Searching for the Astrophysical Gravitational-Wave Background and Prompt Radio Emission from Compact Binaries

Thesis by

Thomas A. Callister, III

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2020
Defended August 19, 2019

Thomas A. Callister, III
ORCID: 0000-0001-9892-177X

Thank you to everyone who made this work possible.

Thank you to my fellow LIGO graduate students, Craig, Max, and Surabhi, for their comradery and constant support.

Thank you to all of my mentors, Joel, Jon, Eric, Jonah, Gregg, and Alan, for the time and energy you invested in me. I would not have arrived here without your friendship, confidence, and support.

Thank you to Jared, Elvira, Sylvia, Serena, and Simona, who taught me how to be a mentor myself.

Thank you to the faculty, staff, and all of my classmates at Carleton College. You sparked my interest in physics and astronomy, and taught me the importance of creativity, curiosity, and morality in science.

Thank you to my family: Karen, Helo, Wilson, Talbot, Mom, Dad, Bunner, and Daddygum. I love you.

Abstract

Gravitational-wave astronomy is now a reality. During my time at Caltech, the Advanced LIGO and Virgo observatories have detected gravitational waves from dozens of compact binary coalescences. All of these gravitational-wave events occurred in the relatively local Universe. In the first part of this thesis, I will instead look towards the *remote* Universe, investigating what LIGO and Virgo may be able to learn about cosmologically-distant compact binaries via observation of the stochastic gravitational-wave background. The stochastic gravitational-wave background is composed of the incoherent superposition of all distant, individually-unresolvable gravitational-wave sources. I explore what we learn from study of the gravitational-wave background, both about the astrophysics of compact binaries and the fundamental nature of gravitational waves. Of course, before we can study the gravitational-wave background we must first detect it. I therefore present searches for the gravitational-wave background using data from Advanced LIGO's first two observing runs, obtaining the most stringent upper limits to date on strength of the stochastic background. Finally, I consider how one might validate an apparent detection of the gravitational-wave background, confidently distinguishing a true astrophysical signal from spurious terrestrial artifacts.

The second part of this thesis concerns the search for electromagnetic counterparts to gravitational-wave events. The binary neutron star merger GW170817 was accompanied by a rich set of electromagnetic counterparts spanning nearly the entire electromagnetic spectrum. Beyond these counterparts, compact binaries may additionally generate powerful radio transients at or near their time of merger. First, I consider whether there is a plausible connection between this so-called "prompt radio emission" and fast radio bursts – enigmatic radio transients of unknown origin. Next, I present the first direct search for prompt radio emission from a compact binary merger using the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA). While no plausible candidates are identified, this effort successfully demonstrates the prompt radio follow-up of a gravitational-wave source, providing a blueprint for LIGO and Virgo follow-up in their O3 observing run and beyond.

Published Content and Contributions

T. A. Callister, M. M. Anderson, G. Hallinan, *et al.*, *A First Search for Prompt Radio Emission from a Gravitational-Wave Event*, Astrophys. J. Letters **877**, L39 (2019). T.C. produced all results and led preparation of the manuscript. doi:10.3847/2041-8213/ab2248.

LIGO Scientific Collaboration & Virgo Collaboration, *A Search for the Isotropic Stochastic Background Using Data from Advanced LIGO's Second Observing Run*, Phys. Rev. D **100**, 061101 (2019). T.C. contributed results and co-wrote the published manuscript. doi:10.1103/PhysRevD.100.061101.

T. A. Callister, M. W. Coughlin, J. B. Kanner, *Gravitational-wave geodesy: A new tool for validating detection of the stochastic gravitational-wave background*, Astrophys. J. Letters **869**, L28 (2018). T.C. produced all results and led preparation of the manuscript. doi:10.3847/2041-8213/aaf3a5.

LIGO Scientific Collaboration & Virgo Collaboration, *Search for Tensor, Vector, and Scalar Polarizations in the Stochastic Gravitational-Wave Background*, Phys. Rev. Letters **120**, 201102 (2018). T.C. led this study and wrote the published manuscript. doi:10.1103/PhysRevLett.120.201102.

P. B. Covas *et al.*, *Identification and Mitigation of Narrow Spectral Artifacts that Degrade Searches for Persistent Gravitational Waves in the First Two Observing Runs of Advanced LIGO*, Phys. Rev. D **97**, 082002 (2018). T.C. developed stochastic search data quality tools presented here and aided in the manuscript preparation. doi:10.1103/PhysRevD.97.082002.

T. Callister, S. A. Biscoveanu, N. Christensen *et al.*, *Polarization-Based Tests of Gravity with the Stochastic Gravitational-Wave Background*, Phys. Rev. X **7**, 041058 (2017). T.C. led this study and wrote the published manuscript. doi:10.1103/PhysRevX.7.041058.

T. A. Callister, J. B. Kanner, T. J. Massinger, *et al.*, *Observing Gravitational Waves with a Single Detector*, Class. Quantum Grav. **34**, 155007 (2017). T.C. led this study and wrote the published manuscript. doi:10.1088/1361-6382/aa7a76.

T. Callister, L. Sammut, S. Qiu *et al.*, *Limits of Astrophysics with Gravitational-Wave Backgrounds*, Phys. Rev. X **6**, 031018 (2016). T.C. produced all results and led preparation of the manuscript. doi:10.1103/PhysRevX.6.031018.

T. Callister, J. Kanner, and A. Weinstein, *Gravitational-Wave Constraints on the Progenitors of Fast Radio Bursts*, Astrophys. J. Letters **825**, L12 (2016). T.C. produced all results and led preparation of the manuscript. doi:10.3847/2041-8205/825/1/L12.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1
Gravitational Waves: Then and Now

Gravitational-wave astronomy is not a new idea. As early as 1960, only forty-four years after Einstein's prediction of gravitational waves, Joseph Weber was already constructing resonant "Weber bar" detectors in an effort to experimentally observe these infinitesimal ripples in spacetime [1]. At the time, consensus was that gravitational waves were little more than mathematical curiosities. It was only in the 1950s that physicists satisfied themselves that gravitational waves were even *real* (rather than spurious mathematical artifacts of General Relativity, as Einstein himself once espoused). And the physical effects of gravitational waves – their microscopic stretching and squeezing of space and time – were deemed so insignificant as to make their direct observation impossible.

Weber rejected this consensus. He realized that extreme astrophysical objects like the (then) newly discovered pulsars [2] might produce gravitational waves strong enough to be detected on Earth [3]. He even sent a device aboard Apollo 17, the final human expedition to the Moon, in an attempt to detect low-frequency gravitational radiation via the slight geological vibrations they might excite within the Moon [4]. Weber believed his efforts successful, announcing in 1969 an apparent detection of gravitational waves [5, 6]. Although these claims were later discredited, Weber *had* succeeded in an arguably more important task – articulating a new kind of **gravitational-wave astronomy**, in which we rely not on light but on gravitational waves to examine the Universe around us.

In 1972, while Weber was busy defending his claims of gravitational-wave detection, Rainer Weiss described in an underwhelmingly-titled yet now-famous *Quarterly Progress Report No. 105* an alternative kind of gravitational-wave detector, one that relied on laser measurements of the precise distance between distant mirrors rather than the vibrations of Weber bars [7]. This idea would form the basis for the LIGO (Laser Interferometer Gravitational-Wave Observatory) project. Comprising three detectors – two in Hanford, Washington and one in Livingston, Louisiana – the LIGO experiment was formally launched in 1984 under the joint leadership of Weiss and Caltech's Kip Thorne and Ronald Drever. Despite a rocky and somewhat scandal-clad beginning (see, for instance, Caltech's collection of oral histories from LIGO's early leaders [8–12]) the project successfully secured funding from the National Science Foundation, and the instruments began their first scientific observing run August 2002 [13, 14].

This first scientific observing run lasted for two weeks. And in these two weeks, LIGO detected precisely nothing [15–19]. Initial LIGO (as the experiment's first incarnation is now known) observations continued in a series of additional observing runs spanning the 2000s. Despite orders of magnitude improvement in LIGO's sensitivity [20], as well as the addition of the Virgo gravitational-wave detector in Italy [21], the Initial LIGO project continued to hear only a gravitational-wave "silence." This was perhaps not unexpected. LIGO's initial proposal to the NSF very carefully promises only "to build and operate" the observatory while continuing relevant research and development. It is fairly directly acknowledged that the successful *detection* of gravitational waves may only be made by a much-improved future "Advanced" LIGO detector[1] [13].

From Einstein's first prediction of gravitational waves through the experiments of Weber to the construction and operation of Initial LIGO, gravitational-wave astronomy is a surprisingly old idea. It was not until four years ago, however, that it became a reality. In 2015, during the second year of my graduate studies, the search for gravitational waves was just concluding a five year hiatus. Final touches were being placed on the LIGO detectors, which had since 2010 been undergoing upgrades to vastly improve their sensitivity [22–24]. In fall 2015 these upgraded Advanced LIGO detectors were first switched on for their first "O1" observing run. Then, only days into their operation, the Advanced LIGO detectors recorded the first unambiguous gravitational-wave signal. Termed GW150914, this signal was identified with violent collision of two black holes, roughly 1.3 billion years ago in a galaxy 400 Mpc away [24–33]. This direct detection of gravitational waves would win Weiss, Thorne, and Initial LIGO director Barry Barish the 2017 Nobel Prize in Physics. It is interesting to note that more time elapsed between LIGO's successful detection and Weber's first attempts than between Weber and Einstein himself.

[1]The projections listed on p. 10 of the proposal remain surprisingly accurate today, roughly 30 years later.

In the few years since, gravitational-wave astronomy has transformed radically. The number of operational gravitational-wave detectors grew from two to three, with the first operation of the European Advanced Virgo detector [34] alongside Advanced LIGO during the 2016-2017 O2 observing run. The Japanese KAGRA detector [35–37] may be brought online in the coming months during LIGO and Virgo's current O3 observing run [37, 38], and construction is beginning on the LIGO-India instrument, adding yet another instrument to the world-wide network of gravitational-wave detectors [38, 39].

Accompanying this growing number of detectors has been an ever-increasing number of events. Since GW150914, gravitational waves have been detected from nine additional binary black hole mergers – two more in O1 [40, 41] and seven in O2 [42–45]. O2 also marked the first observation of a second class of gravitational-wave source. In August 2017, LIGO and Virgo witnessed the collision of two neutron stars via the gravitational-wave signal GW170817 [46–52]. Unlike black holes, which are ultimately nothing but empty space, neutron stars are composed of matter. And when matter is present, we can usually expect light. Indeed, GW170817 was observed not just in gravitational waves, but also with telescopes spanning the electromagnetic spectrum [53, 54]. In perhaps the largest coordinated scientific effort in history, the publication describing the "multi-messenger" gravitational and electromagnetic observations of GW170817 listed 3529 authors and 61 scientific collaborations [54].

In these short four years, we have learned an immense amount of information about the gravitational universe. We have learned that stellar-mass black holes collide far more frequently than once surmised, and we have begun to study their characteristic masses and spins. We have learned that at least some fraction of neutron star mergers are accompanied by electromagnetic emission, and have used the gravitational and electromagnetic signals from GW170817 to jointly probe the nature of the extraordinarily dense matter hidden within neutron star interiors. And yet, as an observational field in its infancy, gravitational-wave astronomy is presently dominated by innumerably many more questions than answers.

In this thesis, I will work towards answering two of these many questions:

1. What can we learn from the stochastic gravitational-wave background?

Binary neutron star and black hole mergers are but one piece of the gravitational-wave sky. Beyond compact binaries, the Advanced LIGO and Virgo experiments are used to search for a host of other gravitational-wave sources, including the **stochastic gravitational-wave background**. For every loud gravitational-wave source in the relatively local Universe, there are innumerably many more that are too distant and/or too weak to individually detect. The stochastic gravitational-wave background is the random gravitational-wave "static" that arises from the combination of all of these quiet sources in the Universe.

With present-day detectors, we are largely limited to directly observing compact binaries in the relatively local Universe. As we will explore in Ch. 4, the stochastic gravitational-wave background offers a glimpse of compact binaries at truly cosmological distances, far earlier in the Universe's history. We have good reason to believe that this early population of black hole and neutron star binaries may be very different from those we see today. We live in an aging Universe – the majority of stars that will ever form have already done so [55], and the accelerating expansion of the Universe means that space will increasingly become a darker and emptier place [56]. The young Universe was very different. Stars were born and consequentially died with a much greater frequency. Heavy elements were scarcer, allowing stars (and their black hole descendants) to sustain much larger masses. Measuring the properties of the stochastic background – the volume and statistical character of the gravitational-wave static – may allow us to infer just how compact binary masses and merger rates have evolved over cosmic time.

The gravitational-wave background will serve not only as a tool with which to learn about distant gravitational-wave sources, but one that can additionally be used to study gravity itself. Our best current understanding of gravity is encapsulated in Einstein's general relativity. The predictions of general relativity have been borne out by every astronomical observation to date – the precession of Mercury, the gravitational deflection of light, and the orbital decay of binary pulsars, to name but a few [57]. So far, though, general relativity has been subjected only to fairly forgiving "weak field" tests, involving small masses moving at relatively slow speeds through weak gravitational fields. With their discovery, gravitational waves offer a means of testing general relativity in truly extreme environments – namely the extraordinarily strong and dynamic gravitational fields in the immediate vicinity of black holes and neutron stars.

Beyond their role as messengers, gravitational waves themselves offer a new arena for testing general relativity. General relativity predicts that gravitational waves travel at exactly the speed of light, a fact that the joint electromagnetic and gravitational-wave observations of GW170817 have confirmed to one part in 10^{15} [53]. General relativity also predicts that gravitational waves take on very particular shapes (or polarizations). In general, the polarizations of gravitational waves are difficult to directly characterize [44, 58]. We will see in Ch. 5 that the stochastic gravitational-wave background can be utilized to measure the polarizations of gravitational waves, allowing us to directly determine whether gravitational waves are consistent (or not) with predictions from general relativity.

Of course, before we can utilize the stochastic gravitational-wave background to study distant compact binaries or constrain deviations from general relativity, we must first succeed in detecting it. The signals comprising the gravitational-wave background are generally orders of magnitude weaker than those from GW150914 or GW170817. Additionally, a host of other terrestrial effects can conspire to mimic this tiny gravitational-wave "static." Together, these facts make the confident detection of the gravitational-wave background a formidable challenge. In Chs. 6 and 7, we will seek to detect the stochastic background with Advanced LIGO, and additionally explore techniques to differentiate a true stochastic signal from terrestrial "mimickers."

2. Do compact binary mergers give rise to prompt radio counterparts?

The binary neutron star merger GW170817 taught us that at least some fraction of gravitational waves come accompanied by light. GW170817 was observed in nearly every electromagnetic band [54]. The gravitational-wave signal was followed seconds later by a burst of high-energy gamma rays generated in the neutron stars' cataclysmic collision [59]. Soon after came the visible light of a "kilonova" – literally the radioactive glow of heavy elements being forged in the collision's aftermath [60, 61]. Finally, GW170817 launched a jet of particles moving near the speed of light, yielding a long-lasting radio signal that faded away only months later [62].

This confluence of gravitational waves and electromagnetic signals proved invaluable. By synthesizing the electromagnetic and gravitational-wave signatures of GW170817, we learned far more than we could have with either signal alone. Analysis of GW170817's optical afterglow revealed that binary neutron stars act as cosmic forges, very likely producing most of the gold found in the present-day Universe [63]. As mentioned above, the relative arrival times of GW170817's gravitational and gamma-ray signals gave us an extraordinarily precise measurement of the speed of gravitational waves [53]. And the gravitational-wave and radio analyses of GW170817, combined with optical observation of its host galaxy, even provided a direct measurement of the expansion rate of the Universe [64].

We do not yet understand if these so-called "multi-messenger" events, appearing to us both via gravitational waves and light (and one day, perhaps, neutrinos), are a rare occurrence or the new norm. Were we simply exceedingly lucky with GW170817, or will a large fraction of neutron stars generically yield observable electromagnetic counterparts? Moreover, beyond gamma-rays, kilonovae, and radio jets, are there additional yet-to-be-observed classes of electromagnetic signals that can accompany compact binary mergers?

In Chs. 8-10, we will switch gears and focus on one such theorized counterpart: **prompt radio emission**. A sharp pulse of radio waves, prompt emission is predicted to be generated in the last instant of a binary neutron star's life, as the two objects collide at nearly the speed of light. Prompt emission is quite speculative. It has not yet been observed, and it is by no means certain that it even exists. If detected, however, prompt emission would provide an invaluable snapshot of a binary's final moments, allowing us to study the conditions at the very heart of the neutron star merger.

Chapter 2
The Basics of Gravitational-Wave Astronomy

I'll start by reviewing the basics of gravitational-wave astronomy – the properties of gravitational-waves, their generation by various astrophysical sources, and the data analysis schemes with which we can detect gravitational waves in otherwise noisy data. My goal is to provide enough background to contextualize and prepare readers for subsequent chapters, and to build intuition with examples that I've found pedagogically useful. In general I will *not* derive results from first physical principles; good introductions to the basic physics of gravitational waves appear in Refs. [65–68]. I will, though, try motivate first-principles results with physical arguments and order-of-magnitude estimates.

2.1 The Physical Effects of Gravitational Waves

Just as electromagnetic waves are vacuum solutions to Maxwell's equations, gravitational waves are vacuum solutions to Einstein's equations. In an otherwise flat spacetime, gravitational waves manifest as wavelike perturbations $h_{\alpha\beta}$ to the ordinary Minkowski metric $\eta_{\alpha\beta}$:

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta} \qquad (h_{\alpha\beta} \ll 1) \tag{2.1}$$

The perturbation $h_{\alpha\beta}$ is known as the **gravitational-wave strain**. Although the components are $h_{\alpha\beta}$ are coordinate-dependent, gravitational waves are most commonly discussed in the **transverse-traceless gauge**, in which both the strain tensor's trace h^α_α and the contraction $h_{\alpha\beta}k^\alpha$ with its wavevector k vanish. In this coordinate system, the strain tensor of a gravitational wave traveling in the z-direction has components

$$h_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{2.2}$$

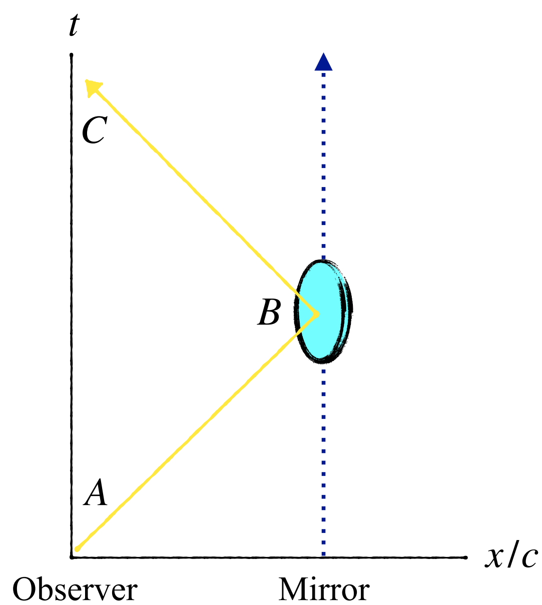

Figure 2.1: Spacetime diagram of the time-of-flight experiment carried out by our ambitious physics student. The student, at coordinate position $x = 0$, hangs a mirror at position $x = L_0$. At time $t = 0$ she then shoots a pulse of light at the mirror, recording the time Δt at which the pulse returns to her.

Equation (2.2) is expressed entirely in terms of two physical quantities: h_+ and $h_\times$. These represent the two polarizations, "plus" (+) and "cross" (×), available to gravitational waves in Einstein's general relativity.[1] A generic gravitational wave can be decomposed into the sum $h_{\alpha\beta} = h_+ \hat{\mathbf{e}}^+_{\alpha\beta} + h_\times \hat{\mathbf{e}}^\times_{\alpha\beta}$, where $\hat{\mathbf{e}}^+_{\alpha\beta}$ and $\hat{\mathbf{e}}^\times_{\alpha\beta}$ are the **basis tensors** for plus and cross modes:

$$\hat{\mathbf{e}}^+_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \hat{\mathbf{e}}^\times_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{2.3}$$

The physical effect of a gravitational wave is to vary the proper distance between freely-falling objects. Consider an enterprising physics student who wishes to measure the distance between herself (at the origin of her coordinate system) and a mirror at position $x = L_0$. She decides to do so via the time-of-flight measurement sketched in Fig. 2.1, shooting a pulse of light towards the mirror at time $t = 0$ (event A), letting the pulse bounce off the mirror (event

[1]We will explore additional polarizations allowed by alternative theories in Ch. 5.

B), and finally recording the time $\Delta t_0 = 2L_0/c$ at which the pulse returns to her (event C).

The student decides to repeat her experiment again, but this time in the presence of a passing gravitational wave moving along her z-axis, into the plane of Fig. 2.1. Assume this gravitational wave has a period much longer than the light's travel time to and from the mirror, so that $h_{\alpha\beta}$ is approximately constant for the duration of the experiment. The round-trip time Δt measured in the presence of this gravitational wave can be computed using the fact that events A and B are null-separated, as are B and C. If the coordinate separation between events A and B is $\Delta x^\alpha = (\frac{1}{2}c\Delta t, L_0, 0, 0)$, then we can use Eq. (2.1) to write

$$\begin{aligned} 0 &= g_{\alpha\beta}\Delta x^\alpha \Delta x^\beta \\ &= -\frac{1}{4}\left(c\Delta t\right)^2 + \left(1 + h_+\right) L_0^2. \end{aligned} \tag{2.4}$$

Solving for Δt and using the fact that $h_+ \ll 1$, the new round trip time is:

$$\begin{aligned} \Delta t &= 2\sqrt{1 + h_+} L_0/c \\ &\approx 2\left(1 + \frac{1}{2}h_+\right) L_0/c. \end{aligned} \tag{2.5}$$

The proper distance between the student and her mirror has increased by $\delta L = h_+ L_0/2$! This is the origin of the term "gravitational-wave strain." Just like a mechanical strain exerted on a material, gravitational waves stretch and shrink the proper distance between freely-falling objects by an amount proportional to the objects' initial separation.

In transverse-traceless coordinates, the effects of gravitational waves are confined entirely to the metric. Initially stationary, freely-falling objects (i.e. the student and her mirror) remain motionless at their initial coordinates, while the space between them expands and contracts to yield the additional time delay Δt. An alternative way to describe gravitational waves involves treating them like a mechanical force. To do this, we adopt local Lorentz coordinates, defined such that $g_{ij} = \eta_{ij}$ and all first derivatives of g_{ij} vanish. In these coordinates, the metric is fixed by design, and it is instead the mirror that moves in response to incident gravitational waves. To see this, we can return to the above example and compute the geodesic deviation of the mirror relative to our student. Remember that the geodesic deviation of a particle at position x^α relative to an observer with four-velocity u^α and proper time τ is (see e.g.

Ref. [66])

$$\frac{D^2x^\alpha}{d\tau^2} = R^\alpha{}_{\beta\mu\nu}u^\beta u^\mu x^\nu, \tag{2.6}$$

temporarily switching to geometrized units in which $c = 1$. Here, $R^\alpha{}_{\beta\mu\nu}$ is the Riemann tensor and $D/d\tau$ represents the covariant derivative along u^α. When applied in the rest frame of our student, this equation tells us that she measures the mirror's acceleration in the x-direction to be

$$\begin{aligned} \frac{d^2x}{dt^2} &\approx R^x{}_{ttx}L_0 \\ &= -R_{xtxt}L_0. \end{aligned} \tag{2.7}$$

In linearized gravity, the $(i0j0)$ components of the Riemann tensor are $R_{i0j0} = -\frac{1}{2}\ddot{h}_{ij}$ (dots denote coordinate time derivatives) and are invariant to gauge transformations [66]. So, despite working in local Lorentz coordinates, we can substitute into Eq. (2.7) the transverse-traceless expression for h_{ij} above. We find that, in these coordinates, gravitational waves serve to accelerate the mirror by $\ddot{x} = \frac{1}{2}\ddot{h}_+L_0$. Integrating twice, at leading order the student measures the mirror's position to be

$$x(t) = \left(1 + \frac{1}{2}h_{xx}\right)L, \tag{2.8}$$

consistent with Eq. (2.5) above.

These different *coordinate-dependent* descriptions – a fluctuating metric vs. an effective force on a fixed background – can be the source of much confusion, offering seemingly contradictory descriptions of how gravitational waves interact with laboratory experiments. When in doubt, it's generally best to revert to thinking in terms of time-of-flight measurements [69, 70].

Gravitational waves carry energy. To calculate the energy carried by gravitational waves, it is necessary to expand Einstein's equations to second order in $h_{\alpha\beta}$; see e.g. Refs. [65–67]. At leading order, one finds that the effective stress-energy tensor of a gravitational wave is

$$T^{\text{GW}}_{\mu\nu} = \frac{1}{32\pi G}\left\langle\left(\partial_\mu h^{\text{TT}}_{\alpha\beta}\right)\left(\partial_\nu h^{\alpha\beta}_{\text{TT}}\right)\right\rangle, \tag{2.9}$$

where the "TT" label denotes transverse-traceless coordinates and the brackets $\langle ... \rangle$ indicate an average taken over several wavelengths. The energy density of a gravitational wave is the T^{GW}_{00} component of Eq. (2.9). Restoring factors of c,

$$\boxed{\rho_{\text{GW}} = \frac{c^2}{32\pi G}\left\langle\dot{h}^{\text{TT}}_{\alpha\beta}\dot{h}^{\alpha\beta}_{\text{TT}}\right\rangle.} \tag{2.10}$$

2.2 Gravitational-Wave Generation

Just as electromagnetic waves are generated by accelerating charges, gravitational waves are sourced by accelerating mass and/or energy. The generation of gravitational-waves is a well-understood problem in general relativity. Before presenting the exact equation governing gravitational-wave generation, though, it is instructive to explore what *form* this equation must take, based only on dimensional analysis and symmetry arguments [65, 71].

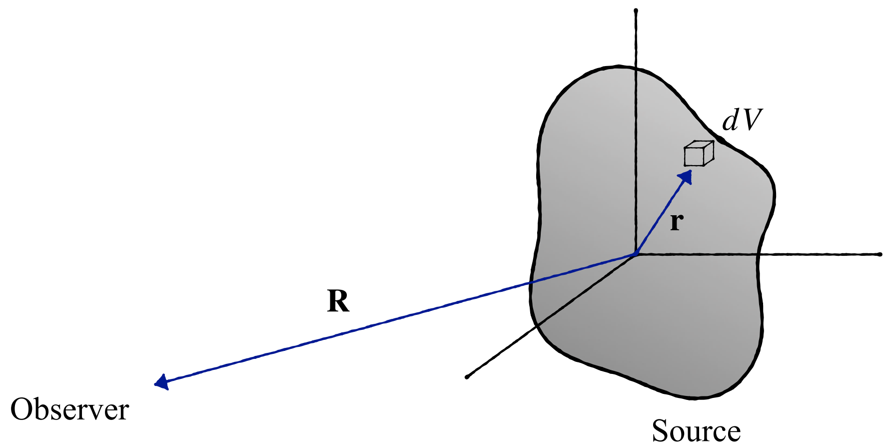

Figure 2.2: A generic gravitational-wave source, a distance R away from an observer. As argued in this section, at leading order only the acceleration of the source's quadrupole moment contributes to gravitational radiation.

The Newtonian gravitational potential is

$$\Phi(\mathbf{R}, t) = G \int \frac{\rho(\mathbf{r}, t_r)}{|\mathbf{R} - \mathbf{r}|} dV, \tag{2.11}$$

where $\mathbf{R}$ is the fixed vector from the origin to the observer and the integral is taken over $\mathbf{r}$. Here, $\rho(\mathbf{r}, t_r)$ is the source's mass-energy density at position $\mathbf{r}$, evaluated at the retarded time $t_r = t - |\mathbf{R} - \mathbf{r}|/c$. If the size of the source is small compared to the distance R to the observer, then we can take a Taylor expansion of $1/|\mathbf{R} - \mathbf{r}|$ in powers of (r/R):

$$\frac{1}{|\mathbf{R} - \mathbf{r}|} = \frac{1}{R} + \frac{\hat{\mathbf{R}} \cdot \hat{\mathbf{r}}}{R} \left(\frac{r}{R}\right) + \frac{3(\hat{\mathbf{R}} \cdot \hat{\mathbf{r}})^2 - 1}{2R} \left(\frac{r}{R}\right)^2 + ... \tag{2.12}$$

where $\hat{\mathbf{R}}$ and $\hat{\mathbf{r}}$ are unit vectors in the direction of $\mathbf{R}$ and $\mathbf{r}$, respectively. Substituting Eq. (2.12) into Eq. (2.11) and rearranging, we can expand the gravitational potential as

$$\Phi = \frac{G}{R}\int \rho\, dV + \frac{GR^i}{R^3}\int \rho r_i\, dV + \frac{3GR^iR^j}{2R^5}\left(\delta_i^k\delta_j^l - \frac{1}{3}\delta_{ij}\delta^{kl}\right)\int \rho r_k r_l dV + ... \tag{2.13}$$

This is a standard **multipole expansion**. The leading term involves the mass monopole $\int \rho\, dV$, while the second includes the mass dipole $\int \rho \mathbf{r}\, dV$. The third term is proportional to the quadrupole moment $\int \rho r_i r_j\, dV$, while the next term (not shown) would be the octopole, involving integration over *three* powers of r. To the extent that gravitational-waves resemble wave-like oscillations in the Newtonian potential, it stands to reason that their generation must also involve some function f of the source's multipole moments:

$$h \sim f\left(\int \rho\, dV,\ \int \rho \mathbf{r}\, dV,\ \int \rho r_i r_j\, dV\right). \tag{2.14}$$

First consider the monopole moment. A gravitational wave's strain tensor is dimensionless. Accordingly, we must construct some dimensionless quantity using only the monopole moment and the relevant physics: the constants G and c, the distance R to the source, and (possibly) time derivatives $\frac{d}{dt}$. We already know from Eq. (2.10) above that the energy of a gravitational wave scales as $\rho_{\text{GW}} \propto h^2$. If a gravitational wave represents true radiation, with energy escaping to infinity, then it must be the case that $\rho_{\text{GW}} 4\pi R^2 \propto h^2 R^2 = \text{constant}$, or $h \propto R^{-1}$. Given this dependence on distance, the only dimensionless quantity we can build with the monopole moment is $h \sim \frac{G}{Rc^2}\int \rho dV$. Note that the monopole moment is simply the total mass M of the source. So we have

$$h_{\text{monopole}} \sim \frac{GM}{Rc^2}. \tag{2.15}$$

Can this term describe gravitational-wave generation? No – mass conservation of an isolated system tells us that the right-hand side of this relation is *constant* in time. In contrast, we expect the left-hand side to be varying sinusoidally, with non-zero time derivatives $\frac{dh}{dt}$ and $\frac{d^2h}{dt^2}$. So a source's monopole moment cannot contribute to gravitational radiation.

Let's move to the dipole moment. The only dimensionless quantity we can build is $h_{\text{dipole}} \sim \frac{G}{Rc^3}\frac{d}{dt}\int \rho \mathbf{r} dV$. Looking more closely, the integral over $\rho\mathbf{r}$ gives us $M\mathbf{r}_c$, where $\mathbf{r}_c$ is the source's center of mass. This implies that $\frac{d}{dt}\int \rho \mathbf{r} dV = M\mathbf{v}_c$, the source's total linear momentum. We now have

$$h_{\text{dipole}} \sim \frac{GM\mathbf{v}_c}{Rc^3}. \tag{2.16}$$

This term also cannot describe an oscillatory gravitational wave. The linear momentum of an isolated system is also conserved, which implies that $\frac{dh}{dt} = \frac{d^2h}{dt^2} = 0$. So, although dipole radiation is the leading-order contributor to *electromagnetic* wave emission, the extra symmetries present here prohibit dipole *gravitational-wave* radiation.

Finally, we arrive at the quadrupole moment, from which we can construct the dimensionless quantity

$$h_{\text{quadrupole}} \sim \frac{G}{Rc^4}\frac{d^2}{dt^2}\int \rho r_i r_j dV. \tag{2.17}$$

We've run out of symmetries – there are no further conservation laws prohibiting time derivatives of this h. This suggests that **quadrupole radiation** represents the leading-order gravitational-wave emission.

We arrived at the quadrupole nature of gravitational waves using purely physical reasoning and dimensional analysis. A rigorous derivation confirms exactly this result – at leading order, gravitational-waves are generated via quadrupole radiation. In particular, the gravitational-wave strain h_{ij} at a large distance R from the source is [65–67]

$$\boxed{h_{ij} = \frac{2G}{Rc^4}\frac{d^2}{dt^2}\bar{I}_{ij}(t_r),} \tag{2.18}$$

where

$$\boxed{\begin{aligned} \bar{I}_{ij} &= I_{ij} - \frac{1}{3}\delta_{ij}\delta^{kl}I_{ij} \\ &= \int \rho \left(r_i r_j - \frac{1}{3}\delta_{ij} r^2 \right) dV, \end{aligned}} \tag{2.19}$$

called the **reduced quadrupole moment**, is the trace-free part of I_{ij}. Recall that it was $\bar{I}_{ij}$, not I_{ij}, that formally appeared in our multipole expansion in Eq. (2.13) above. Equation (2.18) is called the **quadrupole formula**, describing gravitational-wave generation at the lowest (non-zero) order.

As described in Ch. 2.1, gravitational waves are most commonly described via transverse-traceless (TT) gauge. Generically, the strain tensor given by Eq. (2.18) is *not* automatically in these desired coordinates. Thus computing the gravitational wave signal from a given source is generally a two-step process: first we compute h_{ij} according to Eq. (2.18), then we perform a coordinate transformation to obtain h_{ij}^{TT}.

First we make h_{ij} transverse. This is done using the **projection tensor**

$$P_{ij} = \delta_{ij} - n_i n_j. \tag{2.20}$$

Note that $P_{ij} n^i = 0$. Meanwhile, given a vector m^i orthogonal to n^i, $P_{ij} m^i = m_j$. Thus, as its name might suggest, P_{ij} projects vectors onto the plane perpendicular to n_i. Given a gravitational-wave traveling in the k^i direction, we can therefore obtain the transverse strain tensor via $P_i^k P_j^l h_{kl}$. All that's left is to subtract off the trace. All together, we can convert an arbitrary h_{ij} to the transverse-traceless gauge using

$$h_{ij}^{\mathrm{TT}} = \left(P_i^k P_j^l - \frac{1}{2} P_{ij} P^{kl} \right) h_{kl}. \tag{2.21}$$

2.3 Gravitational-Wave Sources

Although any system with a time-varying quadrupole moment yields gravitational radiation, we can only expect detectable gravitational waves from the most extreme systems in the Universe – large masses moving close to the speed of light.

To see this, let's return to the quadrupole formula above and use the approximation that $\frac{d^2}{dt^2} I\!\!\!-_{ij} = \frac{d^2}{dt^2} \int \rho (r_i r_j - \frac{1}{3}\delta_{ij} r^2) dV \sim M v^2$, where v is the system's characteristic velocity and M its total mass. From Eq. (2.18), then, a source with mass M and speed v yields gravitational waves of amplitude

$$\begin{aligned} h &\sim \frac{G}{Rc^4} M v^2 \\ &\sim \frac{GM/c^2}{R} \left(\frac{v}{c} \right)^2 . \end{aligned} \tag{2.22}$$

Note that the first term is just the ratio between the source's Schwarzschild radius and our distance to the source. Plug in some characteristic numbers:

$$h \sim 10^{-24} \left(\frac{100\,\mathrm{Mpc}}{R} \right) \left(\frac{M}{1\,M_\odot} \right) \left(\frac{0.1c}{c} \right)^2 . \tag{2.23}$$

Thus a solar-mass moving at one-tenth of the speed of light 100 Mpc away generates strain of only $h \sim 10^{-24}$. As we will see below, strain of this order is essentially the limit of what we can detect with present-day detectors. Eq. (2.23) also makes it clear why the generation of detectable gravitational waves on Earth is basically impossible. Say we constructed some gravitational-wave emission apparatus (maybe a rotating dumbbell) on Earth, placing it

100 km from our detectors. Even if we could move the dumbbell at speeds of 0.1 c, Eq. (2.23) implies that we'd need the dumbbell's mass to be of order 10^{10} kg, equivalent to roughy one hundred aircraft carriers.

Prime sources of detectable gravitational waves are compact binary mergers, rotating neutron stars, and stellar supernovae. Together, these comprise the canonical sources hunted by present-day gravitational-wave detectors.

