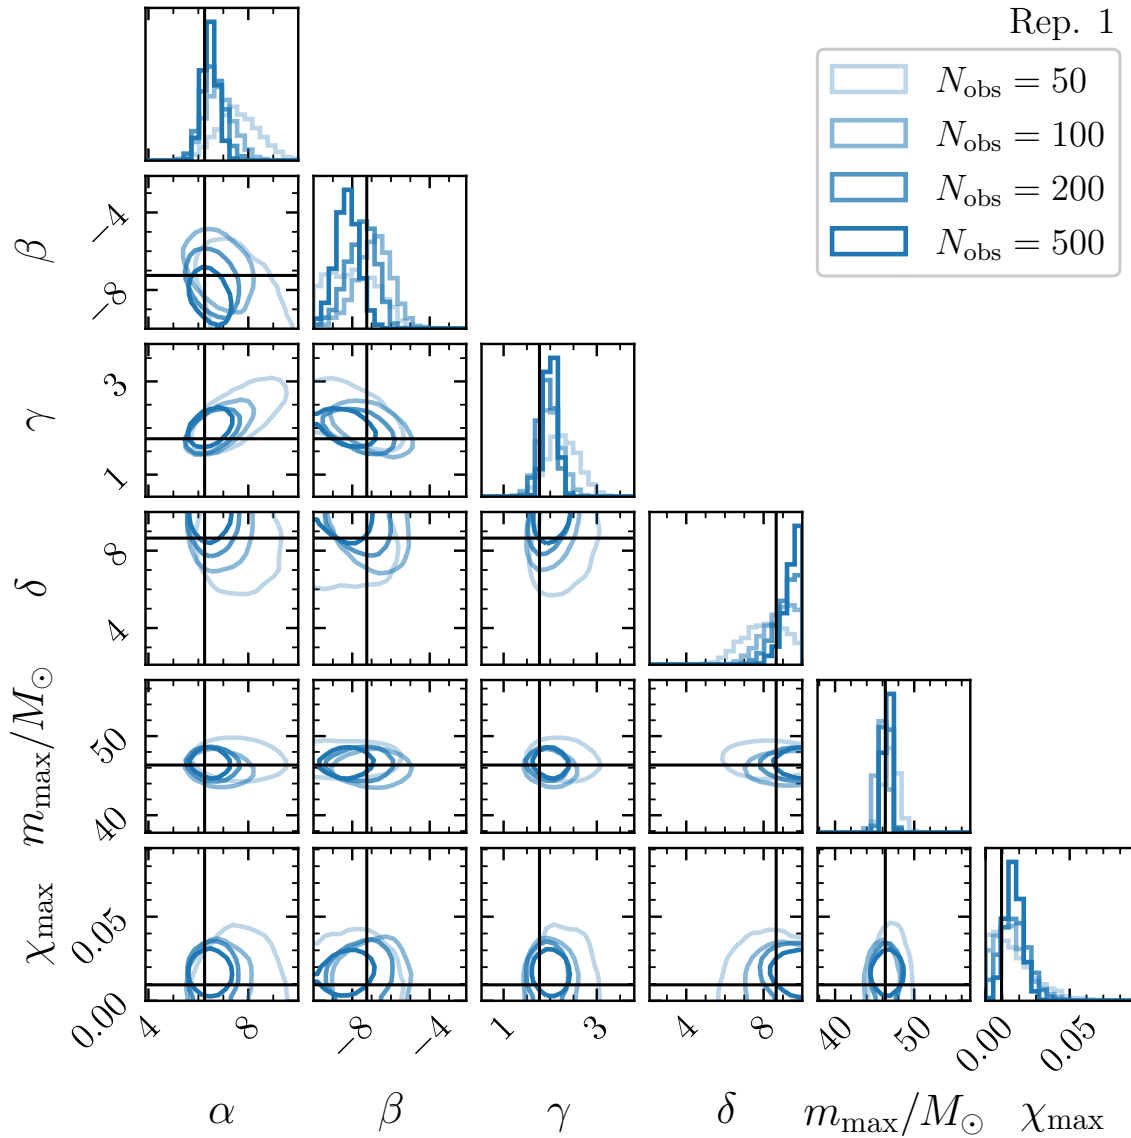


Figure 5.11: One- and two-dimensional marginalized posteriors of the population-level parameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$, corresponding to the $\max d_{\text{H}}$ simulation of Fig. 5.5, as measured from inference runs without measurement errors and selection biases (corresponding to the high SNR limit), and systematics from the DNN population model by drawing mock GW catalogues directly from $p'(\theta|\lambda)$. For the joint two-dimensional panels, each contour encloses the 90% credible region for a single analysis. Injected values are marked with black lines. Mock catalogues drawn for Rep. 1 are presented for an increasing number of observed events, $n = N_{\text{obs}} = 50, 100, 200, 500$ (light to dark shading).

one- and two-dimensional marginal posterior distributions of the hyperparameters are plotted, where the two-dimensional panels display the 90% contours. The solid black lines denote the true values listed above.

In Fig. 5.10 we fix the number of observations in the catalogue to $n = 100$ and perform five independent repetitions of the analysis with five different mock catalogues, given by the different coloured curves. Each run is consistent with both the injected hyperparameter values and each other at the 90% level, though there are significant fluctuations between realizations. Recall that the single-event likelihoods neglect measurement errors; relaxing this assumption and including nonzero widths in those posteriors would decrease the overall accuracy of the hyperparameter measurements and thus blend the results from independent realizations to distributions with greater consistency. Increasing the number of observations in the catalogue improves the hyperparameter measurement error and reduces the statistical fluctuations between realizations; we take $n = 100$ here to approximate the current size of real catalogues [64].

The impact of the growing size of the catalogue is illustrated in the right panel of Fig. 5.11. Here, we choose one particular realization and analyse the catalogue as increasing numbers of events are added incrementally (light- to dark-blue curves). We recover the expected result: The posterior constraints become tighter as n increases from 50 to 500 while remaining consistent with the true hyperparameter values at the 90% level. Larger catalogue sizes also break degeneracies between parameter pairs, e.g., the β - γ correlations, and remove posterior support in regions far from the truth, e.g., in the column for α .

If the events from the mock catalogues are instead drawn from the simulated populations used as validation samples when training the population model, one may expect

a systematic bias in the recovered hyperposteriors as the number of observations increases due to mismodelling in the trained DNN. Indeed, when repeating the above analysis but injecting from simulated validation data while recovering with the DNN, we find hyperposteriors that can exclude the injected values at 90% confidence for the lowest accuracy ($\max d_{\text{H}}$) simulation considered above when $n \geq 100$. However, this point is a considerable outlier in terms of accuracy (see Fig. 5.4). For most regions in the hyperparameter space the mismodelling between injection and recovery remains consistent at 90% confidence. In particular, we verify that this is true for the validation simulation whose hyperparameters are closest to the recovered medians in Sec. 5.5, suggesting our inference on the real catalogue below is robust within the measurement uncertainties. While the tests performed here are admittedly limited in scope, they allowed us to assess the renormalization and sampling capabilities of the pipeline.

5.5 Hierarchical-merger inference with gravitational-wave observations

In the following, we infer the population properties of the binary BHs detected by the LVC given our deep-learned population model of hierarchical mergers. In Fig. 5.12, we present the result of our population inference — the posterior distribution of the hyperparameters λ . Along the diagonal is the one-dimensional marginalization of each hyperparameter, while the other panels display the 50% and 90% confidence intervals of each two-dimensional distribution.

We will also refer to the implied source-parameter PPDs given by

$$p(\theta|\{d\}) = \int p'(\theta|\lambda)p(\lambda|\{d\})d\lambda, \quad (5.12)$$

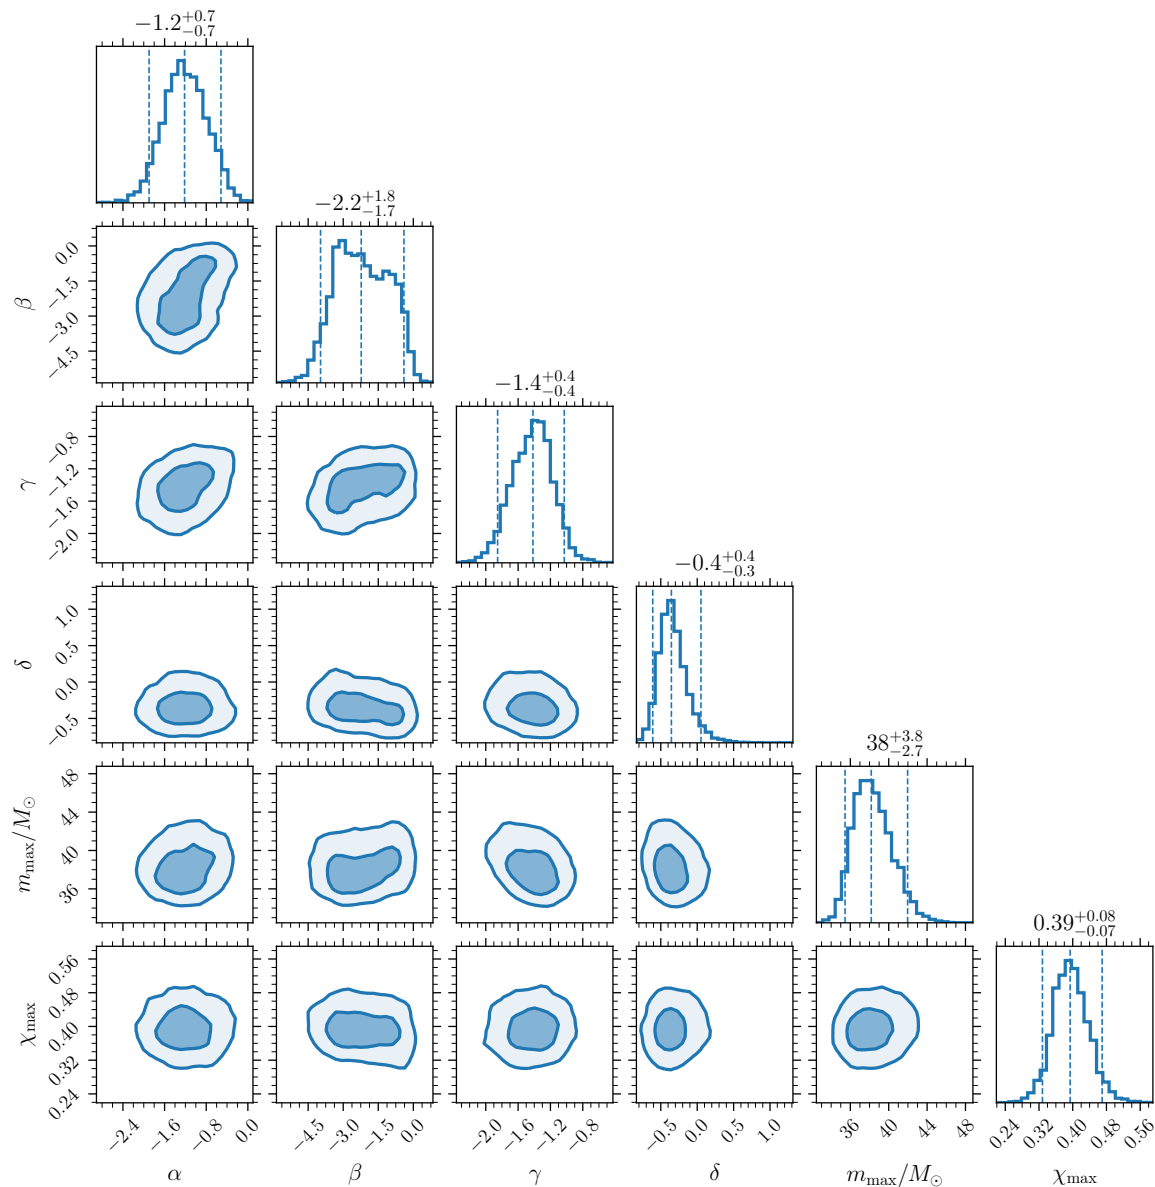


Figure 5.12: One- and two-dimensional marginal distributions of the population-level parameters $\lambda = \{\alpha, \beta, \gamma, \delta, m_{\max}, \chi_{\max}\}$ in our model of hierarchical mergers as measured using the real GW data from the confident ($\text{FAR} < 1 \text{ yr}^{-1}$) binary BH events. In each two-dimensional distribution, the contours enclose the 50% (dark shading) and 90% (light shading) confidence regions. The one-dimensional median and symmetric 90% intervals are reported above each diagonal and are plotted as vertical dashed lines in the corresponding panels.

such that the astrophysical distribution of each source parameter is given by the one-dimensional marginalizations of $p(\theta|\{d\})$.

5.5.1 Host escape speeds

We begin with the properties of host environments. We consider cluster-like hosts — simple collections of individual BHs that may be paired to form binaries — which are solely characterized by their escape speeds v_{esc} . Recall that the simulated clusters are distributed according to a truncated power law $p(v_{\text{esc}}|\delta) \propto v_{\text{esc}}^\delta$, with $0 \text{ km s}^{-1} < v_{\text{esc}} < 500 \text{ km s}^{-1}$. The value of each cluster escape speed controls whether repeated mergers take place since sources receiving larger gravitational kicks are ejected. Though a power-law distribution is a simplified model, it is indicative of a preference (or lack thereof) towards either edge of the domain.

The marginal distribution of the escape-speed index δ is displayed in the fourth diagonal entry in Fig. 5.12. Negative (positive) values of δ indicate an escape-speed distribution favouring lower (higher) values, while $\delta = 0$ corresponds to a uniform distribution in v_{esc} . We report a median and symmetric 90% interval of $\delta = -0.4_{-0.3}^{+0.4}$, corresponding to an escape-speed distribution biased toward smaller values though consistent with uniformity within the 90% credible bounds.

In Fig. 5.13, we display the distribution of escape speeds reconstructed from the marginal hyperposterior $p(\delta|\{d\})$. Marginalizing the escape-speed model $p(v_{\text{esc}}|\delta)$ over the uncertainty in the hyperparameter δ returns the PPD

$$p(v_{\text{esc}}|\{d\}) = \int p(v_{\text{esc}}|\delta)p(\delta|\{d\})d\delta, \quad (5.13)$$

whereas the posterior uncertainty is displayed interior to the 5% and 95% quantiles




Figure 5.13: PPDs of cluster escape velocities v_{esc} (blue), and of gravitational recoils v_{kick} for binaries consisting of two first-generation BHs (orange). The orange distributions are normalized over a range extending beyond the upper v_{esc} limit, as indicated by the arrow. The coloured shaded bands contextualize the velocity scale by denoting the typical escape speeds of globular clusters (green) and nuclear star clusters (purple).

of $p(v_{\text{esc}}|\delta)$, with $\delta \sim p(\delta|d)$. For context, order-of-magnitude estimates of the escape speeds of globular clusters (GCs; $\lesssim 100 \text{ km s}^{-1}$) and nuclear star clusters (NSCs; $\lesssim 500 \text{ km s}^{-1}$) are shown as horizontal coloured bands [292, 557, 558]. Additionally, we display the PPD of gravitational kicks v_{kick} received by the 1g+1g sources implied by our population model and hyperparameter constraints. Recall that, since the first-generation and binary-pairing distributions have parametric forms (see Sec. 5.2), the 1g+1g distribution does also (i.e., power-law mass distributions, uniform dimensionless spin magnitudes, isotropic spin directions; we present the measurements of the population-level parameters governing this distribution in the following sections). Though GW kicks peak at about 200 km s^{-1} , the distribution of escape speeds features support across the defined range up to $v_{\text{esc}} = 500 \text{ km s}^{-1}$. For these 1g+1g sources, we find that $P(v_{\text{kick}} < 500 \text{ km s}^{-1}) = 0.85_{-0.08}^{+0.06}$ and $P(v_{\text{kick}} < v_{\text{esc}}) = 0.37_{-0.12}^{+0.13}$, implying that host environments can retain the kicked remnants of a portion of first-generation mergers and support a population of hierarchical BHs.

5.5.2 Mass distribution

First-generation BHs — those born in stellar collapse — are drawn according to $p(m_{1g}|\gamma, m_{\text{max}}) \propto m_{1g}^{\gamma}$, with $5M_{\odot} < m_{1g} < m_{\text{max}}$ providing a mass limit corresponding to the lower (upper) edge of the purported upper (lower) mass gap. We recover $\gamma = -1.4_{-0.4}^{+0.4}$, implying lighter BHs closer in mass to $5M_{\odot}$ (chosen here to conservatively rule out NS and ambiguous source classifications) preferentially populate the underlying population. Negative 1g mass power-law exponents are expected as they reflect the stellar IMF [588].

If first-generation BHs are drawn from this single power-law prescription, we measure a first-generation upper mass limit of $m_{\text{max}} = 38_{-2.7}^{+3.8}M_{\odot}$. The presence of events

in the GW catalogue with component masses greater than $50M_\odot$ already points to the possibility of hierarchical mergers. Theoretical and simulated estimates of a mass gap location due to PISN typically predict $m_{\max} \sim 50M_\odot$, but they range within $40M_\odot \lesssim m_{\max} \lesssim 70M_\odot$ (or even higher [589]) due to various assumptions on key uncertain parameters [325, 327]. For comparison, taking the POWER LAW + PEAK analysis of Ref. [85] — which features a Gaussian peak with mean μ_m to model mass buildup, potentially due to PISN — we find consistency within 90% credible bounds between the inferred $\mu_m = 34^{+2.6}_{-4.0}M_\odot$ and m_{\max} . Note, however, that while m_{\max} is a sharp cut specifically characterizing the first-generation BH mass limit, the model of Ref. [85] parametrizes all BHs in a single distribution and μ_m is only the mean of a broadened feature, such that m_{\max} and μ_m are not directly equivalent. In our case, BH with masses larger than m_{\max} are accommodated with hierarchical mergers.

Here, we point out a key distinction of our modelling procedure: We assume all first-generation component BHs are drawn from a shared distribution (above) and then binary formation is separately modelled with component pairing probabilities $p(m_1|\alpha) \propto m_1^\alpha$ and $p(m_2|\beta, m_1) \propto m_2^\beta$ ($m_2 < m_1$). This choice differs from, e.g., POWER LAW + PEAK [85], which models each component mass distribution with multiple features superimposed on a power-law distribution. One may be tempted to think that, e.g., the primary mass distribution is equivalent to $p(m_1|\alpha)p(m_1|\gamma, m_{\max}) \propto m_1^{\alpha+\gamma}$ (and similarly for secondary masses), however this applies only to 1g+1g binaries. Our DNN population model additionally captures the interdependence between binary pairing and remnant retention. In short, the power-law indices parametrizing the distributions in this work are not directly comparable to such models. We infer $\alpha = -1.2^{+0.7}_{-0.7}$ and $\beta = -2.2^{+1.8}_{-1.7}$, such that both component pairing probabilities are bottom heavy with positive power-law indices ruled out at the 90% confidence level.