2.3.1 Compact binary mergers

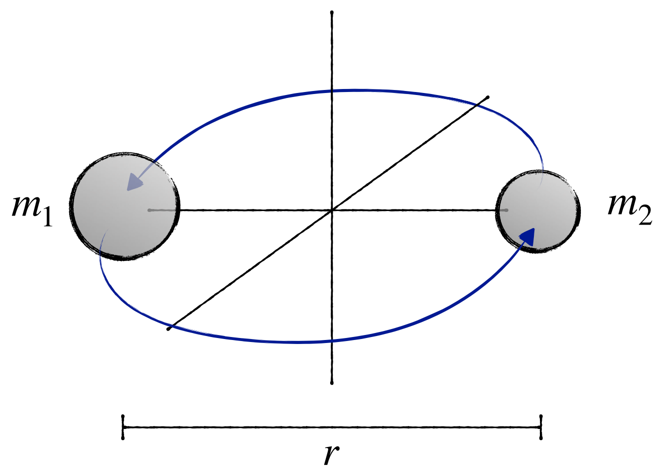

Figure 2.3: An illustration of a compact binary, composed of black holes and/or neutron stars of masses m_1 and m_2 orbiting with separation r.

The most well-studied gravitational-wave source is a **compact binary** composed of two stellar remnants – neutron stars, black holes, and/or white dwarfs (together known as "compact objects"). Figure 2.3 shows a cartoon illustration of a compact binary – two point masses m_1 and m_2 with orbital frequency f and separation r. Using the quadrupole formula [Eq. (2.18)] we can derive some characteristics of the gravitational waves emitted by compact binary sources.

We'll consider the effective one-body problem: a reduced mass $\mu = m_1 m_2/(m_1 + m_2)$ orbiting the system's total mass $M = m_1 + m_2$ at distance r. In this one-body picture, the reduced mass' Cartesian coordinates at some time t are

$$\mathbf{r}_\mu = \Big(r\cos(2\pi f t),\, r\sin(2\pi f t),\, 0\Big), \tag{2.24}$$

and the binary's mass-density is

$$\rho(\mathbf{r}) = M\delta^3(\mathbf{r}) + \mu\,\delta^3(\mathbf{r}-\mathbf{r}_\mu). \tag{2.25}$$

Its reduced quadrupole moment is therefore

$$\begin{aligned}
I\!\!\!-_{ij} &= \int \rho(\mathbf{r}) \left(r_i r_j - \frac{1}{3}\delta_{ij} r^2 \right) dV \\
&= \mu \left(r_{\mu,i} r_{\mu,j} - \frac{1}{3}\delta_{ij} r_\mu^2 \right) \\
&= \frac{\mu r^2}{3} \begin{pmatrix} 3\cos^2(2\pi f t) - 1 & \frac{3}{2}\sin(4\pi f t) & 0 \\ \frac{3}{2}\sin(4\pi f t) & 3\sin^2(2\pi f t) - 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}
\end{aligned} \tag{2.26}$$

with second time derivatives

$$\ddot{I\!\!\!-}_{ij} = -8\pi^2 \mu r^2 f^2 \begin{pmatrix} \cos(4\pi f t) & \sin(4\pi f t) & 0 \\ \sin(4\pi f t) & -\cos(4\pi f t) & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.27}$$

Using Eq. (2.18), the gravitational-wave strain generated by the compact binary is

$$h_{ij} = -\frac{16 G \pi^2 \mu r^2 f^2}{R c^4} \begin{pmatrix} \cos(4\pi f t) & \sin(4\pi f t) & 0 \\ \sin(4\pi f t) & -\cos(4\pi f t) & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.28}$$

Let's simplify Eq. (2.28) by using Kepler's Third Law[2] to replace $r = (GM/4\pi^2 f^2)^{1/3}$. Then

$$\boxed{h_{ij} = -\frac{4}{R c^4} \left(G\mathcal{M}_c\right)^{5/3} \left(2\pi f\right)^{2/3} \begin{pmatrix} \cos(4\pi f t) & \sin(4\pi f t) & 0 \\ \sin(4\pi f t) & -\cos(4\pi f t) & 0 \\ 0 & 0 & 0 \end{pmatrix},} \tag{2.29}$$

where $\mathcal{M}_c = \eta^{3/5} M$ is the binary's **chirp mass**. Here, $\eta = \mu/M = m_1 m_2/(m_1 + m_2)^2$ is the system's reduced mass ratio (also known as the symmetric mass ratio). At leading order, we see that the amplitude of gravitational radiation from a compact binary depends only on this particular combination $\mathcal{M}_c$ of component masses. Equation (2.29) demonstrates another important feature

[2]Kepler's Third Law remains valid for circular orbits around non-rotating black holes [68].

of gravitational radiation – the gravitational waves are generated at *twice* the binary's orbital frequency.

What does our result imply about the polarization of gravitational waves from compact binaries? First consider an observer positioned on the z-axis, such that they view the binary face-on. Fortuitously for this observer, Eq. (2.29) is already in the appropriate transverse-traceless gauge. They see a *circularly-polarized* gravitational-wave signal, with equal amplitude "plus" and "cross" modes exactly 90° out of phase.

What about an observer on the x-axis, who sees the binary "edge on?" First, rotate the strain Eq. (2.29) to the primed coordinates adopted by this observer as shown in Fig. 2.4:

$$h_{i'j'} = h_0 \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\cos(4\pi f t) & -\sin(4\pi f t) \\ 0 & -\sin(4\pi f t) & \cos(4\pi f t) \end{pmatrix}, \tag{2.30}$$

where $h_0 = -\frac{4}{Rc^4}\left(G\mathcal{M}_c\right)^{5/3}\left(2\pi f\right)^{2/3}$. Next, we need to transform to transverse traceless coordinates. We're interested in waves traveling out along the z' axis.

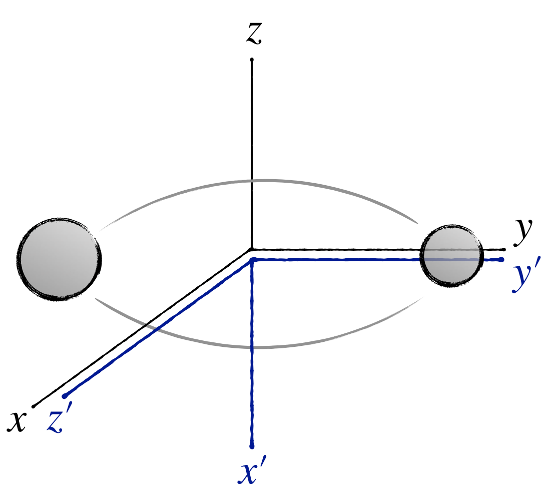

Figure 2.4: Primed coordinates for an observer viewing the binary edge-on. This edge-on observer measures the gravitational waves propagating out along the x-axis, defined as the z'-axis in their frame.

In primed coordinates, the corresponding projection tensor is

$$P_{i'j'} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.31}$$

Substituting Eqs. (2.30) and (2.31) into Eq. (2.21), the transverse-traceless strain measured by our observer on the x-axis is

$$h^{TT}_{i'j'} = \frac{1}{2} h_0 \cos(4\pi f t) \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.32}$$

The edge-on observer, then, sees a purely linear "plus" polarized signal.

2.3.2 Rotating neutron stars

Another canonical gravitational-wave source is a rotating neutron star. A perfectly spherical star has no time-varying quadrupole moment, and so will emit no gravitational radiation. However, neutron stars are predicted to be slightly deformed due to their rapid rotation and/or their intense internal magnetic fields. If these deformations are not perfectly axisymmetric, then they will result in a varying quadrupole moment, and the neutron star will effectively become a source of steady, monochromatic gravitational radiation (called "continuous gravitational waves").

As an example, consider the deformed neutron star in Fig. 2.5, whose rotation axis lies parallel to one of its principal directions. We'll start in a frame co-rotating with the neutron star, with coordinate axes aligned with the neutron star's principal axes, and assume that the star rotates about the z-axis with angular velocity ω. In our co-rotating frame, the star's quadrupole moment is of the form

$$I_{\hat{i}\hat{j}} = \begin{pmatrix} I_{\hat{x}\hat{x}} & 0 & 0 \\ 0 & I_{\hat{y}\hat{y}} & 0 \\ 0 & 0 & I_{\hat{z}\hat{z}} \end{pmatrix}, \tag{2.33}$$

for some non-zero $I_{\hat{x}\hat{x}}$, $I_{\hat{y}\hat{y}}$, and $I_{\hat{z}\hat{z}}$ (I'll use "hats" to denote quantities in the corotating frame). These three moments can be used to define the neutron star's characteristic **ellipticity**:

$$\epsilon = \frac{I_{\hat{x}\hat{x}} - I_{\hat{y}\hat{y}}}{I_{\hat{z}\hat{z}}}. \tag{2.34}$$

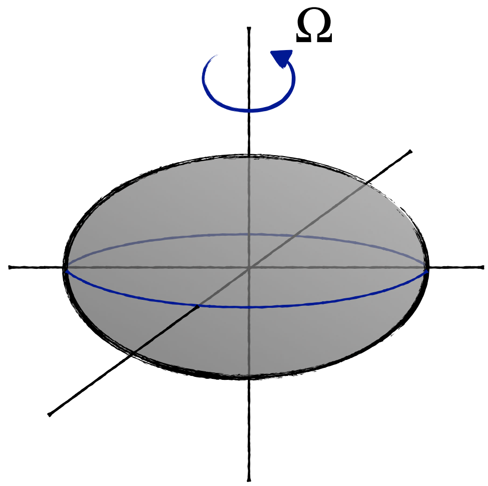

Figure 2.5: Illustration of a deformed, rotating neutron star. For simplicity, we consider the case in which neutron star's rotation is aligned with one of its principal axes.

Given a particular neutron star ellipticity, what is the resulting gravitational radiation? First, we need to transform the quadrupole tensor from co-rotating coordinates to inertial non-rotating coordinates. If the neutron star rotates at angular velocity ω, then the quadrupole tensor $I_{ij}(t)$ measured by an inertial observer at time t is found by rotating $I_{\hat{i}\hat{j}}$ through an angle ωt about the z-axis, giving

$$I_{ij} = \begin{pmatrix} I_{\hat{x}\hat{x}}\cos^2\omega t + I_{\hat{y}\hat{y}}\sin^2\omega t & \frac{1}{2}(I_{\hat{x}\hat{x}} - I_{\hat{y}\hat{y}})\sin 2\omega t & 0 \\ \frac{1}{2}(I_{\hat{x}\hat{x}} - I_{\hat{y}\hat{y}})\sin 2\omega t & I_{\hat{x}\hat{x}}\sin^2\omega t + I_{\hat{y}\hat{y}}\cos^2\omega t & 0 \\ 0 & 0 & \hat{I}_{zz} \end{pmatrix}. \tag{2.35}$$

This expression becomes much nicer if we use the definition of ellipticity to replace $I_{\hat{y}\hat{y}} = I_{\hat{x}\hat{x}} - I_{\hat{z}\hat{z}}\epsilon$. Substituting this into Eq. (2.35) and simplifying, the quadrupole tensor becomes

$$I_{ij} = \begin{pmatrix} I_{\hat{x}\hat{x}} - \epsilon I_{\hat{z}\hat{z}}\sin^2\omega t & \frac{1}{2}\epsilon I_{\hat{z}\hat{z}}\sin 2\omega t & 0 \\ \frac{1}{2}\epsilon I_{\hat{z}\hat{z}}\sin 2\omega t & I_{\hat{x}\hat{x}} - \epsilon I_{\hat{z}\hat{z}}\cos^2\omega t & 0 \\ 0 & 0 & I_{\hat{z}\hat{z}} \end{pmatrix} \tag{2.36}$$

with second time derivatives

$$\ddot{I}_{ij} = \ddot{\text{\it I}\!\!\!-}_{ij} = -2\epsilon I_{\hat{z}\hat{z}}\omega^2 \begin{pmatrix} \cos 2\omega t & \sin 2\omega t & 0 \\ \sin 2\omega t & -\cos 2\omega t & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.37}$$

Finally, the strain measured at distance r is

$$\boxed{h_{ij} = -\frac{16\pi^2 G}{c^4}\frac{\epsilon I_{\hat{z}\hat{z}} f^2}{r}\begin{pmatrix}\cos 2\omega t & \sin 2\omega t & 0\\ \sin 2\omega t & -\cos 2\omega t & 0\\ 0 & 0 & 0\end{pmatrix},} \tag{2.38}$$

expressed in terms of the neutron star's physical rotation frequency f. Just like compact binaries, we again see that rotating neutron stars generate gravitational waves at $2f$.

We've limited ourselves to the fairly simple situation in which the neutron star rotates about one of its principal axes. In reality this is not guaranteed to be the case. A more realistic system rotating about an arbitrary axis will additionally exhibit spin precession. The resulting gravitational waves will be comparable in amplitude to our estimate in Eq. (2.38), but will be radiated at both once and twice the neutron star's rotational frequency; see Ref. [72] for a pedagogical derivation of this effect.

We can use Eq. (2.38) to estimate the ellipticities ϵ that might yield detectable gravitational-wave signals. First, what gravitational-wave strain h might we be able to successfully detect from a continuously rotating neutron star? To help us estimate this, look to the *successfully* detected signal GW150914. The binary black hole merger GW150914 had peak strain $h_{\mathrm{BBH}} \sim 10^{-21}$ and was observed for approximately $0.1\,\mathrm{s}$ [25]. We will see that the **signal-to-noise ratio** (SNR) of a gravitational-wave signal scales as $\mathrm{SNR} \propto h\sqrt{T}$, where T is the signal's duration. If we were able to detect strain 10^{-21} in $0.1\,\mathrm{s}$ of observation, then a rotating neutron star observed for a *year* should be detectable when its strain is of order

$$h_{\mathrm{NS}} \sim h_{\mathrm{BBH}}\sqrt{\frac{0.1\,\mathrm{s}}{10^7\,\mathrm{s}}} \sim 10^{-25}. \tag{2.39}$$

Sure enough, recent Advanced LIGO and Virgo searches for persistent gravitational waves from rotating neutron stars achieve upper limits of order 10^{-25} [73–75].

What neutron star ellipticity does this correspond to? The canonical neutron star moment of inertia is $I_{zz} \sim M_{\mathrm{NS}} R_{\mathrm{NS}}^2 \sim 10^{38}\,\mathrm{kg\,m^2}$, where $M_{\mathrm{NS}} \sim 1\,M_\odot$ and $R_{\mathrm{NS}} \sim 10\,\mathrm{km}$ are typical neutron star masses and radii. Let's also consider a neutron star rotating at $f = 200\,\mathrm{Hz}$ at a distance $r = 1\,\mathrm{kpc}$. The resulting ellipticity is of order

$$\epsilon \sim 10^{-7}. \tag{2.40}$$

This is equal in magnitude to the latest constraints on neutron star ellipticities obtain in LIGO and Virgo's O2 observing run [73, 74].

2.3.3 Stellar core collapse

A final canonical gravitational-wave source is the core collapse of massive stars. Massive stars support themselves through nuclear burning, counteracting their own gravitational pull with the light and heat generated by nuclear fusion. When such a star exhausts its fuel reserves, it can no longer withstand its own gravitational attraction, imploding inward upon itself. The infalling matter strikes a proto-neutron star forming at the star's center and, in what is known as the "core bounce," reverses course, generating an explosion that, with the help of an immense neutrino outflow, disrupts the star in a supernova.

Our previous examples – compact binaries and rotating neutron stars – are, at first order, very simple systems that lend themselves to analytic expressions for the expected gravitational-wave emission. In contrast, core-collapse supernova are inherently complex and chaotic. In fact, at lowest order stellar core-collapse should not radiate gravitational-waves at all! A collapsing and/or bouncing spherically-symmetric shell has no time-varying quadrupole moment. Any gravitational waves emitted in stellar collapses must therefore be associated with higher-order complexities – asymmetry in the collapsing core, turbulence in the ejecta, etc. Unfortunately for us, these complexities do not lend themselves to pencil-and-paper estimates of gravitational-wave production.

Instead, we can *very* crudely estimate the gravitational-wave strain we might expect from a core-collapse supernova; this argument follows Ch. 3 of Ref. [76]. Given a collapsing star of mass M, then we'll start by assuming that some fraction ϵ of the star's mass-energy is converted to gravitational waves:

$$E_{\mathrm{GW}} \sim \epsilon M c^2. \tag{2.41}$$

Next, assume that the gravitational-wave burst is of some characteristic duration τ, so that the luminosity of the burst is $L \sim \frac{E_{\mathrm{GW}}}{\tau}$ and the gravitational-wave flux measured at Earth is $F \sim \frac{L}{4\pi R^2}$, or

$$F \sim \frac{\epsilon M c^2}{4\pi R^2 \tau}. \tag{2.42}$$

Note that gravitational-wave detectors do not measure a gravitational wave's energy, but rather its amplitude. Above we saw that the energy-density of a gravitational wave is $\rho_{\rm GW} \sim \frac{c^2 \dot{h}^2}{32\pi G}$. Then a gravitational wave's flux is related to its strain via

$$\begin{aligned} F &\sim \frac{c^3 \dot{h}^2}{32\pi G} \\ &\sim \frac{\pi c^3}{8G} f^2 h^2, \end{aligned} \tag{2.43}$$

where in the second line we used $\dot{h} \sim 2\pi f h$ for a wave of frequency f. By combining Eqs. (2.41) and (2.43), we can solve for the characteristic amplitude h corresponding to a gravitational-wave burst of efficiency ϵ, frequency f, and duration τ:

$$h \sim \left(\frac{2G}{\pi^2 c} \frac{\epsilon M}{f^2 R^2 \tau} \right)^{1/2}. \tag{2.44}$$

We still need to do some more (even less defensible) work before we can obtain a quantitative estimate for h. First, a burst of duration τ will have peak frequency of order $f \sim (2\pi\tau)^{-1}$. Meanwhile, we might identify τ with the light crossing time of our source. If our source is compact, then its size is of order its Schwarzschild radius $R_S = \frac{2GM}{c^2}$, and so $\tau \sim \frac{2GM}{c^3}$. Alternatively, we might associate τ with the free-fall time for an object of mass M and size $\frac{2GM}{c^2}$; this gives the same result. Substituting these relations into Eq. (2.44) above,

$$h \sim \frac{GM}{\pi^2 c^2} \frac{\epsilon^{1/2}}{R}. \tag{2.45}$$

Note that this result can be expressed even more simply as $h \sim \epsilon^{1/2} \frac{R_S}{R}$.

Let's plug in some typical numbers, assuming a $M = 10 M_\odot$ star at a distance of $R = 10\,\text{kpc}$. State-of-the-art simulations, meanwhile, suggest conversion efficiencies of order $\epsilon \sim 10^{-9}$ between a collapsing progenitor's mass energy and gravitational-wave generation [77–79]. Then we might expect gravitational waves of amplitude

$$\boxed{h \sim 1.0 \times 10^{-22} \left(\frac{M}{10 M_\odot} \right) \left(\frac{\epsilon}{10^{-9}} \right)^{1/2} \left(\frac{R}{10\,\text{kpc}} \right)^{-1},} \tag{2.46}$$

radiated at a characteristic frequency

$$\begin{aligned} f &\sim \frac{1}{2\pi}\frac{c^3}{2GM} \\ &\sim 1.6\,\mathrm{kHz}\left(\frac{M}{10M_\odot}\right)^{-1} \end{aligned} \tag{2.47}$$

over time

$$\begin{aligned} \tau &\sim \frac{2GM}{c^3} \\ &\sim 0.1\,\mathrm{ms}\left(\frac{M}{10M_\odot}\right). \end{aligned} \tag{2.48}$$

Equation (2.46) suggests that we might be able to detect gravitational-wave bursts from core-collapse supernova within the Milky Way. What about more distant sources? The Andromeda Galaxy sits at a distance of approximately 800 kpc. A supernova in Andromeda might therefore yield strain of amplitude $10^{-22}(10\,\mathrm{kpc}/800\,\mathrm{kpc}) \sim 10^{-24}$. This is likely undetectable with current instruments. Therefore, barring exotic emission mechanisms, only supernova in the Milky Way are likely to be accessible to present-day gravitational-wave detectors.

2.4 The Advanced LIGO Detectors

Next I'll briefly describe the Advanced LIGO detectors. This discussion will be largely qualitative; more detailed descriptions of the instrumental and engineering bases of interferometric gravitational-wave detection appear in Refs. [67, 70]. Also see the more technical summaries presented in Refs. [22, 23, 80].

Advanced LIGO comprises two near-identical instruments in Hanford, Washington and Livingston, Louisiana. The Hanford instrument is shown in Figure 2.6. Just like our hypothetical graduate student from Ch. 2.1 above, each LIGO instrument detects passing gravitational waves by measuring apparent changes in length of its two 4 km arms.

Very broadly, each detector operates like a Michelson interferometer. Laser light is sent through a central beamsplitter, directed down each arm, and recombined, directed onto a photodiode at the interferometer's "output port." The interference between the two recombined beams (and hence the total power measured by the photodiode) encodes the relative phases acquired by

Figure 2.6: The LIGO detector in Hanford, Washington. Photo credit: Caltech/MIT/LIGO Laboratory.

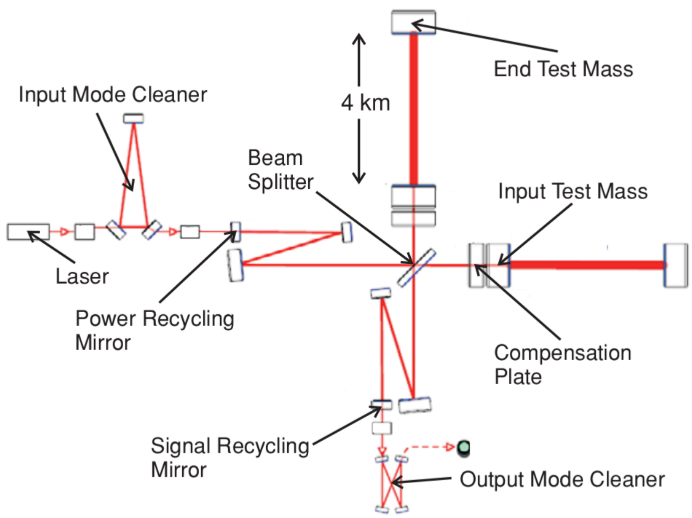

Figure 2.7: Schematic of the optical layout of the Advanced LIGO detectors. Each instrument measures the relative lengths of two 4 km Fabry-Perot arms. Figure from Ref. [22].

the laser light during propagation along each arm, which in turn depends on the arms' relative lengths. In practice, the interferometers are considerably more complex than a simple Michelson. Figure 2.7 illustrates the optical layout of the Advanced LIGO detectors. The interferometers' arms are themselves Fabry-Perot cavities, bounded by two 40 kg mirrors, or "test masses." Each test mass is suspended from a sequence of four pendula, achieving extraordinary mechanical isolation from the surrounding environment and ensuring that the test masses are effectively freely-falling at the frequencies of interest.

One of the main effects that hinders Advanced LIGO's sensitivity is quantum mechanical "shot noise," random Poissonian fluctuations due to the discrete nature of photons. The relative amplitude of shot noise is minimized by *maximizing* the number of photons – in other words, by maximizing the power circulating in the instrument. Advanced LIGO achieves this in two ways. First, high-reflectivity mirror coatings ensure that light remains trapped in the Fabry-Perot arms, undergoing many round trips before escaping and thereby increasing the power contained in each cavity. Second, the detectors additionally employ a so-called power recycling mirror. In the absence of a gravitational wave, recombined laser light is sent back towards the input laser; the power recycling mirror redirects this light once more into the instrument, increasing the effective laser power.

Increased cavity power comes at a cost, however. If $L = 4\,\mathrm{km}$ is the interferometer's arm length and N is the number of round trips along the arms completed by a typical photon, LIGO's sensitivity to gravitational waves falls off rapidly at frequencies $f \gtrsim \frac{c}{2NL}$, where the period of gravitational waves is comparable to or larger than the time a photon remains trapped in the detector arms. Increasing N therefore *decreases* the bandwidth over which LIGO is sensitive to gravitational waves. A final mirror, the signal recycling mirror, is introduced at LIGO's output ports in order compensate for this effect; its precise tuning allows the interferometer to maintain a broad sensitive bandwidth. Alternatively, other signal recycling tunings enable narrowband observations that sacrifice bandwidth in favor of heightened sensitivity across a narrow range of frequencies (envisioned, say, to allow targeted study of the monochromatic radiation from a known neutron star).

Unfortunately, gravitational waves are not the only phenomena that induce a signal in Advanced LIGO. Astrophysical gravitational waves must compete against a relative cacophony of noise sources. We will quantify the strength of noise in a given detector via its **noise power spectral density** (PSD) $P(f)$, defined by

$$\boxed{\langle \tilde{n}(f)\tilde{n}^*(f')\rangle = \frac{1}{2}\delta(f-f')P(|f|).} \tag{2.49}$$

Here, $\tilde{n}(f)$ is the Fourier transform of the detector's noise time series $n(t)$; a star (*) denotes complex conjugation. In practice, since the total strain $s(t)$ measured by a detector is dominated by its noise, we can approximate the

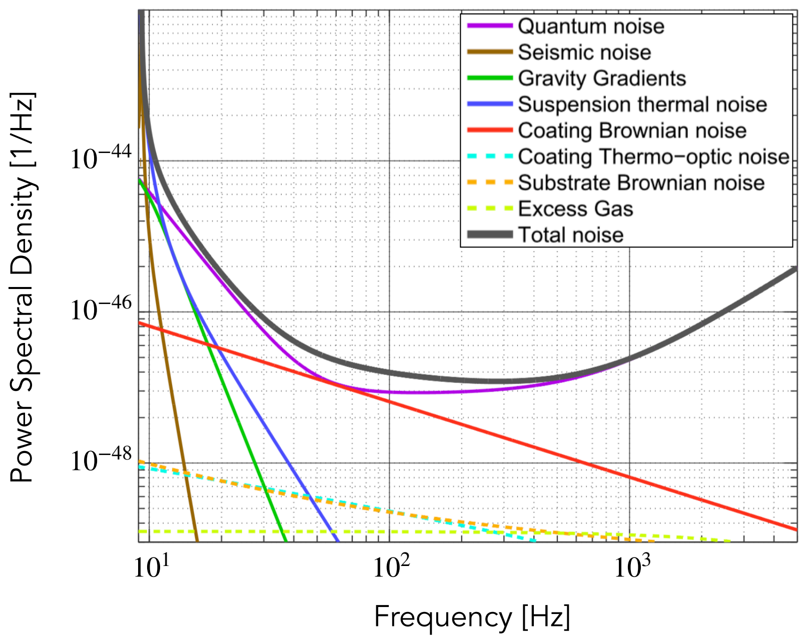

Figure 2.8: The design sensitivity noise power spectral density of Advanced LIGO. Also shown are the individual contributions from known and/or expected sources of instrumental or terrestrial noise. Figure adapted from Ref. [23].

power spectral density as $\langle \tilde{s}(f)\tilde{s}^*(f')\rangle \approx \frac{1}{2}\delta(f-f')P(|f|)$. A related quantity is the **noise amplitude spectral density** (ASD), defined as

$$\mathrm{ASD}(f) = \sqrt{P(f)}. \tag{2.50}$$

The ultimate power spectral density targeted by Advanced LIGO's "design sensitivity" is shown in Fig. 2.8, along with the myriad contributions from individual sources of noise [23]. Some noises sources are environmental – due to the seismic motion of the Earth or the minute variations in the local gravitational fields around the detectors (called gravity gradient noise). Others are instrumental. These include the thermal motion both in the fibers suspending Advanced LIGO's mirrors and in the coatings covering the mirrors themselves, as well as residual gas left behind in the imperfect vacuum tubes housing LIGO's optical systems. Advanced LIGO is even limited by quantum mechanical noise. At high frequencies, this takes the form of shot noise – Poissonian fluctuations in the number of photons emerging from the interferometer to arrive at the photodiode. At low frequencies, meanwhile, quantum mechanical noise manifests as radiation pressure – random fluctuations in the mirrors' positions as they are "kicked" by incident photons.

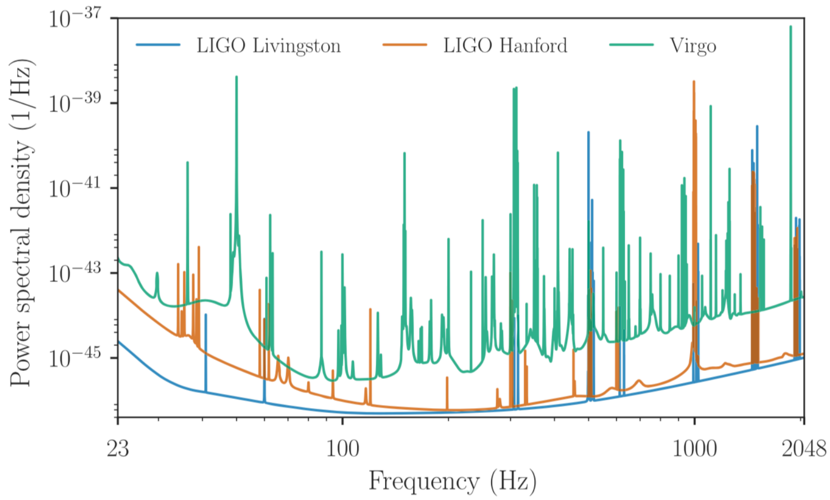

Figure 2.9: The Hanford, Livingston, and Virgo noise power spectral densities near the time of the binary neutron star GW170817. This figure is reproduced from Ref. [47].

For comparison, Fig. 2.9 shows the actual power spectral densities of the Hanford, Livingston, and Virgo detectors at the time of the binary neutron star merger GW170817 during the O2 observing run [47]. Note that Advanced LIGO's O2 PSDs are roughly an order of magnitude larger than the projected design-sensitivity PSDs in Fig. 2.8, indicating the sensitivity gains still possible from future development and commissioning. Second, the O2 PSDs are riddled with sharp narrowband features. These "lines" are due to a variety of sources, including power mains, electronics that sense and control the state of the interferometer, "violin mode" vibrations in the filaments suspending LIGO's optics, and deliberately-injected calibration signals [81].

It is worth taking an additional moment to discuss units. The PSD has what might appear to be fairly odd units: $[P(f)] = [\mathrm{strain}^2\,\mathrm{Hz}^{-1}]$. Correspondingly, the ASD has units $[\mathrm{ASD}(f)] = [\mathrm{strain}\,\mathrm{Hz}^{-1/2}]$. These units reflect the fact that the PSD defines the strain power of random detector noise *per unit frequency bandwidth*. Put another way, the PSD's units should remind us that the amplitude of random noise adds in *quadrature* as we integrate across frequencies [70]. This means that, when estimating noise amplitudes in gravitational wave detectors, we need to specify not only a frequency f but also a bandwidth Δf; the characteristic noise amplitude across this band is then given by $n_0 \sim \sqrt{P(f)\Delta f}$.

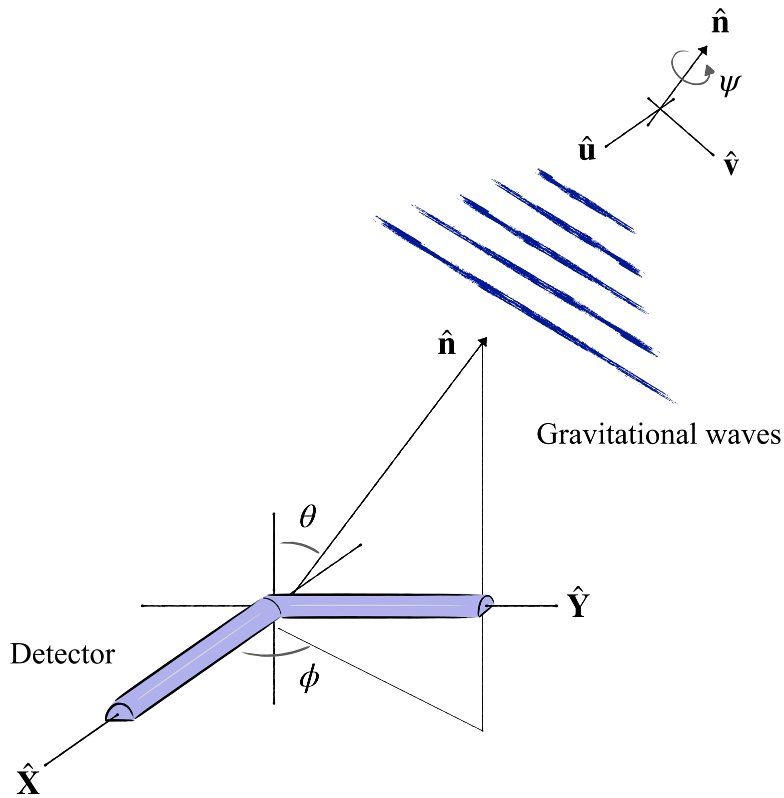

Figure 2.10: A schematic illustrating the frame of a gravitational-wave detector, as well as orthonormal coordinates describing incoming gravitational waves. The detector's arms are oriented along the $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ axes. The unit vector $\hat{\mathbf{n}}$ points towards the gravitational-wave source, and $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ are orthonormal vectors in the plane perpendicular to $\hat{\mathbf{n}}$; the polarization angle ψ describes the rotation of $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ about $\hat{\mathbf{n}}$.

2.5 Identification of Gravitational Waves in Noisy Data

In the previous sections we saw an overview of the properties and generation of gravitational waves, as well as a qualitative overview of the interferometers used to record them. Now we'll turn to the methods used to actually *identify* gravitational waves in the noisy data recorded by gravitational-wave detectors.

Figure 2.10 shows an illustration of a gravitational wave incident upon a LIGO-like interferometric detector from direction $\hat{\mathbf{n}}$. If the gravitational wave is of the form $h_{ij} = h_+\hat{\mathbf{e}}^+_{ij} + h_\times\hat{\mathbf{e}}^\times_{ij}$, then the resulting signal s output by the detector is

$$s = D_{ij}\left(h_+\hat{\mathbf{e}}^{ij}_+(\theta,\phi,\psi) + h_\times\hat{\mathbf{e}}^{ij}_\times(\theta,\phi,\psi)\right). \tag{2.51}$$

The **detector tensor** D_{ij} acts as the transfer function between the incident gravitational-wave strain and our measured signal. Interferometric gravitational-wave detectors like Advanced LIGO and Virgo have detector tensors of the form

$$D_{ij} = \frac{1}{2}\left(X_i X_j - Y_i Y_j\right), \tag{2.52}$$

where X^i and Y^i are unit vectors along each of the detector's arms. Equation (2.51) can alternatively be written

$$s = F^{+}(\theta,\phi,\psi)h_{+} + F^{\times}(\theta,\phi,\psi)h_{\times}, \tag{2.53}$$

where we've defined **antenna patterns** $F^{+}(\theta,\phi,\psi) = D_{ij}\hat{\mathrm{e}}^{ij}_{+}(\theta,\phi,\psi)$ and $F^{\times}(\theta,\phi,\psi) = D_{ij}\hat{\mathrm{e}}^{ij}_{\times}(\theta,\phi,\psi)$. The Advanced LIGO and Virgo antenna patterns are given analytically by[3,4]

$$\boxed{\begin{aligned} F^{+}(\theta,\phi,\psi) &= \frac{1}{2}\left(1+\cos^2\theta\right)\cos 2\phi\cos 2\psi - \cos\theta\sin 2\phi\sin 2\psi \\ F^{\times}(\theta,\phi,\psi) &= \frac{1}{2}\left(1+\cos^2\theta\right)\cos 2\phi\sin 2\psi + \cos\theta\sin 2\phi\cos 2\psi. \end{aligned}} \tag{2.54}$$

Strictly speaking, Eq. (2.54) is only correct in the long-wavelength regime, in which the wavelength of a gravitational wave is much longer than our detector. At smaller wavelengths (higher frequencies) the antenna patterns have non-negligible frequency dependence [70]. For our purposes, however, we will treat the antenna patterns as frequency-independent.

As we saw in Ch. 2.4, astrophysical gravitational-wave signals compete against a chorus of terrestrial and instrumental noise sources. How does the noise power in present-day detectors compare with the strength of astrophysical gravitational-wave signals? The binary neutron star merger GW170817, for example, had peak strain of order $h_0 \sim 10^{-22}$ at frequency $f \sim 2\,\mathrm{kHz}$. We can see in Fig. 2.9, meanwhile, that the corresponding noise power in both Hanford and Livingston is approximately $P(2\,\mathrm{kHz}) \sim 10^{-45}\,\mathrm{Hz}^{-1}$. If we assume a bandwidth $\Delta f \sim f \sim 2\,\mathrm{kHz}$, then the instrumental noise competing with GW170817 has amplitude $n_0 \sim \sqrt{P(2\,\mathrm{kHz})\Delta f} \sim 10^{-21}$. This is an order of magnitude larger than the signal itself! Seemingly paradoxically, though, GW170817 was a *very* loud detection, with signal-to-noise ratios of approximately 20 in each Advanced LIGO detector.

[3] Caution! Sign errors in the analytic antenna patterns abound in the literature.

[4] In Ch. 5 we will also see the analogous antenna patterns for additional polarizations permitted by alternative theories of gravity.

The detection of gravitational-wave signals like GW170817 that are far weaker than instrumental noise is made possible by **matched filtering**. Matched filtering involves searching for excess correlation $C = \int s(t)h(t)dt$ between instrumental output $s(t)$ and a pre-computed **template waveform** $h(t)$ that matches the gravitational-wave signal we hope to detect. We will rigorously work through the mathematics of matched filtering below. First, though, let's derive some back-of-the-envelope results to illustrate how matched filtering enables the detection of very weak gravitational-wave signals.

Consider the case where a gravitational wave is present, and we miraculously happen to have a template that perfectly matches the gravitational-wave signal. Define $\mathcal{S}$ to be the correlation measured in this case (this discussion follows that in Ch. 7 of Ref. [67]):

$$\begin{aligned}\mathcal{S} &= \int s(t)h(t)dt \\ &= \int \left[h(t) + n(t)\right] h(t)dt,\end{aligned} \tag{2.55}$$

where $n(t)$ is the noise present alongside the gravitational-wave signal. Terrestrial noise and the astrophysical gravitational-wave signal have nothing to do with one another, and so are uncorrelated; at leading order $\mathcal{S}$ is then just

$$\begin{aligned}\mathcal{S} &\approx \int h^2(t)dt \\ &\sim h_0^2 N\tau,\end{aligned} \tag{2.56}$$

where h_0 is the scale amplitude of the gravitational-wave, τ is the period of the the wave, and N is the total number of waveform cycles contained in the signal. In contrast, define $\mathcal{N}$ to be the correlation (using this same template) when a gravitational-wave signal is *not* present. Since random detector noise and our template are uncorrelated (and can both be positive or negative), our cross-correlation grows only as a random walk in N:

$$\begin{aligned}\mathcal{N} &= \int s(t)h(t)dt \\ &\sim n_0 h_0 \sqrt{N}\tau,\end{aligned} \tag{2.57}$$

where n_0 is the scale amplitude of random detector noise. The ratio $\mathcal{S}/\mathcal{N}$ now gives an estimate of a resulting **signal-to-noise ratio** (SNR). If $f = 1/\tau$ is the frequency of the gravitational-wave signal and T is the signal's duration, then we find [67]

$$\boxed{\text{SNR} \sim \sqrt{fT}\,\frac{h_0}{n_0}.} \tag{2.58}$$

Equation (2.58) implies that we can detect an arbitrarily weak signal as long as we observe it for a time

$$T \gg \frac{1}{f}\left(\frac{n_0}{h_0}\right)^2. \tag{2.59}$$

Let's return to the example of GW170817. According to Eq. (2.59), in order to detect GW170817-like signals we must observe them for times much longer than $\frac{1}{2\,\mathrm{kHz}}\left(\frac{10^{-21}}{10^{-22}}\right)^2 \sim 0.05\,\mathrm{s}$. In comparison, GW170817 lasted roughly $\sim 100\,\mathrm{s}$ in the band of Advanced LIGO. It's therefore not surprising that GW170817 was an incredibly confident detection, despite being an order of magnitude smaller in amplitude than instrumental noise.

Having derived a few useful rules-of-thumb, we can now more carefully work through the details of matched filtering. More formally, matched filtering is performed by computing $C = \int s(t)h(t)Q(t)dt$, where $Q(t)$ is a yet-unspecified real-valued filter. Ideally, we will choose $Q(t)$ to be the **optimal filter** that maximizes the SNR of a signal matching our template. To identify the optimal filter for a given template, it is easier to work in the frequency domain, re-expressing C as an integral over frequencies. This can be quickly done using the convolution theorem, but it is instructive to work this out by hand – in doing so we'll see several tricks that will reappear when discussing searches for the gravitational-wave background in Ch. 3.

To rewrite C in the frequency domain, first we'll use the fact that $h(t)$ and $Q(t)$ are both real to write

$$\begin{aligned} C &= \int_{-\infty}^{\infty} dt\, s(t)h(t)Q(t) \\ &= \int_{-\infty}^{\infty} dt\, s(t)\left[h(t)Q(t)\right]^* . \end{aligned} \tag{2.60}$$

Now explicitly express both $s(t)$ and $h(t)Q(t)$ as *inverse* Fourier transforms from their frequency domain representations:

$$\begin{aligned} C &= \int_{-\infty}^{\infty} dt \left[\int_{-\infty}^{\infty} df e^{2\pi i f t}\tilde{s}(f)\right]\left[\int_{-\infty}^{\infty} df' e^{2\pi i f' t}\tilde{h}(f')\tilde{Q}(f')\right]^* \\ &= \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df'\, \tilde{s}(f)\tilde{h}^*(f')\tilde{Q}^*(f') \int_{-\infty}^{\infty} dt\, e^{2\pi i (f-f')t}, \end{aligned} \tag{2.61}$$

where in the last line we've done some suggestive rearranging.

Let's take a careful look at the time integral. In the limit that t ranges between $\pm\infty$, it's not clear how exactly to calculate this integral. Instead, we will integrate over the bounded range $(-T/2, T/2)$ and then let $T \to \infty$.

$$\begin{aligned}
\int_{-\infty}^{\infty} dt\, e^{2\pi i(f-f')t} &= \lim_{T\to\infty} \int_{-T/2}^{T/2} dt\, e^{2\pi i(f-f')t} \\
&= \lim_{T\to\infty} \frac{1}{2\pi i(f-f')} \left(e^{2\pi i(f-f')t} \right) \Big|_{-T/2}^{T/2} \\
&= \lim_{T\to\infty} \frac{\sin \pi(f-f')T}{\pi(f-f')T}\, T \\
&= \lim_{T\to\infty} T \operatorname{sinc}\left(\pi(f-f')T\right).
\end{aligned} \tag{2.62}$$

The result is a sinc function, with a peak value of T at $f - f' = 0$, and whose central width is roughly $\Delta(f-f') = 1/T$. As we let $T \to \infty$, the sinc function becomes vanishingly narrow and infinitely tall. In other words, $\int_{-\infty}^{\infty} dt e^{2\pi i(f-f')t}$ behaves like a Dirac delta function:

$$\boxed{\lim_{T\to\infty} \int_{-T/2}^{T/2} dt\, e^{2\pi i(f-f')t} = \delta(f-f').} \tag{2.63}$$

When T is left finite, the integral is called a **finite time delta function**:

$$\boxed{\int_{-T/2}^{T/2} dt\, e^{2\pi i(f-f')t} \equiv \delta_T(f-f'),} \tag{2.64}$$

where

$$\delta_T(f-f') = \begin{cases} T & (f = f') \\ 0 & (\text{otherwise}) \end{cases}. \tag{2.65}$$

Returning to Eq. (2.61) and replacing the time integral with a Dirac delta function, we have

$$\begin{aligned}
C &= \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df'\, \tilde{s}(f) \tilde{h}^*(f') \tilde{Q}^*(f') \delta(f-f') \\
&= \int_{-\infty}^{\infty} df\, \tilde{s}(f) \tilde{h}^*(f) \tilde{Q}^*(f).
\end{aligned} \tag{2.66}$$

As above, we will define define $\mathcal{S}$ as the expectation value $\langle C \rangle$ when a signal matching our template is present in the data. In this case, $\langle \tilde{s}(f) \rangle = \tilde{h}(f)$, giving

$$\begin{aligned}\mathcal{S} &= \left\langle \int_{-\infty}^{\infty} df\, \tilde{s}(f)\tilde{h}^*(f)\tilde{Q}^*(f) \right\rangle \\ &= \int_{-\infty}^{\infty} |\tilde{h}(f)|^2 \tilde{Q}^*(f) df.\end{aligned} \tag{2.67}$$

Similarly, define $\mathcal{N}^2$ as the variance of C in the absence of a signal, such that the detectors measure pure noise: $\tilde{s}(f) = \tilde{n}(f)$. Since $\langle \tilde{n}(f) \rangle = 0$, the expectation value $\langle C \rangle$ vanishes in the noise-only case, and so the variance of C is

$$\begin{aligned}\mathcal{N}^2 &= \langle C^2 \rangle \\ &= \langle CC^* \rangle \\ &= \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} df' \left\langle \tilde{n}(f)\tilde{n}^*(f') \right\rangle \tilde{h}^*(f)\tilde{h}(f')\tilde{Q}^*(f)\tilde{Q}(f') \\ &= \frac{1}{2}\int_{-\infty}^{\infty} df |\tilde{h}(f)|^2 |\tilde{Q}(f)|^2 P(f),\end{aligned} \tag{2.68}$$

where in the second line we used the fact that C is real-valued to write $C = C^*$, and in the final line we used the definition of the power-spectral density $P(f)$. Together, Eqs. (2.67) and (2.68) let us to define the squared signal-to-noise ratio (SNR) of our matched filtering search:

$$\begin{aligned}\mathrm{SNR}^2 &= \frac{\mathcal{S}^2}{\mathcal{N}^2} \\ &= 2\frac{\left(\int_{-\infty}^{\infty} |\tilde{h}(f)|^2 \tilde{Q}^*(f) df\right)^2}{\int_{-\infty}^{\infty} df |\tilde{h}(f)|^2 |\tilde{Q}(f)|^2 P(f)}.\end{aligned} \tag{2.69}$$

We're now finally in a position to identify the filter $\tilde{Q}(f)$ that maximizes SNR^2. The entirely unobvious "trick" is to realize (i.e. be told) that we can rewrite Eq. (2.69) in the form

$$\mathrm{SNR}^2 = \frac{\left\langle h \,\middle|\, h\, Q\, P \right\rangle^2}{\left\langle h\, Q\, P \,\middle|\, h\, Q\, P \right\rangle}, \tag{2.70}$$

where we've defined an *inner product* (not to be confused with an expectation value)

$$\begin{aligned}\left\langle a \,\middle|\, b \right\rangle &= 2\int_{-\infty}^{\infty} df\, \frac{\tilde{a}(f)\tilde{b}^*(f)}{P(f)} \\ &= 4\int_{0}^{\infty} df\, \frac{\tilde{a}(f)\tilde{b}^*(f)}{P(f)}.\end{aligned} \tag{2.71}$$

between between real time-series $a(t)$ and $b(t)$. When dealing with ordinary vectors $\vec{A}$ and $\vec{B}$, the quantity $\frac{(\vec{A}\cdot\vec{B})^2}{\vec{B}\cdot\vec{B}}$ is maximized when the two vectors are parallel, such that $\vec{B} \propto \vec{A}$. This exact same idea holds true here: the squared SNR is maximized when $\tilde{h}(f)\tilde{Q}(f)P(f) \propto \tilde{h}(f)$, or when

$$\tilde{Q}(f) = \frac{1}{P(f)}, \tag{2.72}$$

setting the arbitrary constant of proportionality to one. Then Eq. (2.69) becomes

$$\boxed{\begin{aligned} \mathrm{SNR}^2_{\mathrm{opt}} &= 4\int_0^\infty df\, \frac{|\tilde{h}(f)|^2}{P(f)} \\ &= \langle h \,|\, h \rangle^{1/2}. \end{aligned}} \tag{2.73}$$

This is the **optimal SNR** of a gravitational-wave signal. More generally, the SNR of data $s(t)$ given template $h(t)$ is

$$\boxed{\mathrm{SNR}^2 = \frac{\langle s \,|\, h \rangle^2}{\langle h \,|\, h \rangle}.} \tag{2.74}$$

This reduces to Eq. (2.73) when $h(t) = s(t)$.

2.6 Looking Ahead: The Stochastic Gravitational-Wave Background

In this chapter we reviewed the properties, generation, and detection of gravitational waves. The principles we've seen so far are broadly applicable to virtually all areas of gravitational-wave astronomy. In the next chapter, we'll narrow our focus to one particular class of gravitational-wave signal: the stochastic gravitational-wave background.

Chapter 3
The Stochastic Gravitational-Wave Background

To date, Advanced LIGO and Virgo have published the detection of gravitational waves from eleven compact binary mergers [25, 40, 42–46]. The most distant of these, GW170729, likely occurred at a luminosity distance of 2.8 Gpc [45]; its gravitational-wave signal traveled for roughly six billion years before arriving at Earth. Although this may seem like a vast distance, today's gravitational-wave detectors generally probe only the relatively local Universe. Advanced LIGO, in its present O3 sensitivity, can detect $10+10\,M_\odot$ black hole mergers out to a redshift of $z \sim 0.2$, a distance that pales in comparison to the redshifts $z \sim 10$ that can be studied using electromagnetic telescopes.

For every single gravitational-wave source detected by Advanced LIGO and Virgo, there are a multitude of other sources spread throughout the more distant Universe. Some of these sources are astrophysical, like remote compact binary mergers. Others may be *cosmological*, like the gravitational reverberations left over from the very birth of the Universe. By the time they reach Earth, the gravitational waves from these distant sources are too weak to be detected. The net sum of all such signals, however, may nevertheless be *collectively* detectable in the form of a **stochastic gravitational-wave background** – the random gravitational-wave signal created by the superposition of all distant gravitational-wave sources.

3.1 Observing the Gravitational-Wave Background: A Toy Picture

In Ch. 2.5 above, we saw how matched filtering enables the detection of potentially weak gravitational wave signals. Matched filtering, however, requires a template waveform – we must know *a priori* the exact form of the signal we seek to detect. Unlike the deterministic gravitational-wave signals from single compact binary mergers, the stochastic gravitational-wave background is a random signal, formed from the incoherent superposition of a large number of individual events. The random, unpredictable nature of the gravitational-wave background means that it cannot be detected via standard matched filtering. Instead, we will attempt to detect the gravitational-wave background in the form excess cross-power between widely-separated detectors.

It is useful to first sketch how this **cross-correlation search** for the gravitational-wave background works; we'll go through this more carefully in Ch. 3.6 below. Consider two identical and co-located gravitational-wave detectors, each of which measure random independent noise realizations n_1 and n_2. Both detectors also measure a common signal h (also random) due to the presence of the gravitational-wave background. Together, the signals s_1 and s_2 measured by the two detectors are

$$\begin{aligned} s_1 &= n_1 + h \\ s_2 &= n_2 + h. \end{aligned} \tag{3.1}$$

If we knew n_1 or n_2 perfectly, we could detect the gravitational-wave background with one detector alone, measuring the total signal power $\langle s_1^2 \rangle = \langle n_1^2 \rangle + \langle h^2 \rangle$ and computing the difference $\langle h^2 \rangle = \langle s_1^2 \rangle - \langle n_1^2 \rangle$ between this total power and what we expect from noise alone. In practice, however, this approach is virtually impossible. From Fig. 2.9 above, we can estimate that the Advanced LIGO's Hanford and Livingston detectors have noise power $P(100\,\mathrm{Hz}) \sim 10^{-46}\,\mathrm{Hz}^{-1}$ at $100\,\mathrm{Hz}$. In comparison, we'll demonstrate below that gravitational-wave background may be composed of individual events with strain power $h_0^2 \lesssim 10^{-49}$ So the detection of the gravitational-wave background in this way would require us to know the noise power in our detector to better than one part in 10^3!

Instead of looking for excess power in a single detector, we must instead look for excess *cross-power* $C = \langle s_1 s_2 \rangle$ between our two hypothetical detectors:

$$\begin{aligned} C &= \big\langle (n_1 + h)(n_2 + h) \big\rangle \\ &\boxed{\approx \langle h^2 \rangle.} \end{aligned} \tag{3.2}$$

Note that, since n_1, n_2, and h are mutually independent and uncorrelated, all terms involving instrumental noise vanish at leading order, leaving us with a direct estimate of the strain power $\langle h^2 \rangle$ of the stochastic gravitational-wave background.

Similar to matched filtering, cross-correlation enables the detection of the gravitational-wave background even when it is far weaker than instrumental noise. Like we did in Ch. 2.5, we can build intuition by defining a back-of-the-envelope expression for the signal-to-noise ratio of a cross-correlation search for the gravitational-wave background. Let $\mathcal{S}$ and $\mathcal{N}$ be the expected values of C in the presence and absence of a gravitational-wave background, respectively. If the gravitational waves comprising the stochastic background have characteristic amplitude h_0 and period τ, then at leading order

$$\begin{aligned} \mathcal{S} &= \langle h^2 \rangle \\ &= \frac{1}{T} \int h^2 dt \\ &\sim \frac{1}{T} h_0^2 N \tau, \end{aligned} \tag{3.3}$$

where T is our total observation time and $N \sim fT \sim T/\tau$ is the number of gravitational-wave cycles (of frequency f) completed in this time. Meanwhile, in the presence of noise alone, we expect to measure cross-correlation

$$\begin{aligned} \mathcal{N} &= \langle n_1 n_2 \rangle \\ &= \frac{1}{T} \int n_1 n_2 dt. \end{aligned} \tag{3.4}$$

If the noise in our two detectors is truly independent, then this cross-correlation will vanish as we integrate for infinite time T. Sadly, though, we're limited to finite integration times. In this case, since the product $n_1 n_2$ is a random variable (that can be both positive and negative), the integral in Eq. (3.4) grows as a random walk with the number of cycles N.

$$\mathcal{N} \sim \frac{1}{T} n_0^2 \sqrt{N} \tau. \tag{3.5}$$

Together, the signal-to-noise ratio of the gravitational-wave background is

$$\begin{aligned} \mathrm{SNR} &\sim \frac{\mathcal{S}}{\mathcal{N}} \\ &\boxed{\sim \sqrt{fT} \left(\frac{h_0}{n_0} \right)^2 .} \end{aligned} \tag{3.6}$$

Even if the stochastic gravitational-wave background is much weaker than intrinsic detector noise, we can nevertheless confidently detect it via a cross-correlation search, provided we integrate for a time

$$T \gg \frac{1}{f} \left(\frac{n_0}{h_0} \right)^4 . \tag{3.7}$$

Compare Eq. (3.7) with Eq. (2.59), the integration time required for ordinary matched filtering. In a matched filter search, the integration time required to detect a signal is $T_{\text{matched}} \propto (n_0/h_0)^2$. In contrast, the integration time required for an excess cross-power search is $T_{\text{cross}} \propto (n_0/h_0)^4 \propto T^2_{\text{matched}}$. Compared to a matched filter search, cross-correlation searches evidently take much longer to detect a signal with the same characteristic amplitude h_0. This increased time is the price we unfortunately must pay in order to detect a *random* gravitational-wave signal versus a deterministic one.

3.2 Why the Gravitational-Wave Background?

Using Eq. (3.7), we can estimate the time required to detect a stochastic signal due to distant compact binary mergers. Let's use the binary neutron star GW170817 to estimate the characteristic amplitude h_0 of the gravitational-wave background. As we will see below, the gravitational-wave background is dominated by compact binary mergers at redshifts $z \gtrsim 1$, corresponding to luminosity distances of 7 Gpc and beyond. If GW170817 had strain of order 10^{-22} at a distance of 40 Mpc, then binary neutron stars at $z \sim 1$ have strain amplitudes $h_0 \sim 10^{-22} \left(\frac{40\,\text{Mpc}}{7\,\text{Gpc}}\right) \sim 6 \times 10^{-25}$. Meanwhile, assume that we are searching in Advanced LIGO's most sensitive band, at $f \sim 100\,\text{Hz}$ where $P(f) \sim 10^{-46}\,\text{Hz}^{-1}$ (see Fig. 2.9). Assuming a $\Delta f = 100\,Hz$ bandwidth, then the amplitude of competing instrumental noise is $n_0 \sim \sqrt{P(100\,\text{Hz})\Delta f} \sim 10^{-22}$. Substituting these estimates into Eq. (3.7) above, we find that detection of the gravitational-wave background requires us to integrate for time $T \gg 0.3\,\text{yr}$.

Based on this back-of-the-envelope calculation, it seems that detection of the stochastic gravitational-wave background might entail observation times on the order of *years*! This is a rather significant amount of time, requiring a large investment of both person-power and patience. Detection and analysis of the gravitational-wave background will nevertheless prove valuable for two reasons.

First, the gravitational-wave background is dominated by compact binaries at cosmological distances, well beyond the horizon of current gravitational-wave detectors.

One might expect that the gravitational-wave background is dominated by sources "just beyond earshot" – objects that LIGO and Virgo can't quite de-

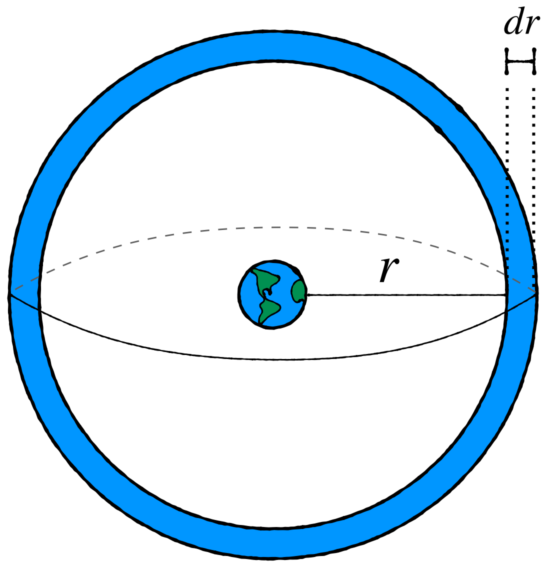

Figure 3.1: A spherical shell of thickness dr and radius r centered at the Earth. In Eqs. (3.8) and (3.9) we demonstrate that, in an isotropic Universe, every such shell contributes an equal SNR to a cross-correlation search for the gravitational-wave background, regardless of the shell's radius.

tect but that still sit relatively nearby in the local Universe. This isn't the case. The stochastic gravitational-wave signal measured in cross-correlation searches is dominated by the vast population of truly distant sources. The gravitational-wave background therefore offers a glimpse of a distinct population of gravitational-wave sources, complementary to those nearby sources that we can observe directly.

To understand why this is true, picture a static, isotropic Universe in which gravitational-wave sources are distributed uniformly in time and volume. Next, imagine dividing the Universe into a set of concentric spherical shells centered around the Earth, each shell with thickness dr; see Fig. 3.1. What is the stochastic signal-to-noise ratio $d(\mathrm{SNR})$ collectively contributed by all the sources in a given shell of radius r (i.e. events occurring at distances between r and $r + dr$)?

If individual sources at distance r each have a characteristic signal-to-noise ratio SNR_0, then the total SNR from the shell is just

$$d(\mathrm{SNR}) = \mathrm{SNR}_0 \, dN, \tag{3.8}$$

where dN is the total number of sources in the shell. From Eq. (3.6), we know that the characteristic stochastic signal-to-noise ratio of a single source scales quadratically with strain. Gravitational-wave strain, meanwhile, scales inversely with distance, and so we have $\mathrm{SNR}_0 \propto h_0^2 \propto r^{-2}$. Meanwhile, in our toy isotropic Universe the number of binaries within a shell is just proportional to the shell's volume: $dN \propto r^2 dr$. Together,

$$\begin{aligned} d(\mathrm{SNR}) &= \mathrm{SNR}_0\, dN \\ &\propto \left(r^{-2}\right)\left(r^2 dr\right) \\ &\boxed{\propto dr.} \end{aligned} \tag{3.9}$$

All factors of r have canceled! In this (naively simple) toy Universe, the stochastic signal-to-noise ratio of a shell does not depend on its radius; a shell of radius $r = 1\,\mathrm{Gpc}$ and a shell of radius $r = 100\,\mathrm{Gpc}$ will contribute equally to the stochastic gravitational-wave signal measured with a cross-correlation search.

An immediate corollary is that the gravitational-wave background is dominated by distant sources. For any specific radius r_0, we can ask whether shells *inside* or *outside* r_0 contribute more to the gravitational-wave background. As long as r_0 is finite, then there are infinitely many more shells beyond r_0 than within it; gravitational-wave sources at radii $r > r_0$ will therefore dominate, regardless of our choice for r_0 (this is essentially Olber's paradox in gravitational-wave form).

In actuality, of course, our Universe is neither static nor isotropic. Due to the cosmological expansion of the Universe, the true signal-to-noise ratio of single events falls more rapidly than $\mathrm{SNR}_0 \propto r^{-2}$ [67], and the Universe's evolving star formation rate causes the expected numbers of sources to peak and subsequently decrease at redshifts $z \gtrsim 1$ [82]. Nevertheless, we will see in more detail below that detection of the gravitational-wave background offers a glimpse of a truly distinct source population at far earlier times in the Universe's evolutionary history.

Second, cross-correlation searches for the stochastic background are model-agnostic, and so are potentially sensitive to new gravitational-wave sources and/or new physics.

In contrast to matched filtered searches for compact binary coalescences, the cross-correlation search for the gravitational-wave background is an *unmod-*

eled search. We search only for excess correlations between detectors, with minimal assumptions about the exact nature of the signal(s) comprising the gravitational-wave background. We just saw above that unmodeled searches can carry a penalty in the time required to detection. Still, unmodeled searches play a crucial role in gravitational-wave astronomy. By virtue of making minimal assumptions about the morphology or statistical character of a gravitational-wave signal, unmodeled searches set us up to discover the unexpected.

Stochastic searches, in particular, might serve to alert us to new types of gravitational-wave sources. As we will see below, we have reasonable predictions for the amplitude and shape of the gravitational-wave background due to standard astrophysical sources like compact binaries or rotating neutron stars. Should the gravitational-wave background be louder than expected, this would alert us to the presence of something unexpected. Alternatively, constraints on the gravitational-wave background provide a simple but powerful sanity check on newly-proposed and/or speculative gravitational-wave sources. It is relatively easy to explain why some new theorized source may so far have avoided direct detection – individual sources may be too quiet or we may lack the sufficiently precise templates needed to detect them via matched filtering. But any source population will generically produce a gravitational-wave background, detectable via cross-correlation of our detectors given sufficient observing time.

Unmodeled searches for the gravitational-wave background can also reveal unforeseen physical effects. We will see below that the exact features of the correlated signal encode a wide range of information about the gravitational-wave sky – its angular distribution, its frequency dependence, and the polarization properties of gravitational waves themselves. The gravitational-wave background can therefore be used as a tool with which to examine the very basic properties of gravitational waves, allowing us to search for and constrain deviations from the firm predictions made by general relativity.

3.3 Characterizing the Stochastic Background

So far our discussion of the gravitational-wave background has been purposefully qualitative. In this section, we'll make things more concrete, reviewing the tools used to quantitatively characterize the stochastic background.

Let $h_A(t, \hat{\mathbf{n}})$ be the net gravitational-wave signal with polarization A arriving at Earth at time t from direction $\hat{\mathbf{n}}$ on the sky. We can equivalently consider the corresponding Fourier domain quantity $\tilde{h}_A(f, \hat{n})$. This signal is due to the incoherent superposition of *all gravitational-wave sources in the Universe* in direction $\hat{\mathbf{n}}$, and so is effectively random. This randomness means that $\tilde{h}_A(f, \hat{\mathbf{n}})$ has no preferred complex phase – at any given instant we are equally likely to measure sine or cosine waves (of positive or negative amplitude). Hence the expectation value of $\tilde{h}_A(f, \hat{\mathbf{n}})$ is necessarily $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \rangle = 0$ for all nonzero f.

Since the expectation value of $\tilde{h}_A(f, \hat{\mathbf{n}})$ vanishes, we'll instead need to characterize the gravitational-wave background via the correlation $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}^*_{A'}(f', \hat{\mathbf{n}}') \rangle$ between the signals from directions $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$ with polarizations A and A' and frequencies f and f'. We'll assume that $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}^*_{A'}(f', \hat{\mathbf{n}}') \rangle$ is separable:

$$\boxed{\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}^*_{A'}(f', \hat{\mathbf{n}}') \rangle = \mathcal{A}(A, A') \Theta(\hat{\mathbf{n}}, \hat{\mathbf{n}}') \mathcal{H}(f, f'),} \tag{3.10}$$

with functions $\mathcal{A}(A, A')$, $\Theta(\hat{\mathbf{n}}, \hat{\mathbf{n}}')$, and $\mathcal{H}(f, f')$ independently characterizing correlations between polarization, sky direction, and frequency.

First consider the function $\mathcal{A}(A, A')$. As we saw in Ch. 2.3, the polarization of a gravitational-wave signal depends on extrinsic parameters like the angle at which we view the source. If gravitational-wave sources are randomly distributed and oriented across the sky, then it must be the case that (i) the strain power of the stochastic background is distributed equally between $+$ and $\times$ polarizations and (ii) that the signals measured in each polarization are uncorrelated. These two facts imply

$$\mathcal{A}(A, A') = \frac{1}{2} \delta_{A,A'}. \tag{3.11}$$

The leading factor of $\frac{1}{2}$ is chosen such that $\sum_{A,A'} \mathcal{A}(A, A') = 1$. Together, these two properties constitute the assumption that the gravitational-wave background is **unpolarized**.

Next, since the gravitational-wave signals from sufficiently different directions $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$ originate from causally disconnected sources, we will assume that the gravitational-wave background is **isotropic**, with

$$\Theta(\hat{\mathbf{n}}, \hat{\mathbf{n}}') = \frac{1}{4\pi} \delta^2(\hat{\mathbf{n}}, \hat{\mathbf{n}}'), \tag{3.12}$$

where $\delta^2(\hat{\mathbf{n}}, \hat{\mathbf{n}}') = \delta(\cos\theta - \cos\theta')\delta(\phi - \phi')$ is the Dirac delta function defined on the sphere.

Finally, consider the term $H(f, f')$. If the stochastic background comprises a large number of incoherent sources each instantaneously radiating at a different frequency f, then we should expect no correlations between the gravitational-wave signals at different frequencies. We'll therefore write

$$\mathcal{H}(f, f') = \frac{1}{2}\delta(f - f')\mathcal{H}(f). \tag{3.13}$$

All together, we have

$$\boxed{\langle \tilde{h}_A(f, \hat{\mathbf{n}})\tilde{h}^*_{A'}(f', \hat{\mathbf{n}}')\rangle = \frac{\delta_{A,A'}}{2}\frac{\delta^2(\hat{\mathbf{n}}, \hat{\mathbf{n}}')}{4\pi}\frac{\delta(f - f')}{2}\mathcal{H}(f).} \tag{3.14}$$

We can integrate Eq. (3.14) over all sky directions, polarizations, and frequencies to obtain

$$\int_0^\infty \mathcal{H}(f)\, df = \sum_{A,\, A'} \iint_{\text{Sky}} d\hat{\mathbf{n}}\, d\hat{\mathbf{n}}' \iint_{-\infty}^{\infty} df\, df' \,\langle \tilde{h}_A(f, \hat{\mathbf{n}})\tilde{h}_{A'}(f', \hat{\mathbf{n}}')\rangle. \tag{3.15}$$

So $\mathcal{H}(f)$ is the **strain power** in the gravitational-wave background – the net power received from all sky directions and polarizations per unit frequency. $\mathcal{H}(f)$ is both real and positive, obeying $\mathcal{H}(-f) = \mathcal{H}(f)$. The factor of $\frac{1}{2}$ in Eq. (3.13) means that $\mathcal{H}(f)$ is a one-sided power, normalized over positive frequencies as in Eq. (3.15).

If we wished, we could stop here and work entirely in terms of the strain power spectrum $\mathcal{H}(f)$ of the gravitational-wave background. Conventionally, however, the gravitational-wave background is described not in terms of its strain power, but by its **dimensionless energy density**

$$\boxed{\Omega(f) = \frac{1}{\rho_c}\frac{d\rho_{\text{GW}}}{d\ln f}.} \tag{3.16}$$

Here $d\rho_{\text{GW}}$ is the energy-density in gravitational-waves per logarithmic frequency interval $d\ln f$, and $\rho_c = 3c^2H_0^2/8\pi G$ is the critical energy density required to close the Universe. H_0 is the Hubble constant.

We can find the relationship between the gravitational-wave background's energy density $\Omega(f)$ and its strain power $\mathcal{H}(f)$ using Eq. (2.10) for the energy density of a gravitational wave:

$$\rho_{\text{GW}} = \frac{c^2}{32\pi G}\left\langle \dot{h}_{\alpha\beta}\dot{h}^{\alpha\beta}\right\rangle. \tag{3.17}$$

In the Fourier domain, $\dot{h}_{\alpha\beta} = 2\pi i f h_{\alpha\beta}$, and so

$$\begin{aligned} \rho_{\text{GW}} &= \frac{c^2}{32\pi G}\left\langle \dot{h}_{\alpha\beta}\dot{h}^{\alpha\beta}\right\rangle \\ &= \frac{\pi c^2 f^2}{8G}\left\langle h_{\alpha\beta}h^{\alpha\beta}\right\rangle \\ &= \frac{\pi c^2 f^2}{8G}\left\langle 2h_+^2 + 2h_\times^2\right\rangle \\ &= \frac{\pi c^2 f^2}{4G}\mathcal{H}(f). \end{aligned} \tag{3.18}$$

Note that we've used our assumption that the background is unpolarized to write $\left\langle h_+^2\right\rangle = \left\langle h_\times^2\right\rangle = \frac{1}{2}\mathcal{H}(f)$. Substituting Eq. (3.18) into Eq. (3.16), we obtain[1]

$$\Omega(f) = \frac{2\pi^2}{3H_0^2} f^3 \mathcal{H}(f), \tag{3.19}$$

or

$$\boxed{\langle \tilde{h}_A(f,\hat{\mathbf{n}})\tilde{h}_{A'}(f',\hat{\mathbf{n}}')\rangle = \frac{3H_0^2}{32\pi^3}\delta_{A,A'}\delta^2(\hat{\mathbf{n}},\hat{\mathbf{n}}')\delta(f-f')f^{-3}\Omega(f).} \tag{3.20}$$

Equation (3.20) has two features worth mentioning. First, in converting from gravitational-wave strain power to energy density $\Omega(f)$, we've acquired a factor of f^{-3}. Two powers of f are due to the time derivatives in the definition Eq. (3.17) of a gravitational wave's energy density. The remaining power of f is purely conventional, due to our choice to work with the energy density $\frac{d\rho_{\text{GW}}}{d\ln f}$ per *logarithmic* frequency interval.

Second, Eq. (3.20) depends on the Hubble constant H_0. The appearance of the Hubble constant is a not-infrequent source of some consternation – why should cosmology have anything to do with purely *local* measurements of the gravitational-wave strain power at Earth? This, again, is simply a widespread convention. Historically, the primary target of stochastic searches by ground-based detectors was the cosmological background due to fluctuations and/or phase transitions in the very early Universe. In this context it is sensible to describe the stochastic background using the language of cosmology. We now know that the gravitational-wave background in the frequency band of Advanced LIGO and Virgo is instead likely dominated by *astrophysical* sources,

[1] Note that different authors use different definitions when defining $\mathcal{H}(f)$. In the standard Ref. [83], for instance, $\mathcal{H}(f)$ is defined via $\langle \tilde{h}_A(f,\hat{\mathbf{n}})\tilde{h}^*_{A'}(f',\hat{\mathbf{n}}')\rangle = \delta_{A,A'}\delta^2(\hat{\mathbf{n}},\hat{\mathbf{n}}')\delta(f-f')\mathcal{H}(f)$, giving $\Omega(f) = \frac{32\pi^3}{3H_0^2}f^3\mathcal{H}(f)$. Every author should agree on Eq. (3.20), however, regardless of convention.

not cosmological effects [26, 48, 84]. The conventional parametrization in terms of $\Omega(f)$, however, has persisted.

There are two additional assumptions that are typically made regarding the gravitational-wave background. First, we will assume that the gravitational-wave background is *Gaussian*, such that the random time series $h_A(t, \hat{\mathbf{n}})$ is characterized completely by its mean (zero) and variance. Second, the gravitational-wave background is assumed to be *stationary*, with statistical properties that do not change over time.

To summarize everything in one place, we have assumed the following about the gravitational-wave background:

- Unpolarized: Power is distributed equally among gravitational-wave polarizations; different polarizations are uncorrelated.
- Isotropic: Average strain power is identical in all directions.
- No spectral correlations: The strain received at frequency f is independent of the instantaneous strain at some other $f' \neq f$.
- Gaussianity: The gravitational-wave background is completely characterized by a mean and variance.
- Stationarity: The statistical properties of the background are unchanging with time.

Most of these are very reasonable. The lack of net large-scale polarization and the absence of spectral and/or polarization correlations are likely very good assumptions. It is difficult to imagine mechanisms that would introduce correlations between frequencies. And only fairly exotic mechanisms, like parity violation in the early Universe [85] or birefringence in gravitational-wave propagation [86–90], can impart the background with significant correlations and/or asymmetry between polarizations. Our assumption of stationarity is similarly robust. The duration of our observations (years) is vanishingly small compared to the age of the Universe. Any evolution in the properties of the gravitational-wave background are therefore imperceptible to us.