In Fig. 5.14, we present the inferred source distributions of the modelled mass parameters — chirp mass M_c and mass ratio q — and the implied distributions of primary and secondary masses, $m_1 = M_c(1+q)^{1/5}/q^{3/5}$ and $m_2 = qm_1$, respectively. Each PPD is plotted as a bold solid line, while the symmetric 90% confidence region of each marginal $p'(\theta|\lambda)$ with $\lambda \sim p(\lambda|d)$ is represented by shaded bands. The chirp mass distribution peaks at the minimum value $5M_\odot$ allowed by our model before an approximately exponential decline, with $M_c \lesssim 40M_\odot$. Equal-mass binaries are preferred in the underlying population, the mass-ratio distribution having a peak at $q = 1$ but with a broader linear decline down to $q \gtrsim 0.1$.

Substructure is apparent in the distributions of component source masses, corroborating the findings of Refs. [85, 590]. Tighter constraints at $M_c \approx 13M_\odot$ and $q \approx 0.6$ result in a cusp in the primary (secondary) mass distribution around $m_1 \approx 20M_\odot$ ($m_2 \approx 12M_\odot$) between two features: the peak of the distribution at $m_1 \approx 12M_\odot$ ($m_2 \approx 8M_\odot$) and a buildup-following decline at the first-generation mass limit $m_{\max} \approx 40M_\odot$. This suggests two contributions to the mass distribution in the range $20M_\odot \lesssim m_1 \lesssim 40M_\odot$: (1) first-generation BHs with masses above the peak of the distribution, and (2) higher-generation BHs with masses still smaller than m_{\max} but whose parents originally had masses in the peak 10–20 M_\odot region. While high-mass outliers above m_{\max} might be considered as clear indicators of repeated mergers, the bottom-heavy nature of the stellar IMF implies that hierarchical mergers may be prominent also for sources with masses below m_{\max} .

The first-generation and combined component mass distributions are compared in Fig. 5.15. In purple, we show the reconstructed distribution of first-generation masses,

$$p(m_{1g}|\{d\}) = \int p(m_{1g}|\gamma, m_{\max})p(\gamma, m_{\max}|d)d\gamma dm_{\max}, \quad (5.14)$$

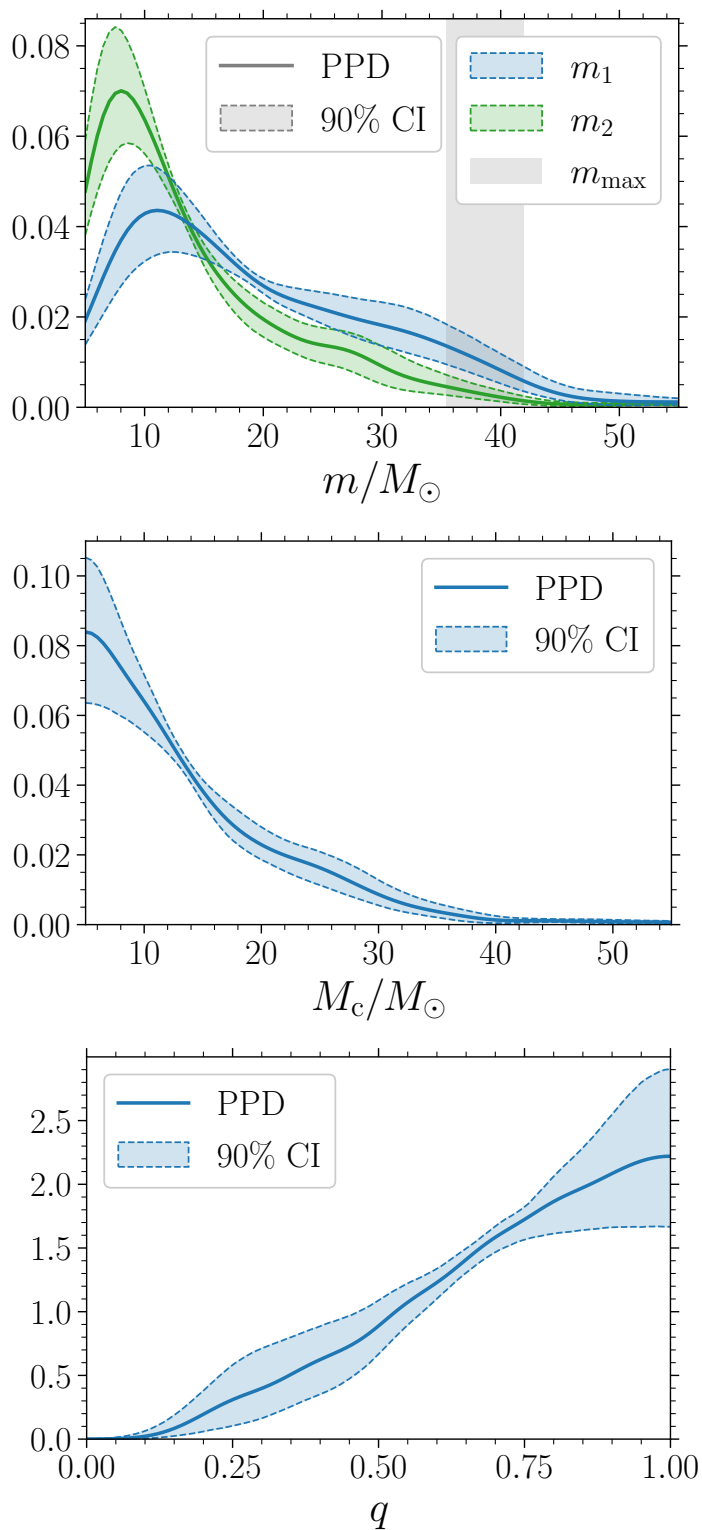


Figure 5.14: Astrophysical distributions of the modelled chirp mass M_c (middle panel) and mass ratio q (bottom panel), as well as the implied distributions of primary and secondary masses, m_1 and m_2 , respectively (top panel), as determined by our DNN population model and Bayesian analysis. The solid blue lines represent the PPDs, while the dashed lines enclose the 90% symmetric confidence intervals (shaded). In the left panel, the vertical grey band encloses the 90% confidence interval for the maximum mass of first-generation BHs, m_{\max} .




Figure 5.15: PPDs (in logarithmic scale) of the first-generation BH masses, m_{1g} , (purple) and the combined distribution of all components masses, m_1 and m_2 (blue). The solid lines denote the means, while the dashed lines bound the shaded 90% symmetric confidence regions. The vertical grey band encloses the 90% constraint on the maximum first-generation BH mass, m_{\max} .

and in blue, we show the joint distribution of all primary and secondary masses. The grey shaded band represents the 90% constraint on the mass limit of first-generation BHs, m_{\max} . Note the logarithmic scale, and that the PPD is a set of expectation values (i.e., means) and, as such, can lie outside the region bounded by given quantiles.

Though declining above the first-generation cutoff, the mass distribution features an extended spectrum above m_{\max} which cannot result from 1g+1g mergers in our model. We find that 99% of all BHs have masses less than $59_{-6.5}^{+7.8} M_{\odot}$. The spectrum ultimately abates at $m_1 > 80 M_{\odot}$ — roughly $2m_{\max}$, implying a lack of greater-than-2g mergers with parent components from the upper end of the 1g mass spectrum — and features multiple small-scale modes in the intervening region (that may be statistical

fluctuations). These observations again point to hierarchical mergers in the underlying population.

5.5.3 Spin distribution

Moving to binary BH spins, recall that the first generation of BHs are modelled with isotropic spins whose dimensionless magnitudes are distributed uniformly up to a maximum $\chi_{\max} \in (0, 1)$, representing the maximum natal spin a BH may be born with in stellar collapse. We infer a value $\chi_{\max} = 0.39^{+0.08}_{-0.07}$. With limited constraining power in the spin observables, the precise constraints reported here are likely to be very model dependent. We opt for a uniform distribution of 1g spin magnitudes because of the large uncertainties surrounding the spin of compact objects following core collapse (e.g., Refs. [329, 330, 470, 473, 474]); this is an area where more accurate observations and more constraining predictions are very much needed. The overall distribution of spins is determined jointly by the first-generation distribution, the binary pairing procedure (as inferred above), the GR mapping of binary to remnant properties, and the ejection or retention of merger remnants in host environments. While we account for the dimensionless spin magnitudes of higher-generation binaries in our population modelling, the spin directions are resampled isotropically.

A more solid finding we report is that the dimensionless spins of 1g+1g binary BHs are limited below those of merger remnants, typically around 0.7 [332]. Hierarchical BHs with much lower spins are extremely rare [541], yet another indication that some higher-generation binary BHs are required to fit the data with our model (cf. Sec. 5.5.4). We measure spins using two effective parameters: the effective aligned spin χ_{eff} measures the binary spin component parallel to the orbital plane [346], and the effective precessing spin χ_{p} measures the in-plane, two-spin projection [426]. For

sources with negligible, misaligned, or (equal-mass) oppositely aligned spins we have $\chi_{\text{eff}} \approx 0$, while large positive (negative) values indicate high aligned (antialigned) spins. Similarly, $\chi_{\text{p}} \approx 0$ for spins that are small, aligned with the orbital angular momentum, or oppositely aligned in the orbital plane. Nonzero values of χ_{p} indicate the presence of spin precession, with $\chi_{\text{p}} > 1$ being a region exclusively occupied by binaries with precessing spin contributions from both BH components.

Figure 5.16 displays the PPDs of these two modelled effective spin parameters. In the top panel, we show the distribution of effective aligned spins χ_{eff} . Here, the assumption of isotropic spins leads to an overly tight constraint. This mismodelling enforces a distribution that is symmetric about and centred on $\chi_{\text{eff}} = 0$, in contrast with more generic spin models that infer asymmetric distributions skewed to positive χ_{eff} [85] (and thus favouring alignment) or those that rule out negative χ_{eff} [455, 456]. However, we find that, typically, $|\chi_{\text{eff}}| \lesssim 0.4$, in agreement with the results of Ref. [85] (GAUSSIAN SPIN model); in particular, we report $|\chi_{\text{eff}}| < 0.46^{+0.04}_{-0.06}$ for 99% of the population.

On the other hand, the bottom panel of Fig. 5.16 shows the distribution of precessing spins measured with χ_{p} , where, unlike Ref. [85], we observe substructure; note that, although they use the earlier χ_{p} definition of Ref. [591], for the majority of events, the two measurements are indistinguishable [426, 592]. We note that, like for χ_{eff} , the uncertainty is likely also underestimated here due to our modelling assumptions. The distribution features two prominent modes. The primary one appears at $\chi \approx 0.2$. A peak at $\chi_{\text{p}} > 0$ is determined by the model, given isotropic spin directions (as is the case for all merger generations in our model) and uniform non-zero spin magnitudes (as for the first-generation binaries). A single-peaked distribution essentially corresponds to the implied χ_{p} prior used in PE analyses [426]. If this feature is astrophysical in

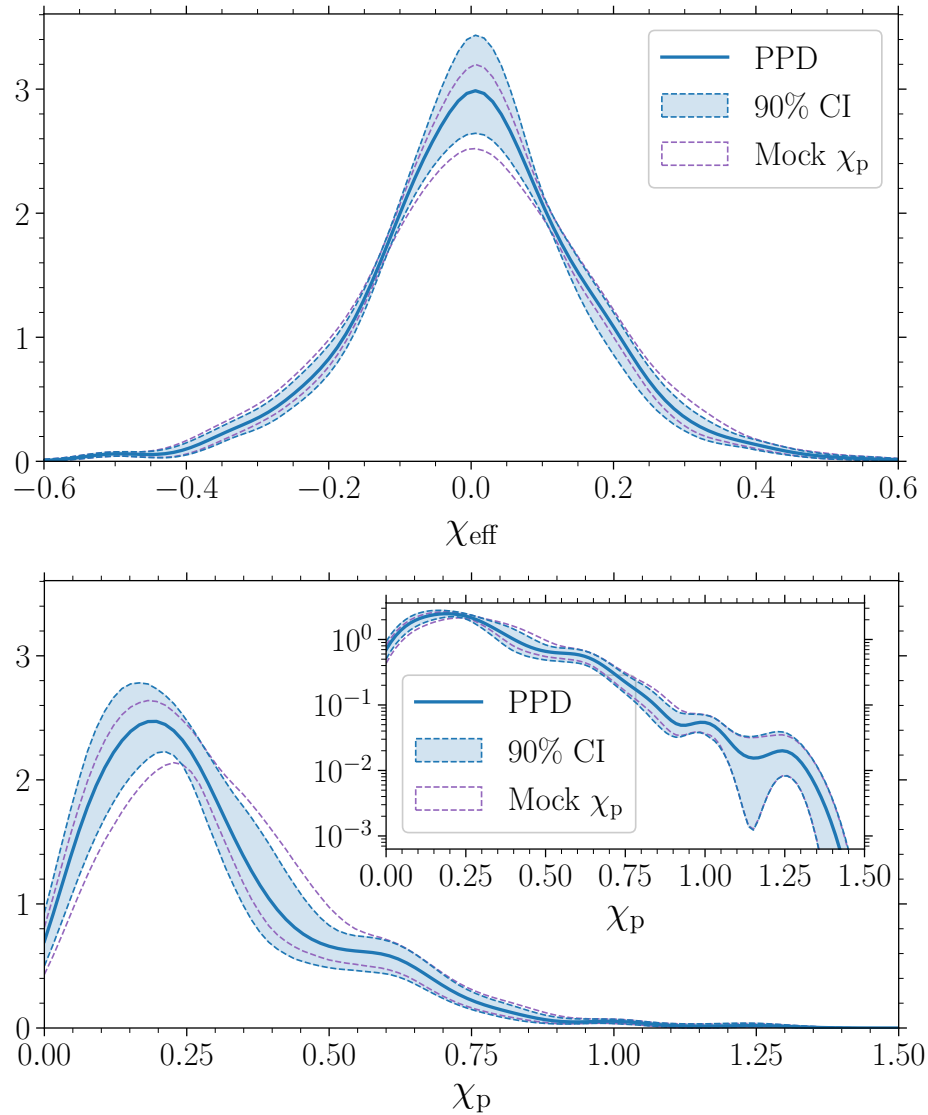


Figure 5.16: PPDs of the effective aligned (χ_{eff} , top panel) and precessing (χ_p , bottom panel) spins derived from our DNN population inference. The means of the distributions are plotted with solid blue lines, and the symmetric 90% confidence intervals are given by the shaded bands bounded by dashed lines. The inset for the χ_p panel shows the same distribution with logarithmic scaling to highlight smaller-scale features. The dashed purple lines bound the 90% confidence region of the distributions that are measured when replacing the true PE results for χ_p with mock samples from the prior for each event in the catalogue.

origin rather than due to our model choices, however, it may imply that sources with at least moderately misaligned spins — and thus undergoing spin precession — make up a sizeable portion of the population. The shape and location of this mode are in broad agreement with the results of Ref. [85]; see their Fig. 16.

However, in contrast to their finding that χ_p measurements can be explained by either a narrow distribution with peak $\chi_p \approx 0.2$ or a broad distribution centred on $\chi_p = 0$ (which results in multimodality when marginalized over the posterior uncertainty), we find that individual distributions drawn according to the hyperposterior always decrease at $\chi_p = 0$, peak at $\chi_p \approx 0.2$, and feature a secondary mode around $\chi_p \approx 0.6$. While our population model naturally accommodates such multimodal structure, the GAUSSIAN SPIN model employed in Ref. [85] only allows for a single peak and is thus unable to jointly capture the narrow $\chi_p \approx 0.2$ peak in addition to the extended distribution above $\chi_p \gtrsim 0.5$, instead favouring one or the other. Indeed, a single Gaussian distribution cannot fit the distribution of χ_p within the 90% credible bounds.

The inset in the χ_p panel of Fig. 5.16 shows the same distribution with logarithmic scaling to highlight smaller-scale features. The distribution falls off above the feature at $\chi_p \approx 0.6$ before a tertiary buildup at $\chi_p \approx 1$ and a final minor mode at $\chi \approx 1.25$, with a large decline in between and eventual declivity beyond. We find minor evidence for a population of sources occupying the exclusive two-spin region $\chi_p > 1$; the 99% quantile lies at $\chi_p = 0.95_{-0.13}^{+0.07}$ while $P(\chi_p > 1) = 0.8_{-0.4}^{+0.6}\%$. There is no support in the population for $\chi_p \gtrsim 1.5$.

Turning to the origins of these spin features, the precessing spin posterior we measure differs from a population prior with uniform spin magnitudes and directions due to the inferred constraint on the maximum natal spin χ_{\max} being less than the prior Kerr limit, leading to a shift towards lower values and, more importantly, the feature

at $\chi_p \approx 0.6$, which is not explainable with such a model. To test whether the posterior constraints are really due to measurements of precession or correlations with other parameters — primarily the best-measured spin χ_{eff} and the mass ratio q — we repeat the hierarchical inference but replace the χ_p posterior for each event in Eq. (5.8) with samples from the PE prior. The measured 90% confidence intervals for the effective spin posteriors are shown by the dashed purple lines in Fig. 5.16. The constraints are qualitatively the same, with only small differences in the 90% credible regions. The purple χ_p distribution favours slightly larger values in the region $\chi_p < 0.6$, suggesting the real precession measurements from GW data offer some information beyond the PE prior, but the differences are minor. The most informative constraints originate from the aforementioned better-measured parameters.

In our single-channel model, a BH with mass above m_{max} is necessarily a merger remnant. Since merger remnants have large spins, this model requires heavy BHs to have large spins if there are masses above m_{max} in the catalogue and natal spins are small, as inferred above. Our DNN model naturally allows for correlations between parameters, unlike simple phenomenological priors, so we can assess which masses contribute to the $\chi_p \approx 0.6$ feature. In the top panel of Fig. 5.17, we split the inferred χ_p population posterior into contributions from primary masses $m_1 \leq m_{\text{max}}$, which can be both first- or higher-generation BHs, and definitely higher-generation sources with $m_1 > m_{\text{max}}$. Though the latter, heavier population of sources necessarily has a preference for larger spins, the contribution to the χ_p distribution from sources with $m_1 \leq m_{\text{max}}$ still contains the feature at $\chi_p \approx 0.6$, implying that this inference is not solely driven by the requirement for heavy BHs to have large spins. This is a consequence of the previous conclusion that hierarchical mergers in our model also populate the region $m_1 \leq m_{\text{max}}$ due to the bottom-heavy mass function.