Isotropy is also reasonable. The solar system moves at speed $v_\oplus \approx 370\,\mathrm{km/s}$ with respect to the cosmic microwave background [91], and hence we might expect large-scale dipole anisotropy across the sky that is a factor of $v_\oplus/c \sim 10^{-3}$

weaker than the overall isotropic monopole moment. Structure formation in the Universe additionally means that astrophysical gravitational-wave sources are not *quite* truly isotropically distributed [92–97]. Although significant disagreement exists between different authors regarding the exact strength of this astrophysical anisotropy [98, 99], all authors agree that it is small – likely no more than a 10% effect for Advanced LIGO and Virgo. Equally significant is anisotropy due to astrophysical "shot noise." Because we can observe the gravitational-wave sky for only a finite period of time, we see only some small, random fraction of all astrophysical sources. The random placement of these sources is subject to considerable variance, which manifests as an apparent anisotropy that masks true anisotropy due to cosmological structure [100, 101].

In contrast, the gravitational-wave background is likely not very Gaussian. The Gaussianity of the gravitational-wave background depends on the duration of individual signals relative to the average time *between* signals. If the signals' durations are much longer than the time between them, then at any given instant many individual sources overlap; the central limit theorem then implies that their collective strain is a Gaussian random variable. If, on the other hand, the signals' durations are much *shorter* than the time between them, then on average no more than one source is active at any given instant. The resulting strain will be decidedly non-Gaussian, characterized non-trivially by higher order moments beyond a mean and variance. As we will see below, the astrophysical background due to binary black holes very likely falls in this non-Gaussian regime. While cross-correlation searches of the type described in Sects. 3.1 and 3.6 are optimal for a truly Gaussian signal, they are *non-optimal* for non-Gaussian backgrounds; the development of more efficient methods for detecting non-Gaussian backgrounds is therefore currently the subject of much discussion [102–108].

3.4 Modeling the Gravitational-Wave Energy Density

Next we'll review how to actually calculate $\Omega(f)$ for a given gravitational-wave source. Since the stochastic gravitational-wave background is due to the superposition of sources throughout the Universe's history, we will need to account for the effects of cosmology. It is crucial to do this carefully – historically the exact relationship between present-day observables and source-frame emission has been the cause of confusion [109].

Let $\frac{dE_s}{df_s}$ be the energy per unit frequency radiated by a single gravitational-wave source, as viewed in the source's frame. If $\frac{dN}{dV_c dz}$ is the number of such sources per comoving volume and unit redshift, then the net gravitational-wave energy density in the present-day Universe is given by the integral

$$\frac{d\rho_{\mathrm{GW}}}{df}(f) = \int dz \, \frac{dN}{dV_c dz} \frac{dE_s}{df_s}\Big|_{f(1+z)}. \tag{3.21}$$

Note that the energy spectrum $\frac{dE_s}{df_s}$ is evaluated at frequency $f(1+z)$; this accounts for the cosmological redshifting of a gravitational wave's frequency between its emission (at redshift z) and its detection (today at redshift $z=0$).

The exact dependence of Eq. (3.21) on $\frac{dE_s}{df_s}$ is a frequent source of confusion, and so it is worth explaining this further. Consider an infinitesimal packet of source-frame energy dE_s radiated between source-frame frequencies f_s and $f_s + df_s$ at redshift z. As this packet propagates over cosmological distances, how is it affected by the expansion of the Universe? First, the packet is redshifted in frequency – observers today will measure the packet's frequency range to span $\frac{f_s}{1+z}$ and $\frac{f_s+df_s}{1+z}$. Thus the observed bandwidth of the packet is $df = \frac{df_s}{1+z}$. Meanwhile, the packet's energy is also redshifted, with observers today measuring energy $dE = \frac{dE_s}{1+z}$. This implies, however, the energy $\frac{dE}{df}$ per unit frequency is *invariant* – observers today measure

$$\begin{aligned} \frac{dE}{df} &= \frac{dE_s/(1+z)}{df_s/(1+z)} \\ &= \frac{dE_s}{df_s}. \end{aligned} \tag{3.22}$$

Hence it is the *source-frame* energy spectrum that appears in Eq. (3.21). We cannot completely disregard cosmological redshifting, though. Although our hypothetical packet's energy per unit frequency has remained constant, its central frequency has still been redshifted down to $\frac{f_s}{1+z}$. So if we wish to know the energy radiated at (present-day) frequency f, we must therefore evaluate $\frac{dE_s}{df_s}$ at source-frame frequency $f_s = f(1+z)$.

Next, we can recast the number $\frac{dN}{dV_c dz}$ of sources per comoving volume and redshift in terms of $\mathcal{R}(z) = \frac{dN}{dV_c dt_s}$, the source-frame merger rate per comoving volume (also called the source's **rate density**):

$$\begin{aligned}\frac{dN}{dV_c dz} &= \frac{dN}{dV_c dt_s}\frac{dt_s}{dz} \\ &= \frac{dN}{dV_c dt_s}\frac{1}{(1+z)H(z)} \\ &\equiv \mathcal{R}(z)\frac{1}{(1+z)H(z)},\end{aligned} \tag{3.23}$$

where $H(z)$ is the Hubble parameter at redshift z:

$$H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_r(1+z)^4 + \Omega_\Lambda}. \tag{3.24}$$

Ω_m, Ω_r, and Ω_Λ are the dimensionless energy densities of matter, radiation, and dark energy, respectively. For most applications we'll encounter in later chapters, we are interested only in redshifts $z \lesssim 10$, and so in practice we can usually neglect the radiation term. Collecting and combining Eqs. (3.16), (3.21), and (3.23), we have [110]

$$\boxed{\Omega(f) = \frac{f}{\rho_c}\int dz \frac{\mathcal{R}(z)\frac{dE_s}{df_s}\big|_{f(1+z)}}{(1+z)H(z)}.} \tag{3.25}$$

Given the energy-density spectrum $\frac{dE_s}{df_s}$ and merger rate $\mathcal{R}(z)$ for a particular source, we can use Eq. (3.25) to very straightforwardly calculate the resulting gravitational-wave background.

3.5 Sources Contributing to the Gravitational-Wave Background

Here I'll give a brief overview of sources that potentially contribute to the stochastic gravitational-wave background in the frequency band of Advanced LIGO and Virgo.

3.5.1 Compact Binary Mergers

Although the gravitational-wave background contains contributions from every class of gravitational-wave source in the Universe, we expect it to be dominated primarily by distant compact binary mergers [26, 48, 84]. The stochastic backgrounds due to distant binary black holes and binary neutron stars are predicted to have energy-densities $\Omega_{\rm BBH}(25\,\mathrm{Hz}) = 5.3^{+4.2}_{-2.5}\times 10^{-10}$ and $\Omega_{\rm BNS}(25\,\mathrm{Hz}) = 3.6^{+8.4}_{-3.1}\times 10^{-10}$ [26, 48, 84].

These predictions each carry a near-100% uncertainty. This considerable uncertainty originates from two sources. The first is simple Poisson uncertainty in the local merger rates of compact binaries. With the detection of only ten binary black holes, their merger rate remains uncertain to an order of magnitude, with coalescence rate estimates ranging between approximately $10 - 110\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ [45, 111]. Binary neutron star rates are even more uncertain, with estimates between $110 - 3840\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ [45, 46]. The second source of error is the considerable systematic uncertainty involved in the extrapolation of observed *local* merger rates to far earlier times in the Universe [26, 112–116]. This extrapolation requires knowledge of the underlying formation rate of stellar progenitors [82, 117], the dependence of compact object formation on stellar metallicities, and the distribution of time delays between a compact object's birth and its eventual gravitational-wave driven merger. At present these systematic uncertainties are smaller than the simple Poisson uncertainties above [26, 48], but will become dominant as our knowledge of local merger rates improves with future detections.

Note that binary neutron stars and binary black holes are predicted to contribute roughly equal energy densities to the gravitational-wave background. Although binary neutron stars are lighter and hence weaker sources of gravitational radiation than binary black holes, they merge roughly ten times as frequently – coincidentally these effects almost exactly cancel out [48].

While the energy densities due to binary neutron stars and black holes are comparable, the gravitational-wave backgrounds due to each object have very different *statistical* natures. A binary neutron star system merges *somewhere* in the Universe roughly every 10 s [48]. The gravitational-wave signal from each such system, meanwhile, persists in the Advanced LIGO frequency band for approximately 200 s. Thus the binary neutron star background is highly *Gaussian* – at any instant, the data measured by Advanced LIGO and Virgo contain overlapping signals from tens of individual binary neutron star mergers. In contrast, binary black holes merge less frequently (between one and five mergers every ten minutes at frequencies accessible to LIGO and Virgo) and have short duration signals lasting only seconds. The result is a so-called **popcorn background** – long durations of gravitational-wave "silence" punctuated every $\sim 200\,\mathrm{s}$ by a lone binary black hole merger. The distinction between the statistically-Gaussian and "popcorn" backgrounds is illustrated in

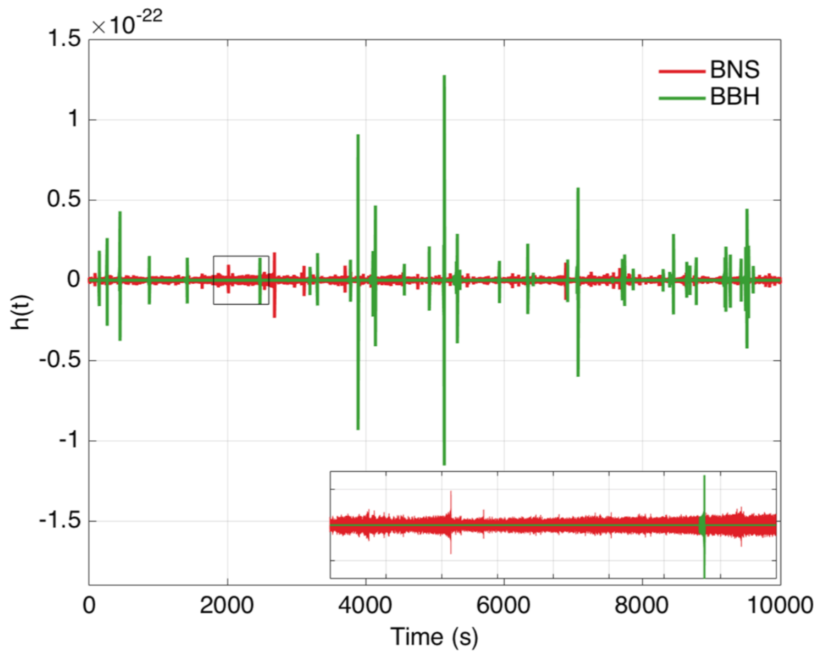

Figure 3.2: The net gravitational-wave strain at Earth due to all binary neutron star (red) and binary black hole (green) mergers. The inset shows a magnified view of the strain time series near 2000 s. Given the large rate and long duration of binary neutron star signals, the binary neutron star background is highly Gaussian, with many signals overlapping at any given instant. In contrast, binary black holes give rise to a non-Gaussian "popcorn" background in which each individual event is well-separated in time. Figure reproduced from Ref. [48].

Fig. 3.2, which shows the net strain at Earth due to neutron star (red) and black hole (green) mergers.

We also expec a background from neutron star-black hole mergers. Although no neutron star-black hole mergers have been detected, the upper limit of 610 $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ on their merger rate [45, 118] enables an upper limit of $\Omega(25\,\mathrm{Hz}) \lesssim 9\times10^{-10}$ on their energy density [84].

We will see in Ch. 4 a detailed calculation of the gravitational-wave background from compact binary mergers. Here, though, we can use Eq. (3.25) to heuristically compute the spectral *shape* of the energy-density spectrum from compact binaries. In the Fourier domain, the strain due to a compact binary merger is $\tilde{h}(f) \propto f^{-7/6}$ [67]. Hence a compact binary's energy spectrum is $\frac{dE}{df} \propto \dot{h}^2 \propto f^2 h^2 \propto f^{-1/3}$, and the resulting gravitational-wave background, at leading order, takes the form of a power-law:

$$\boxed{\Omega_{\mathrm{CBC}}(f) \propto f^{2/3}.} \tag{3.26}$$

This power law extends up to the frequency at which the compact binaries (at $z = 0$) merge, given approximately by [67]

$$f_{\text{Merge}} \approx 2.2\,\text{kHz}\left(\frac{M}{M_\odot}\right)^{-1}, \tag{3.27}$$

where M is the total mass of a binary.

3.5.2 Isolated Neutron Stars

Just as isolated rotating neutron stars generate sources of continuous gravitational waves, so too will they give rise to a gravitational-wave background [112, 119–131]. The exact strength of the predicted background depends strongly on many poorly-understood properties of neutron stars, including the configuration and strength of their internal magnetic fields [123, 125, 127] and the neutron star equation of state [131]. Generally, the most optimistic predictions correspond to gravitational-wave emission by young magnetars – neutron stars with internal fields of up to $10^{16} - 10^{17}$ G that might sustain ellipticities of $\epsilon \sim 10^{-4}$. Precise estimates vary from author to author, but the most optimistic magnetar scenarios generally correspond to a gravitational-wave background of amplitude $\Omega \lesssim 10^{-10}$ at 100 Hz [112, 123, 125, 129].

We can again use Eq. (3.25) to estimate the approximate form of the gravitational-wave background due to rotating neutron stars. First, the spin-down of a rotating neutron star can be written [131]

$$\begin{aligned}\dot{\omega} &= -\frac{B^2R^6\omega^3}{6Ic^3} - \frac{32G\epsilon^2I\omega^5}{5c^5} \\ &\equiv -\dot{\omega}_{\text{EM}} - \dot{\omega}_{\text{GW}}.\end{aligned} \tag{3.28}$$

The first term represents energy loss due to magnetic dipole radiation; B is the neutron star's dipole field, R is its radius, and I is its moment of inertia. The second term, meanwhile, represents energy carried away by gravitational-wave emission; ϵ is the star's rotationally or magnetically-induced quadrupole deformation (this expression assumes deformation along a principal axis normal to the star's spin axis [129]). Next, the gravitational-wave energy spectrum radiated by a single neutron star can be written

$$\frac{dE}{df} = \frac{dE}{dt}\frac{dt}{d\omega}\frac{d\omega}{df}, \tag{3.29}$$

where $\frac{dE}{dt}$ is the energy per unit time carried away via gravitational waves. If $E = \frac{1}{2}I\omega^2$ is the neutron star's total rotational energy, then the energy lost in the form of gravitational waves is just $\frac{dE}{dt} = I\omega\dot{\omega}_{\rm GW}$, where $\dot{\omega}_{\rm GW}$ is the second term in Eq. (3.28). We then have

$$\frac{dE}{df} = \frac{32\pi G}{5c^5}\frac{\epsilon^2 I^2 \omega^6}{|\dot{\omega}|}. \tag{3.30}$$

If the neutron star primarily loses energy through magnetic dipole radiation, then the first term in Eq. (3.28) dominates the star's overall spindown, giving $\dot{\omega} \propto \omega^3$. Then the energy spectrum of a single neutron star is $\frac{dE}{df} \propto \omega^3 \propto f^3$, and, from Eq. (3.25), the corresponding stochastic background has the approximate shape [112]

$$\boxed{\Omega_{\rm NS}(f) \propto f^4.} \tag{3.31}$$

On the other hand, if gravitational radiation dominates, then the overall spindown rate is $\dot{\omega} \propto \omega^5$, and the spectrum for a single source is $\frac{dE}{df} \propto f$. This is the case typically presumed for magnetars, which are thought to have (relatively) large magnetically-induced quadrupole deformations. The resulting gravitational-wave background in this case has shape [112]

$$\boxed{\Omega_{\rm Magnetar}(f) \propto f^2.} \tag{3.32}$$

Beyond gravitational waves from continuous solid-body rotation, a gravitational-wave background can in principle arise from other neutron star mechanisms, including abrupt spin "glitches" [126] and internal instabilities excited by gravitational-wave emission [119, 124, 130].

3.5.3 Core-Collapse Supernovae

The gravitational-wave background will also contain a contribution from distant core-collapse supernova [132–150]. Recall that in Ch. 2.3.3 we assumed that core collapse converts some fraction ϵ of a star's mass energy to gravitational radiation. Early studies presumed that this efficiency factor could be quite large, perhaps $\epsilon \sim 0.1$, and hence yield incredibly strong gravitational-wave backgrounds with $\Omega \sim 0.1$ [151–153]. We now know that gravitational-wave production is *much* weaker, with some simulations suggesting $\epsilon \sim 10^{-9}$ [77–79]. Hence the gravitational-wave background from stellar core collapse is likely not relevant for the current generation of ground-based detectors. Recent estimates project that the energy-density from core-collapse supernova might lie in the range $\Omega(f) \sim 10^{-10} - 10^{-9}$ between 10-100 Hz [147, 149]. These studies, however, presumed gravitational-wave production efficiencies in the range $\epsilon \sim 10^{-5}$; if ϵ is in fact closer to 10^{-9} then these projections could be four orders of magnitude too high.

At lowest order, we can think of the gravitational-wave emission from core-collapse as a white-noise burst, with a flat energy spectrum $\frac{dE}{df}$ = constant. Then, from Eq. (3.25), the resulting energy-density spectrum will scale as $\Omega(f) \propto f$. In practice, though, this is not a very good approximation. Simulated gravitational waveforms from supernova show many additional complex features that cause dE/df to deviate significantly from a flat spectrum, including a high-frequency steepening and exponential cutoff as well as amplified low-frequency emission [150].

3.5.4 Cosmological Gravitational Waves

Beyond stochastic signals from unresolvable astrophysical sources, so-called **cosmological backgrounds** are expected to arise from a variety of early Universe effects. Given that astrophysical backgrounds are due to a multitude of individual independent signals, it's clear why such backgrounds are necessarily stochastic. It's less clear why we should expect a gravitational-wave signal of cosmological origin to be stochastic in nature.

A nice argument explaining why a cosmological signal is inevitably stochastic is given in Ref. [154]. Consider an arbitrary process acting at redshift z that generates gravitational-wave emission, and let this gravitational-wave emission be spatially correlated across some distance l. This distance is necessarily smaller than the cosmological horizon at redshift z: $l \leq c/H(z)$. As viewed by an observer today, the angular extent of one such causally-connected "patch" is $\delta\Theta = l/D_A(z)$, where

$$D_A(z) = \frac{c}{1+z}\int_0^z \frac{dz'}{H(z')} \tag{3.33}$$

is the *angular diameter distance* [155]. Then the total number of such patches over the entire sky is $N \sim 4\pi\delta\Theta^{-2}$, or

$$N \sim 4\pi \frac{H^2(z)}{(1+z)^2}\left(\int_0^z \frac{dz'}{H(z')}\right)^2 . \tag{3.34}$$

Following Ref. [154], we might specifically imagine a hypothetical process active at the time of photon decoupling, at $z \approx 1090$. From Eq. (3.34), today we would see this process occurring across $N \sim 6 \times 10^4$ causally disconnected patches, each with an angular size of approximately 0.8 deg.

How does this compare to the angular resolution of the $\Delta x = 3000\,\mathrm{km}$ Advanced LIGO baseline? Choose a reference frequency of 1 kHz, corresponding to wavelengths $\lambda(1\,\mathrm{kHz}) \approx 300\,\mathrm{km}$. Then the maximum resolution achievable at this frequency is $\frac{\lambda}{\Delta x} \sim 6\,\mathrm{deg}$, far larger than the size of our causally-connected patches; every sky "pixel" resolvable by Advanced LIGO would contain emission from $\left(\frac{5.7\,\mathrm{deg}}{0.8\,\mathrm{deg}}\right) \sim 60$ different patches. At 100 Hz, meanwhile, Advanced LIGO's angular resolution is approximately 60 deg, and every pixel contains the superposition of over 5500 independent patches. By the central limit theorem, any such gravitational-wave signal generated in the early Universe is therefore necessarily stochastic.

There are three commonly-discussed cosmological sources of gravitational waves; see Refs. [154, 156, 157] for excellent reviews. First, inflation may have given rise to a cosmological background via the amplification of pre-inflation quantum fluctuations, or via the excitation of gravitational waves as inflation comes to a halt [158–161]. Gravitational waves may also be generated by various processes associated with a first-order phase transition, including sound waves and collisions between "bubbles" of high or low-temperature phases [162–167]. Finally, a gravitational-wave background may arise due to a population of gravitational-wave bursts from cosmic strings [84, 168–174].

3.6 Measuring the Gravitational-Wave Background

In Ch. 3.4 above, we characterized the stochastic gravitational-wave background in terms of the correlations $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}_{A'}(f', \hat{\mathbf{n}}') \rangle$ between gravitational-waves from different directions, of different polarizations, and at different frequencies. Under a fairly extensive set of assumptions, we also saw how these correlations could be expressed in terms of the energy-density $\Omega(f)$ of gravitational-waves. In this section, we will describe in greater detail how to actually *measure* the stochastic background with gravitational-wave detectors.

If gravitational-wave detectors were strongly directional – that is, if they could be precisely *pointed* towards specific directions on the sky – then we could measure $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}_{A'}(f', \hat{\mathbf{n}}') \rangle$ directly. Given two detectors, we would simply point them at directions $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$, rotate them such that they measure polarizations A and A', and then cross-correlate their outputs. Real gravitational-wave detectors, however, are decidedly *not* directional. On the one hand this is a blessing – a single detector can measure signals nearly omni-directionally across the sky. On the other hand, at any instant a single detector receives a superposition of signals from many different directions and polarizations. This makes the measurement of the gravitational-wave background a bit more complicated; we will need to do some work to relate the cross-power $\langle s_1 s_2 \rangle$ between the outputs of two gravitational-wave detectors to $\langle \tilde{h}_A(f, \hat{\mathbf{n}}) \tilde{h}_{A'}(f', \hat{\mathbf{n}}') \rangle$.

First, consider the output $s(t)$ of a single detector located at position $\mathbf{x}(t)$. We can expand $s(t)$ as the sum of instrumental noise $n(t)$ plus a set of plane gravitational waves spanning all directions $\hat{\mathbf{n}}$, polarizations A, and frequencies f (see Ref. [175] for a discussion of different Fourier conventions in stochastic searches):

$$s(t) = n(t) + \sum_A \int_{\mathrm{Sky}} d\hat{\mathbf{n}}\, F^A(\hat{\mathbf{n}}, t) \int_{-\infty}^{\infty} df\, \tilde{h}_A(f, \hat{\mathbf{n}}) e^{2\pi i f [t - \hat{\mathbf{n}} \cdot \mathbf{x}(t)/c]}. \tag{3.35}$$

Note the explicit time dependence in Eq. (3.35). The detector's location $\mathbf{x}(t)$ changes over time with the rotation and movement of the Earth. The Earth's rotation also means that the antenna patterns $F^A(\hat{\mathbf{n}}, t)$ are time-dependent. In principle, the movement of the Earth also Doppler shifts the apparent frequencies of gravitational-wave signals, but this effect is negligible in stochastic searches (in contrast, the correction of Doppler shifts and other higher order effects is crucial in searches for continuous waves from rotating neutron stars).

Ideally, we'd like to rewrite Eq. (3.35) in the frequency domain. Strictly speaking, this is made difficult by the time-dependent detector location and antenna patterns. We'll circumvent this problem altogether by taking the Fourier transform of Eq. (3.35) over a small time interval T; as long as T is much less than 24 hours we can assume that $F^A(\hat{\mathbf{n}}, t)$ and $\mathbf{x}(t)$ are effectively constant over the span of this interval. Current searches, for instance, take Fourier transforms over $T \sim 100\,\mathrm{s}$ windows [84, 176]. Treating the detector's antenna patterns and position as constants, the Fourier transform of its output $s(t)$ centered around time t, is

$$
\begin{aligned}
&\tilde{s}(f,t) \\
&= \int_{t-T/2}^{t+T/2} dt'\, s(t') e^{-2\pi i f t'} \\
&= \int_{t-T/2}^{t+T/2} dt'\, n(t') e^{-2\pi i f t'} \\
&\qquad + \sum_A \int_{t-T/2}^{t+T/2} dt' e^{-2\pi i f t'} \int d\hat{\mathbf{n}}\, F^A(\hat{\mathbf{n}}, t') \\
&\qquad \times \int df'\, \tilde{h}_A(f', \hat{\mathbf{n}}) e^{2\pi i f' [t' - \hat{\mathbf{n}}\cdot\mathbf{x}(t')/c]} \\
&\approx \tilde{n}(f) + \sum_A \int d\hat{\mathbf{n}}\, F^A(\hat{\mathbf{n}}, t) \int df'\, \tilde{h}_A(f, \hat{\mathbf{n}}) e^{-2\pi i f' \hat{\mathbf{n}}\cdot\mathbf{x}(t)/c} \\
&\qquad \times \int_{t-T/2}^{t+T/2} dt' e^{-2\pi i (f-f')t'} \\
&= \tilde{n}(f) + \sum_A \int d\hat{\mathbf{n}}\, F^A(\hat{\mathbf{n}}, t) \int_{-\infty}^{\infty} df'\, \tilde{h}_A(f', \hat{\mathbf{n}}) e^{-2\pi i f' \hat{\mathbf{n}}\cdot\mathbf{x}(t)/c} \delta_T(f - f').
\end{aligned}
\tag{3.36}
$$

In the third line, we've let $F^A(\hat{\mathbf{n}}, t') \approx F^A(\hat{\mathbf{n}}, t)$ and $\mathbf{x}(t') \approx \mathbf{x}(t)$. In the last line, meanwhile, we've used the definition Eq. (2.64) of the finite time delta function. If we were to let $T \to \infty$ then this would become a true delta function; if we further neglected the time dependence of the detector's position and antenna patterns, Eq. (3.36) would become

$$
\tilde{s}(f,t) = \tilde{n}(f) + \sum_A \int d\hat{\mathbf{n}}\, F^A(\hat{\mathbf{n}}, t) \tilde{h}_A(f, \hat{\mathbf{n}}) e^{-2\pi i f \hat{\mathbf{n}}\cdot\mathbf{x}/c}. \tag{3.37}
$$

This is called the **plane-wave expansion** of a measured gravitational-wave signal. We'll want to keep careful track of our finite time delta functions and their normalizing factors of T, though, so we'll continue working with Eq. (3.36).

Given two detectors, we can cross-correlate their outputs $s_1(t)$ and $s_2(t)$ and investigate how the result is related to the underlying gravitational-wave energy density. Multiplying two copies of Eq. (3.36), we have

$$\begin{aligned}
&\langle \tilde{s}_1(f,t)\tilde{s}_2^*(f',t)\rangle \\
&\quad= \Big\langle \Big[\sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t) \int_{-\infty}^{\infty} dk\, \tilde{h}_A(k,\hat{\mathbf{n}}) e^{-2\pi i k \hat{\mathbf{n}}\cdot\mathbf{x}_1(t)/c} \delta_T(f-k)\Big] \\
&\qquad \Big[\sum_{A'} \int d\hat{\mathbf{n}}'\, F_2^{A'}(\hat{\mathbf{n}}',t) \int_{-\infty}^{\infty} dk'\, \tilde{h}_A^*(k',\hat{\mathbf{n}}') e^{2\pi i k' \hat{\mathbf{n}}'\cdot\mathbf{x}_2(t)/c} \delta_T(f'-k')\Big]\Big\rangle \\
&\quad= \sum_A \sum_{A'} \iint d\hat{\mathbf{n}}\, d\hat{\mathbf{n}}' F_1^A(\hat{\mathbf{n}},t) F_2^{A'}(\hat{\mathbf{n}}',t) \iint dk\, dk' \delta_T(f-k)\delta_T(f'-k') \\
&\qquad\qquad \times \Big\langle \tilde{h}_A(k,\hat{\mathbf{n}}) h_A^*(k',\hat{\mathbf{n}}')\Big\rangle e^{-2\pi i k \hat{\mathbf{n}}\cdot\mathbf{x}_1(t)/c} e^{2\pi i k' \hat{\mathbf{n}}'\cdot\mathbf{x}_2(t)/c}.
\end{aligned} \tag{3.38}$$

To keep our notation somewhat manageable, we're using variables k and k' as dummy variables inside the two frequency integrals (as f'' and f''' would get rather unwieldy). Also, we've neglected the instrumental noise terms $\tilde{n}_1(f,t)$ and $\tilde{n}_2(f,t)$, since any terms involving the instrumental noise vanish when we take the expectation value. Substituting Eq. (3.14) for the expectation value $\langle \tilde{h}_A(k,\hat{\mathbf{n}}) h_A^*(k',\hat{\mathbf{n}}')\rangle$, we get

$$\begin{aligned}
&\langle \tilde{s}_1(f,t)\tilde{s}_2^*(f',t)\rangle \\
&\qquad= \frac{1}{16\pi} \sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t) F_2^A(\hat{\mathbf{n}},t) \iint dk\, dk' \delta_T(f-k)\delta_T(f'-k') \\
&\qquad\qquad \times \delta(k-k')\mathcal{H}(k) e^{-2\pi i k \hat{\mathbf{n}}\cdot\mathbf{x}_1(t)/c} e^{2\pi i k' \hat{\mathbf{n}}\cdot\mathbf{x}_2(t)/c} \\
&\qquad= \frac{1}{16\pi} \sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t) F_2^A(\hat{\mathbf{n}},t) \int dk\, \delta_T(f-k)\delta_T(f'-k) \\
&\qquad\qquad \times \mathcal{H}(k) e^{-2\pi i k \hat{\mathbf{n}}\cdot\Delta\mathbf{x}(t)/c},
\end{aligned} \tag{3.39}$$

where we've integrated over the true delta function $\delta(k-k')$ and defined the detector separation vector $\Delta\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_2$.

Now we need to deal with the two remaining finite time delta functions. Remember that $\delta_T(f-k)$ is peaked at $\delta_T(0) = T$, with a width of approximately $1/T$. If $T = 100\,\mathrm{s}$, then $\delta_T(f-k)$ is only $\sim 0.01\,\mathrm{Hz}$ wide. In contrast, the strain power $\mathcal{H}(k)$ of the gravitational-wave background is expected to vary much more slowly, changing significantly only over tens of Hz. What about the exponential? The exponential term changes significantly only over a frequency range $\Delta k \sim c/\Delta x$. If we take Δx to be the radius of the Earth ($6 \times 10^4\,\mathrm{km}$), then the exponential varies only over $\Delta k \sim 5\,\mathrm{Hz}$. Hence for any Earth-based detectors, the exponential too varies much more slowly with frequency than $\delta_T(f-k)$. It's therefore a fairly reasonable approximation to treat one of the finite time delta functions like a true delta function and integrate over k to obtain:

$$\begin{aligned}
&\langle \tilde{s}_1(f,t)\tilde{s}_2^*(f',t)\rangle \\
&\quad \approx \frac{1}{16\pi}\sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t)F_2^A(\hat{\mathbf{n}},t)\int dk\, \delta(f-k)\delta_T(f'-k) \\
&\qquad\qquad \times \mathcal{H}(k)e^{-2\pi i k\hat{\mathbf{n}}\cdot\Delta\mathbf{x}(t)/c} \qquad (3.40) \\
&\quad = \boxed{\frac{\delta_T(f-f')}{16\pi}\mathcal{H}(f)\sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t)F_2^A(\hat{\mathbf{n}},t)e^{-2\pi i f\hat{\mathbf{n}}\cdot\Delta\mathbf{x}(t)/c}.}
\end{aligned}$$

Let's clean this up slightly by defining

$$\boxed{\Gamma(f) = \sum_A \int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}},t)F_2^A(\hat{\mathbf{n}},t)e^{-2\pi i f\hat{\mathbf{n}}\cdot\Delta\mathbf{x}(t)/c},} \qquad (3.41)$$

in terms of which the cross-correlation between our two detectors is

$$\boxed{\langle \tilde{s}_1(f,t)\tilde{s}_2^*(f',t)\rangle = \frac{\delta_T(f-f')}{16\pi}\Gamma(f)\mathcal{H}(f).} \qquad (3.42)$$

$\Gamma(f)$ is called the **overlap reduction function** between our two detectors. The overlap reduction function can be thought of as the "transfer function" between the gravitational-wave background's strain power $\mathcal{H}(f)$ and the cross-power measured between detectors. The overlap reduction function depends only on the detectors' geometry. It is maximized for identical co-located and co-oriented detectors and decreases with other baseline geometries, penalizing detectors that are separated or rotated with respect to one another. Note also that the overlap reduction function is *time-independent*. Because $F_1^A(\hat{\mathbf{n}},t)$, $F_1^A(\hat{\mathbf{n}},t)$, and $\Delta\mathbf{x}(t)$ rotate synchronously, the integral over all sky directions evaluates to the same result regardless of the exact time t at which we calculate $\Gamma(f)$. Hence we can just write our correlation as $\langle \tilde{s}_1 f)\tilde{s}_2^*(f')\rangle$, without any explicit reference to time; in this case, we presume that $F_1^A(\hat{\mathbf{n}})$, $F_1^A(\hat{\mathbf{n}})$, and $\Delta\mathbf{x}$ are mutually evaluated at some arbitrary time.

In the LIGO and Virgo literature it is more common to see the **normalized overlap reduction function** $\gamma(f) = \lambda\Gamma(f)$, where the normalization factor

$$\lambda = \left(\sum_A \int d\hat{\mathbf{n}}\, [F^A(\hat{\mathbf{n}})]^2\right)^{-1} \tag{3.43}$$

is defined such that identical co-located and co-oriented detectors have $\gamma(f) = 1$. For interferometric detectors with 90 deg opening angles like Advanced LIGO and Virgo, $\lambda = 5/8\pi$ giving

$$\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle = \frac{\delta_T(f-f')}{10}\gamma(f)\mathcal{H}(f). \tag{3.44}$$

Expressed in terms of the gravitational-wave energy density,

$$\boxed{\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle = \frac{3H_0^2}{20\pi^2}\delta_T(f-f')f^{-3}\gamma(f)\Omega(f).} \tag{3.45}$$

Let's now define a somewhat tailored cross-correlation statistic

$$\boxed{\hat{C}(f) = \frac{1}{T}\frac{20\pi^2}{3H_0^2}f^3\,\tilde{s}_1(f)\,\tilde{s}_2^*(f)} \tag{3.46}$$

Using Eq. (3.45), we can immediately see that the expectation value of $\hat{C}(f)$ is

$$\boxed{\langle \hat{C}(f)\rangle = \gamma(f)\Omega(f).} \tag{3.47}$$

That is, $\hat{C}(f)$ is an estimator of the energy-density of the gravitational-wave background, multiplied by the overlap reduction function. This is a slightly different statistic than that defined in most older references. Most searches for the gravitational-wave background are presented in terms of a measurement $\hat{Y}(f) = \hat{C}(f)/\gamma(f)$, defined such that $\langle \hat{Y}(f)\rangle = \Omega(f)$ is a direct estimator of the background's energy-density spectrum. In later chapters, though, it will eventually be to our advantage to work in terms of $\hat{C}(f)$, rather than $\hat{Y}(f)$.

What's the variance of $\hat{C}(f)$? Before calculating this, it is useful to take a quick digression and return to the definition of a detector's PSD $P(f)$:

$$\langle \tilde{n}(f)\tilde{n}^*(f')\rangle = \frac{1}{2}\delta(f-f')P(f). \tag{3.48}$$

How is this expression modified if we're considering data Fourier transformed only over a finite time T? First, write the detector's noise time series as an inverse Fourier transform:

$$n(t) = \int df\, \tilde{n}(f) e^{2\pi i f t}. \tag{3.49}$$

Now do a finite-time Fourier transform back into the frequency domain:

$$\begin{aligned}
\tilde{n}(f) &= \int_{-T/2}^{T/2} dt' n(t') e^{-2\pi i f t'} \\
&= \int_{-T/2}^{T/2} dt' \int df' \tilde{n}(f') e^{2\pi i f' t'} e^{-2\pi i f t'} \\
&= \int df' \tilde{n}(f') \int_{t-T/2}^{t+T/2} dt' e^{-2\pi i (f-f')t'} \\
&= \int df' \delta_T(f-f') \tilde{n}(f').
\end{aligned} \tag{3.50}$$

Then, again using k and k' as dummy frequencies,

$$\begin{aligned}
\langle \tilde{n}(f) \tilde{n}^*(f') \rangle &= \int dk \int dk'\, \delta_T(f-k) \delta_T(f'-k') \langle \tilde{n}(k) \tilde{n}^*(k') \rangle \\
&= \frac{1}{2} \int dk \int dk'\, \delta_T(f-k) \delta_T(f'-k') \delta(k-k') P(k) \\
&= \frac{1}{2} \int dk\, \delta_T(f-k) \delta_T(f'-k) P(k) \\
&\approx \frac{1}{2} \delta_T(f-f') P(f),
\end{aligned} \tag{3.51}$$

where in the final line we've again used our trick of replacing one of the finite time delta functions with a true Dirac delta. Comparing this result against Eq. (3.48), accounting for finite Fourier transforms amounts (not surprisingly) to replacing the Dirac delta with a finite time delta function.