Figure 5.17: PPD of the effective precessing spin χ_p as a function of the primary mass m_1 (top panel) and mass ratio q (bottom panel). The mean distributions are given by the solid lines and the symmetric 90% confidence intervals are given by the shaded bands. In the top panel, we split the χ_p posterior for primary masses below (blue) and above (orange) the maximum first-generation mass m_{\max} . In the bottom panel, we split the χ_p posterior for mass ratios $0 < q \leq 1/3$ (blue), $1/3 < q \leq 2/3$ (orange), and $2/3 < q \leq 1$ (green).

In the bottom panel of Fig. 5.17, we similarly observe the population distribution of χ_p as a function of the mass ratio q . Larger precessing spins are suppressed for more unequal masses $q \leq 1/3$ with the peak lowered to $\chi_p \approx 0.15$, while for mass ratios $q > 2/3$ closer to unity, it increases to the slightly larger value $\chi_p \approx 0.25$. Spins around $\chi_p \approx 0.6$ are present for these mass ratios, but the distinct feature is most prominent for $1/3 < q \leq 2/3$. This region of the parameter space is prominently occupied by mixed-generation mergers, e.g., 1g+2g. We verify that this secondary structure is consistent with repeated mergers in our model as follows. Starting with the 1g+1g PPD and binary pairing measurements to compute the distribution of 2g remnant masses and dimensionless spin magnitudes (computed with Refs. [430, 435] as in Sec. 5.2.3), the distributions of χ_p for binaries formed either with a 1g BH and a remnant BH (i.e., 1g+2g) or two remnant BHs (i.e., 2g+2g) both feature peaks at $\chi_p \approx 0.6$. This is because, for 1g+2g sources, the dominant contribution to χ_p is from the primary, which is more likely to be 2g, while for 2g+2g sources, the primary and secondary are more likely to contribute equally such that their average is similar to the 1g+2g case.

5.5.4 Merger generations

Given our DNN population model, the observations of the previous sections suggest the presence of hierarchical mergers in the underlying population of merging stellar-mass binary BHs. Taking samples $\lambda \sim p(\lambda|\{d\})$ from the posterior distribution of population parameters, the corresponding draws from the posterior of merger-generation fractions can be derived as $f_g(\lambda) \sim p(f_g|\{d\})$, where f_g is given by the DNN described in Sec. 5.3.4.

Figure 5.18 presents the posterior distributions of the fractional contributions to

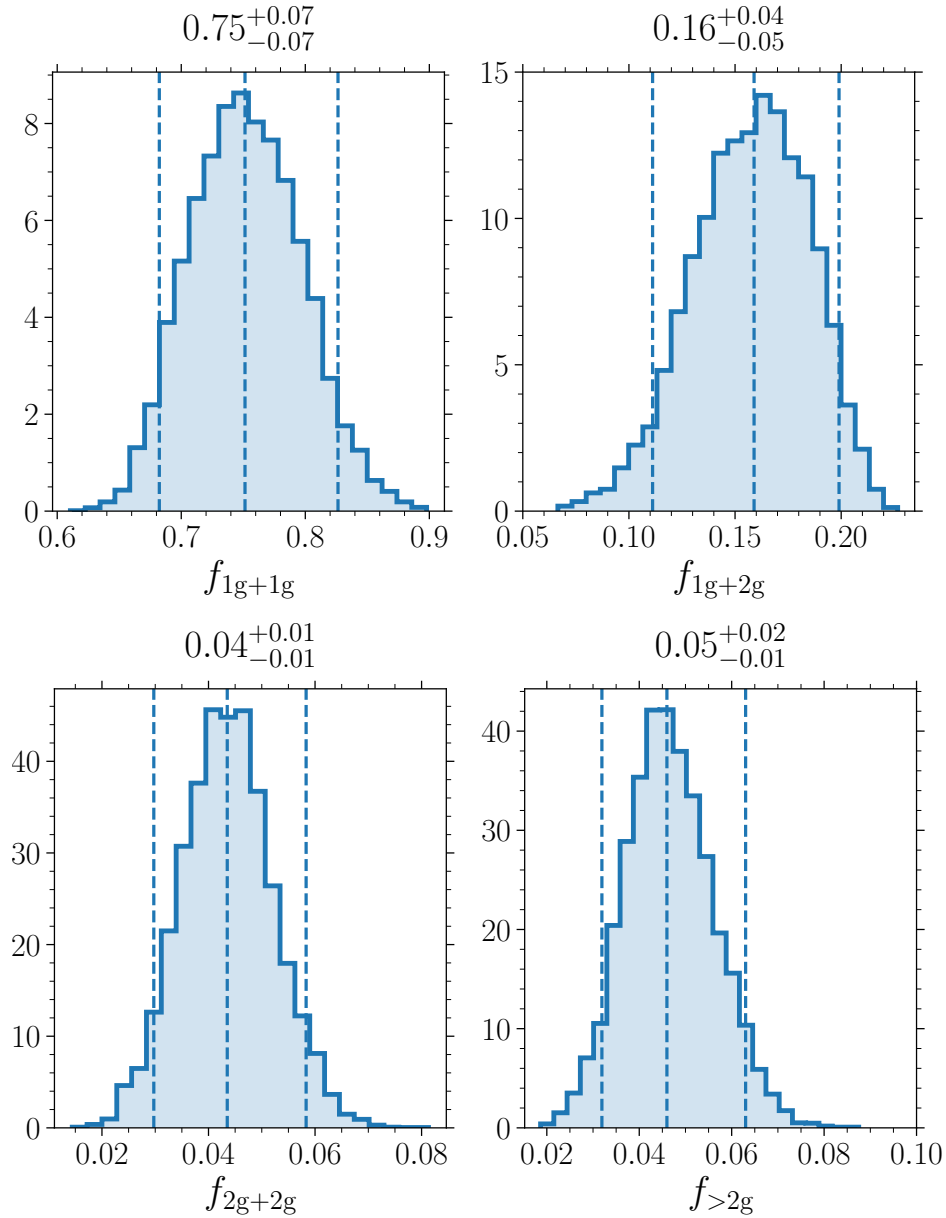


Figure 5.18: Distributions of the branching fractions — left to right: f_{1g+1g} , f_{1g+2g} (top row), and f_{2g+2g} , $f_{>2g}$ (bottom row) — for merger generations in the astrophysical distribution of merging stellar-mass binary BHs, as measured with our deep-learning approach to population inference. The median and symmetric 90% confidence region for each generation fraction is reported above — and plotted as vertical dashed lines within — the corresponding panel.

the population from each merger generation; the medians and 90% symmetric intervals are quoted and indicated as vertical dashed lines. In the underlying distribution, $75_{-7}^{+7}\%$ of sources contain only first-generation BHs (1g+1g), which implies around 25% contain a component that is the remnant of a previous merger, with 90% (99%) one-sided support for $1 - f_{1g+1g} \gtrsim 0.19$ (0.14). Mixed-generation binaries with both a first- and second-generation component make up the second-largest portion of the population, with $f_{1g+2g} = 0.16_{-0.05}^{+0.04}$, while binaries containing two second-generation BHs or any component of even higher generation contribute equally at about the 5% level ($f_{2g+2g} = 0.04_{-0.01}^{+0.01}$ and $f_{>2g} = 0.05_{-0.01}^{+0.02}$, respectively).

Previous studies of older GW catalogues found weak evidence for the presence of hierarchical mergers [593, 594]. However, new detections brought the addition of events whose properties, including higher masses and mass ratios, hinted at higher-generation origins. Ref. [528] presented a population analysis based on a phenomenological model of globular clusters, implying the presence of at least one second-generation BH in the catalogue with greater than 96% probability, rising to greater than 99.99% when considering their highest Bayes-factor model corresponding to an escape speed $v_{\text{esc}} \sim 300\text{km s}^{-1}$. In this case, they found median relative merger rates of 0.15 and 0.01 when comparing 1g+2g and 2g+2g binaries to the 1g+1g case, respectively, with 99% upper limits of 0.29 and 0.04. Equivalently, in our analysis we find broadly consistent relative branching fractions $f_{1g+2g}/f_{1g+1g} = 0.21_{-0.08}^{+0.08}$ and $f_{2g+2g}/f_{1g+1g} = 0.06_{-0.02}^{+0.03}$ (reporting medians and symmetric 90% confidence intervals). Given the disparity of the underlying model assumptions between the two analyses and the addition of new detections used in our study, our results jointly point to the fact that, if admitted in the fitted population, a modest number of binary BHs with hierarchical origin appears to be a likely way to explain the observed GW data.

5.6 Conclusions

While we are able to highlight some key insights into the astrophysics of stellar-mass black-hole mergers, we stress that our DNN population model is based on simulations that are simplified by, e.g., employing various power-law parametrizations. This work serves as a test case to demonstrate the efficacy of the modelling procedure by bridging the gap between phenomenological and accurate simulated models.

The complexity of our simulated populations can be increased in various ways. First, we model the spin distributions of first-generation binary BHs as uniform and isotropic, and while we account for the spin magnitude of merger remnants, we continue the assumption of isotropicity in higher-generation mergers. Capturing more realistic distributions of higher-generation spins requires both retaining information on post-merger spin directions and characterizing any changes in relative orientation during binary formation.

Employing more sophisticated spin magnitude models and adding an additional hyperparameter to control the degree of first-generation spin alignment, we can better capture the behaviour of a wider class of host environments, e.g., isolated evolution [465] or the disks of active galactic nuclei [446]. More generally, allowing for contributions from a mixture of distinct formation channels would lead to a more realistic fit. In particular, χ_{eff} has been shown to favour positive values, which may indicate a significant contribution from isolated binary formation to the merger rate [85] but we only consider a single-channel dynamical-environment formation model. By underestimating the location of the χ_{eff} distribution, the fraction of hierarchical mergers may be overestimated [595]. The added complexity of multichannel modelling is beyond the scope of this first study, and we aim to address it in future work.

Second, we model redshifts with a fixed distribution corresponding to a rate of events that is uniform in comoving volume and source-frame time. A simple extension would be to include a parametrized redshift model [407], though this would also increase the dimensionality of the hyperparameter space. We consider the mergers in our simulated populations as an ensemble and do not account for dynamical assembly of hierarchical mergers, i.e., for the fact that a remnant BH can only form a new binary at later times than its parent system [334]. In practice, one should include the effect of time delays between formation and merger and thus more realistically model the merger rate; we leave such explorations to future work.

Compared to previous work [404, 517], the approach presented here replaces the histogram approximation of simulated binary BH distributions with Gaussian KDEs and the emulation of these distributions across both the source- and population-level parameter spaces via GPR with DNNs. While GPR has been shown to be an ineffective approach in higher dimensions [518], alternative deep-learning techniques such as normalizing flows [521, 522] have proven successful [501, 519, 523].

Rather than training on probability-density evaluations (the required number of which scales exponentially with the dimensionality) as in this work, normalizing flows are trained directly on samples from the true distribution (thus scaling linearly with dimensionality), making the latter more effective in high-dimensional spaces. Furthermore, our methodology requires truncating the population model in the unbounded chirp-mass parameter in order to generate training data and employ numerical normalization, introducing a refactoring term in the population-level likelihood. This issue may be solved with domain compactification by, e.g., modelling the inverse of the mass scale instead. Normalizing flows also have the advantage of being generative models (i.e., from which new predictive samples can be drawn) that additionally pro-

vide density estimation with correctly enforced normalization. However, often either just the forward (density prediction) or inverse (sample generation) model is efficient to evaluate [520, 596]. On the other hand, neither deep-learning approach provides any estimation of modelling uncertainty, whereas GPR does; this is an area where Bayesian deep learning may prove fruitful (see, e.g., Ref. [597]).

That said, our DNN framework has some advantages. The separation of density estimation and emulation adds a level of flexibility not otherwise available. The KDE-first approach allows for sufficient smoothing of the training distributions prior to the learning stage. This distinction is important for simulations with modelling choices that lead to unphysical numerical discretization of outputs or low sampling densities; e.g., despite employing normalizing flows, Ref. [501] required KDE resampling to boost the set of simulated BH mergers.

This work serves as a showcase of the developments made possible by combining advanced techniques from the fields of deep learning and statistical analysis, applied within the context of GW astrophysics. Our deep-learning population pipeline, which we applied to the case study of simple simulations of hierarchical stellar-mass BH mergers, is thus ready to be used in conjunction with more sophisticated simulated populations. Combined with the state of the art in population synthesis, we will be able to constrain the properties of progenitor formation environments by directly comparing GW data with higher-dimensional models of binary evolution.

Conclusions

Phrases to the effect of “GWs have opened a new window onto the universe” have now been heard so often that you’d be forgiven for rolling your eyes at hearing them again — but it’s true; GWs genuinely are an extraordinary compliment to the remit of messengers we can use for observational astronomy. Not only that, but they probe the regions of space and time previously forbidden to us — those dictated by the rules of strong gravity. Several years on from the first observation of GWs, we now almost take for granted the fact that direct measurements of compact objects and their collisions have become routine. This new probe has excited research programs in astrophysics, cosmology, and relativity — with unprecedented returns (Chapter 1).

Of particular interest is the interplay between astrophysics and relativity. The formation of GW sources, including binary BHs, is determined by astrophysical processes that are often uncertain. Their GW-driven inspirals are governed in contrast by the mathematical predictions of GR. If we observe GWs and measure the properties of their sources, can we unravel their histories and demographics? That is the motivational question driving this Thesis.

In Chapter 2, we saw that astrophysical predictions do not necessarily translate directly into GW observables. The up–down spin-precession instability is an extreme

example in which a very singular natal prediction becomes a very different one at detectable GW frequencies. Luckily, we derived that unstable does not imply unpredictable — there is a smoking-gun signature for this peculiarity of the relativistic two-body problem that we can test with continuing GW observations.

Converting these measurements made with GW detectors to a reference suitable for astrophysical modelling is also possible, as we showed in Chapter 3. This conversion is not just for convenience — it also corrects a fundamental conceptual issue that plagues GW population studies. Future observational catalogues may reveal the influence of such systematic biases and it will become crucial to correct them in order to make accurate inferences on the formation of GW populations.

We gave an example of this kind of GW-to-astrophysics inference in Chapter 4. If binary BH formation models specify key evolutionary processes for which astrophysical inputs result in clear outcomes for merging compact binaries, we can conversely leverage the observed GW catalogue — whose sources may have gone through those processes — to constrain their occurrence. Targeted population modelling is crucial in this regard.

The best way to build astrophysics-informed population models for GW sources is to use those very simulations into which we can project all of our prior knowledge. We successfully demonstrated how to transform population simulations — which are inherently discrete — into the accurate and continuous forms required for Bayesian inference (i.e., measurement instead of comparison) in Chapter 5. We predict that, backed by powerful emulation tools such as deep learning, simulation-based inference will be a major influence on the future of GW population studies, enabling hitherto inaccessible astrophysical and cosmological constraints.

It remains to perform a comprehensive astrophysically-motivated population analysis of GW sources. This will require a full and up-to-date combined model of all the

major competitor formation pathways, their subchannels, and selection biases, in tandem with reliable hierarchical Bayesian inference — no mean feat, and no surprise such a study is so far lacking in the literature. Nonetheless, advanced probabilistic learning methods [395, 598, 599, 600, 601] may prove fruitful (as suggested by our preliminary investigations, not presented in this Thesis) in the not-too-distant future.

Astrophysics-agnostic inference is another avenue that we have yet to explore. Rather than replacing overly parametrized population models with highly informative astrophysical models, one could instead perform inference over a class of distributions that is flexible enough to capture any realistic population, independent of astrophysical parameters. Though this removes interpretability, the key advantage is that we are less limited by our prior assumptions, explicit or otherwise. Suggested approaches include data-driven statistical methods [457], regression over symbolically parametrized function spaces [602], flexible semi-parametric models [603, 604], and normalizing flows [601] — in our view this last option appears most promising as it readily extends to high-dimensional distributions, thus allowing for complex correlations in the binary BH parameters that naturally result from astrophysical formation; outstanding issues are the inclusion of selection biases and robust estimation of measurement uncertainties, however. Our future work will chase down these loose ends.

Astrophysics governs the births and natal development of BH binaries, while relativity governs their remaining lives and deaths. As astronomers, we intrude on their working biographies via the messages carried by GWs. Despite the uncertainties at either extrema of these inspiral timelines, one thing is for certain: we will only get more data. It is our imperative to maximize the scientific outcome of these GW observations. I look forward to the future in this GW astronomy.

Bibliography

- [1] O. Heaviside, *Electromagnetic Theory* **1**, 455 (1893).
- [2] H. Poincaré, *Rendiconti del Circolo matematico di Palermo* **21**, 129 (1906).
- [3] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* , 844 (1915).
- [4] A. Einstein, *Annalen der Physik* **354**, 769 (1916).
- [5] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* , 688 (1916).
- [6] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* , 154 (1918).
- [7] D. Kennefick, *Traveling at the Speed of Thought: Einstein and the Quest for Gravitational Waves* (2007).
- [8] H. Collins, *Gravity's Ghost and Big Dog: Scientific Discovery and Social Analysis in the Twenty-First Century*, EBL-Schweitzer (University of Chicago Press, 2014).
- [9] J. L. Cervantes-Cota, S. Galindo-Uribarri, and G. F. Smoot, *Universe* **2**, 10.3390/universe2030022 (2016).
- [10] H. Collins, *Gravity's Kiss: The Detection of Gravitational Waves*, The MIT Press (MIT Press, 2017).
- [11] A. S. Eddington, *Proceedings of the Royal Society of London Series A* **102**, 268 (1922).
- [12] A. Einstein and N. Rosen, *Journal of The Franklin Institute* **223**, 43 (1937).
- [13] F. A. E. Pirani, *Acta Physica Polonica* **15**, 389 (1956).
- [14] J. Weber, *Phys. Rev. Lett.* **18**, 498 (1967).
- [15] J. Weber, *Phys. Rev. Lett.* **20**, 1307 (1968).