Now we're ready to go back and calculate the variance of $\hat{C}(f)$:

$$\begin{aligned}
&\langle \hat{C}(f) \hat{C}^*(f') \rangle - \langle \hat{C}(f) \rangle \langle \hat{C}^*(f') \rangle \\
&= \frac{1}{T^2} \left(\frac{20\pi^2}{3H_0^2} \right)^2 f^6 \left\langle \tilde{s}_1(f)\, \tilde{s}_2^*(f)\, \tilde{s}_1^*(f')\, \tilde{s}_2(f') \right\rangle - \gamma(f)\Omega(f)\gamma(f')\Omega(f') \\
&= \frac{1}{T^2} \left(\frac{20\pi^2}{3H_0^2} \right)^2 f^6 \left\langle \tilde{s}_1(f)\, \tilde{s}_1^*(f') \right\rangle \left\langle \tilde{s}_2(f')\, \tilde{s}_2^*(f) \right\rangle - \gamma(f)\Omega(f)\gamma(f')\Omega(f') \\
&= \frac{1}{T^2} \left(\frac{20\pi^2}{3H_0^2} \right)^2 f^6 \left(\frac{1}{2} \delta_T(f-f') P_1(f) \right) \left(\frac{1}{2} \delta_T(f-f') P_2(f) \right) \\
&\qquad - \gamma(f)\Omega(f)\gamma(f')\Omega(f') \\
&= \delta(f-f') \sigma^2(f) - \gamma(f)\Omega(f)\gamma(f')\Omega(f'),
\end{aligned} \tag{3.52}$$

where

$$\boxed{\sigma^2(f) = \frac{1}{T}\left(\frac{10\pi^2}{3H_0^2}\right)^2 f^6\, P_1(f)P_2(f).} \tag{3.53}$$

Note that one of the finite time delta-functions has been converted to a true Dirac delta; the other evaluated to $\delta_T(f-f) = T$. The first term is proportional to the product $P_1(f)P_2(f)$ between the detectors' PSDs; in general this is orders of magnitude larger than the background's strain power $H^2(f)$. We'll therefore usually approximate the variance of $\hat{C}$ as $\langle\hat{C}\hat{C}^*\rangle - \langle\hat{C}\rangle\langle\hat{C}^*\rangle \approx \langle\hat{C}\hat{C}^*\rangle = \delta(f-f')\sigma^2(f)$.

Current LIGO and Virgo searches for the gravitational-wave background operate via measuring $\hat{C}(f)$, searching for statistically significant correlations between instruments. Given a measured cross-correlation spectrum $\hat{C}(f)$ (and associated uncertainty $\sigma^2(f)$), how do we gauge if it is, in fact, statistically significant?

Define a *broadband* detection statistic

$$\hat{C}_B = \frac{\int df\, \hat{C}(f)\lambda(f)}{\int df'\, \lambda(f')}, \tag{3.54}$$

where $\lambda(f)$ is some weighting function. The expectation value and variance of $\hat{C}$ are, respectively,

$$\begin{aligned} \langle\hat{C}_B\rangle &= \frac{\int df\, \langle\hat{C}(f)\rangle\lambda(f)}{\int df'\, \lambda(f')} \\ &= \frac{\int df\, \gamma(f)\Omega(f)\lambda(f)}{\int df'\, \lambda(f')} \end{aligned} \tag{3.55}$$

and

$$\begin{aligned} \sigma_B^2 &= \langle\hat{C}_B^2\rangle - \langle\hat{C}_B\rangle^2 \\ &= \frac{\int df \int df' \left[\langle\hat{C}(f)\hat{C}(f')\rangle - \langle\hat{C}(f)\rangle\langle\hat{C}(f')\rangle\right]\lambda(f)\lambda(f')}{\left(\int df''\, \lambda(f'')\right)^2} \\ &\approx \frac{\int df \int df' \delta(f-f')\sigma^2(f)\lambda(f)\lambda(f')}{\left(\int df''\, \lambda(f'')\right)^2} \\ &= \frac{\int df\, \sigma^2(f)\lambda^2(f)}{\left(\int df'\, \lambda(f')\right)^2}. \end{aligned} \tag{3.56}$$

The expected broadband signal-to-noise ratio is

$$\begin{aligned}\langle \mathrm{SNR} \rangle &= \frac{\langle \hat{C}_B \rangle}{\sqrt{\sigma_B^2}} \\ &= \frac{\int df\, \gamma(f)\Omega(f)\lambda(f)}{\sqrt{\int df\, \sigma^2(f)\lambda^2(f)}}.\end{aligned} \tag{3.57}$$

We can now choose $\lambda(f)$ to maximize this signal-to-noise ratio. We'll use the same trick that we saw when choosing the optimal filter in Ch. 2.5. If we define the inner product

$$\begin{aligned}(A \,|\, B) &= \int_{-\infty}^{\infty} df\, \frac{\tilde{A}(f)\tilde{B}(f)}{\sigma^2(f)} \\ &= \left(\frac{3H_0^2}{10\pi^2}\right)^2 T \int_{-\infty}^{\infty} df\, \frac{\tilde{A}(f)\tilde{B}(f)}{f^6 P_1(f) P_2(f)},\end{aligned} \tag{3.58}$$

then Eq. (3.57) can be written in the form

$$\langle \mathrm{SNR} \rangle = \frac{(\gamma\Omega \,|\, \lambda\sigma^2)}{\sqrt{(\lambda\sigma^2 \,|\, \lambda\sigma^2)}}. \tag{3.59}$$

Using the same argument as in Ch. 2.5, the SNR is maximized when $\lambda(f)\sigma^2(f)$ and $\gamma(f)\Omega(f)$ are "parallel" – i.e. when

$$\lambda(f) = \frac{\gamma(f)\Omega(f)}{\sigma^2(f)}. \tag{3.60}$$

Thus the **optimal SNR** of a stochastic signal is

$$\boxed{\begin{aligned}\mathrm{SNR_{opt}} &= \sqrt{(\gamma\Omega \,|\, \gamma\Omega)} \\ &= \left[\left(\frac{3H_0^2}{10\pi^2}\right)^2 T \int_{-\infty}^{\infty} df\, \frac{\gamma^2(f)\Omega^2(f)}{f^6 P_1(f) P_2(f)}\right]^{1/2}.\end{aligned}} \tag{3.61}$$

Just like optimal matched filtering required us to know in advance the specific signal for which we were searching, here too we need to know the spectral shape of $\Omega(f)$ in order to optimally choose the filter function $\lambda(f)$. If we've yet to actually *detect* the gravitational-wave background, then obviously we can't already know $\Omega(f)$. In practice, we therefore have to choose some model $\Omega_M(f)$ for the background's energy-density spectrum. Given cross-correlation measurements $\hat{C}(f)$ and a model spectrum $\Omega_M(f)$, the measured signal-to-noise ratio is

$$\boxed{\mathrm{SNR} = \frac{\left(\hat{C}\,|\,\gamma\Omega_M\right)}{\sqrt{(\gamma\Omega_M\,|\,\gamma\Omega_M)}}.} \tag{3.62}$$

This analysis has so far relied on an implicit assumption regarding the stability of instrumental noise – namely, that it is **stationary**, with PSDs that do not evolve in time. While stationarity is a very good assumption regarding the power of the gravitational-wave background (see Ch. 3.3), it is a particularly poor assumption for detector noise, which can vary significantly over days, hours, and even minutes. We cannot, therefore, perform a single cross-correlation with time T equal to, say, the length of an observing run. Instead, we must measure a *sequence* of many cross-correlations $\hat{C}(f,t)$ (with associated variances $\sigma^2(f,t)$), each taken over a sufficiently short time T that the instrumental noise can be considered approximately stationary. In the O1 and O2 observing runs, for instance, cross-correlations were measured over individual segments of duration $T = 192\,\mathrm{s}$ [84, 176].

Given multiple measurements $\hat{C}(f,t_i)$ of a correlation spectrum at multiple times t_i, how do we combine them to into a single spectrum? We've inadvertently already answered this question. Look back to Eqs. (3.54), (3.56), and (3.60). Together, these equations told us how to optimally combine cross-correlations measured at different frequencies. An exactly analogous argument holds when combining measurements at different *times*. Assuming a temporally "flat" model spectrum $\Omega(t) = \text{constant}$, the optimal combination of different time segments is given by

$$\boxed{\hat{C}(f) = \frac{\sum_i \hat{C}(f,t_i)/\sigma^2(f,t_i)}{\sum_i 1/\sigma^2(f,t_i)},} \tag{3.63}$$

with variance

$$\boxed{\frac{1}{\sigma^2(f)} = \sum_i \frac{1}{\sigma^2(f,t_i)}.} \tag{3.64}$$

3.7 Looking Ahead

This section introduced us to the stochastic gravitational-wave background. I argued qualitatively that the gravitational-wave background offers a glimpse at the most distant gravitational-wave sources – remote astrophysical sources and cosmological gravitational waves from the early Universe. In Ch. 3.3 I formally defined the stochastic signal in terms of its power spectrum, or, equivalently, its dimensionless gravitational-wave energy density. In Ch. 3.4 I derived a generic prescription for modeling the energy-density of the gravitational-wave background, and in Ch. 3.5 we saw an overview of the astrophysical and cosmological sources that may contribute to the stochastic background. Finally, in Ch. 3.6 we learned how cross-correlation measurements between gravitational-wave detectors are used in an attempt to detect and characterize a stochastic gravitational-wave signal.

In the next several chapters I will apply these principles to further investigate several aspects of the stochastic gravitational-wave background, its interpretation, and its detection. First, in Ch. 4, I will quantify how astrophysical sources at different redshifts contribute to the stochastic signal measured by Advanced LIGO. I will also examine exactly which spectral properties of the background's energy-density spectrum are measurable with present-day instruments. Ch. 5, meanwhile, will explore what the gravitational-wave background might teach us about fundamental physics – namely, the polarizations of gravitational waves. In Ch. 6 I will then present searches for the stochastic gravitational-wave background, using data from Advanced LIGO's O1 and O2 observing run. Finally, in Ch. 7 I will consider how we might, one day, seek to unambiguously *validate* a tentative detection of the stochastic background.

Chapter 4
Astrophysical Information Contained in the Gravitational-Wave Background

∾

This chapter contains work published in:

> T. Callister, L. Sammut, S. Qiu *et al.*, *Limits of Astrophysics with Gravitational-Wave Backgrounds*, Phys. Rev. X **6**, 031018 (2016).

I produced all results and wrote the majority of the published manuscript.

At the time this work was published, Advanced LIGO had detected only one unambiguous gravitational-wave event: GW150914. Besides this one event, we knew little about the broader population of binary black holes, and so in Callister (2016) we had to simply presume that all binary black hole mergers had masses comparable to GW150914's. We now know, of course, that this is not true. However, when presenting the results below, I'll temporarily neglect the existence of these other gravitational-wave detections and, as in Callister (2016), restrict purely to GW150914. The main findings of this chapter, though, will nevertheless remain valid when we again consider the broader population of black holes.

∾

The detection of an astrophysical stochastic background would be a major accomplishment for the gravitational-wave community, providing us a glimpse of sources at truly cosmological distances. As we saw in Ch. 3, the energy-density spectrum $\Omega(f)$ of the gravitational-wave background depends on the detailed astrophysics underlying gravitational-wave sources – their intrinsic properties and their distribution across redshifts.

Even though the gravitational-wave background depends on this rich astrophysics, it is not clear how well we can actually *extract* this information given a measurement of the gravitational-wave background. In this chapter I will explore this problem, investigating the following three questions:

First, how does the information contained in the gravitational-wave background compare to what we learn from resolvable binaries in the local Universe? In Ch. 4.1, we will find that the stochastic signal due to binary black hole mergers is dominated by sources between redshifts $z \approx 0.1 - 3.5$. Observations of the stochastic background will therefore probe a binary black hole population that is truly distinct from directly resolvable sources in the local Universe.

Second, which spectral features of the gravitational-wave background can we expect to successfully measure? In Ch. 4.2, I will demonstrate that, while second generation gravitational wave detectors may successfully measure the amplitude of the stochastic background, it is difficult to further distinguish higher-order features of the background's energy density spectrum, such as its spectral index. This means that, unfortunately, our ability to study detailed astrophysics and cosmology using the gravitational-wave background *alone* is somewhat limited.

Third, how does the presence of binary black holes affect our ability to measure other potentially interesting contributions to the gravitational-wave background? In Ch. 4.3, we will see that the astrophysical background due to binary black hole mergers acts as a limiting foreground, significantly obscuring the presence of other astrophysical or cosmological sources.

4.1 Information contained in the astrophysical gravitational-wave background

In Ch. 3.4 above, we derived a general prescription for modeling the energy-density spectrum of the gravitational-wave background:

$$\Omega(f) = \frac{f}{\rho_c} \int dz \frac{\mathcal{R}(z) \frac{dE_s}{df_s}\big|_{f(1+z)}}{(1+z)H(z)}, \tag{4.1}$$

where $\mathcal{R}(z)$ is the source-frame event rate per comoving volume and $\frac{dE_s}{df_s}$ is the source-frame energy spectrum of a single event. Let's begin by applying this prescription to the specific case of binary black hole mergers. Specifically, we will follow the so-called FIDUCIAL model for the binary black hole background described in Ref. [26].

We will describe the gravitational-wave signals from individual binary black holes using the phenomenological model of Ref. [177], which provides a semi-analytic description of a binary black hole's inspiral, merger, and final ring-down. The resulting energy spectrum for a single binary is [114, 177]

$$\frac{dE_s}{df_s} = \frac{1}{3}\left(G\pi\right)^{2/3}\mathcal{M}_c^{5/3}e(f), \tag{4.2}$$

where

$$e(f) = \begin{cases} f^{-1/3} & (f < f_{\mathrm{merge}}) \\ \frac{f^{2/3}}{f_{\mathrm{merge}}} & (f_{\mathrm{merge}} \leq f < f_{\mathrm{ring}}) \\ \frac{1}{f_{\mathrm{merge}} f_{\mathrm{ring}}^{4/3}}\left(\frac{f}{1+\left(\frac{f-f_{\mathrm{ring}}}{\sigma/2}\right)^2}\right)^2 & (f_{\mathrm{ring}} \leq f < f_{\mathrm{cutoff}}) \\ 0 & (f \geq f_{\mathrm{cutoff}}) \end{cases}. \tag{4.3}$$

The merger, ringdown, and cutoff frequencies f_{merge}, f_{ring}, and f_{cutoff} as well as the ringdown bandwidth σ found by fitting to results from numerical relativity; numerical definitions for these fitted quantities are given in Sect. IV of Ref. [177].

We also need a prescription for $\mathcal{R}(z)$, the binary black hole merger rate as a function of redshift. Assuming that the black holes observed with Advanced LIGO are the remnants of stellar progenitors (rather than primordial black holes), then the binary merger rate is likely to trace the Universe's star formation rate, modulo a time delay between a main sequence binary's formation and the eventual gravitational-wave driven merger of its black hole remnants. The Fiducial gravitational-wave background model of Ref. [26], for example, adopts the star formation rate [117]

$$R_*(z) = \nu \frac{a\, e^{b(z-z_m)}}{a-b+b\, e^{a(z-z_m)}}\frac{M_\odot}{\mathrm{Mpc}^3\,\mathrm{yr}}, \tag{4.4}$$

with $\nu = 0.145$, $z_m = 1.86$, $a = 2.80$, and $b = 2.62$. Equation (4.4) is calibrated to the observed distribution of gamma-ray bursts [178].

It may also be the case that binary black holes are born preferentially in low-metallicity environments. Metal-rich stars are known to lose significant mass to stellar winds. Metal-poor stars, in contrast, retain a greater fraction of their initial mass over their lifetime, and hence may more readily yield massive black holes [179–182]. As in the Fiducial model, we will reweight Eq. (4.4) by the fraction $F(z)$ of stars formed redshift z with metallicities below $Z_\odot/2$, where $Z_\odot = 0.02$ is the solar metallicity.

As a function of redshift, the mean stellar metallicity is [82]

$$\overline{\log Z(z)} = 0.5 + \log\left(\frac{y(1-R)}{\rho_b}\int_z^{20}\frac{R_*^{\rm MD}(z')dz'}{H(z')(1+z')}\right), \tag{4.5}$$

with stellar metal yield $y = 0.019$, return fraction $R = 0.27$, baryon density $\rho_b = 2.77\times10^{11}\,\Omega_b\,h_0^2\,M_\odot\,\mathrm{Mpc}^{-3}$, and $\Omega_b = 0.045$. The star formation rate used in calibrating y and R is [82]

$$R_*^{\rm MD}(z) = 0.015\frac{(1+z)^{2.7}}{1+\left(\frac{1+z}{2.9}\right)^{5.6}}\frac{M_\odot}{\mathrm{Mpc}^3\,\mathrm{yr}}. \tag{4.6}$$

Assuming that stellar metallicities are log-normally distributed with a standard deviation of 0.5, the fraction of stars with metallicities below $Z_\odot/2$ is

$$F(z) = \sqrt{\frac{2}{\pi}}\int_{-\infty}^{\log Z_\odot/2} e^{-2\left(\log Z - \overline{\log Z(z)}\right)^2} d\log Z. \tag{4.7}$$

We have one final ingredient to consider – the time delay between a binary's formation and eventual merger. Let $p(t_d)$ be the probability distribution of time delays t_d. Then the merger rate per comoving volume is the convolution of the metallicity-weighted stellar formation rate $R_*(z)F(z)$ with $p(t_d)$:

$$\mathcal{R}(z) = \mathcal{R}_0\frac{\int_{t_{\rm min}}^{t_{\rm max}} R_*(z_f(t_d,z))\,F(z_f(t_d,z))\,p(t_d)\,dt_d}{\int_{t_{\rm min}}^{t_{\rm max}} R_*(z_f(t_d,0))\,F(z_f(t_d,0))\,p(t_d)\,dt_d}. \tag{4.8}$$

Here $z_f(t_d,z)$ is the formation redshift of a binary that later merges at redshift z after a delay t_d. Equation (4.8) is normalized such that the constant $\mathcal{R}_0$ is the local merger rate per comoving volume at redshift $z = 0$. The integration limits $t_{\rm min}$ and $t_{\rm max}$ are the minimum and maximum delay times over which we believe compact binaries can evolve to merger. We take $t_{\rm min} = 50\,\mathrm{Myr}$ and $t_{\rm max} = 13.5\,\mathrm{Gyr}$, the present-day age of the Universe. The shape of $p(t_d)$ itself is not very certain; we assume a power-law form [183]:

$$p(t_d) \propto \begin{cases} t_d^{-1} & (t_{\rm min}\leq t_d\leq t_{\rm max}) \\ 0 & (\text{else}) \end{cases}. \tag{4.9}$$

Figure 4.1 shows examples of energy-density spectra created using Eqs. (4.1), (4.2), and (4.8). In particular, we vary the average chirp mass $\mathcal{M}_c$ of the presumed binary black hole population. Also shown are power-law integrated (PI)

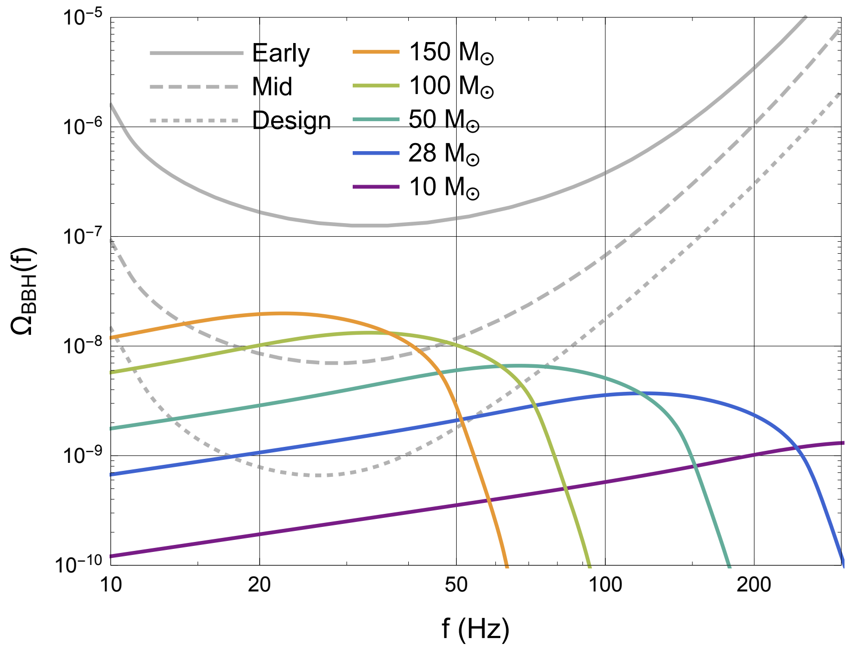

Figure 4.1: Energy-density spectra corresponding to binary black holes of various chirp masses, assuming a local coalescence rate of $\mathcal{R}_0 = 16\,\mathrm{Gpc}^{-3}\mathrm{yr}^{-1}$. Power-law integrated curves for one year of integration with Advanced LIGO at early, middle, and design-sensitivity are shown for comparison [184]. Approximately 95% of the signal-to-noise ratio comes from a band spanning $15-45\,\mathrm{Hz}$.

curves [184] indicating the sensitivity of the stochastic search after one year of integration with Advanced LIGO at early, middle, and design-sensitivity.[1] From Fig. 4.1 we can immediately see that the stochastic energy-density spectrum depends on black hole masses in two ways. First, since $\frac{dE_s}{df_s} \propto \mathcal{M}_c^{5/3}$, increasing $\mathcal{M}_c$ (at fixed R_0) causes the peak value of $\Omega(f)$ to correspondingly grow like $\mathcal{M}_c^{5/3}$. Second, since more massive binaries merge at lower frequencies, increasing chirp mass shifts the binary black hole energy-density leftward. Specifically, the knee frequency f_{max} at which the energy-density spectrum peaks scales as $f_{\mathrm{max}} \propto 1/\mathcal{M}_c$.

We can use this newly constructed model of the binary black hole background to carefully answer the following question:

If we detect the gravitational-wave background, exactly which systems in the Universe have we actually detected?

In Ch. 3.2 above, we argued very qualitatively that the gravitational-wave background contains information about truly distant gravitational-wave sources,

[1] Note that these were *projected* sensitivities, estimated in 2016 when these results were published. With the benefit of hindsight, these sensitivity projections were rather optimistic.

rather than the closest "sub-threshold" events just beyond the range of direct detection. In this argument we appealed to a toy Universe that was static and isotropic. Of course, the Universe is neither static nor isotropic. The Universe itself evolves due to cosmological expansion, and the distribution of binary black hole mergers is expected to trace the star formation rate via Eq. (4.8).

Recall that the optimal signal-to-noise ratio of a gravitational-wave background $\Omega(f)$ is

$$\mathrm{SNR} = \left[2T\left(\frac{3H_0^2}{10\pi^2}\right)^2 \int_0^\infty \frac{\gamma(f)^2\Omega(f)^2}{f^6 P_1(f)P_2(f)}df\right]^{1/2}. \tag{4.10}$$

To more rigorously quantify the SNR contribution from binaries at different redshifts, we can define an "SNR density"

$$\boxed{\frac{d(\mathrm{SNR})}{dz} = \frac{2T}{\mathrm{SNR}}\left(\frac{3H_0^2}{10\pi^2}\right)^2 \int_0^\infty \frac{\gamma^2(f)\Omega(f)\frac{d\Omega}{dz}(f,z)}{f^6 P_1(f)P_2(f)}df,} \tag{4.11}$$

with

$$\frac{d\Omega}{dz}(f,z) = \frac{f}{\rho_c}\frac{\frac{dE_s}{df_s}(f(1+z))\mathcal{R}(z)}{(1+z)H(z)}. \tag{4.12}$$

Equation (4.11) quantifies the stochastic signal-to-noise ratio due to sources between redshifts z and $z+dz$.

The stochastic SNR density for design-sensitivity Advanced LIGO is plotted as a function of z in Fig. 4.2 for several choices of chirp mass, assuming the FIDUCIAL binary black hole model described above. For purposes of comparison, each curve is normalized to a total signal-to-noise ratio SNR $= 1$. Also shown is the *cumulative* SNR, obtained by integrating $d(\mathrm{SNR})/dz$ up to some cutoff z. Within each subplot, the dashed vertical lines indicate threshold redshifts $z_{50\%}$ beyond which BBHs of each chirp mass (indicated by the respective colors) are individually resolvable less than 50% of the time. These threshold redshifts therefore indicate the typical range of a direct search for compact binary coalescences – binaries at redshifts $z < z_{50\%}$ are, on average, directly resolvable, while those at $z > z_{50\%}$ are not. In general, $z_{50\%}$ increases with binary chirp mass, as more massive systems are more readily detectable at greater distances. Note, however, that extremely massive $\mathcal{M}_c = 150\,M_\odot$ systems have a threshold redshift *lower* than that for $\mathcal{M}_c = 100\,M_\odot$ objects; this is due to the fact that gravitational-waves from very massive systems are rapidly redshifted out of Advanced LIGO's sensitivity band.

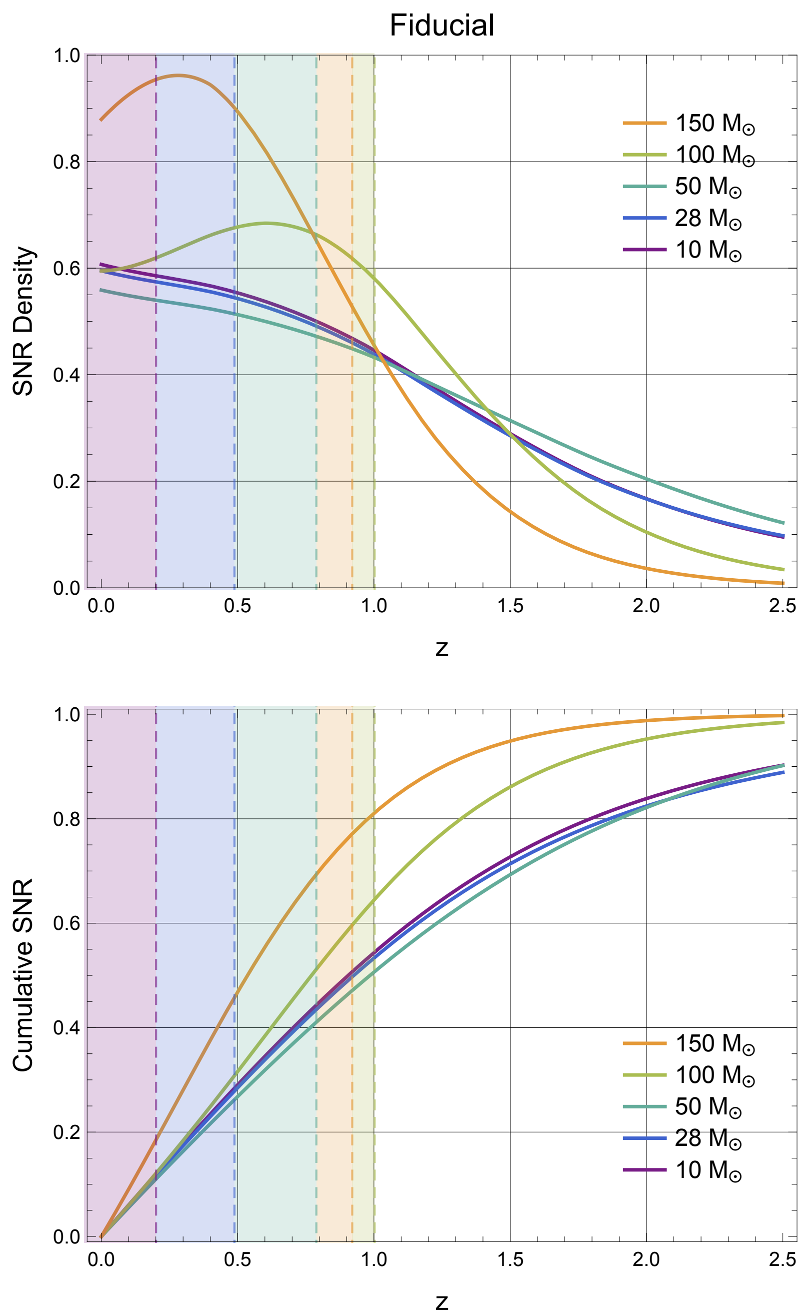

Figure 4.2: *Top*: SNR density $d(\mathrm{SNR})/dz$ for various choices of chirp mass, assuming the FIDUCIAL model for the gravitational-wave background due to binary black hole mergers. Each curve is normalized to a total signal-to-noise ratio of 1. *Bottom*: The cumulative SNR as a function of maximum redshift.

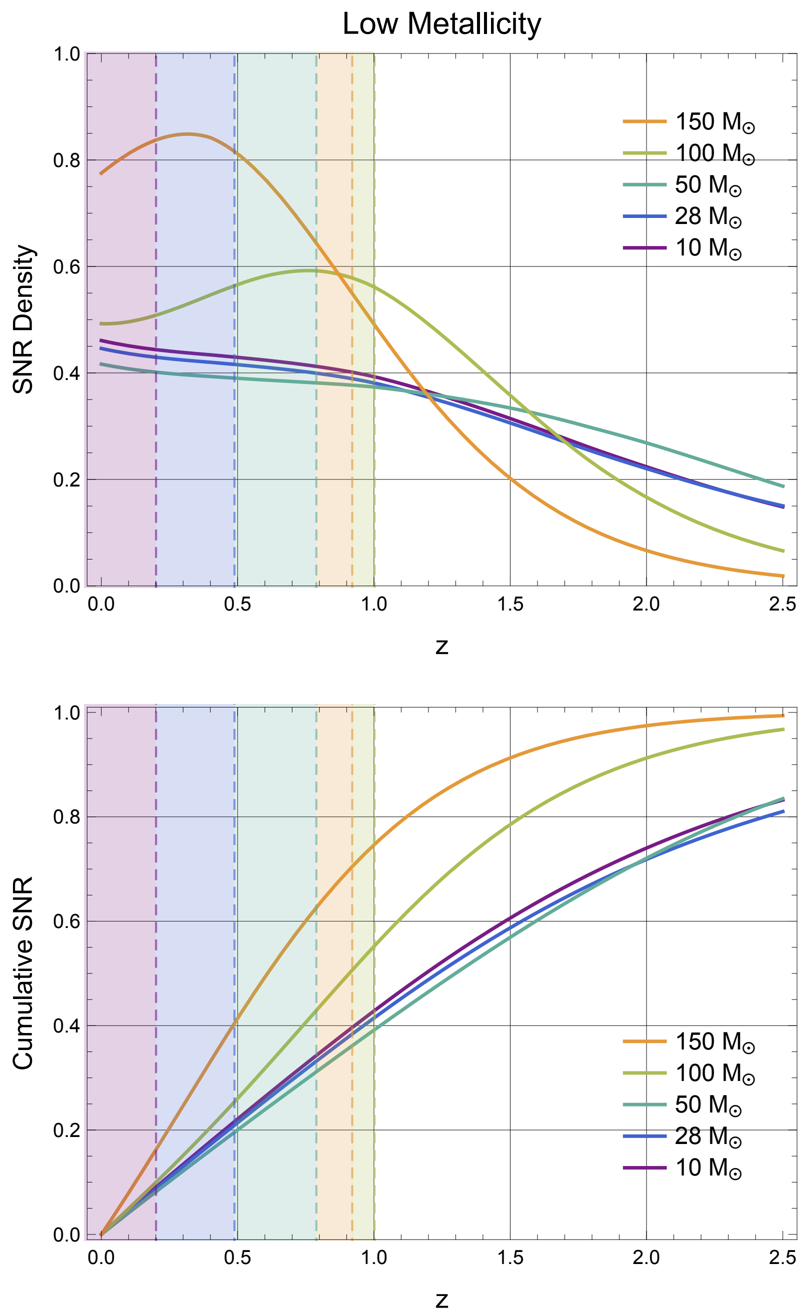

Figure 4.3: Stochastic SNR density $d(\mathrm{SNR})/dz$, assuming that compact binaries are due to stellar progenitors with metallicities $Z < Z_\odot/10$.

For binaries like GW150914, with $\mathcal{M}_c \approx 28\,M_\odot$ and $z_{50\%} \approx 0.5$, approximately 70% of the stochastic SNR is due to unresolvable binaries when assuming the Fiducial model. Sources between redshifts 0.1 and 3.5 contribute the central 90% of the total signal-to-noise ratio, and half of the total SNR is due to binaries beyond $z \approx 1.2$.

It is also interesting to see consider how SNR density changes with average chirp mass. For chirp masses $\mathcal{M}_c \lesssim 50\,M_\odot$, all SNR density curves appear similar; this is because the "knee frequency" of $\Omega(f)$ sits beyond the sensitive part of the Advanced LIGO band (see Fig. 4.1 above). At $\mathcal{M}_c \approx 100\,M_\odot$, however, the SNR density instead shows a peak at moderate z. This peak corresponds to the redshift at which the binary's final merger is redshifted into Advanced LIGO's most sensitive band. Finally, as $\mathcal{M}_c$ increases further to $150\,M_\odot$, the merger from high-redshift signals is shifted below the LIGO band entirely, leaving mostly signal from low-z sources.

The details do, of course, depend on our exact prescription for $\Omega(f)$. Varying the precise astrophysics (the time delay distribution, star formation rate, etc.) assumed in our model for the gravitational-wave background does inevitably alter the inferred SNR density. One of the largest sources of systematic uncertainties is the dependence of compact binary formation on metallicity. In the Fiducial model described above, we assumed that compact binary formation is dominated by stellar progenitors with metallicities $Z < Z_\odot/2$. Alternatively, if we assume extremely metal-poor stellar progenitors with $Z < Z_\odot/10$ (the Low Metallicity model of Ref. [26]), then we instead obtain the results in Fig. 4.3.

Requiring very metal-poor progenitors implies a much larger number of compact binary mergers at high redshift, when the Universe was younger and largely metal-free. Compared to the SNR densities in Fig. 4.2, Fig. 4.3 therefore shows stochastic SNR densities shifted towards higher redshifts. In the case of GW150914-like binaries, 80% of the stochastic SNR is due to unresolvable binaries, with 90% of the total signal contributed by binaries between redshifts 0.1 and 4.2 (in contrast to redshifts 0.1 to 3.5 in the Fiducial model above).

4.2 Extracting Astrophysical Information from the Stochastic Background

Our results in Ch. 4.1 should suggest that valuable astrophysical information is contained in the binary black hole background. Figure 4.1, for instance, demonstrated that the amplitude and spectral shape of the background's energy-density spectrum depends on the masses and coalescence rates of the binary black hole population. Furthermore, this population of binary black holes is *not* the same population directly detected in the local Universe, but a distinct population comprising sources at cosmological distances, extending up to redshifts $z \approx 3$ or 4 (depending on our prescription for progenitor metallicity). The degree to which this information can be *extracted*, however, is an altogether separate question, depending on our ability to perform model selection and parameter estimation with measurements of the gravitational-wave background.

Before proceeding quantitatively, we can gain some intuition by qualitatively reexamining Fig. 4.1 above. In Ch. 3.5.1, we argued that the energy-density due to compact binaries should scale approximately as $\Omega(f) \propto f^{2/3}$. Figure 4.1 confirms exactly this. At low frequencies, all example energy-density spectra trace $f^{2/3}$ power laws, deviating only at relatively high-frequencies, where the slope of $\Omega(f)$ increases slightly before turning over and going to zero. It is precisely this departure from a power law that will break degeneracies between the mass and rate of binary black holes. If this departure is undetectable, then we will be limited in our ability to independently measure the population of distant compact binaries using the gravitational-wave background alone.

Parameter estimation has already been shown to be difficult for gravitational-wave backgrounds dominated by compact binaries of several solar masses [185], which only depart from $\Omega(f) \propto f^{2/3}$ power laws at frequencies above $\sim 1\,\text{kHz}$. The discovery of GW150914, though, has since taught us that the gravitational-wave background will have a significant contribution from *high-mass* binary black holes. The energy-density spectrum due to these more massive black holes will be shifted to lower frequencies, where departures from power-law spectra are increasingly visible to ground-based detectors. This tentatively suggests that binary black hole backgrounds may be more promising targets for model selection and parameter estimation.

To generically evaluate our prospects for extracting astrophysical information from the binary black hole background, we will investigate at which point an astrophysical background (following the FIDUCIAL model) can be distinguished from a simple power-law spectrum

$$\Omega_{\text{PL}}(f) = \Omega_0 \left(\frac{f}{f_0}\right)^{2/3}, \tag{4.13}$$

where Ω_0 is the background's amplitude at an arbitrary reference frequency f_0.