-
- [16] J. Weber, *Phys. Rev. Lett.* **22**, 1320 (1969).
- [17] J. Weber, *Phys. Rev. Lett.* **24**, 276 (1970).
- [18] V. B. Braginskii, A. B. Manukin, E. I. Popov, V. N. Rudenko, and A. A. Khorev, *Soviet Journal of Experimental and Theoretical Physics Letters* **16**, 108 (1972).
- [19] J. A. Tyson, C. G. MacLennan, and L. J. Lanzerotti, *Phys. Rev. Lett.* **30**, 1006 (1973).
- [20] R. L. Garwin and J. L. Levine, *Phys. Rev. Lett.* **31**, 176 (1973).
- [21] J. L. Levine and R. L. Garwin, *Phys. Rev. Lett.* **31**, 173 (1973).
- [22] J. A. Tyson, *Phys. Rev. Lett.* **31**, 326 (1973).
- [23] R. W. P. Drever, J. Hough, R. Bland, and G. W. Lessnoff, *Nature* **246**, 340 (1973).
- [24] V. B. Braginskii, A. B. Manukin, E. I. Popov, V. N. Rudenko, and A. A. Khorev, *Soviet Journal of Experimental and Theoretical Physics* **39**, 387 (1974).
- [25] J. L. Levine and R. L. Garwin, *Phys. Rev. Lett.* **33**, 794 (1974).
- [26] H. Billing, P. Kafka, K. Maischberger, F. Meyer, and W. Winkler, *Nuovo Cimento Lettere* **12**, 111 (1975).
- [27] D. H. Douglass, R. Q. Gram, J. A. Tyson, and R. W. Lee, *Phys. Rev. Lett.* **35**, 480 (1975).
- [28] H. Hirakawa and K. Narihara, *Phys. Rev. Lett.* **35**, 330 (1975).
- [29] J. L. Logan, *Physics Today* **26**, 44 (1973).
- [30] J. A. Tyson and R. P. Giffard, *ARA&A* **16**, 521 (1978).
- [31] A. Blum, R. Lalli, and J. Renn, *Nature Astronomy* **2**, 534 (2018).
- [32] D. Rickles and C. M. DeWitt, in *The Role of Gravitation in Physics: Report from the 1957 Chapel Hill Conference* (2011).
- [33] H. Bondi, *Nature* **179**, 1072 (1957).
- [34] J. Weber and J. A. Wheeler, *Reviews of Modern Physics* **29**, 509 (1957).
- [35] J. Weber, *Physical Review* **117**, 306 (1960).
- [36] R. Weiss, *General Relativity and Gravitation* **54**, 153 (2022).

-
- [37] LIGO Scientific Collaboration, J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, et al., *Classical and Quantum Gravity* **32**, 074001 (2015), 1411.4547 [gr-qc].
- [38] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, et al., *Classical and Quantum Gravity* **32**, 024001 (2015), 1408.3978 [gr-qc].
- [39] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Phys. Rev. Lett.* **116**, 061102 (2016), 1602.03837 [gr-qc].
- [40] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Phys. Rev. Lett.* **116**, 241102 (2016), 1602.03840 [gr-qc].
- [41] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *ApJ* **818**, L22 (2016), 1602.03846 [astro-ph.HE].
- [42] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Annalen der Physik* **529**, 1600209 (2017), 1608.01940 [gr-qc].
- [43] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Phys. Rev. Lett.* **116**, 221101 (2016), 1602.03841 [gr-qc].
- [44] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *Phys. Rev. Lett.* **119**, 161101 (2017), 1710.05832 [gr-qc].
- [45] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *Physical Review X* **9**, 011001 (2019), 1805.11579 [gr-qc].
- [46] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *ApJ* **850**, L40 (2017), 1710.05838 [astro-ph.HE].
- [47] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *ApJ* **848**, L12 (2017), 1710.05833 [astro-ph.HE].

- [48] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *ApJ* **848**, L13 (2017), 1710.05834 [astro-ph.HE].
- [49] A. Goldstein, P. Veres, E. Burns, M. S. Briggs, R. Hamburg, D. Kocevski, C. A. Wilson-Hodge, R. D. Preece, S. Poolakkil, O. J. Roberts, et al., *ApJ* **848**, L14 (2017), 1710.05446 [astro-ph.HE].
- [50] V. Savchenko, C. Ferrigno, E. Kuulkers, A. Bazzano, E. Bozzo, S. Brandt, J. Chenevez, T. J. L. Courvoisier, R. Diehl, A. Domingo, et al., *ApJ* **848**, L15 (2017), 1710.05449 [astro-ph.HE].
- [51] D. A. Coulter, R. J. Foley, C. D. Kilpatrick, M. R. Drout, A. L. Piro, B. J. Shappee, M. R. Siebert, J. D. Simon, N. Ulloa, D. Kasen, et al., *Science* **358**, 1556 (2017), 1710.05452 [astro-ph.HE].
- [52] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *Phys. Rev. Lett.* **121**, 161101 (2018), 1805.11581 [gr-qc].
- [53] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *ApJ* **850**, L39 (2017), 1710.05836 [astro-ph.HE].
- [54] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., *Nature* **551**, 85 (2017), 1710.05835 [astro-ph.CO].
- [55] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **915**, L5 (2021), 2106.15163 [astro-ph.HE].
- [56] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, et al., *Phys. Rev. D* **102**, 043015 (2020), 2004.08342 [astro-ph.HE].
- [57] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, et al., *ApJ* **896**, L44 (2020), 2006.12611 [astro-ph.HE].
- [58] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, et al., *Phys. Rev. Lett.* **125**, 101102 (2020), 2009.01075 [gr-qc].

- [59] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, et al., *ApJ* **900**, L13 (2020), 2009.01190 [astro-ph.HE].
- [60] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Physical Review X* **6**, 041015 (2016), arXiv:1606.04856 [gr-qc].
- [61] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Physical Review X* **9**, 031040 (2019), 1811.12907 [astro-ph.HE].
- [62] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Physical Review X* **11**, 021053 (2021), 2010.14527 [gr-qc].
- [63] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, et al., arXiv e-prints , arXiv:2108.01045 (2021), 2108.01045 [gr-qc].
- [64] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, et al., arXiv e-prints , arXiv:2111.03606 (2021), 2111.03606 [gr-qc].
- [65] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, *ApJ* **872**, 195 (2019), 1811.01921 [gr-qc].
- [66] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, *ApJ* **891**, 123 (2020), 1910.05331 [astro-ph.HE].
- [67] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, *ApJ* **922**, 76 (2021), 2105.09151 [astro-ph.HE].
- [68] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, *ApJ* **946**, 59 (2023).
- [69] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019), 1902.10341 [astro-ph.IM].
- [70] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **101**, 083030 (2020), 1904.07214 [astro-ph.HE].
- [71] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, *Phys. Rev. D* **106**, 043009 (2022), 2201.02252 [astro-ph.HE].

-
- [72] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, et al., *SoftwareX* **13**, 100658 (2021), 1912.11716 [gr-qc].
- [73] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, R. Abbott, H. Abe, F. Acernese, K. Ackley, S. Adhicary, N. Adhikari, R. X. Adhikari, et al., *arXiv e-prints* , arXiv:2302.03676 (2023), 2302.03676 [gr-qc].
- [74] L. S. Collaboration and V. Collaboration, 10.5281/zenodo.5117702 (2022).
- [75] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration, 10.5281/zenodo.5546662 (2021).
- [76] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Phys. Rev. D* **100**, 104036 (2019), 1903.04467 [gr-qc].
- [77] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Phys. Rev. D* **103**, 122002 (2021), 2010.14529 [gr-qc].
- [78] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, et al., *arXiv e-prints* , arXiv:2112.06861 (2021), 2112.06861 [gr-qc].
- [79] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **909**, 218 (2021), 1908.06060 [astro-ph.CO].
- [80] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, et al., *arXiv e-prints* , arXiv:2111.03604 (2021), 2111.03604 [astro-ph.CO].
- [81] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **923**, 14 (2021), 2105.06384 [gr-qc].
- [82] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, R. Abbott, H. Abe, F. Acernese, K. Ackley, S. Adhicary, N. Adhikari, R. X. Adhikari, et al., *arXiv e-prints* , arXiv:2304.08393 (2023), 2304.08393 [gr-qc].

-
- [83] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **882**, L24 (2019), 1811.12940 [astro-ph.HE].
- [84] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **913**, L7 (2021), 2010.14533 [astro-ph.HE].
- [85] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Physical Review X* **13**, 011048 (2023).
- [86] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Phys. Rev. D* **94**, 102001 (2016), 1605.01785 [gr-qc].
- [87] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos, et al., *Phys. Rev. D* **101**, 084002 (2020), 1908.03584 [astro-ph.HE].
- [88] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **104**, 122004 (2021), 2107.03701 [gr-qc].
- [89] C. L. Fryer and K. C. B. New, *Living Reviews in Relativity* **14**, 1 (2011).
- [90] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Phys. Rev. D* **103**, 064017 (2021), 2012.12128 [gr-qc].
- [91] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **105**, 022002 (2022), 2109.09255 [astro-ph.HE].
- [92] R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **106**, 062002 (2022), 2201.10104 [gr-qc].
- [93] R. Abbott, H. Abe, F. Acernese, K. Ackley, S. Adhicary, N. Adhikari, R. X. Adhikari, V. K. Adkins, V. B. Adya, C. Affeldt, et al., *ApJ* **941**, L30 (2022), 2209.02863 [astro-ph.HE].
- [94] J. T. Whelan, R. Tenorio, J. K. Wofford, J. A. Clark, E. J. Daw, E. Goetz, D. Keitel, A. Neunzert, A. M. Sintes, K. J. Wagner, et al., *arXiv e-prints*, arXiv:2302.10338 (2023), 2302.10338 [astro-ph.HE].

-
- [95] T. L. Killestein, M. Mould, D. Steeghs, J. Casares, D. K. Galloway, and J. T. Whelan, *MNRAS* **520**, 5317 (2023), 2302.00018 [astro-ph.HE].
- [96] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Phys. Rev. D* **104**, 082004 (2021), 2107.00600 [gr-qc].
- [97] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *ApJ* **932**, 133 (2022), 2112.10990 [gr-qc].
- [98] R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **106**, 102008 (2022), 2201.00697 [gr-qc].
- [99] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *ApJ* **921**, 80 (2021), 2105.11641 [astro-ph.HE].
- [100] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **105**, 082005 (2022), 2111.15116 [gr-qc].
- [101] R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **106**, 042003 (2022), 2204.04523 [astro-ph.HE].
- [102] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, O. Behnke, et al., *ApJ* **909**, 79 (2021), 2009.12260 [astro-ph.HE].
- [103] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, et al., *Phys. Rev. D* **104**, 022005 (2021), 2103.08520 [gr-qc].
- [104] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Phys. Rev. D* **105**, 122001 (2022), 2110.09834 [gr-qc].
- [105] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binétruy, M. Born, D. Bortoluzzi, et al., *arXiv e-prints*, arXiv:1702.00786 (2017), 1702.00786 [astro-ph.IM].
- [106] P. Amaro-Seoane, S. Aoudia, S. Babak, P. Binétruy, E. Berti, A. Bohé, C. Caprini, M. Colpi, N. J. Cornish, K. Danzmann, et al., *GW Notes* **6**, 4 (2013), 1201.3621 [astro-ph.CO].

-
- [107] E. Barausse, E. Berti, T. Hertog, S. A. Hughes, P. Jetzer, P. Pani, T. P. Sotiriou, N. Tamanini, H. Witek, K. Yagi, et al., *General Relativity and Gravitation* **52**, 81 (2020), 2001.09793 [gr-qc].
- [108] S. A. Hughes, *MNRAS* **331**, 805 (2002), astro-ph/0108483 [astro-ph].
- [109] A. Vecchio, *Phys. Rev. D* **70**, 042001 (2004), astro-ph/0304051 [astro-ph].
- [110] A. Sesana, F. Haardt, P. Madau, and M. Volonteri, *ApJ* **623**, 23 (2005), astro-ph/0409255 [astro-ph].
- [111] A. Sesana, M. Volonteri, and F. Haardt, *MNRAS* **377**, 1711 (2007), astro-ph/0701556 [astro-ph].
- [112] K. G. Arun, S. Babak, E. Berti, N. Cornish, C. Cutler, J. Gair, S. A. Hughes, B. R. Iyer, R. N. Lang, I. Mandel, et al., *Classical and Quantum Gravity* **26**, 094027 (2009), 0811.1011 [gr-qc].
- [113] G. Pratten, A. Klein, C. J. Moore, H. Middleton, N. Steinle, P. Schmidt, and A. Vecchio, *arXiv e-prints*, arXiv:2212.02572 (2022), 2212.02572 [gr-qc].
- [114] P. Amaro-Seoane, J. R. Gair, M. Freitag, M. C. Miller, I. Mandel, C. J. Cutler, and S. Babak, *Classical and Quantum Gravity* **24**, R113 (2007), astro-ph/0703495 [astro-ph].
- [115] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017), 1703.09722 [gr-qc].
- [116] A. Stroeer and A. Vecchio, *Classical and Quantum Gravity* **23**, S809 (2006), astro-ph/0605227 [astro-ph].
- [117] A. J. Ruiter, K. Belczynski, M. Benacquista, S. L. Larson, and G. Williams, *ApJ* **717**, 1006 (2010), 0705.3272 [astro-ph].
- [118] V. Korol, S. Toonen, A. Klein, V. Belokurov, F. Vincenzo, R. Buscicchio, D. Gerosa, C. J. Moore, E. Roebber, E. M. Rossi, et al., *A&A* **638**, A153 (2020), 2002.10462 [astro-ph.GA].
- [119] E. Roebber, R. Buscicchio, A. Vecchio, C. J. Moore, A. Klein, V. Korol, S. Toonen, D. Gerosa, J. Goldstein, S. M. Gaebel, et al., *ApJ* **894**, L15 (2020), 2002.10465 [astro-ph.GA].
- [120] E. Finch, G. Bartolucci, D. Chucherko, B. G. Patterson, V. Korol, A. Klein, D. Bandopadhyay, H. Middleton, C. J. Moore, and A. Vecchio, *arXiv e-prints*, arXiv:2210.10812 (2022), 2210.10812 [astro-ph.SR].

-
- [121] K. Breivik, C. L. Rodriguez, S. L. Larson, V. Kalogera, and F. A. Rasio, *ApJ* **830**, L18 (2016), 1606.09558 [astro-ph.GA].
- [122] D. Gerosa, S. Ma, K. W. K. Wong, E. Berti, R. O’Shaughnessy, Y. Chen, and K. Belczynski, *Phys. Rev. D* **99**, 103004 (2019), 1902.00021 [astro-ph.HE].
- [123] R. Busicchio, A. Klein, E. Roebber, C. J. Moore, D. Gerosa, E. Finch, and A. Vecchio, *Phys. Rev. D* **104**, 044065 (2021), 2106.05259 [astro-ph.HE].
- [124] A. Klein, G. Pratten, R. Busicchio, P. Schmidt, C. J. Moore, E. Finch, A. Bonino, L. M. Thomas, N. Williams, D. Gerosa, et al., *arXiv e-prints*, arXiv:2204.03423 (2022), 2204.03423 [astro-ph.HE].
- [125] J. Crowder and N. J. Cornish, *Phys. Rev. D* **72**, 083005 (2005), gr-qc/0506015 [gr-qc].
- [126] S. Kawamura, T. Nakamura, M. Ando, N. Seto, K. Tsubono, K. Numata, R. Takahashi, S. Nagano, T. Ishikawa, M. Musha, et al., *Classical and Quantum Gravity* **23**, S125 (2006).
- [127] S. Kawamura, M. Ando, N. Seto, S. Sato, T. Nakamura, K. Tsubono, N. Kanda, T. Tanaka, J. Yokoyama, I. Funaki, et al., *Classical and Quantum Gravity* **28**, 094011 (2011).
- [128] M. Arca Sedda, C. P. L. Berry, K. Jani, P. Amaro-Seoane, P. Auclair, J. Baird, T. Baker, E. Berti, K. Breivik, A. Burrows, et al., *Classical and Quantum Gravity* **37**, 215011 (2020), 1908.11375 [gr-qc].
- [129] P. R. Saulson, *American Journal of Physics* **65**, 501 (1997).
- [130] D. Reitze, P. Saulson, and H. Grote, *Advanced Interferometric Gravitational-Wave Detectors. Volume I: Essentials of Gravitational-Wave Detectors* (2019).
- [131] M. Tinto and S. V. Dhurandhar, *Living Reviews in Relativity* **8**, 4 (2005).
- [132] R. W. Hellings and G. S. Downs, *ApJ* **265**, L39 (1983).
- [133] R. S. Foster and D. C. Backer, *ApJ* **361**, 300 (1990).
- [134] G. Desvignes, R. N. Caballero, L. Lentati, J. P. W. Verbiest, D. J. Champion, B. W. Stappers, G. H. Janssen, P. Lazarus, S. Osłowski, S. Babak, et al., *MNRAS* **458**, 3341 (2016), 1602.08511 [astro-ph.HE].
- [135] B. C. Joshi, P. Arumugasamy, M. Bagchi, D. Bandyopadhyay, A. Basu, N. Dhanda Batra, S. Bethapudi, A. Choudhary, K. De, L. Dey, et al., *Journal of Astrophysics and Astronomy* **39**, 51 (2018).

- [136] S. Ransom, A. Brazier, S. Chatterjee, T. Cohen, J. M. Cordes, M. E. DeCesar, P. B. Demorest, J. S. Hazboun, M. T. Lam, R. S. Lynch, et al., in *Bulletin of the American Astronomical Society*, Vol. 51 (2019) p. 195, 1908.05356 [astro-ph.IM].
- [137] B. B. P. Perera, M. E. DeCesar, P. B. Demorest, M. Kerr, L. Lentati, D. J. Nice, S. Osłowski, S. M. Ransom, M. J. Keith, Z. Arzoumanian, et al., *MNRAS* **490**, 4666 (2019), 1909.04534 [astro-ph.HE].
- [138] M. Kerr, D. J. Reardon, G. Hobbs, R. M. Shannon, R. N. Manchester, S. Dai, C. J. Russell, S. Zhang, W. van Straten, S. Osłowski, et al., *PASA* **37**, e020 (2020), 2003.09780 [astro-ph.IM].
- [139] A. Sesana, A. Vecchio, and C. N. Colacino, *MNRAS* **390**, 192 (2008), 0804.4476 [astro-ph].
- [140] Z. Arzoumanian, P. T. Baker, H. Blumer, B. Bécsy, A. Brazier, P. R. Brook, S. Burke-Spolaor, S. Chatterjee, S. Chen, J. M. Cordes, et al., *ApJ* **905**, L34 (2020), 2009.04496 [astro-ph.HE].
- [141] Kagra Collaboration, T. Akutsu, M. Ando, K. Arai, Y. Arai, S. Araki, A. Araya, N. Aritomi, H. Asada, Y. Aso, et al., *Nature Astronomy* **3**, 35 (2019), 1811.08079 [gr-qc].
- [142] T. Akutsu, M. Ando, K. Arai, Y. Arai, S. Araki, A. Araya, N. Aritomi, Y. Aso, S. Bae, Y. Bae, et al., *Progress of Theoretical and Experimental Physics* **2021**, 05A101 (2021), 2005.05574 [physics.ins-det].
- [143] T. Akutsu, M. Ando, K. Arai, Y. Arai, S. Araki, A. Araya, N. Aritomi, H. Asada, Y. Aso, S. Bae, et al., *Progress of Theoretical and Experimental Physics* **2021**, 05A103 (2021), 2008.02921 [gr-qc].
- [144] R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. X. Adhikari, V. K. Adkins, V. B. Adya, C. Affeldt, D. Agarwal, et al., *Progress of Theoretical and Experimental Physics* **2022**, 063F01 (2022), 2203.01270 [gr-qc].
- [145] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos, et al., *Living Reviews in Relativity* **23**, 3 (2020).
- [146] V. Kalogera, B. S. Sathyaprakash, M. Bailes, M.-A. Bizouard, A. Buonanno, A. Burrows, M. Colpi, M. Evans, S. Fairhurst, S. Hild, et al., *arXiv e-prints*, arXiv:2111.06990 (2021), 2111.06990 [gr-qc].
- [147] M. Punturo, M. Abernathy, F. Acernese, B. Allen, N. Andersson, K. Arun, F. Barone, B. Barr, M. Barsuglia, M. Beker, et al., *Classical and Quantum Gravity* **27**, 194002 (2010).