Recall that, in cross-correlation searches for the stochastic background, our measurement noise is assumed to be Gaussian. Given a model $\Omega_M(f)$ for the background's energy-density spectrum, we'll therefore define a Gaussian likelihood for the cross-correlation $\hat{C}(f)$ measured within a single frequency bin of width df: [83, 185]

$$\mathcal{L}_f\left(\hat{C}(f)\,|\,\Omega_M(f)\right) \propto \exp\left(-\frac{\left[\hat{C}(f) - \gamma(f)\Omega_M(f)\right]^2}{2\sigma^2(f)}df\right), \tag{4.14}$$

where $\hat{C}(f)$ and $\sigma^2(f)$ are given by Eqs. (3.46) and (3.53), respectively. Crucially, here we do not have actual measurements $\hat{C}(f)$, but will instead be using Eq. (4.14) to *forecast* what future observations might reveal. Thus we're interested not in $\mathcal{L}_f$, but its expectation value:

$$\left\langle\mathcal{L}_f\left(\hat{C}(f)\,|\,\Omega_M(f)\right)\right\rangle \propto \exp\left(-\frac{\left[\gamma(f)\Omega(f) - \gamma(f)\Omega_M(f)\right]^2}{4\sigma^2(f)}df\right). \tag{4.15}$$

Equation (4.15) is obtained by substituting $\hat{C}(f) = \gamma(f)\Omega(f) + \delta C(f)$ into Eq. (4.14). $\delta C(f)$ is the random error associated with a single measurement; it is Gaussian distributed with a mean of zero and a variance of $\sigma^2(f)$. Marginalizing over all possible noise instantiations $\delta C(f)$ gives Eq. (4.15).

Equation (4.15) is almost identical to Eq. (4.14), but with an extra factor of 1/2 in the exponential. Thus the proper consideration of measurement error is crucial when forecasting future observational prospects; studies that use Eq. (4.14) to forecast future results implicitly assume *zero* measurement error, and hence obtain overly-precise results. Put another way, one should take care to use the average likelihood, $\langle\mathcal{L}_f(\hat{C}(f)\,|\,\Omega_M(f))\rangle$, and *not* the likelihood of the average, $\mathcal{L}_f(\langle\hat{C}(f)\rangle\,|\,\Omega_M(f))$.

Of course, we don't measure the cross-correlation between detectors in a single frequency bin, but measure an entire spectrum of cross-correlations. The full (ensemble-averaged) likelihood is the product of Eq. (4.15) across all frequencies:

$$\begin{aligned}\mathcal{L}(\Omega \,|\, \Omega_M) &\propto \prod_f \langle \mathcal{L}_f \rangle \\ &= \mathcal{N} \exp\left[-\frac{1}{4}\left(\Omega - \Omega_M \,|\, \Omega - \Omega_M\right)\right],\end{aligned} \tag{4.16}$$

where $\mathcal{N}$ is a normalization factor and we have defined the inner product

$$(A \,|\, B) = 2T \left(\frac{3H_0^2}{10\pi^2}\right)^2 \int_0^\infty \frac{\gamma(f)^2 \tilde{A}(f) \tilde{B}(f)}{f^6 P_1(f) P_2(f)} df. \tag{4.17}$$

Note that Ω, not $\hat{C}$, appears on the left-hand side of Eq. (4.16), since this ensemble-averaged likelihood depends only on the expectation value $\langle \hat{C}(f) \rangle = \gamma(f)\Omega(f)$.

Given a *true* gravitational-wave background described by the FIDUCIAL model, we can now compute the **maximum-likelihood ratio**

$$\mathcal{R} = \frac{\mathcal{L}^{\text{ML}}(\Omega_{\text{BBH}} \,|\, \Omega_M = \Omega_{\text{BBH}})}{\mathcal{L}^{\text{ML}}(\Omega_{\text{BBH}} \,|\, \Omega_M = \Omega_{\text{PL}})} \tag{4.18}$$

between the astrophysical and power-law models, defined by Eqs. (4.1) and (4.13), respectively. Large values of $\mathcal{R}$ will indicate that the astrophysical model is (correctly) preferred over the power-law background model. Meanwhile, values close to $\mathcal{R} = 1$ indicate that the background's energy-density spectrum is indistinguishable from a simple power law.

The maximum likelihood when correctly assuming the astrophysical FIDUCIAL model is simply

$$\mathcal{L}^{\text{ML}}(\Omega_{\text{BBH}} \,|\, \Omega_M = \Omega_{\text{BBH}}) = \mathcal{N}; \tag{4.19}$$

since the true underlying background is contained within the space of our model, the difference $\Omega - \Omega_M$ vanishes in Eq. (4.16). The maximum likelihood when (incorrectly) assuming a power law model can also be derived analytically. The power law model has a single free parameter: its amplitude Ω_0. The amplitude maximizing Eq. (4.16) satisfies

$$\begin{aligned}0 &= \frac{d \log \mathcal{L}(\Omega_{\text{BBH}} \,|\, \Omega_M = \Omega_{\text{PL}})}{d\Omega_0} \\ &= -\frac{1}{4} \frac{d}{d\Omega_0} \left(\Omega_{\text{BBH}} - \Omega_0 (f/f_0)^{2/3} \,\middle|\, \Omega_{\text{BBH}} - \Omega_0 (f/f_0)^{2/3}\right) \\ &= -\frac{1}{4}\left[-2\left(\Omega_{\text{BBH}} \,|\, \omega\right) + 2\Omega_0 \left(\omega \,|\, \omega\right)\right]\end{aligned} \tag{4.20}$$

or

$$\Omega_0^{\rm ML} = \frac{(\omega\,|\,\Omega_{\rm BBH})}{(\omega\,|\,\omega)} \tag{4.21}$$

where $\omega(f) = (f/f_0)^{2/3}$. The corresponding maximum likelihood is

$$\begin{aligned}\mathcal{L}^{\rm ML}&(\Omega_{\rm BBH}\,|\,\Omega_M = \Omega_{\rm PL}) \\ &= \mathcal{N}\exp\left\{-\frac{1}{4}\left((\Omega_{\rm BBH}\,|\,\Omega_{\rm BBH}) - \frac{(\omega\,|\,\Omega_{\rm BBH})^2}{(\omega\,|\,\omega)}\right)\right\}.\end{aligned} \tag{4.22}$$

Thus

$$\boxed{\mathcal{R} = \exp\left\{\frac{1}{4}\left((\Omega_{\rm BBH}\,|\,\Omega_{\rm BBH}) - \frac{(\omega\,|\,\Omega_{\rm BBH})^2}{(\omega\,|\,\omega)}\right)\right\}.} \tag{4.23}$$

Today, maximum likelihood ratios are relatively rare in the gravitational-wave data analysis community. Instead, they've been supplanted by more formally-correct (and correspondingly more difficult to compute) **Bayes factors** between competing hypotheses. As will be discussed in Ch. 5 below, there exists a simple conceptual relationship between Bayes factors and maximum likelihood ratios. In particular, a Bayes factor may be approximated by an ordinary maximum likelihood ratio, multiplied by an additional "Occam's factor" that further penalizes the more complex of the two models under consideration [106]. The inclusion of this Occam's factor here would only serve to penalize the complex astrophysical model for the binary black hole background. By neglecting the Occam's factor, we are effectively showing the *most optimistic* prospects for discerning the form of an astrophysical BBH background.

Figure 4.4 shows contours of the maximum log likelihood ratio $\ln\mathcal{R}$, as a function of the local coalescence rate $\mathcal{R}_0$ and presumed average chirp mass $\mathcal{M}_c$ of binary black holes, after three years of observation with design-sensitivity Advanced LIGO. For reference, the solid black curve indicates the merger rates above which a binary black hole background is detectable with optimal SNR $= 3$ after three years, when correctly assuming an astrophysical model for $\Omega(f)$. The dashed black curve similarly indicates rates above which binary black hole backgrounds are detectable with SNR $= 3$ when assuming a power-law model that does *not* turn over at high frequencies.[2] The chirp mass and local merger rate inferred from GW150914 [26–28] are indicated by a star.

[2]Note that this *not* an optimal SNR, since the space of power law models does not contain the true BBH spectrum.

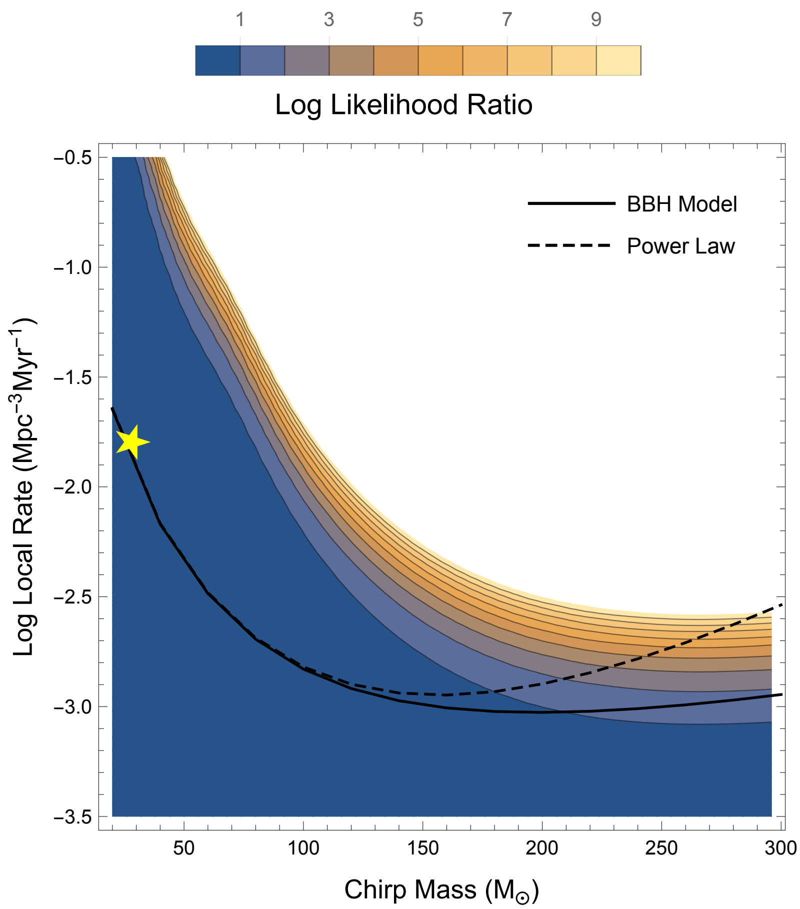

Figure 4.4: Contours of maximum log likelihood ratios $\ln \mathcal{R}$ between the astrophysical and power-law background models (Eqs. (4.1) and (4.13), respectively) given three years of observation with design-sensitivity Advanced LIGO. The solid and dashed black curves indicate the local coalescence rates at which the background is *detectable* with $\text{SNR} = 3$ when using astrophysical and power-law models, respectively. The star indicates the background associated with GW150914 [26]. Although the background inferred from GW150914 may be marginally detectable with Advanced LIGO after three years of observation, it is indistinguishable from a simple power-law model. The background remains indistinguishable from a power-law even for co-located detectors, which are predicted to make a strong detection of the BBH background.

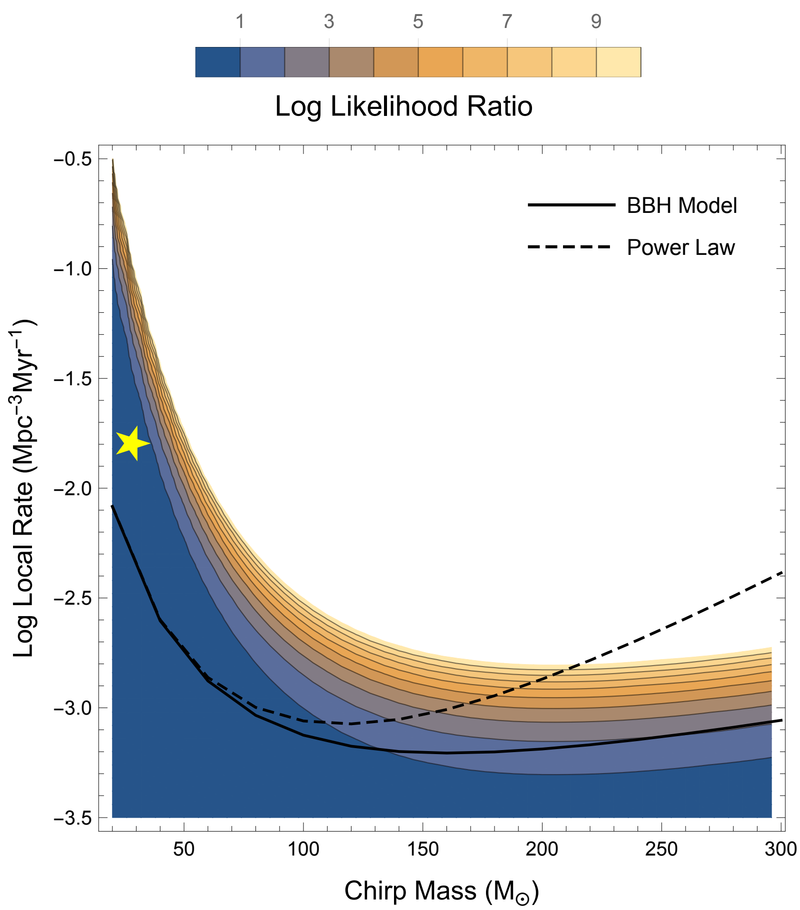

Figure 4.5: As in Fig. 4.4 above, but assuming *co-located* detectors, which are predicted to make a strong detection of the binary black hole background.

Over a majority of the plotted parameter space, we see that $\ln \mathcal{R} \lesssim 1$; in this region the power-law and astrophysical models cannot be distinguished. Only for chirp masses and local rates much larger than those implied by GW150914 is $\ln \mathcal{R} > 1$. Thus, while Advanced LIGO may possibly detect the stochastic background associated with GW150914, such a background is indistinguishable from a simple power-law. Quantitatively, approximately 6000 years of observation at design sensitivity are required to attain $\ln \mathcal{R} = 3$!

Recall that the sensitivity of Advanced LIGO's Hanford-Livingston (H1-L1) network to a stochastic background is limited at high frequencies by the overlap reduction function, which rapidly approaches zero at frequencies $f \gtrsim 60$ Hz [186]. Hence it is the *overlap reduction function* that effectively prevents us from detecting departures in $\Omega(f)$ from ordinary power laws. The critical frequency above which the overlap reduction function goes to zero is inversely proportional to the separation between detectors. Searches for the isotropic gravitational-wave background would therefore benefit from *smaller* baselines between detectors. There is some historical precedent for this – during Initial LIGO, a third interferometer (H2) was present at Hanford, co-located and co-oriented with H1 [187]. With a frequency-independent overlap reduction function $\gamma_{\text{H1-H2}}(f) = 1$, the H1-H2 pair is *significantly* more sensitive (geometrically speaking) to the high-frequency gravitational-wave background than H1-L1. Of course, in reality there are substantial practical barriers to successfully operating two co-located detectors, namely strongly correlated noise due to common seismic, magnetic, and anthropogenic environments [187].

While there are currently no plans to reinstall a second interferometer at Hanford during Advanced LIGO, it is interesting to consider how Fig. 4.4 might change given a hypothetical H1-H2 network of co-located 4 km Advanced LIGO interferometers. Figure 4.5 illustrates the maximum likelihood ratios between our astrophysical and power law models for this hypothetical H1-H2 network. Once again, while the binary black hole background implied by GW150914 is detectable with H1-H2, it remains indistinguishable from a power law. Approximately 50 years of observation with design-sensitivity co-located detectors are required to reach $\ln\mathcal{R} = 3$ in favor of the astrophysical model. Although this is a significant improvement (by a factor ≈ 140) over the H1-L1 performance above, it nevertheless remains an impractically long time.

The behavior of Figs. 4.4 and 4.5 can be better understood by re-plotting contours of $\ln\mathcal{R}$ as functions of the background's energy-density $\Omega(10\,\text{Hz})$ at 10 Hz and the frequency f_{max} at which the energy-density is at a maximum (recall that $\Omega \sim \mathcal{M}_c^{5/3} R_0$ and $f_{\text{max}} \sim 1/\mathcal{M}_c$; see Fig. 4.1). The result is shown in Figs. 4.6 and 4.7 for the H1-L1 and H1-H2 baselines, respectively. From Fig. 4.6, it is apparent that the only backgrounds distinguishable from power laws using H1-L1 are those whose masses place f_{max} between $\sim 10 - 50$ Hz (chirp masses between approximately $70 - 300\,M_\odot$), corresponding to the most sensitive frequency band for the isotropic stochastic search. The H1-H2 network avoids the penalty associated with the overlap reduction function and so shows sensitivity across a broader f_{max} range, limited only by the detectors' own PSDs at high and low frequencies.

Any configuration of advanced detectors appears unlikely to differentiate an astrophysical gravitational-wave background model from a simple power-law.

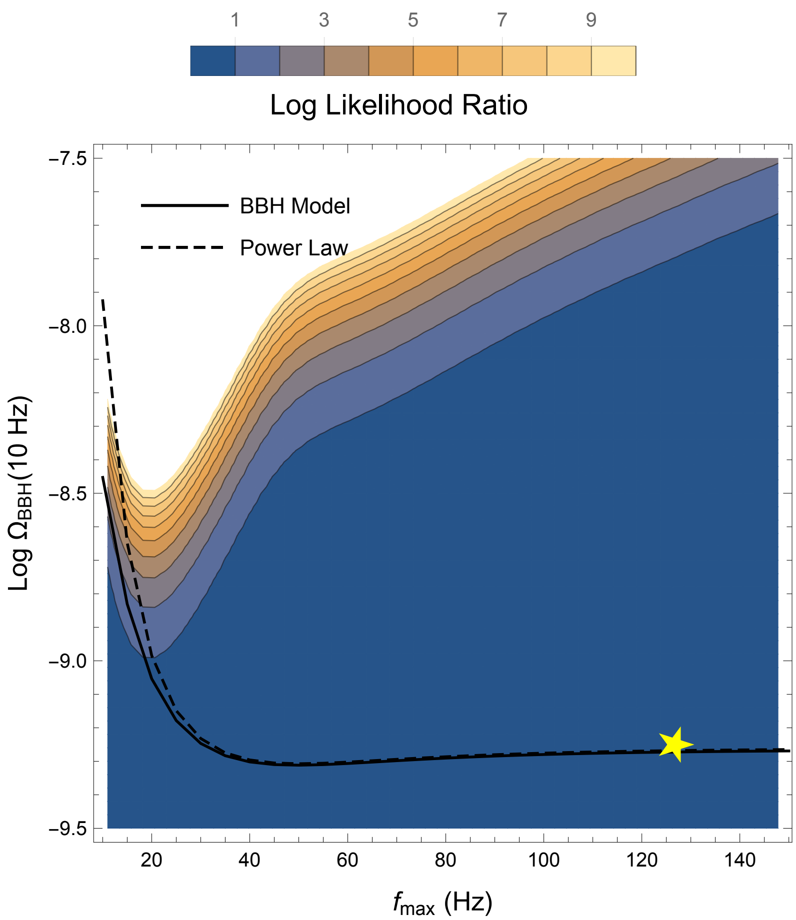

Figure 4.6: Maximum log-likelihood contours between the astrophysical and power-law models, shown as a function $f_{\rm max}$ (see Fig. 4.1) and the background's amplitude at 10 Hz. The results shown assume three years of integration with design-sensitivity Advanced LIGO. As in Fig. 4.4, solid and dashed black curves show the amplitudes at which a background is detectable when assuming astrophysical and power-law models, respectively, while the star indicates the astrophysical background associated with GW150914. Advanced LIGO is best able to distinguish realistic background models from power laws for frequencies $f_{\rm max}$ between 10-50 Hz, corresponding to the most sensitive frequency band for the stochastic search.

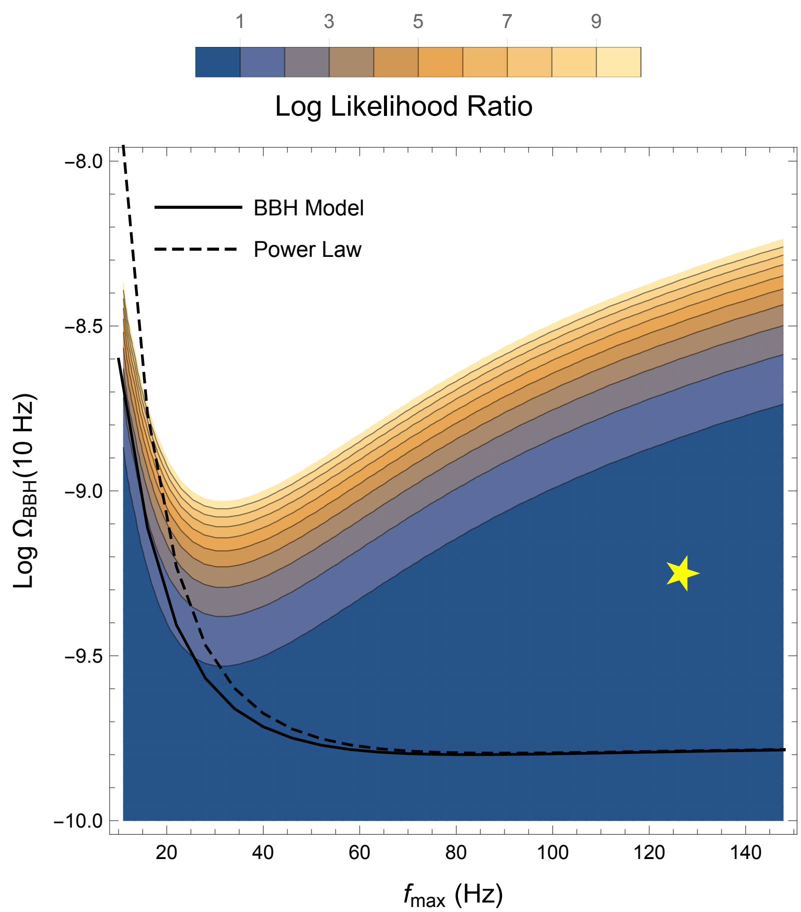

Figure 4.7: As in Fig. 4.6, but assuming a hypothetical H1-H2 network of co-located and co-oriented Advanced LIGO detectors.

Hence astrophysical inference on the gravitational-wave background will likely be limited to a single piece of information – its amplitude Ω_0 – rather than its spectral shape.

As a rule of thumb, it should be possible to distinguish two models for the background's energy density only when their *difference* $\Delta\Omega(f)$ is itself detectable. For reference, a gravitational-wave background with energy-density $\Omega(f) = 10^{-9}\,(f/10\,\mathrm{Hz})^{2/3}$ has expected signal-to-noise ratio SNR ≈ 3 after one year of integration with Advanced LIGO. Therefore, given integration time T, it will be possible to select between two models for the gravitational-wave background only if their difference is of order

$$\Delta\Omega(f) \gtrsim 10^{-9}\left(\frac{f}{10\,\mathrm{Hz}}\right)^{2/3}\left(\frac{1\,\mathrm{yr}}{T}\right)^{1/2}. \tag{4.24}$$

4.3 The binary black hole background as a limiting foreground

So far in this chapter, we have considered only the gravitational-wave background's ability to teach us about distant compact binaries. The gravitational-wave background, of course, is not only due to compact binary mergers alone. It will invariably comprise signals from a multitude of other sources – rotating neutron stars, core-collapse supernovae, gravitational waves of cosmological origin, and maybe even exotic or as-of-yet unforeseen sources of gravitational radiation. Once a stochastic gravitational-wave signal is observed by advanced detectors, a natural question will therefore be:

Is the observed signal consistent with expectations from binary black hole mergers alone, or is there a contribution from something else?

In this sense, the binary black hole background now becomes a *limiting foreground*, possibly obscuring the presence of additional, weaker contributions to the net gravitational-wave background.

As a simple scenario, consider a combined signal composed of an astrophysical background due to GW150914-like black holes (chirp mass $\mathcal{M}_c = 28\,M_\odot$ and local rate $\mathcal{R}_0 = 16\,\mathrm{Gpc}^{-3}\mathrm{yr}^{-1}$) and a flat energy-density spectrum with amplitude Ω_c of cosmological origin. How loud must Ω_c be in order to be detectable against the binary black hole background $\Omega_{\mathrm{BBH}}(f)$? An equivalent question is: how loud must the stochastic signal be in order to detect a spectral index that is inconsistent with the binary black hole scenario? Since the energy-density spectrum due to compact binaries is known to follow $f^{2/3}$, only the measurement of a spectral index inconsistent with 2/3 can provide direct evidence of a distinct additional source population in the stochastic background.

We can recast this question as another model selection problem. In Ch. 4.2, we just demonstrated that the astrophysical background due to binary black holes is likely indistinguishable from a power law. Hence we will assume that a background composed of binary black holes alone (the "BBH-only" hypothesis) is described by

$$\Omega_{\mathrm{BBH-}}(f) = \Omega_0 \left(\frac{f}{f_0}\right)^{2/3}. \tag{4.25}$$

Meanwhile, when allowing for an additional contribution to the gravitational-wave background (the "BBH+" hypothesis), we'll use the energy-density spectrum

$$\Omega_{\text{BBH}+}(f) = \Omega_1 \left(\frac{f}{f_0}\right)^{2/3} + \Omega_2, \tag{4.26}$$

where Ω_2 is a constant.

Once again, given a true gravitational-wave background with energy density $\Omega(f)$, we can form a maximum likelihood ratio

$$\mathcal{R} = \frac{\mathcal{L}^{\text{ML}}(\Omega \,|\, \Omega_M = \Omega_{\text{BBH}+})}{\mathcal{L}^{\text{ML}}(\Omega \,|\, \Omega_M = \Omega_{\text{BBH}-})} \tag{4.27}$$

quantifying the evidence that the background contains an additional contribution beyond binary black holes alone. The "BBH-only" likelihood is maximized by the same amplitude Ω_0^{ML} given in Eq. (4.21), provided we replace $\Omega_{\text{BBH}}(f)$ with the combined background $\Omega(f) = \Omega_{\text{BBH}}(f) + \Omega_c$ considered here. The "BBH+" likelihood is in turn maximized by solving

$$0 = \frac{d \log \mathcal{L}(\Omega \,|\, \Omega_M = \Omega_{\text{BBH}+})}{d\Omega_1} \tag{4.28}$$

and

$$0 = \frac{d \log \mathcal{L}(\Omega \,|\, \Omega_M = \Omega_{\text{BBH}+})}{d\Omega_2}, \tag{4.29}$$

which together give

$$\begin{aligned}
\Omega_1^{\text{ML}} &= \frac{(\omega \,|\, 1)(\Omega \,|\, 1) - (1 \,|\, 1)(\Omega \,|\, \omega)}{(\omega \,|\, 1)^2 - (\omega \,|\, \omega)(1 \,|\, 1)} \\
\Omega_2^{\text{ML}} &= \frac{(\omega \,|\, 1)(\Omega \,|\, \omega) - (\omega \,|\, \omega)(\Omega \,|\, 1)}{(\omega \,|\, 1)^2 - (\omega \,|\, \omega)(1 \,|\, 1)}.
\end{aligned} \tag{4.30}$$

Figure 4.8 shows contours in $\ln \mathcal{R}$ as a function of the cosmological background amplitude Ω_c and the total integration time, assuming the design-sensitivity H1-L1 detector network. The solid and dashed black curves indicate the observation times necessary to *detect* the combined astrophysical and cosmological backgrounds with optimal signal-to-noise ratios of SNR = 3 and 5, respectively. Note that these curves diverge downwards as Ω_c approaches zero, corresponding to the fixed detection time required for the binary black hole background alone (which is held constant in this exercise). The solid and dashed grey curves, meanwhile, indicate that cosmological background amplitudes that *would* have been detectable with SNR = 3 and 5 if there existed no binary black holes.

The fact that the grey curves lie deep within the $\ln \mathcal{R} \simeq 0$ region implies that the presence of the binary black hole background serves to largely obscure

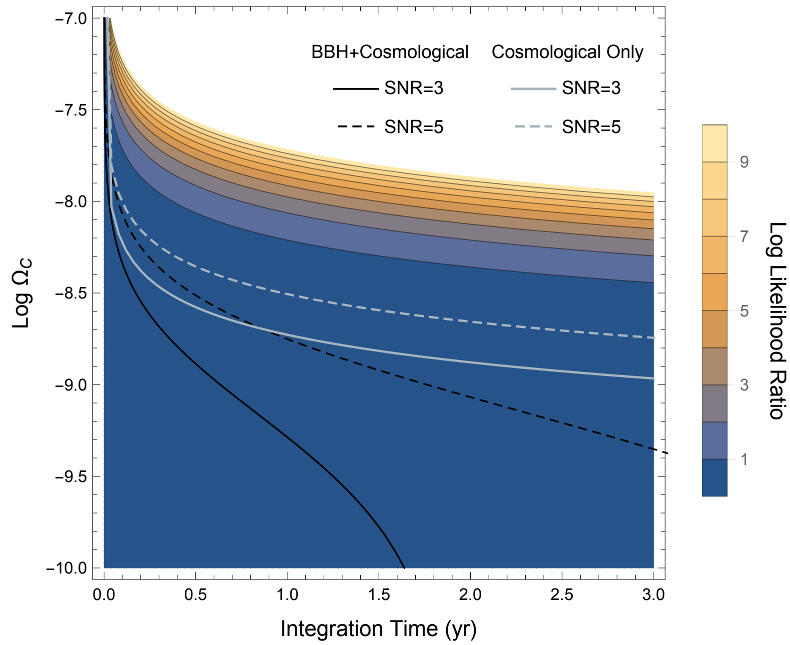

Figure 4.8: Contours of the maximum likelihood ratio between the "BBH+" and "BBH-only" models, as a function of the cosmological background's amplitude Ω_c and total integration time with design-sensitivity Advanced LIGO. Black curves indicate observation times required to detect the combined astrophysical and cosmological background with a given optimal SNR, while grey curves show the amplitudes Ω_c that *alone* would be detectable.

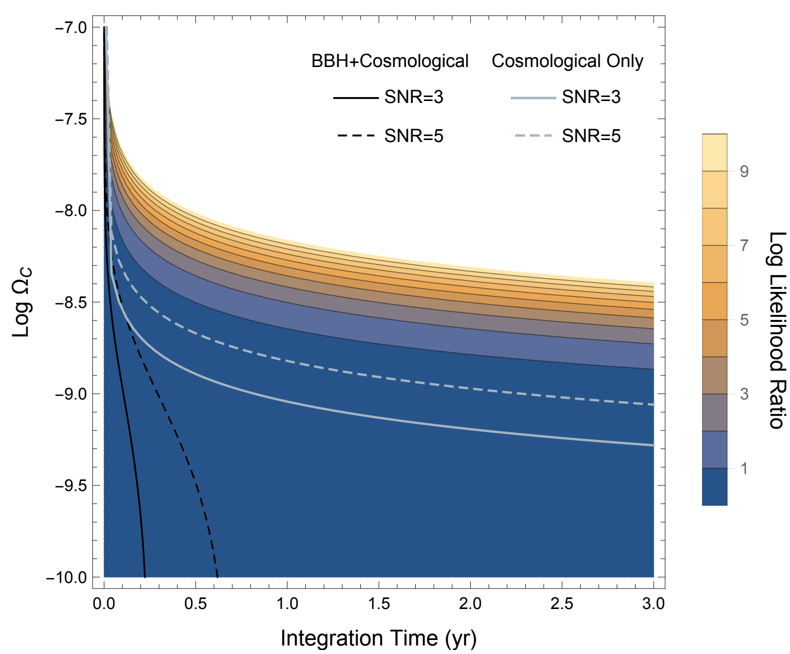

Figure 4.9: As in Fig. 4.9, but assuming a hypothetical H1-H2 network of co-located Advanced LIGO detectors.

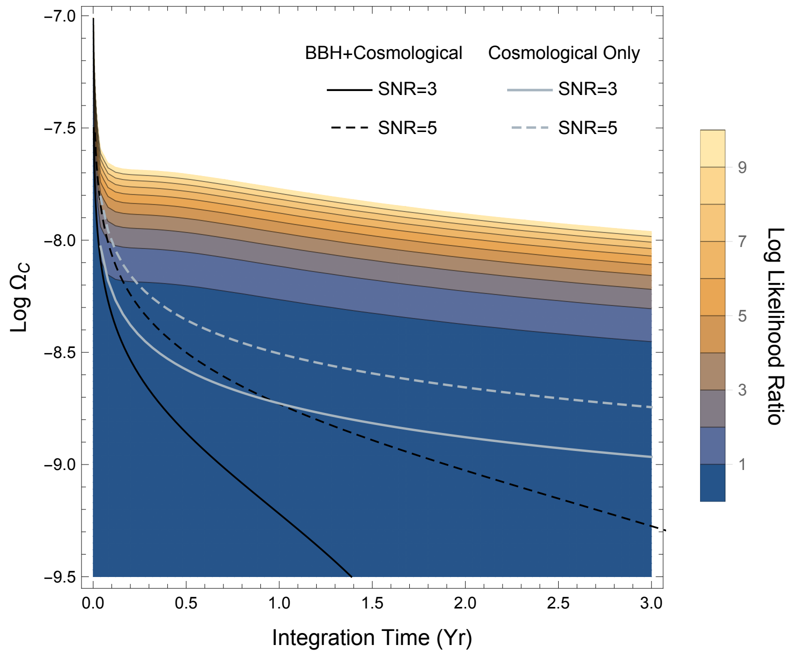

Figure 4.10: As in Fig. 4.9 above, but assuming that the amplitudes Ω_0 and Ω_1 (in the "BBH-only" and "BBH+" models, respectively) of the binary black hole background are known *a priori* to within a factor of two.

any cosmological background that might otherwise be detectable. If no binary black hole background were present, for instance, Advanced LIGO could detect a cosmological background of amplitude $\Omega_c \approx 10^{-9.0}$ with SNR $= 3$ after three years of observation. When the binary black hole background is present, however, a much larger $\Omega_c \approx 10^{-8.2}$ (corresponding to $\ln \mathcal{R} = 3$) is required to both detect and *resolve* the additional presence of the cosmological signal. After three years of observation at design sensitivity, Advanced LIGO will therefore be able to constrain the amplitudes of additional background components to $\Omega_c \lesssim 10^{-8.2}$.

A network of co-located detectors performs somewhat better. As shown in Fig. 4.9, our hypothetical H1-H2 network could place constraints $\Omega_c \lesssim 10^{-8.4}$ after one year of observation and $\Omega_c \lesssim 10^{-8.6}$ after three.

To obtain Figs. 4.8 and 4.9, we've treated the power-law amplitudes Ω_0 and Ω_1 of Eqs. (4.25) and (4.26) as entirely free parameters. In reality, we will likely be able to place a strong prior on these parameters, leveraging the direct detection of nearby compact binaries to estimate of the average binary chirp mass and local coalescence rate as well as improving our understanding of systematic uncertainties in the formation history of compact binaries [26]. However, even if we assume tight *a priori* constraints on the amplitudes Ω_0 and Ω_1, we find little change in our results. Fig. 4.10 shows contours of $\ln\mathcal{R}$ when we assume that the amplitude of the binary black hole background is known to within a factor of two, an optimistic assumption given the uncertainty in merger rate evolution with redshift [26]. Even with these optimistic priors, we see that our ability to resolve a cosmological background is notably improved only for $\Omega_c \gtrsim 10^{-8}$ and integration times $T \lesssim 1\,\mathrm{yr}$. After $T \approx 1\,\mathrm{yr}$ of integration, the experimental uncertainty on the gravitational-wave background amplitude has become *smaller* than our presumed prior uncertainty, and so our *a priori* knowledge is no longer useful.

4.4 Takeaways and Prospects

In this chapter, we sought to address three questions concerning what we might learn from detection of the astrophysical gravitational-wave background.

First, how does the information contained in a stochastic background compare with what can be learned from nearby, individually-resolvable binary mergers? We saw that, while direct searches for binary black hole mergers are sensitive to redshifts less than $z_{50\%} \approx 0.5$, the stochastic background is dominated by binary mergers in the far more distant Universe, with 90% of the stochastic SNR due to sources between redshifts $z \approx 0.1-3.5$. The stochastic background therefore encodes astrophysical information about a population of black hole binaries that is distinct from the local population visible to direct matched-filter searches for compact binaries.

Second, what astrophysics can we hope to extract from future observations of the binary black hole background? In principle, the functional form of the background's energy density spectrum depends upon the precise characteristics of the underlying binary black hole population – its mean chirp mass, local coalescence rate, star formation history, and metallicity dependence. We found, however, that for realistic chirp masses and coalescence rates, the form of the stochastic background is indistinguishable from a simple $f^{2/3}$ power-law with Advanced LIGO. In the near future, astrophysical inference using the gravitational-wave background will largely be limited to considering only the overall *amplitude* of the the background and not its shape.

Finally, how is our ability to measure other stochastic backgrounds affected by the presence of an astrophysical binary black hole background? We quantified the degree to which the binary black holes obscure additional contributions to the stochastic background, like gravitational waves of cosmological origin. For a cosmological background to be resolvable, it must be strong enough to overcome our uncertainty in the amplitude of the binary black hole background. In this sense, the stochastic signal due to compact binaries now acts as a foreground, limiting Advanced LIGO's sensitivity to additional, weaker background components.

At face value, these conclusions are somewhat disheartening. There remains cause for optimism, though. First, our analysis has assumed a standard cross-correlation search for the gravitational-wave background. Cross-correlation searches are optimal provided that the gravitational-wave background is stationary, isotropic, and Gaussian. However, as described in Ch. 3.5.1, we know that the gravitational-wave background is decidedly *not* Gaussian. It may therefore be possible to improve upon our results with future data analysis methods optimized for non-Gaussian backgrounds.

In the more distant future, third generation detectors like the Einstein Telescope (ET) [188, 189] and Cosmic Explorer (CE) will be able to probe black hole binaries at cosmological distances. ET, for instance, is projected to resolve individual GW150914-like events to redshifts of $z \sim 15$ [189], allowing for precision observation of the binary black hole population over the entire history of star-formation. The ability of ET to resolve such events raises the exciting possibility of the individual identification and subtraction of each BBH coalescence from the data, opening the way for the detection of weaker, underlying stochastic backgrounds of astrophysical or even cosmological origin.

Even in the present, there is much to be learned by synthesizing the results from *both* direct searches for nearby binary mergers and cross-correlation measurements of the gravitational-wave background. By acting as a measure of the *integrated* merger history of compact binaries, knowledge of the gravitational-wave background's amplitude alone, when combined with knowledge from the local Universe, may allow us to make powerful statements about the formation and merger history of compact binaries.

Chapter 5
Measuring Gravitational-Wave Polarizations with the Stochastic Background

∾

This chapter contains work published in:

> T. Callister, S. A. Biscoveanu, N. Christensen *et al.*, *Polarization-Based Tests of Gravity with the Stochastic Gravitational-Wave Background*, Phys. Rev. X **7**, 041058 (2017).

I conceived of this project, produced most of the results in the published text (all of the results shown here), and wrote the majority of the manuscript.

∾

5.1 Gravitational-Wave Polarizations

In this section we'll adopt geometrized units in which $c = 1$. The speed of light will return in Ch. 5.2.

When introducing gravitational waves in Ch. 2, we declared that they are described entirely by two polarizations, the plus ($+$) and cross ($\times$) modes. These two gravitational-wave polarizations have basis tensors

$$\hat{\mathbf{e}}^{+}_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{5.1}$$

and

$$\hat{\mathbf{e}}^{\times}_{ij} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{5.2}$$

and deform rings of freely-falling test masses as shown in Fig. 5.1.

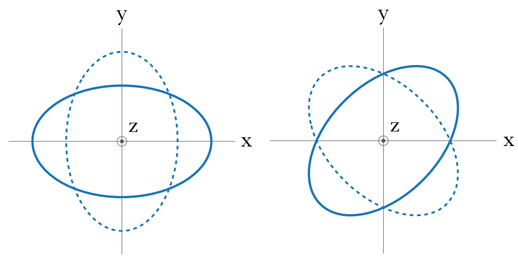

Figure 5.1: The effect of $+$ and $\times$ polarizations on rings of freely-falling test masses, assuming a gravitational-wave propagating in the $\hat{z}$ direction.

It is worth investigating why we only have two gravitational-wave polarizations. The easiest way to understand this is by counting independent components of the Riemann tensor $R_{\alpha\beta\gamma\delta}$. While the Riemann tensor has 256 components, only 20 of these are independent given the many symmetries that the Riemann tensor must obey. The components of the Riemann tensor must also obey the Bianchi identities $\nabla_{[\lambda} R_{\alpha\beta]\mu\nu} = 0$. If we restrict to a *plane wave* on an otherwise flat background, such that $R_{\alpha\beta\gamma\delta} \equiv R_{\alpha\beta\gamma\delta}(t-z)$, then the Bianchi identities leave us with only *six* independent quantities – the components R_{i0j0}.

So far this counting argument has assumed *only* a metric theory of gravity. Let's now restrict to our favorite metric theory of gravity – general relativity. In vacuum, Einstein's equations provide four more constraints on the Riemann tensor, reducing the number of independent components to two:

$$R_{x0x0} = -R_{y0y0} \tag{5.3}$$

and

$$R_{x0y0} = R_{y0x0}. \tag{5.4}$$

How do these remaining Riemann tensor components relate to gravitational waves? Recall that the Riemann tensor in linearized gravity is

$$R_{\alpha\beta\mu\nu} = \frac{1}{2}\left(\partial_\alpha\partial_\nu h_{\beta\mu} + \partial_\beta\partial_\mu h_{\alpha\nu} - \partial_\alpha\partial_\mu h_{\beta\nu} - \partial_\beta\partial_\nu h_{\alpha\mu}\right), \tag{5.5}$$

neglecting terms of $\mathcal{O}(h^2)$. In the Newtonian limit,

$$\begin{aligned} R_{i0j0} &= \frac{1}{2}\left(\partial_i\partial_0 h_{0j} + \partial_0\partial_j h_{i0} - \partial_i\partial_j h_{00} - \partial_0\partial_0 h_{ij}\right) \\ &\approx -\frac{1}{2}\ddot{h}_{ij}, \end{aligned} \tag{5.6}$$

where, since gravitational waves travel at the speed of light, we have assumed that time derivatives of h_{ij} are much larger than spatial derivatives ($\partial_0 h_{ij} \gg \partial_k h_{ij}$). So the Riemann tensor components R_{i0j0} simply correspond to (time derivatives of) the components of a gravitational wave's strain tensor h_{ij}. The fact that the Riemann tensor has only two independent components therefore implies that gravitational waves have only two independent components as well: our familiar $+$ and $\times$ modes.

If general relativity is *not* the correct description of gravity, though, then the above argument breaks down. In a general metric theory of gravity, we are left with our original six independent components of the Riemann tensor, yielding six independent polarizations available to gravitational waves. The four additional degrees of freedom are most commonly decomposed into the x, y, breathing (b), and longitudinal (l) modes, with polarization tensors

$$\hat{\mathbf{e}}^x_{ij} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \tag{5.7}$$

$$\hat{\mathbf{e}}^y_{ij} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \tag{5.8}$$

$$\hat{\mathbf{e}}^b_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{5.9}$$

and

$$\hat{\mathbf{e}}^l_{ij} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{5.10}$$

The effect of each of these alternative polarizations is shown in Fig. 5.2. In contrast to the $+$ and $\times$ polarizations, which are purely transverse, three of the four alternative polarizations (x, y, and l) have components *along* the gravitational wave's direction of propagation.

The $+$ and $\times$ modes are often referred to as "tensor" polarizations (alternatively "spin-2" modes). Analogously, the x and y modes are referred to as "vector" (spin-1) polarizations while the breathing and longitudinal modes are called "scalar" (spin-0) polarizations. These names are inspired by the

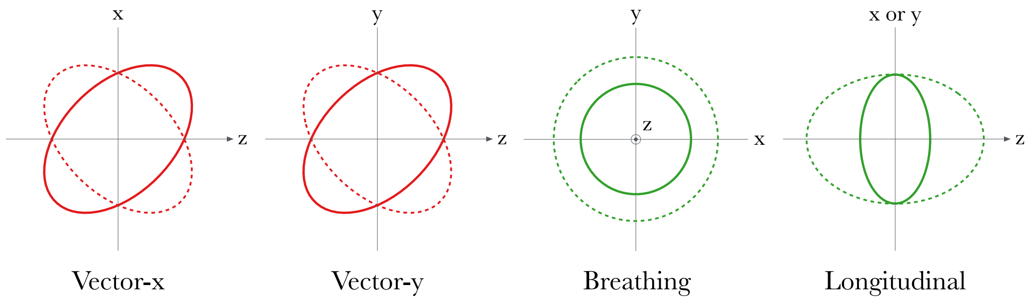

Figure 5.2: The effect of alternative x, y, breathing, and longitudinal polarizations on rings of freely-falling test masses, assuming a gravitational-wave propagating in the $\hat{z}$ direction.

responses of the polarization tensors to rotations. Consider rotating the superposition $h^T_{ij} = a\,\hat{\mathbf{e}}^+_{ij} + b\,\hat{\mathbf{e}}^\times_{ij}$ by an angle ϕ about the $\hat{z}$ axis,

$$h^T_{\hat{i}\hat{j}} = R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi) h^T_{ij}, \tag{5.11}$$

where $R^i{}_{\hat{j}}(\phi)$ is a rotation matrix (not a contraction of the Riemann tensor). Simplifying,

$$\begin{aligned} h^T_{\hat{i}\hat{j}} &= a\big[R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi)\,\hat{\mathbf{e}}^+_{ij}\big] + b\big[R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi)\,\hat{\mathbf{e}}^\times_{ij}\big] \\ &= a\left(\cos 2\phi\,\hat{\mathbf{e}}^+_{ij} + \sin 2\phi\,\hat{\mathbf{e}}^\times_{ij}\right) + b\left(-\sin 2\phi\,\hat{\mathbf{e}}^+_{ij} + \cos 2\phi\,\hat{\mathbf{e}}^\times_{ij}\right) \\ &= \left(a\cos 2\phi - b\sin 2\phi\right)\,\hat{\mathbf{e}}^+_{ij} + \left(a\sin 2\phi + b\cos 2\phi\right)\,\hat{\mathbf{e}}^\times_{ij}. \end{aligned} \tag{5.12}$$

The gravitational-wave returns to its original state after a rotation of only $\phi = \pi\,$rad, just like a spin-2 particle in the language of quantum mechanics.

Next, similarly rotate the superposition of "vector" modes $h^V_{ij} = a\,\hat{\mathbf{e}}^x_{ij} + b\,\hat{\mathbf{e}}^y_{ij}$:

$$\begin{aligned} h^V_{\hat{i}\hat{j}} &= R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi) h^V_{ij} \\ &= a\big[R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi)\,\hat{\mathbf{e}}^x_{ij}\big] + b\big[R^i{}_{\hat{i}}(\phi) R^j{}_{\hat{j}}(\phi)\,\hat{\mathbf{e}}^y_{ij}\big] \\ &= a\left(\cos\phi\,\hat{\mathbf{e}}^x_{ij} + \sin\phi\,\hat{\mathbf{e}}^y_{ij}\right) + b\left(-\sin\phi\,\hat{\mathbf{e}}^x_{ij} + \cos\phi\,\hat{\mathbf{e}}^y_{ij}\right) \\ &= \left(a\cos\phi - b\sin\phi\right)\,\hat{\mathbf{e}}^x_{ij} + \left(a\sin\phi + b\cos\phi\right)\,\hat{\mathbf{e}}^y_{ij}. \end{aligned} \tag{5.13}$$

The x and y components of h^V_{ij} transform like an ordinary vector (or a spin-1 particle).

Finally, the breathing and longitudinal modes are symmetric in the transverse plane, and so do not transform at all under rotations about the $\hat{z}$ axis. Hence these polarizations behave just like scalars (or spin-0 particles) under rotations.

The direct measurement of gravitational-wave polarizations is a particularly clean test of general relativity. General relativity predicts that gravitational waves are described purely by the two tensor modes. The observation of vector or scalar modes would therefore be an immediate and direct indicator of new physics beyond general relativity. Moreover, polarization-based tests of gravity require almost *no* assumptions about a gravitational-wave signal or its source. Whereas many tests of general relativity with gravitational waves are highly model dependent, involving the measurement of coefficients parametrizing the exact phase and amplitude evolution of a gravitational-wave signal [52, 190], here we are instead interested in a local, purely geometric characterization of a gravitational wave's strain tensor.

5.2 Extended Theories of Gravity and Sources of Alternative Polarizations

In the spirit of model-independence, our focus in this chapter is largely phenomenological – do alternative polarizations exist (invalidating general relativity) or not? Nevertheless, there are good theoretical reasons to search for alternative gravitational-wave polarizations. Among the broad range of alternative gravitational theories, particularly well-studied is the class of **scalar-tensor theories**, which contain an additional scalar field that is non-minimally coupled to spacetime curvature [191]. Scalar-tensor theories effectively elevate Newton's constant G to a field that evolves dynamically alongside the spacetime metric and matter fields.

Within the context of scalar-tensor theories, core-collapse supernovae (CCSNe) constitute a potential source of scalar gravitational waves. Spherically-symmetric stellar collapses have no time-varying quadrupole moment, and so CCSNe are expected to be very weak sources of gravitational radiation in general relativity. They do, however, radiate scalar breathing modes in canonical scalar-tensor theories. While the direct observation of gravitational waves from CCSNe could therefore place strong constraints on scalar-tensor theories [192], only supernovae within the Milky Way are likely to be directly detectable using current instruments [193, 194]. Such events are rare, occurring at a rate between $(0.6 - 10.5) \times 10^{-2}\,\mathrm{yr}^{-1}$ [195]. Nevertheless, certain extreme phenomenological supernovae models predict gravitational radiation many orders of magnitude stronger than in more conventional models, possibly allowing their detection at extragalactic distances [193]. Additionally, the superposition of all distant CCSNe could give rise to a gravitational-wave background of breathing modes [147, 150].

Compact binary coalescences may also yield scalar-polarized gravitational waves. In many scalar-tensor theories, bodies may carry a "scalar charge" that sources the emission of scalar gravitational waves [196, 197]. Monopole scalar radiation is suppressed due to conservation of scalar charge (just as conservation of mass suppressed monopole gravitational radiation in Ch. 2 above), but there is in general no conservation law suppressing dipole radiation. Scalar dipole radiation from compact binaries is enhanced by a factor of $(v/c)^{-2}$ relative to ordinary quadrupole tensor radiation (where v is the orbital velocity of the binary) and thus represents a potentially promising source of scalar gravitational waves. Electromagnetic observations of pulsar binaries have placed stringent constraints on anomalous energy loss beyond that predicted by general relativity; these constraints may be translated into a strong limit on the presence of additional scalar-dipole radiation [198, 199]. These limits, though, are strongly model-dependent, assuming *a priori* only small deviations from general relativity. Additionally, pure vacuum solutions like binary black holes are not necessarily subject to these constraints [200–203].

A variety of exotic sources may radiate alternative polarizations as well. Cosmic strings, for instance, generically radiate alternative polarizations in extended theories of gravity [169, 204]. Another potential source of alternative polarizations are the so-called "bubble walls" generated by first order phase transitions in the early Universe [156, 165, 167]. In scalar-tensor theories, bubbles are expected to produce strong scalar emission [196]. Gravitational waves from bubbles are heavily redshifted, though, and today may have frequencies too low for Advanced LIGO to detect [165]. Bubble walls may therefore be a more promising target for future space-based detectors like LISA than for current ground-based instruments.

Finally, it is in principle possible for alternative polarizations to be generated more effectively from sources at very large distances. There are several ways in which this might occur. First, modifications to the gravitational-wave dispersion relation can lead to mixing between different polarizations in vacuum (an effect analogous to neutrino oscillations). This can cause mixing between the usual tensor modes [205], and also between tensor modes and other polarizations [206, 207]. Thus alternative polarizations can be generated during propagation even if tensor modes alone are produced at the source. This kind of behavior appears, for instance, in generic Lorentz-violating theories of gravity [208, 209]. Secondly, as mentioned above alternative theories may promote fundamental constants like Newton's constant G to dynamical fields. If these fields behaved sufficiently differently at earlier stages in the Universe's evolution, local constraints on scalar emission may not apply to emission from remote sources [210]. Finally, it is possible to posit screening mechanisms that suppress the emission of alternative polarizations by local sources but that do not affect more remote sources [211].

5.3 Measuring Gravitational-Wave Polarizations

The crucial reason we are able to measure gravitational-wave polarizations at all is that our detectors have different geometrical responses to different gravitational-wave polarizations. In Ch. 2.5, we defined the antenna patterns $F_+(\hat{\mathbf{n}})$ and $F_\times(\hat{\mathbf{n}})$ characterizing a detector's response s to $+$ and $\times$-polarized waves from direction $\hat{n}$:

$$s = F_+(\hat{\mathbf{n}})h_+ + F_\times(\hat{\mathbf{n}})h_\times, \tag{5.14}$$

where the antenna patterns are themselves the contraction of the detector tensor D_{ij} with the basis tensor $\hat{\mathbf{e}}^+_{ij}(\hat{\mathbf{n}})$ and $\hat{\mathbf{e}}^\times_{ij}(\hat{\mathbf{n}})$. The antenna patterns describing the response to alternative polarizations are defined analogously, e.g. $F_x(\hat{\mathbf{n}}) = D_{ij}e^{ij}_x(\hat{\mathbf{n}})$, $F_y(\hat{\mathbf{n}}) = D_{ij}e^{ij}_y(\hat{\mathbf{n}})$, etc. Absolute values of the LIGO/Virgo antenna patterns for the tensor, vector, and scalar polarizations are shown in Figs. 5.3 – 5.5. In each figure it is assumed that the detector lies in the $z = 0$ plane with arms in the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ direction. Note that only one plot is shown in Fig. 5.5. This is because breathing and longitudinal polarizations induce *identical* (up to a sign) responses in "L"-shaped interferometric detectors like LIGO and Virgo.

The fact that these antenna patterns are (with the exception of F_b and F_l) distinct will allow us to infer the polarization of a gravitational-wave signal, provided that signal is seen in multiple detectors [212]. Consider a generically-polarized gravitational-wave burst arriving from direction $\hat{\mathbf{n}}$:

$$h_{ij} = h_+\hat{\mathbf{e}}^+_{ij}(\hat{\mathbf{n}}) + h_\times\hat{\mathbf{e}}^\times_{ij}(\hat{\mathbf{n}}) + h_x\hat{\mathbf{e}}^x_{ij}(\hat{\mathbf{n}}) + h_y\hat{\mathbf{e}}^y_{ij}(\hat{\mathbf{n}}) + h_b\hat{\mathbf{e}}^b_{ij}(\hat{\mathbf{n}}) + h_l\hat{\mathbf{e}}^l_{ij}(\hat{\mathbf{n}}). \tag{5.15}$$

Measuring the burst's polarization amounts to measuring the six amplitudes h_+, $h_\times$, etc. Given one detector, though, we can only measure the single linear

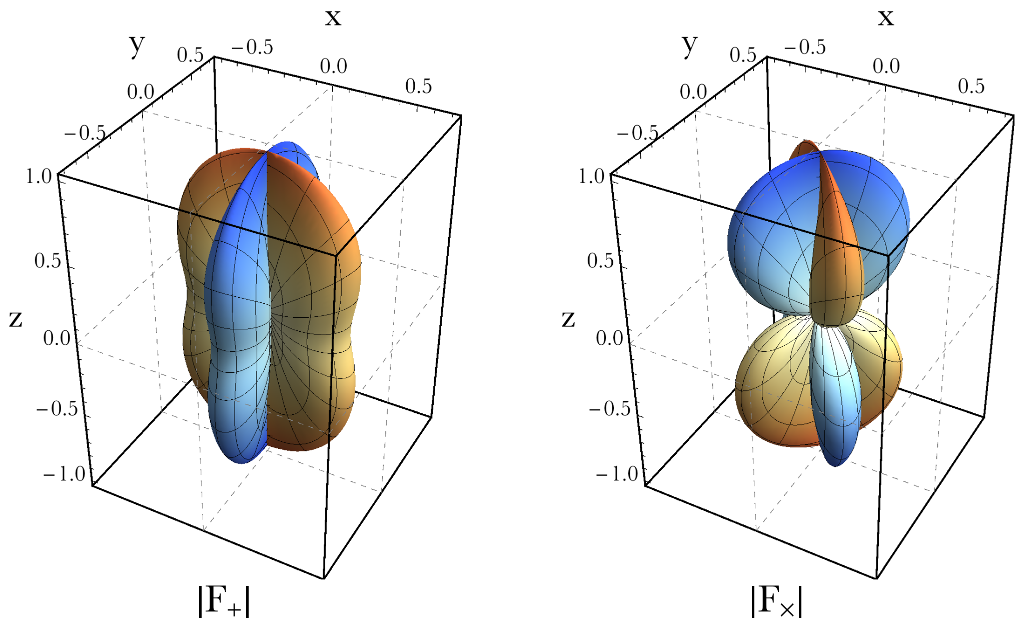

Figure 5.3: Absolute values of the LIGO and Virgo antenna patterns F_+ and $F_\times$ for tensor $+$ and $\times$ modes. The detector is presumed to lie in the $z = 0$ plane, with its vertex at the origin and its arms extending in the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ directions.

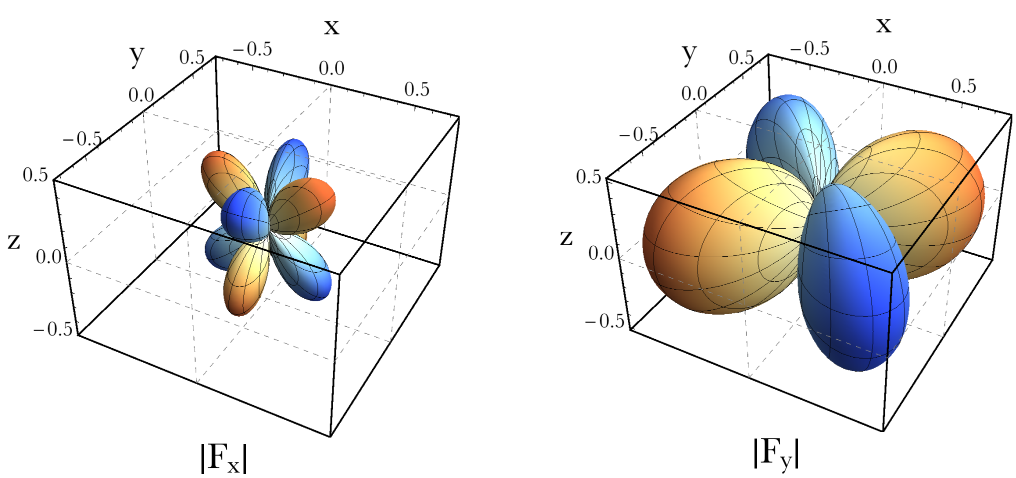

Figure 5.4: Absolute values of the LIGO and Virgo antenna patterns F_x and F_y for vector x and y modes.

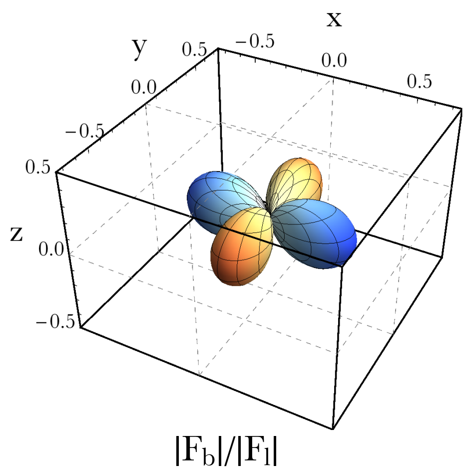

Figure 5.5: Absolute values of the LIGO and Virgo antenna patterns F_b and F_l for scalar b and l modes. The breathing and longitudinal antenna patterns are identical, up to an overall sign.

combination

$$\begin{aligned} s_1 &= F_1^+(\hat{\mathbf{n}})h_+ + F_1^\times(\hat{\mathbf{n}})h_\times + F_1^x(\hat{\mathbf{n}})h_x + F_1^y(\hat{\mathbf{n}})h_y + F_1^b(\hat{\mathbf{n}})h_b + F_1^l(\hat{\mathbf{n}})h_l \\ &= \begin{pmatrix} F_1^+(\hat{\mathbf{n}}) & F_1^\times(\hat{\mathbf{n}}) & F_1^x(\hat{\mathbf{n}}) & F_1^y(\hat{\mathbf{n}}) & F_1^b(\hat{\mathbf{n}}) & F_1^l(\hat{\mathbf{n}}) \end{pmatrix} \cdot \vec{\mathbf{h}}\,, \end{aligned} \tag{5.16}$$

where $\vec{\mathbf{h}} = (h_+\ h_\times\ h_x\ h_y\ h_b\ h_l)^{\mathrm{T}}$ is a vector defining the wave's polarization. This equation clearly isn't invertible. To solve for $\vec{\mathbf{h}}$, in general we need *six* detectors, each measuring signal s_i with antenna patterns F_i^a (with i labeling the detector and a labeling the polarization). If we define a vector $\vec{\mathbf{s}} = (s_1\ s_2\ s_3\ ...)^{\mathrm{T}}$ whose components are these six measurements, then we can write

$$\begin{aligned} \vec{\mathbf{s}} &= \begin{pmatrix} F_1^+(\hat{\mathbf{n}}) & F_1^\times(\hat{\mathbf{n}}) & F_1^x(\hat{\mathbf{n}}) & F_1^y(\hat{\mathbf{n}}) & F_1^b(\hat{\mathbf{n}}) & F_1^l(\hat{\mathbf{n}}) \\ F_2^+(\hat{\mathbf{n}}) & F_2^\times(\hat{\mathbf{n}}) & F_2^x(\hat{\mathbf{n}}) & F_2^y(\hat{\mathbf{n}}) & F_2^b(\hat{\mathbf{n}}) & F_2^l(\hat{\mathbf{n}}) \\ & & & \vdots & & \\ F_6^+(\hat{\mathbf{n}}) & F_6^\times(\hat{\mathbf{n}}) & F_6^x(\hat{\mathbf{n}}) & F_6^y(\hat{\mathbf{n}}) & F_6^b(\hat{\mathbf{n}}) & F_6^l(\hat{\mathbf{n}}) \end{pmatrix} \cdot \vec{\mathbf{h}} \\ &\equiv \mathbf{F} \cdot \vec{\mathbf{h}}. \end{aligned} \tag{5.17}$$

Provided that none of our six detectors are co-located and co-oriented (so that no two rows of $\mathbf{F}$ are linearly-dependent), $\mathbf{F}$ can be inverted and we can solve Eq. (5.17) for $\vec{\mathbf{h}}$.

This isn't quite yet a complete picture, though. There are two additional complicating factors. First, Eq. (5.15) describing our generically polarized gravitational wave has not six but *eight* unknowns: the six polarization amplitudes and two additional parameters describing the event's *a priori* unknown sky location $\hat{\mathbf{n}}$. Given a sufficiently loud and sufficiently brief burst, in principle we can use time-of-flight measurements between detectors to constrain $\hat{\mathbf{n}}$. But in practice localization constraints are non-trivially informed by the relative amplitudes measured between detectors. Thus Eq. (5.17) is optimistic – in reality *at least* six detectors are needed to measure the $6 + 2$ unknowns encapsulated in our hypothetical gravitational-wave burst.

Second, we specified above that we need detectors that are not co-located and co-oriented; two parallel detectors would simply measure the same linear combination s_i of the unknown polarization amplitudes, providing us with no new information. Unfortunately, LIGO's Hanford and Livingston detectors are *very nearly parallel*, with antenna patterns as close to one another's as allowed by the curvature of the Earth. This choice was intentional – LIGO's early researchers wished to maximize the probability that a signal seen in one detector would also appear in the other. Before the first experimental confirmation of gravitational waves, this design choice was quite sensible as a cross-check on any putative gravitational-wave signal. Now, however, this choice means that while we may have a network of three detectors, we can really only measure *two* polarization components.

With these challenges in mind, Advanced LIGO and Virgo have now begun the first direct study of gravitational-wave polarizations. The binary black hole GW170814 was the first event to be observed in all three LIGO and Virgo detectors, allowing the first preliminary analysis of its polarization content [44, 58]. When analyzed with models assuming pure tensor, pure vector, and pure scalar polarizations, GW170814 moderately favored the tensor-only hypothesis, with log-Bayes factors $\ln \mathcal{B}^T_V = 5.3$ and $\ln \mathcal{B}^T_S = 6.9$ (later revised to the less stringent values $\ln \mathcal{B}^T_V = 3.4$ and $\ln \mathcal{B} = 5.4$) relative to the vector-only and scalar-only models [44, 213].

Far more powerful constraints were enabled by the binary neutron star merger GW170817. Above, we saw that, barring independent knowledge of a gravitational wave's source location, not only six but *eight* measurements are needed to fully characterize the gravitational wave's polarization. In the case of

GW170817, this independent knowledge came in the form of a plethora of electromagnetic counterparts [46, 53, 54], which allowed the unambiguous identification of the signal's host galaxy and direction of propagation [60]. GW170817 therefore yielded much more powerful evidence in favor of the tensor-only model, rejecting the vector- and scalar-only hypothesis with $\ln \mathcal{B}^T_V = 47.9$ and $\ln \mathcal{B}^T_S = 53.2$ [52].

These results from GW170817 and GW170814 represent significant first steps in polarization-based tests of gravity. They are, however, somewhat limited – with three detectors, we can say little about the possibility of *mixed* polarizations. We would struggle, for instance, to say anything about the presence of tensor *and* scalar modes within a transient gravitational wave [58]. Future detectors like KAGRA [35, 37] or LIGO-India [39] will therefore be necessary before we can break existing degeneracies and confidently distinguish vector or scalar polarizations in gravitational-wave transients.

5.4 Polarization Measurements with the Gravitational-Wave Background

Above, we argued that (in the best case) six independent measurements are required to characterize the polarization of an arbitrary gravitational-wave. When discussing gravitational-wave transients,[1] six independent measurements means six (non-parallel) detectors. Of course, we expect other classes of gravitational-wave signals beyond transients. In particular, LIGO and Virgo search for a variety of *long-duration* signals, including continuous waves from rotating neutron stars and the stochastic gravitational-wave background. A long-duration signal observed over the course of a sidereal day allows for *multiple* linearly independent measurements even with a single detector – as the Earth rotates, the gravitational-wave signal sweeps through different angles across the detector's antenna patterns. In the language of Eq. (5.17) above, we can construct an invertible $\mathbf{F}$ not by combining antenna patterns $F^a_i(\hat{\mathbf{n}})$ from six different detectors, but by combining the antenna patterns $F^a_1(\hat{\mathbf{n}}_i)$ of a *single* instrument measured in six different directions $\hat{\mathbf{n}}_i$. Unlike gravitational-wave transients, long-lived sources like the stochastic background [214–216] and continuous waves [217, 218] could be leveraged to directly mea-

[1] Here, "transient" means anything whose duration in the LIGO/Virgo frequency band is much smaller than a sidereal day.

sure gravitational-wave polarizations *today*, without the need for additional detectors or electromagnetic counterparts.

Recall that the cross-correlation between measurements from two gravitational-wave detectors gives [83, 219]

$$\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle = \frac{1}{10}\delta(f-f')\gamma(f)\mathcal{H}(f), \tag{5.18}$$

where $\gamma(f)$ is the overlap reduction function between the two detectors and $\mathcal{H}(f)$ is the total strain power of the stochastic gravitational-wave background (see Eq. (3.44)). In terms of the detectors' antenna patterns and separation vector $\Delta\mathbf{x}$, the overlap reduction function is written [186]

$$\gamma(f) = \frac{5}{8\pi}\sum_{A\in\{+,\times\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}, \tag{5.19}$$

normalized so that $\gamma(f) = 1$ for co-incident and co-located detectors. Implicit in Eq. (5.18) is the assumption that *general relativity is correct*; that is, gravitational waves are purely tensor polarized. If we instead step back and allow for the existence of all six possible polarizations, then Eq. (5.18) must be amended.

In general, the output of a single detector is

$$\tilde{s}(f) = \sum_a \int d\hat{\mathbf{n}}\, F^a(\hat{\mathbf{n}})\tilde{h}_a(f,\hat{\mathbf{n}})e^{-2\pi i f\mathbf{x}\cdot\hat{\mathbf{n}}/c}, \tag{5.20}$$

where $a \in \{+,\times,x,y,b,l\}$ indexes the different gravitational-wave polarizations, and the cross-power between two such detectors is

$$\begin{aligned}\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle &= \frac{1}{8\pi}\delta(f-f')\sum_a \mathcal{H}^a(f)\int d\hat{\mathbf{n}}\, F_1^a(\hat{\mathbf{n}})F_2^a(\hat{\mathbf{n}})e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\ &\equiv \frac{1}{8\pi}\delta(f-f')\mathcal{H}^a(f)\int d\hat{\mathbf{n}}\, F_1^a(\hat{\mathbf{n}})F_2^a(\hat{\mathbf{n}})e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}.\end{aligned} \tag{5.21}$$

Rather than explicitly writing sums over polarizations, we'll adopt an Einstein summation-like notation, with repeated indices indicating a sum over a. Also note the leading factor of 8π, rather than 16π that appears in Eq. (3.40). This difference arises because $\mathcal{H}^a(f)$ is the strain power *per polarization*, whereas $\mathcal{H}(f)$ in Eq. (3.40) is the total strain power across plus and cross modes (we will explicitly denote this latter quantity $\mathcal{H}^T(f)$ below).

Buried in Eq. (5.21) are our standard assumptions about the gravitational-wave background, namely that it is isotropic, stationary, and Gaussian. We've also assumed that the various gravitational-wave polarizations are mutually uncorrelated, with $\langle \tilde{h}_a(f,\hat{\mathbf{n}})\tilde{h}^*_{a'}(f',\hat{\mathbf{n}}')\rangle$ vanishing unless $a = a'$, and that tensor and vector sectors are each unpolarized, such that

$$\mathcal{H}^+(f) = \mathcal{H}^\times(f) = \frac{1}{2}\mathcal{H}^T(f) \tag{5.22}$$

and

$$\mathcal{H}^x(f) = \mathcal{H}^y(f) = \frac{1}{2}\mathcal{H}^V(f), \tag{5.23}$$

where $\mathcal{H}^T(f)$ and $\mathcal{H}^V(f)$ are the total strain powers in tensor and vector modes, respectively. This assumption follows from the same argument made in Ch. 3.3 – if gravitational-wave sources are randomly distributed and oriented with respect to the Earth, we have no reason to expect more x-polarized signals than y-polarized ones. The same argument *cannot* be made for the scalar polarizations. Because the scalar breathing and longitudinal polarizations *cannot* be rotated into one another, source isotropy does not imply equal power in each scalar polarization.

The above assumptions are not all equally justifiable, and may be broken by various alternative theories of gravity. For instance, one should not expect an unpolarized background in any parity-violating theory, like Chern-Simons gravity [86–90], even in the absence of non-tensorial modes [220]. Furthermore, different polarizations may not be statistically independent, as is the case for the breathing and longitudinal modes in linearized massive gravity [175]. Finally, we should expect a departure from isotropy in any theory violating Lorentz invariance [205, 208, 209]. These exceptions notwithstanding, for simplicity we'll proceed under the assumptions listed above.

In terms of $\mathcal{H}^T(f)$ and $\mathcal{H}^V(f)$, Eq. (5.21) becomes

$$\begin{aligned}
&\langle \tilde{s}_1(f)\tilde{s}^*_2(f')\rangle \\
&\quad = \frac{1}{8\pi}\delta(f-f') \times \\
&\qquad \Bigg\{ \frac{1}{2}\mathcal{H}^T(f)\int d\hat{\mathbf{n}}\, \left[F^+_1(\hat{\mathbf{n}})F^+_2(\hat{\mathbf{n}}) + F^\times_1(\hat{\mathbf{n}})F^\times_2(\hat{\mathbf{n}})\right] e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&\qquad + \frac{1}{2}\mathcal{H}^V(f)\int d\hat{\mathbf{n}}\, \left[F^x_1(\hat{\mathbf{n}})F^x_2(\hat{\mathbf{n}}) + F^y_1(\hat{\mathbf{n}})F^y_2(\hat{\mathbf{n}})\right] e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&\qquad + \mathcal{H}^b(f)\int d\hat{\mathbf{n}}\, F^b_1(\hat{\mathbf{n}})F^b_2(\hat{\mathbf{n}}) e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&\qquad + \mathcal{H}^l(f)\int d\hat{\mathbf{n}}\, F^l_1(\hat{\mathbf{n}})F^l_2(\hat{\mathbf{n}}) e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \Bigg\}.
\end{aligned} \tag{5.24}$$

Remember, though, that for right-angle interferometric detectors like LIGO and Virgo, the antenna response patterns to breathing and longitudinal polarizations are identical up to a sign: $F_i^b(\hat{\mathbf{n}}) = -F_i^l(\hat{\mathbf{n}})$. So the final two lines in Eq. (5.24) can be written

$$\begin{aligned}\mathcal{H}^b(f)\int d\hat{\mathbf{n}}\, F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} &+ \mathcal{H}^l(f)\int d\hat{\mathbf{n}}\, F_1^l(\hat{\mathbf{n}})F_2^l(\hat{\mathbf{n}})e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\ &= \left[\mathcal{H}^b(f)+\mathcal{H}^l(f)\right]\int d\hat{\mathbf{n}}\, F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\ &= \left[\mathcal{H}^b(f)+\mathcal{H}^l(f)\right]\int d\hat{\mathbf{n}}\, \frac{F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})+F_1^l(\hat{\mathbf{n}})F_2^l(\hat{\mathbf{n}})}{2}e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\ &= \frac{1}{2}\mathcal{H}^S(f)\int d\hat{\mathbf{n}}\, \left[F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})+F_1^l(\hat{\mathbf{n}})F_2^l(\hat{\mathbf{n}})\right]e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}.\end{aligned} \tag{5.25}$$

Specifically, in the second line we replaced $F_1^l(\hat{\mathbf{n}})F_2^l(\hat{\mathbf{n}}) = F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})$ in order to factor out $\mathcal{H}^l(f)$. Then, in the third line, we re-expanded $F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}}) = \frac{1}{2}[F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}}) + F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}})] = \frac{1}{2}[F_1^b(\hat{\mathbf{n}})F_2^b(\hat{\mathbf{n}}) + F_1^l(\hat{\mathbf{n}})F_2^l(\hat{\mathbf{n}})]$. Since the responses of Advanced LIGO and Virgo to breathing and longitudinal modes are completely degenerate, we are sensitive only to the total power $\mathcal{H}^S(f)$ in scalar modes, rather than the individual energies in the breathing and longitudinal polarizations [212, 214].