-
- [148] M. Maggiore, C. Van Den Broeck, N. Bartolo, E. Belgacem, D. Bertacca, M. A. Bizouard, M. Branchesi, S. Clesse, S. Foffa, J. García-Bellido, et al., *J. Cosmology Astropart. Phys.* **2020**, 050 (2020), 1912.02622 [astro-ph.CO].
- [149] M. Branchesi, M. Maggiore, D. Alonso, C. Badger, B. Banerjee, F. Beirnaert, S. Bhagwat, G. Boileau, S. Borhanian, D. D. Brown, et al., *arXiv e-prints*, [arXiv:2303.15923](#) (2023), 2303.15923 [gr-qc].
- [150] D. Reitze, R. X. Adhikari, S. Ballmer, B. Barish, L. Barsotti, G. Billingsley, D. A. Brown, Y. Chen, D. Coyne, R. Eisenstein, et al., in *Bulletin of the American Astronomical Society*, Vol. 51 (2019) p. 35, 1907.04833 [astro-ph.IM].
- [151] M. Evans, R. X. Adhikari, C. Afle, S. W. Ballmer, S. Biscoveanu, S. Borhanian, D. A. Brown, Y. Chen, R. Eisenstein, A. Gruson, et al., *arXiv e-prints*, [arXiv:2109.09882](#) (2021), 2109.09882 [astro-ph.IM].
- [152] V. Baibhav, E. Berti, D. Gerosa, M. Mapelli, N. Giacobbo, Y. Bouffanais, and U. N. Di Carlo, *Phys. Rev. D* **100**, 064060 (2019), 1906.04197 [gr-qc].
- [153] C. M. Will, *Living Reviews in Relativity* **17**, 4 (2014), 1403.7377 [gr-qc].
- [154] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, 831 (1915).
- [155] A. Einstein, *Jahrbuch der Radioaktivität und Elektronik* **4**, 411 (1908).
- [156] A. Einstein, *Annalen der Physik* **340**, 898 (1911).
- [157] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, 1030 (1914).
- [158] F. W. Dyson, A. S. Eddington, and C. Davidson, *Philosophical Transactions of the Royal Society of London Series A* **220**, 291 (1920).
- [159] A. Einstein, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, 142 (1917).
- [160] A. Friedmann, *Zeitschrift für Physik* **10**, 377 (1922).
- [161] A. Friedmann, *Zeitschrift für Physik* **21**, 326 (1924).
- [162] G. Lemaître, *MNRAS* **91**, 483 (1931).
- [163] H. P. Robertson, *ApJ* **82**, 284 (1935).
- [164] H. P. Robertson, *ApJ* **83**, 187 (1936).

-
- [165] H. P. Robertson, *ApJ* **83**, 257 (1936).
- [166] A. G. Walker, *Proceedings of the London Mathematical Society* **42**, 90 (1937).
- [167] E. Hubble, *Proceedings of the National Academy of Science* **15**, 168 (1929).
- [168] G. Lemaître, *Nature* **127**, 706 (1931).
- [169] P. J. E. Peebles, *Principles of Physical Cosmology* (1993).
- [170] I. I. Shapiro, *Phys. Rev. Lett.* **13**, 789 (1964).
- [171] I. I. Shapiro, M. E. Ash, R. P. Ingalls, W. B. Smith, D. B. Campbell, R. B. Dyce, R. F. Jurgens, and G. H. Pettengill, *Phys. Rev. Lett.* **26**, 1132 (1971).
- [172] R. A. Hulse and J. H. Taylor, *ApJ* **195**, L51 (1975).
- [173] J. H. Taylor and J. M. Weisberg, *ApJ* **253**, 908 (1982).
- [174] J. G. Williams, X. X. Newhall, and J. O. Dickey, *Phys. Rev. D* **53**, 6730 (1996).
- [175] J. Müller, T. W. Murphy, U. Schreiber, P. J. Shelus, J.-M. Torre, J. G. Williams, D. H. Boggs, S. Bouquillon, A. Bourgoïn, and F. Hofmann, *Journal of Geodesy* **93**, 2195 (2019).
- [176] C. W. F. Everitt, D. B. Debra, B. W. Parkinson, J. P. Turneare, J. W. Conklin, M. I. Heifetz, G. M. Keiser, A. S. Silbergleit, T. Holmes, J. Kolodziejczak, et al., *Phys. Rev. Lett.* **106**, 221101 (2011), 1105.3456 [gr-qc].
- [177] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, et al., *ApJ* **875**, L1 (2019), 1906.11238 [astro-ph.GA].
- [178] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, et al., *ApJ* **875**, L2 (2019), 1906.11239 [astro-ph.IM].
- [179] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, and J. a. Barrett, *ApJ* **875**, L3 (2019), 1906.11240 [astro-ph.GA].
- [180] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, et al., *ApJ* **875**, L4 (2019), 1906.11241 [astro-ph.GA].

-
- [181] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, et al., *ApJ* **875**, L5 (2019), 1906.11242 [astro-ph.GA].
- [182] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, et al., *ApJ* **875**, L6 (2019), 1906.11243 [astro-ph.GA].
- [183] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, J. C. Algaba, R. Anantua, K. Asada, R. Azulay, U. Bach, A.-K. Baczko, et al., *ApJ* **930**, L12 (2022).
- [184] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, J. C. Algaba, R. Anantua, K. Asada, R. Azulay, U. Bach, A.-K. Baczko, et al., *ApJ* **930**, L13 (2022).
- [185] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, J. C. Algaba, R. Anantua, K. Asada, R. Azulay, U. Bach, A.-K. Baczko, et al., *ApJ* **930**, L14 (2022).
- [186] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, J. C. Algaba, R. Anantua, K. Asada, R. Azulay, U. Bach, A.-K. Baczko, et al., *ApJ* **930**, L15 (2022).
- [187] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, J. C. Algaba, R. Anantua, K. Asada, R. Azulay, U. Bach, A.-K. Baczko, et al., *ApJ* **930**, L16 (2022).
- [188] I. Debono and G. F. Smoot, *Universe* **2**, 23 (2016), 1609.09781 [gr-qc].
- [189] C. Rovelli, *Quantum Gravity* (2004).
- [190] C. Kiefer, *Quantum gravity* (2007).
- [191] D. Oriti, *Approaches to Quantum Gravity* (2009).
- [192] S. Mukhi, *Classical and Quantum Gravity* **28**, 153001 (2011), 1110.2569 [physics.pop-ph].
- [193] A. Ashtekar and E. Bianchi, *Reports on Progress in Physics* **84**, 042001 (2021), 2104.04394 [gr-qc].
- [194] A. Einstein, *Annalen der Physik* **322**, 891 (1905).
- [195] A. Einstein, *Annalen der Physik* **323**, 639 (1905).

-
- [196] A. Einstein, *Annalen der Physik* **328**, 371 (1907).
- [197] S. Weinberg, *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity* (1972).
- [198] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (1973).
- [199] S. W. Hawking and G. F. R. Ellis, *The large-scale structure of space-time*. (1973).
- [200] R. M. Wald, *General Relativity* (1984).
- [201] B. F. Schutz, *A First Course in General Relativity* (1985).
- [202] S. W. Hawking and W. Israel, *Three Hundred Years of Gravitation* (1989).
- [203] J. B. Hartle, *Gravity: an introduction to Einstein's general relativity* (2003).
- [204] S. M. Carroll, *Spacetime and geometry. An introduction to general relativity* (2004).
- [205] E. Poisson, *A relativist's toolkit: the mathematics of black-hole mechanics* (2004).
- [206] T. A. Moore, *A General Relativity Workbook* (2013).
- [207] B. O'Neill, *Semi-Riemannian Geometry With Applications to Relativity* (1983).
- [208] M. P. Do Carmo, *Riemannian Geometry* (1992).
- [209] J. M. Lee, *Introduction to Riemannian Manifolds* (2018).
- [210] A. Einstein, *The meaning of relativity* (1921).
- [211] A. Einstein, *Annalen der Physik* **343**, 355 (1912).
- [212] H. Minkowski, *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* **1908**, 53 (1908).
- [213] K. Schwarzschild, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* , 189 (1916).
- [214] J. Droste, *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences* **19**, 197 (1917).
- [215] M. D. Kruskal, *Physical Review* **119**, 1743 (1960).
- [216] J. T. Jebsen, *Arkiv for Matematik, Astronomi och Fysik* **15**, 18 (1921).
- [217] G. D. Birkhoff and R. E. Langer, *Relativity and modern physics* (1923).

-
- [218] N. V. Johansen and F. Ravndal, *General Relativity and Gravitation* **38**, 537 (2006).
- [219] A. S. Eddington, *Nature* **113**, 192 (1924).
- [220] G. Lemaître, *Annales de la Société Scientifique de Bruxelles* **53**, 51 (1933).
- [221] D. Finkelstein, *Physical Review* **110**, 965 (1958).
- [222] R. Penrose, *Phys. Rev. Lett.* **14**, 57 (1965).
- [223] R. P. Kerr, *Phys. Rev. Lett.* **11**, 237 (1963).
- [224] R. Penrose, *Nuovo Cimento Rivista Serie* **1**, 252 (1969).
- [225] H. Reissner, *Annalen der Physik* **355**, 106 (1916).
- [226] H. Weyl, *Annalen der Physik* **359**, 117 (1917).
- [227] G. Nordström, *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences* **20**, 1238 (1918).
- [228] G. B. Jeffery, *Proceedings of the Royal Society of London Series A* **99**, 123 (1921).
- [229] E. T. Newman and A. I. Janis, *Journal of Mathematical Physics* **6**, 915 (1965).
- [230] E. T. Newman, E. Couch, K. Chinnapared, A. Exton, A. Prakash, and R. Torrence, *Journal of Mathematical Physics* **6**, 918 (1965).
- [231] J. R. Oppenheimer and H. Snyder, *Physical Review* **56**, 455 (1939).
- [232] S. W. Hawking, *Proceedings of the Royal Society of London Series A* **294**, 511 (1966).
- [233] S. W. Hawking, *Proceedings of the Royal Society of London Series A* **295**, 490 (1966).
- [234] S. W. Hawking, *Proceedings of the Royal Society of London Series A* **300**, 187 (1967).
- [235] S. W. Hawking and R. Penrose, *Proceedings of the Royal Society of London Series A* **314**, 529 (1970).
- [236] R. D. Blandford, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (1989) Chap. 8, pp. 277–329.

- [237] A. Celotti, J. C. Miller, and D. W. Sciama, *Classical and Quantum Gravity* **16**, A3 (1999), [astro-ph/9912186](#) [astro-ph].
- [238] F. Hoyle and W. A. Fowler, *MNRAS* **125**, 169 (1963).
- [239] D. Lynden-Bell, *Nature* **223**, 690 (1969).
- [240] R. D. Blandford and R. L. Znajek, *MNRAS* **179**, 433 (1977).
- [241] M. J. Rees, *ARA&A* **22**, 471 (1984).
- [242] M. C. Begelman, R. D. Blandford, and M. J. Rees, *Reviews of Modern Physics* **56**, 255 (1984).
- [243] J. Kormendy and D. Richstone, *ARA&A* **33**, 581 (1995).
- [244] A. Eckart and R. Genzel, *Nature* **383**, 415 (1996).
- [245] A. Eckart and R. Genzel, *MNRAS* **284**, 576 (1997).
- [246] A. M. Ghez, B. L. Klein, M. Morris, and E. E. Becklin, *ApJ* **509**, 678 (1998), [astro-ph/9807210](#) [astro-ph].
- [247] B. L. Webster and P. Murdin, *Nature* **235**, 37 (1972).
- [248] C. T. Bolton, *Nature* **235**, 271 (1972).
- [249] M. Maggiore, *Gravitational Waves: Volume 1: Theory and Experiments* (Oxford University Press, 2007).
- [250] M. Maggiore, *Gravitational Waves: Volume 2: Astrophysics and Cosmology* (2018).
- [251] J. Creighton and W. Anderson, *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis*. (2011).
- [252] É. É. Flanagan and S. A. Hughes, *New Journal of Physics* **7**, 204 (2005), [gr-qc/0501041](#) [astro-ph].
- [253] A. Buonanno, *arXiv e-prints*, [arXiv:0709.4682](#) (2007), [0709.4682](#) [gr-qc].
- [254] P. C. Peters and J. Mathews, *Physical Review* **131**, 435 (1963).
- [255] P. C. Peters, *Physical Review* **136**, 1224 (1964).
- [256] M. G. Haehnelt, M. B. Davies, and M. J. Rees, *MNRAS* **366**, L22 (2006), [astro-ph/0511245](#) [astro-ph].

- [257] S. Komossa, H. Zhou, and H. Lu, *ApJ* **678**, L81 (2008), 0804.4585 [astro-ph].
- [258] D. Batcheldor, A. Robinson, D. J. Axon, E. S. Perlman, and D. Merritt, *ApJ* **717**, L6 (2010), 1005.2173 [astro-ph.CO].
- [259] M. Eracleous, T. A. Boroson, J. P. Halpern, and J. Liu, *ApJS* **201**, 23 (2012), 1106.2952 [astro-ph.CO].
- [260] J. M. Comerford and J. E. Greene, *ApJ* **789**, 112 (2014), 1405.6711 [astro-ph.GA].
- [261] M. Koss, L. Blecha, R. Mushotzky, C. L. Hung, S. Veilleux, B. Trakhtenbrot, K. Schawinski, D. Stern, N. Smith, Y. Li, et al., *MNRAS* **445**, 515 (2014), 1401.6798 [astro-ph.GA].
- [262] D. Lena, A. Robinson, A. Marconi, D. J. Axon, A. Capetti, D. Merritt, and D. Batcheldor, *ApJ* **795**, 146 (2014), 1409.3976 [astro-ph.GA].
- [263] R. S. Barrows, J. M. Comerford, J. E. Greene, and D. Pooley, *ApJ* **829**, 37 (2016), 1606.01253 [astro-ph.GA].
- [264] D. C. Kim, I. Yoon, G. C. Privon, A. S. Evans, D. Harvey, S. Stierwalt, and J. H. Kim, *ApJ* **840**, 71 (2017), 1704.05549 [astro-ph.GA].
- [265] D. W. Pesce, J. A. Braatz, J. J. Condon, and J. E. Greene, *ApJ* **863**, 149 (2018), 1807.04598 [astro-ph.GA].
- [266] C. J. Skipper and I. W. A. Browne, *MNRAS* **475**, 5179 (2018), 1801.03456 [astro-ph.GA].
- [267] J. D. Hogg, L. Blecha, C. S. Reynolds, K. L. Smith, and L. M. Winter, *MNRAS* **503**, 1688 (2021), 2103.00012 [astro-ph.GA].
- [268] D. W. Pesce, A. C. Seth, J. E. Greene, J. A. Braatz, J. J. Condon, B. R. Kent, and D. Krajnović, *ApJ* **909**, 141 (2021), 2101.07932 [astro-ph.GA].
- [269] M. Mapelli, in *Handbook of Gravitational Wave Astronomy* (2021) p. 16.
- [270] I. Mandel and F. S. Broekgaarden, *Living Reviews in Relativity* **25**, 1 (2022), 2107.14239 [astro-ph.HE].
- [271] I. Mandel and A. Farmer, *Phys. Rep.* **955**, 1 (2022), 1806.05820 [astro-ph.HE].
- [272] K. Belczynski, V. Kalogera, and T. Bulik, *ApJ* **572**, 407 (2002), astro-ph/0111452 [astro-ph].
- [273] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *ApJ* **759**, 52 (2012), 1202.4901 [astro-ph.HE].

-
- [274] N. Ivanova, S. Justham, X. Chen, O. De Marco, C. L. Fryer, E. Gaburov, H. Ge, E. Glebbeek, Z. Han, X. D. Li, et al., *A&A Rev.* **21**, 59 (2013), 1209.4302 [astro-ph.HE].
- [275] S. Stevenson, A. Vigna-Gómez, I. Mandel, J. W. Barrett, C. J. Neijssel, D. Perkins, and S. E. de Mink, *Nature Communications* **8**, 14906 (2017), 1704.01352 [astro-ph.HE].
- [276] F. Santoliquido, M. Mapelli, N. Giacobbo, Y. Bouffanais, and M. C. Artale, *MNRAS* **502**, 4877 (2021), 2009.03911 [astro-ph.HE].
- [277] P. Marchant, N. Langer, P. Podsiadlowski, T. M. Tauris, and T. J. Moriya, *A&A* **588**, A50 (2016), 1601.03718 [astro-ph.SR].
- [278] I. Mandel and S. E. de Mink, *MNRAS* **458**, 2634 (2016), 1601.00007 [astro-ph.HE].
- [279] S. E. de Mink and I. Mandel, *MNRAS* **460**, 3545 (2016), 1603.02291 [astro-ph.HE].
- [280] S. F. Portegies Zwart and S. L. W. McMillan, *ApJ* **528**, L17 (2000), astro-ph/9910061 [astro-ph].
- [281] S. Banerjee, H. Baumgardt, and P. Kroupa, *MNRAS* **402**, 371 (2010), 0910.3954 [astro-ph.SR].
- [282] M. Mapelli, Y. Bouffanais, F. Santoliquido, M. Arca Sedda, and M. C. Artale, *MNRAS* **511**, 5797 (2022), 2109.06222 [astro-ph.HE].
- [283] R. M. O’Leary, F. A. Rasio, J. M. Fregeau, N. Ivanova, and R. O’Shaughnessy, *ApJ* **637**, 937 (2006), astro-ph/0508224 [astro-ph].
- [284] J. M. B. Downing, M. J. Benacquista, M. Giersz, and R. Spurzem, *MNRAS* **407**, 1946 (2010), 0910.0546 [astro-ph.SR].
- [285] C. L. Rodriguez, M. Morscher, B. Pattabiraman, S. Chatterjee, C.-J. Haster, and F. A. Rasio, *Phys. Rev. Lett.* **115**, 051101 (2015), 1505.00792 [astro-ph.HE].
- [286] C. L. Rodriguez, S. Chatterjee, and F. A. Rasio, *Phys. Rev. D* **93**, 084029 (2016), 1602.02444 [astro-ph.HE].
- [287] A. Askar, M. Szkudlarek, D. Gondek-Rosińska, M. Giersz, and T. Bulik, *MNRAS* **464**, L36 (2017), 1608.02520 [astro-ph.HE].
- [288] G. Fragione and B. Kocsis, *Phys. Rev. Lett.* **121**, 161103 (2018), 1806.02351 [astro-ph.GA].

-
- [289] J. Samsing, *Phys. Rev. D* **97**, 103014 (2018), 1711.07452 [astro-ph.HE].
- [290] R. M. O’Leary, B. Kocsis, and A. Loeb, *MNRAS* **395**, 2127 (2009), 0807.2638 [astro-ph].
- [291] F. Antonini and H. B. Perets, *ApJ* **757**, 27 (2012), 1203.2938 [astro-ph.GA].
- [292] F. Antonini and F. A. Rasio, *ApJ* **831**, 187 (2016), 1606.04889 [astro-ph.HE].
- [293] C. Petrovich and F. Antonini, *ApJ* **846**, 146 (2017), 1705.05848 [astro-ph.HE].
- [294] A. Rasskazov and B. Kocsis, *ApJ* **881**, 20 (2019), 1902.03242 [astro-ph.HE].
- [295] B. M. Ziosi, M. Mapelli, M. Branchesi, and G. Tormen, *MNRAS* **441**, 3703 (2014), 1404.7147 [astro-ph.GA].
- [296] S. Banerjee, *MNRAS* **467**, 524 (2017), 1611.09357 [astro-ph.HE].
- [297] U. N. Di Carlo, N. Giacobbo, M. Mapelli, M. Pasquato, M. Spera, L. Wang, and F. Haardt, *MNRAS* **487**, 2947 (2019), 1901.00863 [astro-ph.HE].
- [298] G. Fragione and S. Banerjee, *ApJ* **913**, L29 (2021), 2103.10447 [astro-ph.HE].
- [299] Y. Kozai, *AJ* **67**, 591 (1962).
- [300] M. L. Lidov, *Planet. Space Sci.* **9**, 719 (1962).
- [301] T. O. Kimpson, M. Spera, M. Mapelli, and B. M. Ziosi, *MNRAS* **463**, 2443 (2016), 1608.05422 [astro-ph.GA].
- [302] F. Antonini, S. Toonen, and A. S. Hamers, *ApJ* **841**, 77 (2017), 1703.06614 [astro-ph.GA].
- [303] B.-M. Hoang, S. Naoz, B. Kocsis, F. A. Rasio, and F. Dosopoulou, *ApJ* **856**, 140 (2018), 1706.09896 [astro-ph.HE].
- [304] J. M. Bellovary, M.-M. Mac Low, B. McKernan, and K. E. S. Ford, *ApJ* **819**, L17 (2016), 1511.00005 [astro-ph.GA].
- [305] I. Bartos, B. Kocsis, Z. Haiman, and S. Márka, *ApJ* **835**, 165 (2017), 1602.03831 [astro-ph.HE].
- [306] A. Secunda, J. Bellovary, M.-M. Mac Low, K. E. S. Ford, B. McKernan, N. W. C. Leigh, W. Lyra, and Z. Sándor, *ApJ* **878**, 85 (2019), 1807.02859 [astro-ph.HE].
- [307] G. Fabj, S. S. Nasim, F. Caban, K. E. S. Ford, B. McKernan, and J. M. Bellovary, *MNRAS* **499**, 2608 (2020), 2006.11229 [astro-ph.GA].

- [308] N. C. Stone, B. D. Metzger, and Z. Haiman, *MNRAS* **464**, 946 (2017), 1602.04226 [astro-ph.GA].
- [309] N. W. C. Leigh, A. M. Geller, B. McKernan, K. E. S. Ford, M. M. Mac Low, J. Bellovary, Z. Haiman, W. Lyra, J. Samsing, M. O’Dowd, et al., *MNRAS* **474**, 5672 (2018), 1711.10494 [astro-ph.GA].
- [310] H. Tagawa, Z. Haiman, and B. Kocsis, *ApJ* **898**, 25 (2020), 1912.08218 [astro-ph.GA].
- [311] B. McKernan, K. E. S. Ford, J. Bellovary, N. W. C. Leigh, Z. Haiman, B. Kocsis, W. Lyra, M. M. Mac Low, B. Metzger, M. O’Dowd, et al., *ApJ* **866**, 66 (2018), 1702.07818 [astro-ph.HE].
- [312] Y. Yang, I. Bartos, Z. Haiman, B. Kocsis, S. Márka, and H. Tagawa, *ApJ* **896**, 138 (2020), 2003.08564 [astro-ph.HE].
- [313] M. Gröbner, W. Ishibashi, S. Tiwari, M. Haney, and P. Jetzer, *A&A* **638**, A119 (2020), 2005.03571 [astro-ph.GA].
- [314] S. Chandrasekhar, *ApJ* **97**, 255 (1943).
- [315] D. Merritt and M. Y. Poon, *ApJ* **606**, 788 (2004), astro-ph/0302296 [astro-ph].
- [316] L. Mayer, S. Kazantzidis, P. Madau, M. Colpi, T. Quinn, and J. Wadsley, *Science* **316**, 1874 (2007), 0706.1562 [astro-ph].
- [317] G. Lodato, S. Nayakshin, A. R. King, and J. E. Pringle, *MNRAS* **398**, 1392 (2009), 0906.0737 [astro-ph.CO].
- [318] M. Milosavljević and D. Merritt, in *The Astrophysics of Gravitational Wave Sources*, American Institute of Physics Conference Series, Vol. 686, edited by J. M. Centrella (2003) pp. 201–210, astro-ph/0212270 [astro-ph].
- [319] M. Colpi, *Space Sci. Rev.* **183**, 189 (2014), 1407.3102 [astro-ph.GA].
- [320] L. Barack, V. Cardoso, S. Nissanke, T. P. Sotiriou, A. Askar, C. Belczynski, G. Bertone, E. Bon, D. Blas, R. Brito, et al., *Classical and Quantum Gravity* **36**, 143001 (2019), 1806.05195 [gr-qc].
- [321] P. Amaro-Seoane, J. Andrews, M. Arca Sedda, A. Askar, Q. Baghi, R. Balasov, I. Bartos, S. S. Bavera, J. Bellovary, C. P. L. Berry, et al., *Living Reviews in Relativity* **26**, 2 (2023), 2203.06016 [gr-qc].

-
- [322] K. Belczynski, A. Heger, W. Gladysz, A. J. Ruiter, S. Woosley, G. Wiktorowicz, H. Y. Chen, T. Bulik, R. O’Shaughnessy, D. E. Holz, et al., *A&A* **594**, A97 (2016), 1607.03116 [astro-ph.HE].
- [323] S. E. Woosley, *ApJ* **836**, 244 (2017), 1608.08939 [astro-ph.HE].
- [324] S. E. Woosley, *ApJ* **878**, 49 (2019), 1901.00215 [astro-ph.SR].
- [325] R. Farmer, M. Renzo, S. E. de Mink, P. Marchant, and S. Justham, *ApJ* **887**, 53 (2019), 1910.12874 [astro-ph.SR].
- [326] R. Farmer, M. Renzo, S. E. de Mink, M. Fishbach, and S. Justham, *ApJ* **902**, L36 (2020), 2006.06678 [astro-ph.HE].
- [327] S. E. Woosley and A. Heger, *ApJ* **912**, L31 (2021), 2103.07933 [astro-ph.SR].
- [328] Y. Qin, T. Fragos, G. Meynet, J. Andrews, M. Sørensen, and H. F. Song, *A&A* **616**, A28 (2018), 1802.05738 [astro-ph.SR].
- [329] J. Fuller and L. Ma, *ApJ* **881**, L1 (2019), 1907.03714 [astro-ph.SR].
- [330] N. Steinle and M. Kesden, *Phys. Rev. D* **103**, 063032 (2021), arXiv:2010.00078 [astro-ph.HE].
- [331] V. Kalogera, *ApJ* **541**, 319 (2000), astro-ph/9911417 [astro-ph].
- [332] D. Gerosa and M. Fishbach, *Nature Astronomy* **5**, 749 (2021), 2105.03439 [astro-ph.HE].
- [333] M. Mapelli, M. Dall’Amico, Y. Bouffanais, N. Giacobbo, M. Arca Sedda, M. C. Artale, A. Ballone, U. N. Di Carlo, G. Iorio, F. Santoliquido, et al., *MNRAS* **505**, 339 (2021), 2103.05016 [astro-ph.HE].
- [334] D. Gerosa and E. Berti, *Phys. Rev. D* **95**, 124046 (2017), 1703.06223 [gr-qc].
- [335] M. Fishbach, D. E. Holz, and B. Farr, *ApJ* **840**, L24 (2017), 1703.06869 [astro-ph.HE].
- [336] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, *ApJ* **832**, L2 (2016), 1609.05916 [astro-ph.HE].
- [337] W. M. Farr, S. Stevenson, M. C. Miller, I. Mandel, B. Farr, and A. Vecchio, *Nature* **548**, 426 (2017), 1706.01385 [astro-ph.HE].
- [338] S. Stevenson, C. P. L. Berry, and I. Mandel, *MNRAS* **471**, 2801 (2017), 1703.06873 [astro-ph.HE].

-
- [339] F. Pretorius, *Phys. Rev. Lett.* **95**, 121101 (2005), gr-qc/0507014 [gr-qc].
- [340] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, *Phys. Rev. Lett.* **96**, 111101 (2006), gr-qc/0511048 [gr-qc].
- [341] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, *Phys. Rev. Lett.* **96**, 111102 (2006), gr-qc/0511103 [gr-qc].
- [342] M. D. Duez and Y. Zlochower, *Reports on Progress in Physics* **82**, 016902 (2019), 1808.06011 [gr-qc].
- [343] T. W. Baumgarte and S. L. Shapiro, *Numerical Relativity: Solving Einstein's Equations on the Computer* (2010).
- [344] L. Blanchet, *Living Reviews in Relativity* **9**, 4 (2006).
- [345] E. Poisson and C. M. Will, *Gravity* (2014).
- [346] É. Racine, *Phys. Rev. D* **78**, 044021 (2008), 0803.1820 [gr-qc].
- [347] T. Damour, *Phys. Rev. D* **64**, 124013 (2001), gr-qc/0103018 [gr-qc].
- [348] K. S. Thorne and J. B. Hartle, *Phys. Rev. D* **31**, 1815 (1985).
- [349] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [350] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995), gr-qc/9506022 [gr-qc].
- [351] J. D. Schnittman, *Phys. Rev. D* **70**, 124020 (2004), astro-ph/0409174 [astro-ph].
- [352] M. Kesden, D. Gerosa, R. O'Shaughnessy, E. Berti, and U. Sperhake, *Phys. Rev. Lett.* **114**, 081103 (2015), 1411.0674 [gr-qc].
- [353] D. Gerosa, M. Kesden, U. Sperhake, E. Berti, and R. O'Shaughnessy, *Phys. Rev. D* **92**, 064016 (2015), 1506.03492 [gr-qc].
- [354] A. Klein, *arXiv e-prints* , arXiv:2106.10291 (2021), 2106.10291 [gr-qc].
- [355] D. Gerosa, G. Fumagalli, M. Mould, G. Cavallotto, D. Padilla Monroy, D. Gargardt, and V. De Renzi, *arXiv e-prints* , arXiv:2304.04801 (2023), 2304.04801 [gr-qc].
- [356] K. Chatziioannou, A. Klein, N. Cornish, and N. Yunes, *Phys. Rev. Lett.* **118**, 051101 (2017), 1606.03117 [gr-qc].

-
- [357] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish, *Phys. Rev. D* **95**, 104004 (2017), 1703.03967 [gr-qc].
- [358] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions with formulas, graphs, and mathematical tables* (1965).
- [359] N. K. Johnson-McDaniel, S. Kulkarni, and A. Gupta, *Phys. Rev. D* **106**, 023001 (2022), arXiv:2107.11902 [astro-ph.HE].
- [360] J. V. Wall and C. R. Jenkins, *Practical Statistics for Astronomers*, Vol. 3 (2003).
- [361] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis* (2014).
- [362] E. Thrane and C. Talbot, *PASA* **36**, e010 (2019), 1809.02293 [astro-ph.IM].
- [363] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos, et al., *Classical and Quantum Gravity* **37**, 055002 (2020), 1908.11170 [gr-qc].
- [364] M. Bayes and M. Price, *Philosophical Transactions of the Royal Society of London Series I* **53**, 370 (1763).
- [365] M. Pitkin, S. Reid, S. Rowan, and J. Hough, *Living Reviews in Relativity* **14**, 5 (2011), arXiv:1102.3355 [astro-ph.IM].
- [366] C. Bond, D. Brown, A. Freise, and K. A. Strain, *Living Reviews in Relativity* **19**, 3 (2016).
- [367] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992), gr-qc/9209010 [gr-qc].
- [368] P. Jaranowski and A. Królak, *Living Reviews in Relativity* **15**, 4 (2012).
- [369] A. Buikema, C. Cahillane, G. L. Mansell, C. D. Blair, R. Abbott, C. Adams, R. X. Adhikari, A. Ananyeva, S. Appert, K. Arai, et al., *Phys. Rev. D* **102**, 062003 (2020), 2008.01301 [astro-ph.IM].
- [370] F. Acernese, M. Agathos, A. Ain, S. Albanesi, A. Allocca, A. Amato, T. Andrade, N. Andres, M. Andrés-Carcasona, T. Andrić, et al., *arXiv e-prints* , arXiv:2205.01555 (2022), 2205.01555 [gr-qc].
- [371] KAGRA Collaboration, H. Abe, R. X. Adhikari, T. Akutsu, M. Ando, A. Araya, N. Aritomi, H. Asada, Y. Aso, S. Bae, et al., *arXiv e-prints* , arXiv:2203.07011 (2022), 2203.07011 [astro-ph.IM].

- [372] D. A. Brown, *Searching for Gravitational Radiation from Binary Black Hole MA-CHOs in the Galactic Halo*, Ph.D. thesis, - (2007).
- [373] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999), [gr-qc/9808076](#) [gr-qc].
- [374] T. Dal Canton, A. H. Nitz, A. P. Lundgren, A. B. Nielsen, D. A. Brown, T. Dent, I. W. Harry, B. Krishnan, A. J. Miller, K. Wette, et al., *Phys. Rev. D* **90**, 082004 (2014), [1405.6731](#) [gr-qc].
- [375] S. Sachdev, S. Caudill, H. Fong, R. K. L. Lo, C. Messick, D. Mukherjee, R. Magee, L. Tsukada, K. Blackburn, P. Brady, et al., *arXiv e-prints* , [arXiv:1901.08580](#) (2019), [1901.08580](#) [gr-qc].
- [376] F. Aubin, F. Brighenti, R. Chierici, D. Estevez, G. Greco, G. M. Guidi, V. Juste, F. Marion, B. Mours, E. Nitoglia, et al., *Classical and Quantum Gravity* **38**, 095004 (2021), [2012.11512](#) [gr-qc].
- [377] S. Klimenko, G. Vedovato, M. Drago, F. Salemi, V. Tiwari, G. A. Prodi, C. Lazzaro, K. Ackley, S. Tiwari, C. F. Da Silva, et al., *Phys. Rev. D* **93**, 042004 (2016), [1511.05999](#) [gr-qc].
- [378] M. Drago, S. Klimenko, C. Lazzaro, E. Milotti, G. Mitselmakher, V. Necula, B. O'Brian, G. A. Prodi, F. Salemi, M. Szczepanczyk, et al., *SoftwareX* **14**, 100678 (2021), [2006.12604](#) [gr-qc].
- [379] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton, *Classical and Quantum Gravity* **39**, 245013 (2022), [2207.03429](#) [astro-ph.IM].
- [380] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R.olta, M. Halpern, R. S. Hill, N. Odegard, et al., *ApJS* **208**, 19 (2013), [1212.5226](#) [astro-ph.CO].
- [381] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, et al., *A&A* **594**, A13 (2016), [1502.01589](#) [astro-ph.CO].
- [382] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al., *A&A* **641**, A6 (2020), [1807.06209](#) [astro-ph.CO].
- [383] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, *Physical Review Research* **1**, 033015 (2019), [1905.09300](#) [gr-qc].

- [384] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, et al., *Phys. Rev. D* **103**, 104056 (2021), 2004.06503 [gr-qc].
- [385] S. Ossokine, A. Buonanno, S. Marsat, R. Cotesta, S. Babak, T. Dietrich, R. Haas, I. Hinder, H. P. Pfeiffer, M. Pürrer, et al., *Phys. Rev. D* **102**, 044055 (2020), 2004.09442 [gr-qc].
- [386] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *PASP* **125**, 306 (2013), 1202.3665 [astro-ph.IM].
- [387] J. S. Speagle, *MNRAS* **493**, 3132 (2020), 1904.02180 [astro-ph.IM].
- [388] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, et al., *ApJS* **241**, 27 (2019), 1811.02042 [astro-ph.IM].
- [389] J. Lange, R. O’Shaughnessy, and M. Rizzo, *arXiv e-prints*, arXiv:1805.10457 (2018), 1805.10457 [gr-qc].
- [390] J. Skilling, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, American Institute of Physics Conference Series, Vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint (2004) pp. 395–405.
- [391] J. Skilling, *Bayesian Analysis* **1**, 833 (2006).
- [392] J. Goodman and J. Weare, *Communications in Applied Mathematics and Computational Science* **5**, 65 (2010).
- [393] R. J. E. Smith, G. Ashton, A. Vajpeyi, and C. Talbot, *MNRAS* **498**, 4492 (2020), 1909.11873 [gr-qc].
- [394] M. J. Williams, J. Veitch, and C. Messenger, *Phys. Rev. D* **103**, 103006 (2021), 2102.11056 [gr-qc].
- [395] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **127**, 241103 (2021), arXiv:2106.12594 [gr-qc].
- [396] F. Iacovelli, M. Mancarella, S. Foffa, and M. Maggiore, *ApJS* **263**, 2 (2022), 2207.06910 [astro-ph.IM].
- [397] F. Iacovelli, M. Mancarella, S. Foffa, and M. Maggiore, *ApJ* **941**, 208 (2022), 2207.02771 [gr-qc].

-
- [398] K. W. K. Wong, M. Gabri e, and D. Foreman-Mackey, *arXiv e-prints* , arXiv:2211.06397 (2022), 2211.06397 [astro-ph.IM].
- [399] T. D. P. Edwards, K. W. K. Wong, K. K. H. Lam, A. Coogan, D. Foreman-Mackey, M. Isi, and A. Zimmerman, *arXiv e-prints* , arXiv:2302.05329 (2023), 2302.05329 [astro-ph.IM].
- [400] K. W. K. Wong, M. Isi, and T. D. P. Edwards, *arXiv e-prints* , arXiv:2302.05333 (2023), 2302.05333 [astro-ph.IM].
- [401] T. J. Loredo and D. Q. Lamb, *Phys. Rev. D* **65**, 063002 (2002), astro-ph/0107260 [astro-ph].
- [402] I. Mandel, W. M. Farr, and J. R. Gair, *MNRAS* **486**, 1086 (2019), 1809.02063 [physics.data-an].
- [403] S. Vitale, D. Gerosa, W. M. Farr, and S. R. Taylor, in *Handbook of Gravitational Wave Astronomy* (2022) p. 45.
- [404] S. R. Taylor and D. Gerosa, *Phys. Rev. D* **98**, 083017 (2018), 1806.08365 [astro-ph.HE].
- [405] J. Roulet, T. Venumadhav, B. Zackay, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **102**, 123022 (2020), 2008.07014 [astro-ph.HE].
- [406] S. M. Gaebel, J. Veitch, T. Dent, and W. M. Farr, *MNRAS* **484**, 4008 (2019), 1809.03815 [astro-ph.IM].
- [407] M. Fishbach, D. E. Holz, and W. M. Farr, *ApJ* **863**, L41 (2018), arXiv:1805.10270 [astro-ph.HE].
- [408] V. Tiwari, *Classical and Quantum Gravity* **35**, 145009 (2018), 1712.00482 [astro-ph.HE].
- [409] W. M. Farr, *Research Notes of the American Astronomical Society* **3**, 66 (2019), 1904.10879 [astro-ph.IM].
- [410] D. Gerosa, G. Pratten, and A. Vecchio, *Phys. Rev. D* **102**, 103020 (2020), 2007.06585 [astro-ph.HE].
- [411] K. W. K. Wong, K. K. Y. Ng, and E. Berti, *arXiv e-prints* , arXiv:2007.10350 (2020), 2007.10350 [astro-ph.HE].
- [412] C. Talbot and E. Thrane, *ApJ* **927**, 76 (2022).

-
- [413] M. Mould, D. Gerosa, and S. R. Taylor, *Phys. Rev. D* **106**, 103013 (2022), 2203.03651 [astro-ph.HE].
- [414] C. E. A. Chapman-Bird, C. P. L. Berry, and G. Woan, *arXiv e-prints* , arXiv:2212.06166 (2022), 2212.06166 [astro-ph.HE].
- [415] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *ApJS* **227**, 14 (2016), arXiv:1606.03939 [astro-ph.HE].
- [416] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993), gr-qc/9301003 [gr-qc].
- [417] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. E. Holz, T. Bulik, and F. Pannarale, *ApJ* **806**, 263 (2015), 1405.7016 [astro-ph.HE].
- [418] D. Gerosa, *dgerosa/gwdet* (2017).
- [419] M. Mould and D. Gerosa, *Phys. Rev. D* **101**, 124037 (2020), 2003.02281 [gr-qc].
- [420] M. Mould, *arXiv e-prints* , arXiv:2104.15011 (2021), 2104.15011 [gr-qc].
- [421] M. Mould and D. Gerosa, *Phys. Rev. D* **105**, 024076 (2022), 2110.05507 [astro-ph.HE].
- [422] M. Mould, D. Gerosa, F. S. Broekgaarden, and N. Steinle, *MNRAS* **517**, 2738 (2022), 2205.12329 [astro-ph.HE].
- [423] L. Reali, M. Mould, D. Gerosa, and V. Varma, *Classical and Quantum Gravity* **37**, 225005 (2020), 2005.01747 [gr-qc].
- [424] V. Varma, M. Mould, D. Gerosa, M. A. Scheel, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. D* **103**, 064003 (2021), 2012.07147 [gr-qc].
- [425] V. De Renzi, D. Gerosa, M. Mould, R. Busicchio, and L. Zanga, *arXiv e-prints* , arXiv:2304.13063 (2023), 2304.13063 [gr-qc].
- [426] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, *Phys. Rev. D* **103**, 064067 (2021), 2011.11948 [gr-qc].
- [427] V. De Renzi, D. Gerosa, G. Pratten, P. Schmidt, and M. Mould, *Phys. Rev. D* **106**, 084040 (2022), 2207.00030 [gr-qc].
- [428] V. Baibhav, D. Gerosa, E. Berti, K. W. K. Wong, T. Helfer, and M. Mould, *Phys. Rev. D* **102**, 043002 (2020), 2004.00650 [astro-ph.HE].

-
- [429] V. Baibhav, E. Berti, D. Gerosa, M. Mould, and K. W. K. Wong, *Phys. Rev. D* **104**, 084002 (2021), 2105.12140 [gr-qc].
- [430] D. Gerosa, M. Mould, D. Gangardt, and S. Oller, *dgerosa/precession: Precession* (2023).
- [431] M. Mould, *mdmould/qluster* (2023).
- [432] M. Mould, *mdmould/popodds* (2023).
- [433] M. Mould, *mdmould/shallow* (2023).
- [434] D. Gerosa, M. Kesden, R. O’Shaughnessy, A. Klein, E. Berti, U. Sperhake, and D. Trifirò, *Phys. Rev. Lett.* **115**, 141102 (2015), arXiv:1506.09116 [gr-qc].
- [435] D. Gerosa and M. Kesden, *Phys. Rev. D* **93**, 124066 (2016), arXiv:1605.01067 [astro-ph.HE].
- [436] C. O. Lousto and J. Healy, *Phys. Rev. D* **93**, 124074 (2016), arXiv:1601.05086 [gr-qc].
- [437] D. Gerosa, U. Sperhake, and J. Vošmera, *Classical and Quantum Gravity* **34**, 064004 (2017), arXiv:1612.05263 [gr-qc].
- [438] A. Buonanno, Y. Chen, and T. Damour, *Phys. Rev. D* **74**, 104005 (2006), arXiv:gr-qc/0508067 [gr-qc].
- [439] M. Campanelli, C. O. Lousto, H. Nakano, and Y. Zlochower, *Phys. Rev. D* **79**, 084010 (2009), arXiv:0808.0713 [gr-qc].
- [440] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009), arXiv:0907.0700 [gr-qc].
- [441] M. Kesden, U. Sperhake, and E. Berti, *Phys. Rev. D* **81**, 084054 (2010), arXiv:1002.2643 [astro-ph.GA].
- [442] D. Gerosa, M. Kesden, E. Berti, R. O’Shaughnessy, and U. Sperhake, *Phys. Rev. D* **87**, 104028 (2013), arXiv:1302.4442 [gr-qc].
- [443] M. Fishbach and D. E. Holz, *ApJ* **891**, L27 (2020), arXiv:1905.12669 [astro-ph.HE].
- [444] Y. Yang, I. Bartos, Z. Haiman, B. Kocsis, Z. Márka, N. C. Stone, and S. Márka, *ApJ* **876**, 122 (2019), arXiv:1903.01405 [astro-ph.HE].

-
- [445] Y. Yang, I. Bartos, V. Gayathri, K. E. S. Ford, Z. Haiman, S. Klimentko, B. Kocsis, S. Márka, Z. Márka, B. McKernan, et al., *Phys. Rev. Lett.* **123**, 181101 (2019), [arXiv:1906.09281 \[astro-ph.HE\]](#).
- [446] B. McKernan, K. E. S. Ford, R. O’Shaughnessy, and D. Wysocki, *MNRAS* **494**, 1203 (2020), [arXiv:1907.04356 \[astro-ph.HE\]](#).
- [447] A. Klein, E. Barausse, A. Sesana, A. Petiteau, E. Berti, S. Babak, J. Gair, S. Aoudia, I. Hinder, F. Ohme, et al., *Phys. Rev. D* **93**, 024003 (2016), [arXiv:1511.05581 \[gr-qc\]](#).
- [448] V. Varma, M. Isi, S. Biscoveanu, W. M. Farr, and S. Vitale, *Phys. Rev. D* **105**, 024045 (2022), [arXiv:2107.09692 \[astro-ph.HE\]](#).
- [449] K. K. Y. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, *Phys. Rev. D* **98**, 083007 (2018), [arXiv:1805.03046 \[gr-qc\]](#).
- [450] I. M. Romero-Shaw, C. Talbot, S. Biscoveanu, V. D’Emilio, G. Ashton, C. P. L. Berry, S. Coughlin, S. Galaudage, C. Hoy, M. Hübner, et al., *MNRAS* **499**, 3295 (2020), [arXiv:2006.00714 \[astro-ph.IM\]](#).
- [451] T. A. Callister, *arXiv e-prints*, [arXiv:2104.09508 \(2021\)](#), [arXiv:2104.09508 \[gr-qc\]](#).
- [452] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014), [arXiv:1308.3271 \[gr-qc\]](#).
- [453] C. Talbot, R. Smith, E. Thrane, and G. B. Poole, *Phys. Rev. D* **100**, 043030 (2019), [arXiv:1904.02863 \[astro-ph.IM\]](#).
- [454] T. Bogdanović, C. S. Reynolds, and M. C. Miller, *ApJ* **661**, L147 (2007), [arXiv:astro-ph/0703054 \[astro-ph\]](#).
- [455] J. Roulet, H. S. Chia, S. Olsen, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, *Phys. Rev. D* **104**, 083010 (2021), [arXiv:2105.10580 \[astro-ph.HE\]](#).
- [456] S. Galaudage, C. Talbot, T. Nagar, D. Jain, E. Thrane, and I. Mandel, *ApJ* **921**, L15 (2021), [arXiv:2109.02424 \[gr-qc\]](#).
- [457] S. Rinaldi and W. Del Pozzo, *MNRAS* **509**, 5454 (2022), [arXiv:2109.05960 \[astro-ph.IM\]](#).
- [458] I. Romero-Shaw, P. D. Lasky, E. Thrane, and J. Calderón Bustillo, *ApJ* **903**, L5 (2020), [arXiv:2009.04771 \[astro-ph.HE\]](#).

-
- [459] I. Romero-Shaw, P. D. Lasky, and E. Thrane, *ApJ* **921**, L31 (2021), [arXiv:2108.01284 \[astro-ph.HE\]](#).
- [460] I. Romero-Shaw, P. D. Lasky, and E. Thrane, *ApJ* **940**, 171 (2022), [arXiv:2206.14695 \[astro-ph.HE\]](#).
- [461] V. Gayathri, J. Healy, J. Lange, B. O'Brien, M. Szczepańczyk, I. Bartos, M. Campanelli, S. Klimentko, C. O. Lousto, and R. O'Shaughnessy, *Nature Astronomy* **6**, 344 (2022).
- [462] I. M. Romero-Shaw, D. Gerosa, and N. Loutrel, *MNRAS* **519**, 5352 (2023), [arXiv:2211.07528 \[astro-ph.HE\]](#).
- [463] M. Zevin, I. M. Romero-Shaw, K. Kremer, E. Thrane, and P. D. Lasky, *ApJ* **921**, L43 (2021), [arXiv:2106.09042 \[astro-ph.HE\]](#).
- [464] H. Yu, S. Ma, M. Giesler, and Y. Chen, *Phys. Rev. D* **102**, 123009 (2020), [arXiv:2007.12978 \[gr-qc\]](#).
- [465] D. Gerosa, E. Berti, R. O'Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, *Phys. Rev. D* **98**, 084036 (2018), [arXiv:1808.02491 \[astro-ph.HE\]](#).
- [466] N. Steinle and M. Kesden, *Phys. Rev. D* **106**, 063028 (2022), [arXiv:2206.00391 \[astro-ph.HE\]](#).
- [467] F. S. Broekgaarden, S. Stevenson, and E. Thrane, *ApJ* **938**, 45 (2022), [arXiv:2205.01693 \[astro-ph.HE\]](#).
- [468] H. C. Spruit, *A&A* **381**, 923 (2002), [arXiv:astro-ph/0108207 \[astro-ph\]](#).
- [469] J. Fuller, A. L. Piro, and A. S. Jermyn, *MNRAS* **485**, 3661 (2019), [arXiv:1902.08227 \[astro-ph.SR\]](#).
- [470] K. Belczynski, J. Klencki, C. E. Fields, A. Olejak, E. Berti, G. Meynet, C. L. Fryer, D. E. Holz, R. O'Shaughnessy, D. A. Brown, et al., *A&A* **636**, A104 (2020), [arXiv:1706.07053 \[astro-ph.HE\]](#).
- [471] D. Kushnir, M. Zaldarriaga, J. A. Kollmeier, and R. Waldman, *MNRAS* **462**, 844 (2016), [arXiv:1605.03839 \[astro-ph.HE\]](#).
- [472] K. Hotokezaka and T. Piran, *ApJ* **842**, 111 (2017), [arXiv:1702.03952 \[astro-ph.HE\]](#).
- [473] M. Zaldarriaga, D. Kushnir, and J. A. Kollmeier, *MNRAS* **473**, 4174 (2018), [arXiv:1702.00885 \[astro-ph.HE\]](#).

-
- [474] S. S. Bavera, T. Fragos, Y. Qin, E. Zapartas, C. J. Neijssel, I. Mandel, A. Batta, S. M. Gaebel, C. Kimball, and S. Stevenson, *A&A* **635**, A97 (2020), [arXiv:1906.12257 \[astro-ph.HE\]](#).
- [475] S. F. Portegies Zwart and F. Verbunt, *A&A* **309**, 179 (1996).
- [476] V. M. Kaspi, A. G. Lyne, R. N. Manchester, F. Crawford, F. Camilo, J. F. Bell, N. D’Amico, I. H. Stairs, N. P. F. McKay, D. J. Morris, et al., *ApJ* **543**, 321 (2000), [arXiv:astro-ph/0005214 \[astro-ph\]](#).
- [477] T. M. Tauris and T. Sennels, *A&A* **355**, 236 (2000), [arXiv:astro-ph/9909149 \[astro-ph\]](#).
- [478] M. S. Sipior, S. Portegies Zwart, and G. Nelemans, *MNRAS* **354**, L49 (2004), [arXiv:astro-ph/0407268 \[astro-ph\]](#).
- [479] S. Toonen, H. B. Perets, A. P. Igoshev, E. Michaely, and Y. Zenati, *A&A* **619**, A53 (2018), [arXiv:1804.01538 \[astro-ph.HE\]](#).
- [480] V. Venkatraman Krishnan, M. Bailes, W. van Straten, N. Wex, P. C. C. Freire, E. F. Keane, T. M. Tauris, P. A. Rosado, N. D. R. Bhat, C. Flynn, et al., *Science* **367**, 577 (2020), [arXiv:2001.11405 \[astro-ph.HE\]](#).
- [481] M. Zevin and S. S. Bavera, *ApJ* **933**, 86 (2022), [arXiv:2203.02515 \[astro-ph.HE\]](#).
- [482] C. Talbot and E. Thrane, *Phys. Rev. D* **96**, 023012 (2017), [arXiv:1704.08370 \[astro-ph.HE\]](#).
- [483] D. Wysocki, J. Lange, and R. O’Shaughnessy, *Phys. Rev. D* **100**, 043012 (2019), [arXiv:1805.06442 \[gr-qc\]](#).
- [484] L. S. Collaboration, V. C. S. Collaboration, V. Collaboration, and K. Collaboration, [10.5281/zenodo.5650061](#) (2023).
- [485] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration, [10.5281/zenodo.5636815](#) (2021).
- [486] R. Essick and W. Farr, [arXiv e-prints](#) , [arXiv:2204.00461](#) (2022), [arXiv:2204.00461 \[astro-ph.IM\]](#).
- [487] J. Golomb and C. Talbot, *ApJ* **926**, 79 (2022), [arXiv:2106.15745 \[astro-ph.HE\]](#).
- [488] C. Talbot and J. Golomb, [arXiv e-prints](#) , [arXiv:2304.06138](#) (2023), [arXiv:2304.06138 \[astro-ph.IM\]](#).

-
- [489] S. Vitale, S. Biscoveanu, and C. Talbot, *arXiv e-prints*, arXiv:2204.00968 (2022), arXiv:2204.00968 [gr-qc].
- [490] T. A. Callister, S. J. Miller, K. Chatziioannou, and W. M. Farr, *ApJ* **937**, L13 (2022), arXiv:2205.08574 [astro-ph.HE].
- [491] A. Olejak and K. Belczynski, *ApJ* **921**, L2 (2021), arXiv:2109.06872 [astro-ph.HE].
- [492] Y. Shao and X.-D. Li, *ApJ* **930**, 26 (2022), arXiv:2203.14529 [astro-ph.HE].
- [493] Z. Doctor, B. Farr, and D. E. Holz, *ApJ* **914**, L18 (2021), arXiv:2103.04001 [astro-ph.HE].
- [494] J. T. Gálvez Gherzi and L. C. Stein, *Classical and Quantum Gravity* **38**, 045012 (2021), arXiv:2007.11578 [gr-qc].
- [495] S. Biscoveanu, M. Isi, S. Vitale, and V. Varma, *Phys. Rev. Lett.* **126**, 171103 (2021), arXiv:2007.09156 [astro-ph.HE].
- [496] I. M. Romero-Shaw, E. Thrane, and P. D. Lasky, *PASA* **39**, e025 (2022), arXiv:2202.05479 [astro-ph.IM].
- [497] T. A. Callister, C.-J. Haster, K. K. Y. Ng, S. Vitale, and W. M. Farr, *ApJ* **922**, L5 (2021), arXiv:2106.00521 [astro-ph.HE].
- [498] M. Fishbach, Z. Doctor, T. Callister, B. Edelman, J. Ye, R. Essick, W. M. Farr, B. Farr, and D. E. Holz, *ApJ* **912**, 98 (2021), arXiv:2101.07699 [astro-ph.HE].
- [499] S. S. Bavera, M. Fishbach, M. Zevin, E. Zapartas, and T. Fragos, *A&A* **665**, A59 (2022), arXiv:2204.02619 [astro-ph.HE].
- [500] S. Biscoveanu, T. A. Callister, C.-J. Haster, K. K. Y. Ng, S. Vitale, and W. M. Farr, *ApJ* **932**, L19 (2022), arXiv:2204.01578 [astro-ph.HE].
- [501] K. W. K. Wong, K. Breivik, K. Kremer, and T. Callister, *Phys. Rev. D* **103**, 083021 (2021), arXiv:2011.03564 [astro-ph.HE].
- [502] M. Zevin, S. S. Bavera, C. P. L. Berry, V. Kalogera, T. Fragos, P. Marchant, C. L. Rodriguez, F. Antonini, D. E. Holz, and C. Pankow, *ApJ* **910**, 152 (2021), arXiv:2011.10057 [astro-ph.HE].
- [503] J. Stegmann and F. Antonini, *Phys. Rev. D* **103**, 063007 (2021), arXiv:2012.06329 [astro-ph.HE].

-
- [504] J. Kumamoto, M. S. Fujii, and A. Tanikawa, *MNRAS* **486**, 3942 (2019), [arXiv:1811.06726 \[astro-ph.HE\]](#).
- [505] U. N. Di Carlo, M. Mapelli, N. Giacobbo, M. Spera, Y. Bouffanais, S. Rastello, F. Santoliquido, M. Pasquato, A. Ballone, A. A. Trani, et al., *MNRAS* **498**, 495 (2020), [arXiv:2004.09525 \[astro-ph.HE\]](#).
- [506] A. A. Trani, A. Tanikawa, M. S. Fujii, N. W. C. Leigh, and J. Kumamoto, *MNRAS* **504**, 910 (2021), [arXiv:2102.01689 \[astro-ph.HE\]](#).
- [507] H. Tong, S. Galaudage, and E. Thrane, *Phys. Rev. D* **106**, 103019 (2022).
- [508] J. W. Barrett, S. M. Gaebel, C. J. Neijssel, A. Vigna-Gómez, S. Stevenson, C. P. L. Berry, W. M. Farr, and I. Mandel, *MNRAS* **477**, 4685 (2018), [arXiv:1711.06287 \[astro-ph.HE\]](#).
- [509] K. Belczynski, A. Romagnolo, A. Olejak, J. Klencki, D. Chattopadhyay, S. Stevenson, M. Coleman Miller, J. P. Lasota, and P. A. Crowther, *ApJ* **925**, 69 (2022), [arXiv:2108.10885 \[astro-ph.HE\]](#).
- [510] Y. Bouffanais, M. Mapelli, F. Santoliquido, N. Giacobbo, U. N. Di Carlo, S. Rastello, M. C. Artale, and G. Iorio, *MNRAS* **507**, 5224 (2021), [arXiv:2102.12495 \[astro-ph.HE\]](#).
- [511] G. Franciolini and P. Pani, *Phys. Rev. D* **105**, 123024 (2022), [arXiv:2201.13098 \[astro-ph.HE\]](#).
- [512] B. Edelman, Z. Doctor, J. Godfrey, and B. Farr, *ApJ* **924**, 101 (2022), [arXiv:2109.06137 \[astro-ph.HE\]](#).
- [513] V. Tiwari, *Classical and Quantum Gravity* **38**, 155007 (2021), [arXiv:2006.15047 \[astro-ph.HE\]](#).
- [514] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tiño, and J. Veitch, *MNRAS* **465**, 3254 (2017), [arXiv:1608.08223 \[astro-ph.HE\]](#).
- [515] J. Sadiq, T. Dent, and D. Wysocki, *Phys. Rev. D* **105**, 123014 (2022), [arXiv:2112.12659 \[gr-qc\]](#).
- [516] J. W. Barrett, I. Mandel, C. J. Neijssel, S. Stevenson, and A. Vigna-Gómez, in *Astroinformatics*, Vol. 325, edited by M. Brescia, S. G. Djorgovski, E. D. Feigelson, G. Longo, and S. Cavuoti (2017) pp. 46–50, [arXiv:1704.03781 \[astro-ph.HE\]](#).
- [517] K. W. K. Wong and D. Gerosa, *Phys. Rev. D* **100**, 083015 (2019), [arXiv:1909.06373 \[astro-ph.HE\]](#).

-
- [518] D. H. T. Cheung, K. W. K. Wong, O. A. Hannuksela, T. G. F. Li, and S. Ho, *Phys. Rev. D* **106**, 083014 (2022), arXiv:2112.06707 [astro-ph.IM].
- [519] K. W. K. Wong, G. Contardo, and S. Ho, *Phys. Rev. D* **101**, 123005 (2020), arXiv:2002.09491 [astro-ph.IM].
- [520] G. Papamakarios, T. Pavlakou, and I. Murray, *arXiv e-prints*, arXiv:1705.07057 (2017), arXiv:1705.07057 [stat.ML].
- [521] G. Papamakarios, E. Nalisnick, D. Jimenez Rezende, S. Mohamed, and B. Lakshminarayanan, *arXiv e-prints*, arXiv:1912.02762 (2019), arXiv:1912.02762 [stat.ML].
- [522] I. Kobyzev, S. J. D. Prince, and M. A. Brubaker, *arXiv e-prints*, arXiv:1908.09257 (2019), arXiv:1908.09257 [stat.ML].
- [523] K. W. K. Wong, G. Franciolini, V. De Luca, V. Baibhav, E. Berti, P. Pani, and A. Riotto, *Phys. Rev. D* **103**, 023026 (2021), arXiv:2011.01865 [gr-qc].
- [524] D. Gerosa, S. Vitale, and E. Berti, *Phys. Rev. Lett.* **125**, 101103 (2020), arXiv:2005.04243 [astro-ph.HE].
- [525] C. L. Rodriguez, K. Kremer, M. Y. Grudić, Z. Hafen, S. Chatterjee, G. Fragione, A. Lamberts, M. A. S. Martinez, F. A. Rasio, N. Weatherford, et al., *ApJ* **896**, L10 (2020), arXiv:2005.04239 [astro-ph.HE].
- [526] A. S. Hamers and M. Safarzadeh, *ApJ* **898**, 99 (2020), arXiv:2005.03045 [astro-ph.HE].
- [527] M. Safarzadeh and K. Hotokezaka, *ApJ* **897**, L7 (2020), arXiv:2005.06519 [astro-ph.HE].
- [528] C. Kimball, C. Talbot, C. P. L. Berry, M. Zevin, E. Thrane, V. Kalogera, R. Buscicchio, M. Carney, T. Dent, H. Middleton, et al., *ApJ* **915**, L35 (2021), arXiv:2011.05332 [astro-ph.HE].
- [529] G. Fragione, A. Loeb, and F. A. Rasio, *ApJ* **902**, L26 (2020), arXiv:2009.05065 [astro-ph.GA].
- [530] M. Arca-Sedda, F. P. Rizzuto, T. Naab, J. Ostriker, M. Giersz, and R. Spurzem, *ApJ* **920**, 128 (2021), arXiv:2105.07003 [astro-ph.GA].
- [531] J. Samsing, I. Bartos, D. J. D’Orazio, Z. Haiman, B. Kocsis, N. W. C. Leigh, B. Liu, M. E. Pessah, and H. Tagawa, *Nature* **603**, 237 (2022), arXiv:2010.09765 [astro-ph.HE].

-
- [532] M. Fishbach and D. E. Holz, *ApJ* **904**, L26 (2020), [arXiv:2009.05472 \[astro-ph.HE\]](#).
- [533] R. Essick, A. Farah, S. Galaudage, C. Talbot, M. Fishbach, E. Thrane, and D. E. Holz, *ApJ* **926**, 34 (2022), [arXiv:2109.00418 \[astro-ph.HE\]](#).
- [534] B. Liu and D. Lai, *MNRAS* **502**, 2049 (2021), [arXiv:2009.10068 \[astro-ph.HE\]](#).
- [535] H. Tagawa, B. Kocsis, Z. Haiman, I. Bartos, K. Omukai, and J. Samsing, *ApJ* **908**, 194 (2021), [arXiv:2012.00011 \[astro-ph.HE\]](#).
- [536] W. Lu, P. Beniamini, and C. Bonnerot, *MNRAS* **500**, 1817 (2021), [arXiv:2009.10082 \[astro-ph.HE\]](#).
- [537] J. A. González, M. Hannam, U. Sperhake, B. Brügmann, and S. Husa, *Phys. Rev. Lett.* **98**, 231101 (2007), [arXiv:gr-qc/0702052 \[gr-qc\]](#).
- [538] M. Campanelli, C. Lousto, Y. Zlochower, and D. Merritt, *ApJ* **659**, L5 (2007), [arXiv:gr-qc/0701164 \[gr-qc\]](#).
- [539] D. Gerosa, F. Hébert, and L. C. Stein, *Phys. Rev. D* **97**, 104049 (2018), [arXiv:1802.04276 \[gr-qc\]](#).
- [540] D. Gerosa and E. Berti, *Phys. Rev. D* **100**, 041301 (2019), [arXiv:1906.05295 \[astro-ph.HE\]](#).
- [541] D. Gerosa, N. Giacobbo, and A. Vecchio, *ApJ* **915**, 56 (2021), [arXiv:2104.11247 \[astro-ph.HE\]](#).
- [542] M. D. McKay, R. J. Beckman, and W. J. Conover, *Technometrics* **21**, 239 (1979).
- [543] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, *ApJ* **591**, 288 (2003), [arXiv:astro-ph/0212469 \[astro-ph\]](#).
- [544] S. E. Woosley, S. Blinnikov, and A. Heger, *Nature* **450**, 390 (2007), [arXiv:0710.3314 \[astro-ph\]](#).
- [545] P. Marchant and T. J. Moriya, *A&A* **640**, L18 (2020), [arXiv:2007.06220 \[astro-ph.HE\]](#).
- [546] J. R. Rice and B. Zhang, *ApJ* **908**, 59 (2021), [arXiv:2009.11326 \[astro-ph.HE\]](#).
- [547] M. Safarzadeh and Z. Haiman, *ApJ* **903**, L21 (2020), [arXiv:2009.09320 \[astro-ph.HE\]](#).
- [548] Z. Roupas and D. Kazanas, *A&A* **632**, L8 (2019), [arXiv:1911.03915 \[astro-ph.GA\]](#).

- [549] P. Natarajan, *MNRAS* **501**, 1413 (2021), [arXiv:2009.09156 \[astro-ph.GA\]](#).
- [550] L. A. C. van Son, S. E. De Mink, F. S. Broekgaarden, M. Renzo, S. Justham, E. Laplace, J. Morán-Fraile, D. D. Hendriks, and R. Farmer, *ApJ* **897**, 100 (2020), [arXiv:2004.05187 \[astro-ph.HE\]](#).
- [551] K. Belczynski, R. Hirschi, E. A. Kaiser, J. Liu, J. Casares, Y. Lu, R. O’Shaughnessy, A. Heger, S. Justham, and R. Soria, *ApJ* **890**, 113 (2020), [arXiv:1911.12357 \[astro-ph.HE\]](#).
- [552] J. S. Vink, E. R. Higgins, A. A. C. Sander, and G. N. Sabhahit, *MNRAS* **504**, 146 (2021), [arXiv:2010.11730 \[astro-ph.HE\]](#).
- [553] A. Tanikawa, H. Susa, T. Yoshida, A. A. Trani, and T. Kinugawa, *ApJ* **910**, 30 (2021), [arXiv:2008.01890 \[astro-ph.HE\]](#).
- [554] E. Farrell, J. H. Groh, R. Hirschi, L. Murphy, E. Kaiser, S. Ekström, C. Georgy, and G. Meynet, *MNRAS* **502**, L40 (2021), [arXiv:2009.06585 \[astro-ph.SR\]](#).
- [555] T. Kinugawa, T. Nakamura, and H. Nakano, *MNRAS* **501**, L49 (2021), [arXiv:2009.06922 \[astro-ph.HE\]](#).
- [556] G. Costa, A. Bressan, M. Mapelli, P. Marigo, G. Iorio, and M. Spera, *MNRAS* **501**, 4514 (2021), [arXiv:2010.02242 \[astro-ph.SR\]](#).
- [557] D. Merritt, M. Milosavljević, M. Favata, S. A. Hughes, and D. E. Holz, *ApJ* **607**, L9 (2004), [arXiv:astro-ph/0402057 \[astro-ph\]](#).
- [558] O. Y. Gnedin, H. Zhao, J. E. Pringle, S. M. Fall, M. Livio, and G. Meylan, *ApJ* **568**, L23 (2002), [arXiv:astro-ph/0202045 \[astro-ph\]](#).
- [559] E. Barausse, V. Morozova, and L. Rezzolla, *ApJ* **758**, 63 (2012), [arXiv:1206.3803 \[gr-qc\]](#).
- [560] E. Barausse and L. Rezzolla, *ApJ* **704**, L40 (2009), [arXiv:0904.2577 \[gr-qc\]](#).
- [561] F. Hofmann, E. Barausse, and L. Rezzolla, *ApJ* **825**, L19 (2016), [arXiv:1605.01938 \[gr-qc\]](#).
- [562] C. O. Lousto and Y. Zlochower, *Phys. Rev. D* **77**, 044028 (2008), [arXiv:0708.4048 \[gr-qc\]](#).
- [563] C. O. Lousto, Y. Zlochower, M. Dotti, and M. Volonteri, *Phys. Rev. D* **85**, 084015 (2012), [arXiv:1201.1923 \[gr-qc\]](#).

-
- [564] C. O. Lousto and Y. Zlochower, *Phys. Rev. D* **87**, 084027 (2013), arXiv:1211.7099 [gr-qc].
- [565] J. D. Bekenstein, *ApJ* **183**, 657 (1973).
- [566] M. J. Fitchett, *MNRAS* **203**, 1049 (1983).
- [567] M. Zevin and D. E. Holz, *ApJ* **935**, L20 (2022), arXiv:2205.08549 [astro-ph.HE].
- [568] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al., *Nature Methods* **17**, 261 (2020), arXiv:1907.10121 [cs.MS].
- [569] B. W. Silverman, *Density estimation for statistics and data analysis* (1986).
- [570] M. Mould, `mdmould/kaydee` (2023).
- [571] T. Odland, `tommyod/kdepy: Kernel density estimation in python` (2018).
- [572] S. J. Sheather and M. C. Jones, *Journal of the Royal Statistical Society. Series B (Methodological)* **53**, 683 (1991).
- [573] Z. I. Botev, J. F. Grotowski, and D. P. Kroese, *The Annals of Statistics* **38**, 2916 (2010).
- [574] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al., *TensorFlow: Large-scale machine learning on heterogeneous systems* (2015), software available from tensorflow.org.
- [575] T. Developers, `Tensorflow` (2023).
- [576] B. Xu, N. Wang, T. Chen, and M. Li, arXiv e-prints , arXiv:1505.00853 (2015), arXiv:1505.00853 [cs.LG].
- [577] D. P. Kingma and J. Ba, arXiv e-prints , arXiv:1412.6980 (2014), arXiv:1412.6980 [cs.LG].
- [578] E. Hellinger, *Journal für die reine und angewandte Mathematik* **1909**, 210 (1909).
- [579] C. J. Moore and D. Gerosa, *Phys. Rev. D* **104**, 083008 (2021), arXiv:2108.02462 [gr-qc].
- [580] L. S. Finn, *Phys. Rev. D* **53**, 2878 (1996), arXiv:gr-qc/9601048 [gr-qc].
- [581] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, *PASP* **131**, 024503 (2019), arXiv:1807.10312 [astro-ph.IM].

-
- [582] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, C. Capano, T. Dent, L. Pekowsky, S. De, et al., [gwastro/pycbc: v2.1.1 release of pycbc \(2023\)](#).
- [583] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016), [arXiv:1508.07250 \[gr-qc\]](#).
- [584] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016), [arXiv:1508.07253 \[gr-qc\]](#).
- [585] A. Sesana, J. Gair, E. Berti, and M. Volonteri, *Phys. Rev. D* **83**, 044036 (2011), [arXiv:1011.5893 \[astro-ph.CO\]](#).
- [586] Y. Bouffanais, M. Mapelli, D. Gerosa, U. N. Di Carlo, N. Giacobbo, E. Berti, and V. Baibhav, *ApJ* **886**, 25 (2019), [arXiv:1905.11054 \[astro-ph.HE\]](#).
- [587] A. Toubiana, K. W. K. Wong, S. Babak, E. Barausse, E. Berti, J. R. Gair, S. Marsat, and S. R. Taylor, *Phys. Rev. D* **104**, 083027 (2021), [arXiv:2106.13819 \[gr-qc\]](#).
- [588] P. Kroupa, *MNRAS* **322**, 231 (2001), [arXiv:astro-ph/0009005 \[astro-ph\]](#).
- [589] K. Belczynski, *ApJ* **905**, L15 (2020), [arXiv:2009.13526 \[astro-ph.HE\]](#).
- [590] V. Tiwari and S. Fairhurst, *ApJ* **913**, L19 (2021), [arXiv:2011.04502 \[astro-ph.HE\]](#).
- [591] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015), [arXiv:1408.1810 \[gr-qc\]](#).
- [592] C. Henshaw, R. O’Shaughnessy, and L. Cadonati, *Classical and Quantum Gravity* **39**, 125003 (2022), [arXiv:2201.05220 \[gr-qc\]](#).
- [593] Z. Doctor, D. Wysocki, R. O’Shaughnessy, D. E. Holz, and B. Farr, *ApJ* **893**, 35 (2020), [arXiv:1911.04424 \[astro-ph.HE\]](#).
- [594] C. Kimball, C. Talbot, C. P. L. Berry, M. Carney, M. Zevin, E. Thrane, and V. Kalogera, *ApJ* **900**, 177 (2020), [arXiv:2005.00023 \[astro-ph.HE\]](#).
- [595] M. Fishbach, C. Kimball, and V. Kalogera, *ApJ* **935**, L26 (2022), [arXiv:2207.02924 \[astro-ph.HE\]](#).
- [596] A. van den Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. van den Driessche, E. Lockhart, L. C. Cobo, F. Stimberg, et al., [arXiv e-prints](#), [arXiv:1711.10433 \(2017\)](#), [arXiv:1711.10433 \[cs.LG\]](#).

-
- [597] L. Valentin Jospin, W. Buntine, F. Boussaid, H. Laga, and M. Bennamoun, [arXiv e-prints](#) , [arXiv:2007.06823](#) (2020), [arXiv:2007.06823](#) [cs.LG].
- [598] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, [arXiv e-prints](#) , [arXiv:1906.04032](#) (2019), [arXiv:1906.04032](#) [stat.ML].
- [599] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville, [arXiv e-prints](#) , [arXiv:1804.00779](#) (2018), [arXiv:1804.00779](#) [cs.LG].
- [600] N. De Cao, I. Titov, and W. Aziz, [arXiv e-prints](#) , [arXiv:1904.04676](#) (2019), [arXiv:1904.04676](#) [stat.ML].
- [601] D. Ruhe, K. Wong, M. Cranmer, and P. Forré, [arXiv e-prints](#) , [arXiv:2211.09008](#) (2022), [arXiv:2211.09008](#) [astro-ph.IM].
- [602] K. W. K. Wong and M. Cranmer, [arXiv e-prints](#) , [arXiv:2207.12409](#) (2022), [arXiv:2207.12409](#) [astro-ph.IM].
- [603] B. Edelman, B. Farr, and Z. Doctor, *ApJ* **946**, 16 (2023), [arXiv:2210.12834](#) [astro-ph.HE].
- [604] J. Golomb and C. Talbot, [arXiv e-prints](#) , [arXiv:2210.12287](#) (2022), [arXiv:2210.12287](#) [astro-ph.HE].