All together, Eq. (5.24) can be rewritten as [214]

$$\boxed{\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle = \frac{1}{10}\delta(f-f')\gamma_A(f)\mathcal{H}^A(f),} \tag{5.26}$$

where the repeated capitalized index A indicates summation over tensor, vector, and scalar modes $A \in \{T, V, S\}$. We now have not one, but three overlap reduction functions, $\gamma_T(f)$, $\gamma_V(f)$, and $\gamma_S(f)$, which separately quantify the response of the baseline to each class of polarization [186, 214]:

$$\gamma_T(f) = \frac{5}{8\pi}\sum_{A\in\{+,\times\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}, \tag{5.27}$$

$$\gamma_V(f) = \frac{5}{8\pi}\sum_{A\in\{x,y\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}, \tag{5.28}$$

and

$$\gamma_S(f) = \frac{5}{8\pi}\sum_{A\in\{b,l\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}. \tag{5.29}$$

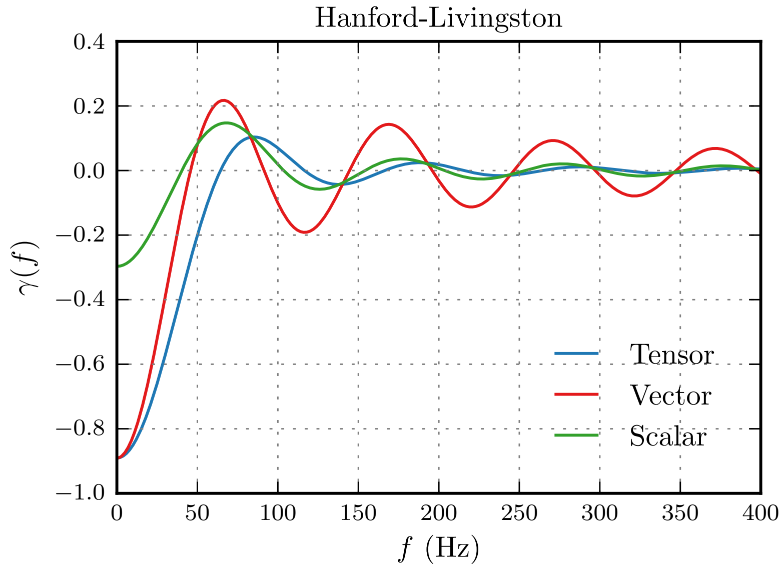

Figure 5.6: Overlap reduction functions quantifying the sensitivity of the Hanford-Livingston baseline to isotropic backgrounds of tensor, vector, and scalar-polarized gravitational waves.

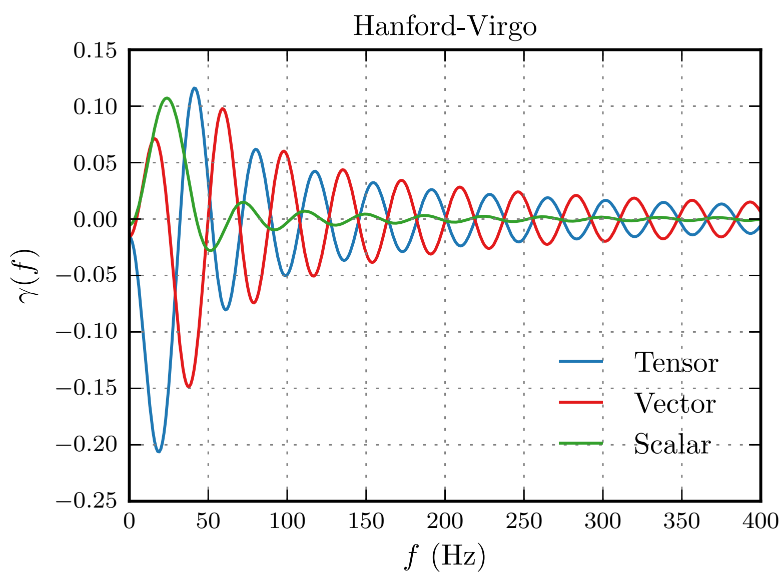

Figure 5.7: As in Fig. 5.6, but for the Hanford-Virgo baseline. Since the distance between Hanford and Virgo is much larger than that between Hanford and Livingston; the Hanford-Virgo overlap reduction functions are smaller in amplitude and more rapidly oscillatory.

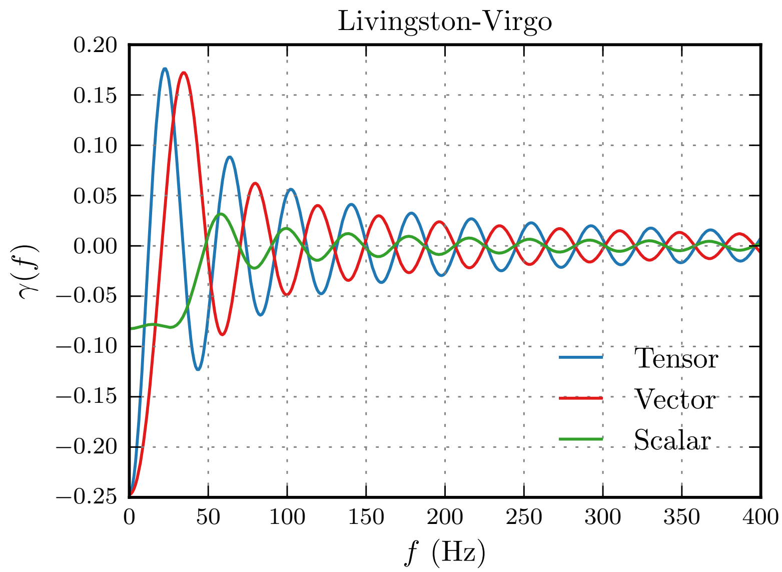

Figure 5.8: As in Fig. 5.6, but for the Livingston-Virgo baseline.

Equations (5.27)-(5.29) are normalized such that co-located and co-oriented right-angle interferometric detectors have $\gamma_T(f) = 1$.

Figure 5.6 shows these three overlap reduction functions for the Hanford-Livingston (H1-L1) Advanced LIGO network. Although the vector and scalar overlap reduction functions are qualitatively similar to $\gamma_T(f)$, there are important quantitative differences between the three curves. First, $\gamma_V(f)$ and $\gamma_T(f)$ are of comparable magnitude at low frequencies, but $\gamma_V(f)$ remains relatively large at frequencies above 64 Hz, where $\gamma_T(f)$ is effectively zero. As a result, we will see that Advanced LIGO is in many cases more sensitive to vector-polarized backgrounds than to standard tensor backgrounds. Second, the scalar overlap reduction function is smallest in magnitude, with $|\gamma_S(0)|$ roughly a factor of three smaller than $|\gamma_T(0)|$ and $|\gamma_V(0)|$. Advanced LIGO is therefore least sensitive to scalar-polarized backgrounds. This reflects a generic feature of quadrupole gravitational-wave detectors, which geometrically have a smaller response to scalar modes than to vector and tensor polarizations [218]. For an extreme example of the opposite case, see pulsar timing arrays, which are orders of magnitude *more* sensitive to longitudinal polarizations than to standard tensor-polarized signals [221–223].

For comparison, Figs. 5.7 and 5.8 show the overlap reduction functions for the

Hanford-Virgo (H1-V1) and Livingston-Virgo (L1-V1) baselines. As the separations between Hanford/Livingston and Virgo are much greater than that between Hanford and Livingston, the H1-V1 and L1-V1 overlap reduction functions are generally much smaller in amplitude and more rapidly oscillatory, translating into weaker sensitivity to the stochastic background. Note, however, that the H1-V1 and L1-V1 tensor overlap reduction functions remain larger in amplitude than H1-L1's at frequencies $f \gtrsim 200\,\mathrm{Hz}$, implying heightened relative sensitivity to tensor backgrounds at high frequencies [224].

The strain powers $\mathcal{H}^A(f)$ appearing in Eq. (5.26) are theory-independent; they are observable quantities that can be directly measured in the detector frame. We conventionally describe gravitational-wave backgrounds not with strain power, though, but by their gravitational-wave energy-density $\Omega(f)$ (see Eq. (3.16)). Within general relativity, the background's energy-density is related to $\mathcal{H}^A(f)$ via [83]

$$\Omega^A(f) = \frac{2\pi^2}{3H_0^2} f^3 \mathcal{H}^A(f). \tag{5.30}$$

As shown in Ch. 3.3, Eq. (5.30) is a consequence of Isaacson's formula for the effective stress-energy of gravitational waves [83, 175, 225]. Alternate theories of gravity, though, can predict different expressions for the stress-energy of gravitational-waves and hence different relationships between strain power $\mathcal{H}(f)$ and energy-density $\Omega(f)$ [175]. For ease of comparison to previous studies, we will use Eq. (5.30) to *define* the **canonical energy-density** $\Omega^A(f)$ in polarization sector A. If we allow Isaacson's formula to hold, then $\Omega^A(f)$ may be directly interpreted as a physical energy density. If not, though, then $\Omega^A(f)$ can instead be understood simply as a derived function of the observable $\mathcal{H}^A(f)$.

5.5 Advanced LIGO's Sensitivity to Backgrounds of Alternative Polarizations

When allowing for the existence of alternative gravitational-wave polarizations, the cross-correlation statistic $\hat{C}(f)$ introduced in Ch. 3.6 now has the expectation value [83, 214]

$$\langle \hat{C}(f) \rangle = \gamma_A(f) \Omega^A(f), \tag{5.31}$$

with variance $\langle C(f)C(f')\rangle = \delta(f-f')\sigma^2(f)$, where

$$\sigma^2(f) = \frac{1}{T}\left(\frac{10\pi^2}{3H_0^2}\right)^2 f^6 P_1(f)P_2(f). \tag{5.32}$$

Here, T is the total coincident observation time between detectors and $P_i(f)$ is the noise power spectral density of detector i.

Remember also that a spectrum of cross-correlation measurements $\hat{C}(f)$ may be combined to obtain a single broadband signal-to-noise ratio, given by

$$\mathrm{SNR}^2 = \frac{\left(\hat{C}\,|\,\gamma_A\Omega_M^A\right)^2}{\left(\gamma_B\Omega_M^B\,|\,\gamma_C\Omega_M^C\right)}, \tag{5.33}$$

defined in terms of the inner product

$$(A\,|\,B) = \left(\frac{3H_0^2}{10\pi^2}\right)^2 2T\int_0^\infty \frac{\tilde{A}^*(f)\tilde{B}(f)}{f^6 P_1(f)P_2(f)}df. \tag{5.34}$$

In Eq. (5.33), $\Omega_M^A(f)$ represents our adopted model for the energy-density spectra of tensor, vector, and scalar modes within the stochastic background. The optimal SNR is obtained when our model matches the background's true energy density, giving

$$\mathrm{SNR}^2_{\mathrm{opt}} = (\gamma_A\Omega^A\,|\,\gamma_B\Omega^B). \tag{5.35}$$

We will continue to model stochastic energy-density spectra as power laws, such that

$$\Omega_M^A(f) = \Omega_0^A\left(\frac{f}{f_0}\right)^{\alpha_A}, \tag{5.36}$$

where Ω_0^A is the amplitude of the gravitational-wave background at frequency f_0 and with polarization A and α_A is the corresponding spectral index. We demonstrated in Ch. 4.2 above that the ordinary (tensorial) stochastic background from compact binary coalescences, for instance, is well-modeled by a power law of slope $\alpha_T = 2/3$ in the sensitivity band of Advanced LIGO [226]. Slopes of $\alpha = 0$ and $\alpha = 3$, meanwhile, correspond to scale-invariant energy and strain spectra, respectively. While we will largely stick to power-law models in our analysis, in Ch. 5.8 we will also explore the potential consequences if this assumption is in fact incorrect (as would be the case, for instance, for a background of unexpectedly massive binary black holes [226]). Throughout this chapter we will use the reference frequency $f_0 = 25\,\mathrm{Hz}$.

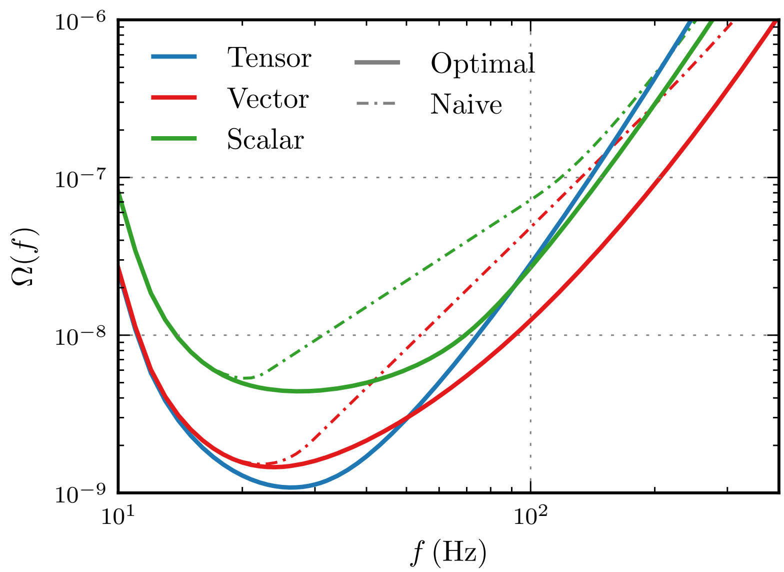

Figure 5.9: Power-law integrated (PI) curves showing the sensitivity of Advanced LIGO to stochastic backgrounds of tensor, vector, and scalar polarizations (solid blue, red, and green, respectively). Power-law energy-density spectra drawn tangent to the PI curves have expected $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 3$ after three years of observation at design sensitivity. Also shown are "naive" PI curves for vector and scalar backgrounds (dashed red and green) illustrating the sensitivity of existing search methods optimized only for tensor polarizations.

With this formalism in hand, we are now equipped to quantify Advanced LIGO's sensitivity to stochastic backgrounds of alternative polarizations. Plotted in Fig. 5.9 are power-law integrated curves (PI) curves representing Advanced LIGO's optimal sensitivity to power-law backgrounds of pure tensor (solid blue), vector (solid red), and scalar (solid green) modes [184]. PI curves are defined by the locus of power-spectra (with slopes ranging from $-\infty$ to $+\infty$) that are individually detectable with $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 3$ after three years of observation with design-sensitivity Advanced LIGO. In quasi-closed form,

$$\mathrm{PI}(f) = \min_{\alpha} \left\{ \Omega_{0,\alpha} \left(f/f_0\right)^{\alpha} \right\}, \tag{5.37}$$

where the amplitudes $\Omega_{0,\alpha}$ are chosen such that each power law $\Omega_{0,\alpha} \left(f/f_0\right)^{\alpha}$ has an optimal SNR of 3. Energy-density spectra lying above and below the PI curves will generally have optimal SNRs greater and less than 3, respectively. The solid curves in Figure 5.10, meanwhile, explicitly show the background amplitudes required for a marginal detection ($\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 3$ after three years of observation) as a function of spectral index.

For spectral indices $\alpha_A \lesssim 0$, we see that Advanced LIGO is approximately

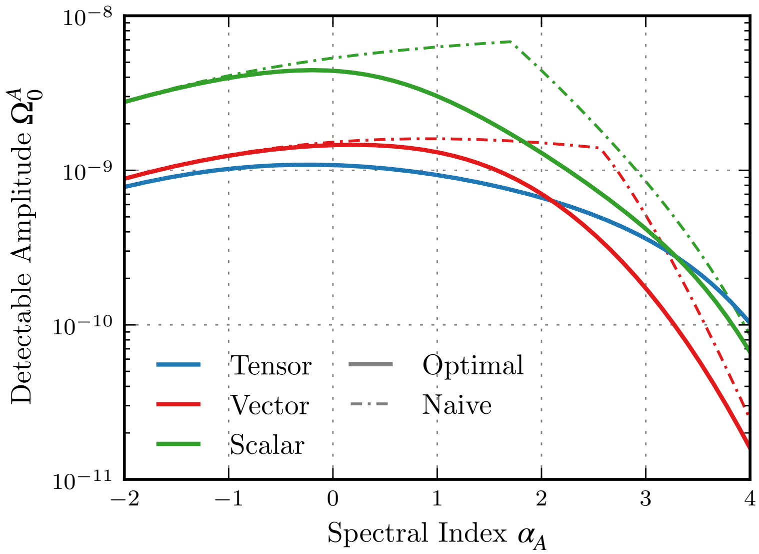

Figure 5.10: Minimum detectable background amplitudes ($\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 3$ after three years of observation at design sensitivity) as a function of spectral index α_A. For small and negative values of α_A, Advanced LIGO is approximately equally sensitive to backgrounds of all three polarizations. For large α_A, Advanced LIGO is instead most sensitive to vector and scalar-polarized backgrounds. The dashed curves show amplitudes detectable with existing "naive" methods. The sensitivity loss between optimal and naive cases is negligible for $\alpha_A \lesssim 0$, but becomes significant at moderate positive slopes (e.g. $\alpha_A \sim 2$). The kinks in the naive curves are due to biases incurred when recovering vector and scalar backgrounds with purely-tensor models; see the text for details.

equally sensitive to tensor and vector-polarized backgrounds, but has reduced sensitivity to scalar signals. When $\alpha_A = 0$, for instance, the minimum optimally-detectable tensor and vector amplitudes are $\Omega_0^T = 1.1 \times 10^{-9}$ and $\Omega_0^V = 1.5 \times 10^{-9}$, while the minimum detectable scalar amplitude is $\Omega_0^S = 4.4 \times 10^{-9}$, a factor of several larger. These values reflect the fact that, as shown in Fig. 5.6, the Hanford-Livingston tensor and vector overlap reduction functions are comparable at low frequencies, while the scalar overlap reduction function is reduced in magnitude.

In contrast, Advanced LIGO's tensor overlap reduction function decays more rapidly at high frequencies than the vector and scalar overlap reduction functions. As a result, Advanced LIGO is more sensitive to vector and scalar backgrounds of large, positive slope than to similarly-shaped tensorial backgrounds. In Fig. 5.9, for instance, the vector and scalar PI curves lie an order of magnitude below the tensor PI curve at frequencies above $f \sim 300\,\mathrm{Hz}$. The constraints that Advanced LIGO can place on positively-sloped vector and scalar backgrounds are therefore as much as an order of magnitude more stringent than those that can be placed on tensor backgrounds of similar slope.

It should be emphasized again that the Hanford-Livingston network's relative sensitivities to tensor, vector, and scalar-polarized backgrounds are due purely to its geometry, rather than properties of the gravitational-waves themselves. If we were instead to consider the Hanford-Virgo baseline, for instance, the right-hand side of Fig. 5.7 shows that at high frequencies the Hanford-Virgo pair is least sensitive to scalar polarizations, whereas the Hanford-Livingston baseline is least sensitive to tensor modes in this same band.

So far we have discussed only Advanced LIGO's *optimal* sensitivity to stochastic backgrounds of alternative polarizations. Existing stochastic searches, though, are *not* optimal for such backgrounds, instead using models $\Omega^A_M(f)$ that allow only for tensor gravitational-wave polarizations.[2] The dashed curves in Figs. 5.9 and 5.10 illustrate Advanced LIGO's "naive" sensitivity to backgrounds of alternative polarizations if we were to *incorrectly* assume a purely-tensor model. The failure to correctly model the polarization of the stochastic background carries two consequences.

The first consequence is a simple reduction in SNR, translating into decreased sensitivity to backgrounds of vector and scalar modes. Sensitivity loss is fairly minimal for slopes $\alpha_A \lesssim 0$. When $\alpha_S = 0$, for example, the minimum detectable scalar amplitude rises from $\Omega^S_0 = 4.4 \times 10^{-9}$ in the optimal case to 5.3×10^{-9} in the naive case, an increase of 20%. The SNR penalty grows more severe, however, for stochastic backgrounds of moderate positive slope. For $\alpha_S = 2$, Advanced LIGO can optimally detect a scalar background of amplitude $\Omega^S_0 = 1.3 \times 10^{-9}$, while "naive" methods would detect only a background of amplitude $\Omega^S_0 = 4.4 \times 10^{-9}$, a factor of 3.4 larger. Note that, since the SNR of the stochastic search accumulates only as SNR $\propto \sqrt{T}$, even a small decrease in sensitivity can result in a somewhat severe increase in the time required to make a detection. To illustrate this, Fig. 5.11 shows the ratio $T_{\text{Naive}}/T_{\text{Optimal}}$ between the observing times required for Advanced LIGO to detect vector (red) and scalar (green) backgrounds using existing "naive" methods and optimal methods. Even at $\alpha_S = 0$, where we saw minimal sensitivity loss in Fig. 5.10, naive methods would require at least 50% more observing time to

[2] Since publication of this work, LIGO and Virgo now perform searches optimized for stochastic backgrounds of alternative polarizations; see Ch. 6.

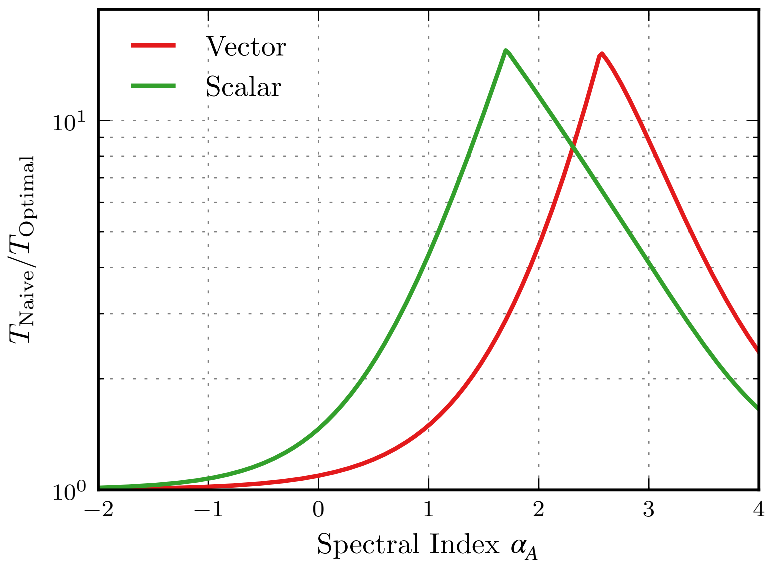

Figure 5.11: The fractional increase in observing time required for Advanced LIGO to make a detection of vector (red) and scalar (green) backgrounds when incorrectly assuming pure tensor polarizations, as a function of the backgrounds' spectral indices α_A. The sharp kinks in each curve are due to biases incurred when fitting vector and scalar backgrounds with a purely-tensor model.

detect a scalar-polarized stochastic background. Since the stochastic background is expected to be optimally detected only after several years, even a 50% increase potentially translates into years of additional observation time, a requirement which may well stress standard experimental lifetimes and operational funding cycles.

The second (arguably more dangerous) consequence is *bias*. If we attempted to fit a sufficiently loud vector or scalar-polarized background with a purely tensorial model, we would invariably recover *something*, but there is no guarantee that the best-fit spectral index α_T under the tensorial model would match the background's true spectral index α_V or α_S. Hence we would suffer from "stealth bias," unknowingly recovering heavily-biased estimates of the amplitude and spectral index of the stochastic background [227, 228]. In general, when fitting a purely tensorial model to vector and scalar backgrounds with slopes $\alpha_{V/S} \gtrsim 3$, our best-fit slope is biased towards large values, such that $\alpha_T > \alpha_{V/S}$. Meanwhile, vector and scalar backgrounds with $\alpha_{V/S} \lesssim 1$ yield best-fit slopes that are biased in the opposite direction, towards smaller values. The sharp kinks in Figs. 5.10 and 5.11 occur at the transitions between these two regimes. Such biases indicate another pitfall of search methods designed only for tensor-polarizations. Even if a vector or scalar-polarized background is recovered with minimal SNR loss, without some independent confirmation we may remain entirely unaware that the detected background indeed violates general relativity.

5.6 Identifying the polarization of the gravitational-wave background

We have seen in Ch. 5.5 that, even when using existing methods assuming only standard tensor polarizations, Advanced LIGO may still be capable of detecting a stochastic background of vector or scalar modes (albeit after potentially much longer observation times). Detection is only the first of two hurdles, though. Once the stochastic background has been detected, we will still need to establish whether it is entirely tensor-polarized, or if it contains vector or scalar-polarized gravitational waves.

Since tensor, vector, and scalar gravitational-wave polarizations each enter into cross-correlation measurements [Eq. (5.31)] with unique overlap reduction functions, the polarization content of a detected stochastic background is in principle discernible from the spectral shape of $\hat{C}(f)$. As an example, Fig. 5.12 shows simulated cross-correlation measurements $\hat{C}(f)$ for both purely tensor (blue) and purely scalar-polarized (green) backgrounds after three years of observation with design-sensitivity Advanced LIGO. Note that these are *extremely* strong backgrounds, with spectra $\Omega^T(f) = 5\times10^{-8}(f/f_0)^{2/3}$ and $\Omega^S(f) = 1.8\times10^{-7}(f/f_0)^{2/3}$; each would be detectable with $\langle \mathrm{SNR}_{\textsc{opt}}\rangle = 150$ after three years. The dashed curves trace the expectation values $\langle\hat{C}(f)\rangle$ of the cross-correlation spectra for each case, while the solid curves show a particular instantiation of measured values. The alternating signs (positive or negative) of each spectrum are determined by the tensor and scalar overlap reduction functions, which have zero-crossings at different characteristic frequencies (see Fig. 5.6). As a result, the tensor and scalar-polarized gravitational waves each impart a unique shape to the cross-correlation spectra, offering a means of visually discriminating between the two cases.

As mentioned above, though, the backgrounds shown in Fig 5.12 are unphysically loud, with $\mathrm{SNR}_{\textsc{opt}} = 150$. A tensor background of this amplitude would *already have been detected* with the standard isotropic search over Advanced LIGO's O1 observing run [176]. Since stochastic searches accumulate SNR

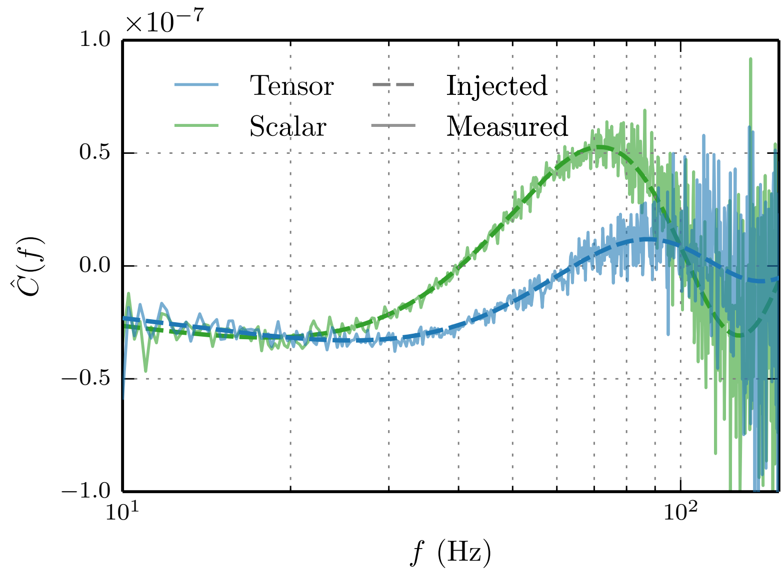

Figure 5.12: Simulated cross-correlation measurements $\hat{C}(f)$ for purely tensor (blue) and purely scalar (green) stochastic backgrounds, recovered after three years of observation with design-sensitivity Advanced LIGO. The backgrounds shown have $\alpha_T = \alpha_S = 2/3$, and have amplitudes chosen such that each is detectable with $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 150$. The measured spectra each show distinct modulations characteristic of the tensor and scalar overlap reduction functions, allowing a clear identification of the polarization in each case.

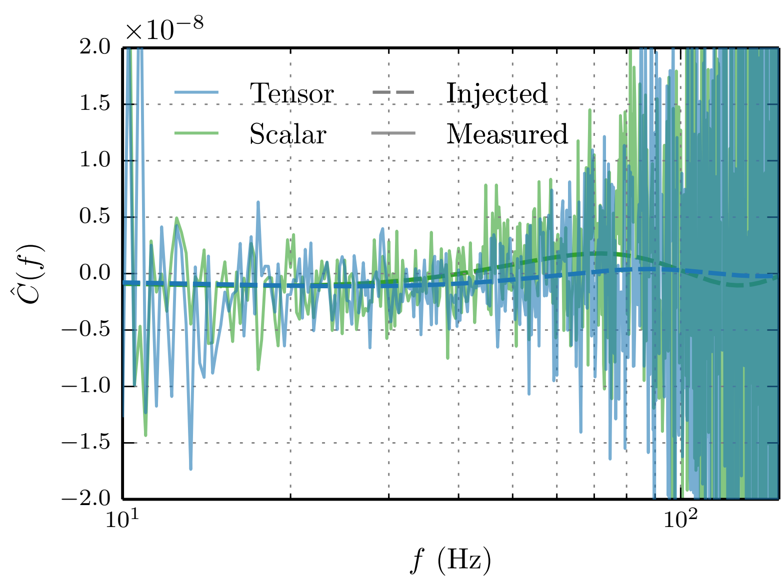

Figure 5.13: As in Fig. 5.12, but now with weaker tensor and scalar backgrounds, detectable with $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 5$ after three years of observation at design sensitivity. While each background is still sufficiently loud to be confidently detected by existing search techniques, the characteristic amplitude modulations and hence the polarization content of each simulated background are no longer evident.

slowly with time, the first detection of the stochastic background will necessarily be marginal; in this case the presence of alternative gravitational-wave polarizations will not be evident. To demonstrate this, Fig. 5.13 shows the simulated recovery of relatively weaker tensor and scalar backgrounds of spectral shape $\Omega^T(f) = 1.7 \times 10^{-9}(f/f_0)^{2/3}$ and $\Omega^S(f) = 6.1 \times 10^{-9}(f/f_0)^{2/3}$, again after three years of observation with Advanced LIGO. Note that these are still considered "loud" signals – both backgrounds would be confidently detected with $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle > 5$ after three years. Despite this, the backgrounds' polarization content is no longer obvious.

Interestingly, even when naively searching for purely-tensor polarized backgrounds, design-sensitivity Advanced LIGO would still detect the quieter scalar background in Fig. 5.13 with SNR = 5.0. This again serves to demonstrate that, when assuming *a priori* that the stochastic background is purely tensor-polarized, any vector or scalar contributions detected with standard search methods may simply be mistaken for ordinary tensor modes. Not only would vector or scalar components fail to be identified, but, as discussed in Ch. 5.5, they would heavily bias parameter estimation of the tensor energy-density spectrum. If we wish to test general relativity with the stochastic background, we will therefore need to develop new tools in order to formally quantify the presence (or absence) of vector or scalar polarizations. Additionally, while we have so far investigated only backgrounds of pure tensor, vector, or scalar polarization, most plausible alternative theories of gravity will predict backgrounds of *mixed* polarization, with vector or scalar components in addition to a tensor component. Any realistic approach must therefore be able to handle a stochastic background of completely generic polarization content.

Our approach will be to detect and classify the stochastic background using Bayesian model selection, adapting the method used in Ref. [218] to study the polarization content of continuous gravitational-wave sources. Specifically, given cross-correlation measurements of the gravitational-wave background, we will calculate **odds ratios** quantifying (*i*) whether a stochastic signal has been detected and (*ii*) whether that stochastic signal contains evidence for alternative gravitational-wave polarizations.

5.6.1 Bayesian Statistics and Odds Ratios

Before proceeding with these calculations, we will first take a brief detour and introduce the fundamentals of **Bayesian inference**.

Consider two experimental outcomes, denoted A and B, and the **joint probability** $P(A \cap B)$ ("the probability of A and B") that *both* outcomes are observed. This joint probability can be rewritten

$$P(A \cap B) = P(A|B)P(B), \tag{5.38}$$

where $P(B)$ is the probability of outcome B (irrespective of A), and $P(A|B)$ ("the probability of A given B") is the **conditional probability** of A if we *a priori* assume B is true. Certainly, though, $P(A \cap B)$ is symmetric in A and B. So we could alternatively have written

$$\begin{aligned} P(A \cap B) &= P(B \cap A) \\ &= P(B|A)P(A). \end{aligned} \tag{5.39}$$

Combing Eqs. (5.38) and (5.39) and solving for $P(A|B)$ gives

$$\boxed{P(A|B) = \frac{P(B|A)P(A)}{P(B)}.} \tag{5.40}$$

Known as **Bayes' theorem**, this simple expression serves as the foundation for most statistical inference across the gravitational-wave community. In essence, Bayes' theorem tells us how to "invert" conditional probabilities – given $P(B|A)$, we can use Eq. (5.40) to compute $P(A|B)$.

The most common application of Bayes' theorem is **parameter estimation**. Consider an Advanced LIGO/Virgo observation of a compact binary merger. Let the variable d symbolically represent the measured strain data, while $\vec{\theta}$ represents the set of (unknown) source parameters – the masses, spins, distance, etc. of the compact binary. Our goal, of course, is to leverage the data d to tell us about the binary's source parameters $\vec{\theta}$. In other words, our aim is to calculate the conditional probability $p(\vec{\theta}|d, \mathcal{S})$ that the source is described by parameters $\vec{\theta}$, given our measurements d and the assumption that a true signal is indeed present (hypothesis $\mathcal{S}$).

Problematically, though, we do not know $p(\vec{\theta}|d, \mathcal{S})$. Instead, what we *do* know (provided that the statistics of instrumental noise are properly understood) is the probability $p(d|\vec{\theta}, \mathcal{S})$ with which we expect to obtain our particular realization of data d, given a source with fixed parameters $\vec{\theta}$. Bayes' theorem provides a prescription for converting between what we know, $p(d|\vec{\theta}, \mathcal{S})$, and what we want:

$$\boxed{p(\vec{\theta}|d, \mathcal{S}) = \frac{p(d|\vec{\theta}, \mathcal{S})p(\vec{\theta}, \mathcal{S})}{p(d, \mathcal{S})}.} \tag{5.41}$$

In practice, the explicit dependence of Eq. (5.41) on hypothesis $\mathcal{S}$ is often suppressed, giving

$$p(\vec{\theta}|d) = \frac{p(d|\vec{\theta})p(\vec{\theta})}{p(d)}, \tag{5.42}$$

where the presence of $\mathcal{S}$ within each term is implied. Different names are conventionally reserved for the different factors in Eq. (5.41). The quantity $p(\vec{\theta}|d)$ is called the **posterior probability** of $\vec{\theta}$, while $p(d|\vec{\theta})$ is the **likelihood**. Meanwhile, $p(\vec{\theta})$ is the **prior** on parameters $\vec{\theta}$ and $p(d)$ is known as the **evidence**.

The prior $p(\vec{\theta})$ is sometimes the source of contention. Unlike the likelihood, which has a single objective definition, there *is no* single correct choice for $p(\vec{\theta})$. Instead, the prior should be chosen to represent one's state of knowledge of $\vec{\theta}$, before the observation in question is carried out. For example, lacking any special prognoses about where in the Universe the next compact binary will occur, we might assume that binary mergers are equally likely to occur anywhere on the sky and hence choose a uniform prior $p(\hat{\mathbf{n}})$ on the sky direction $\hat{\mathbf{n}}$ of a source. Alternatively, when searching for continuous gravitational waves from rotating neutron stars we might choose a prior $p(\hat{\mathbf{n}})$ that is strongly concentrated in the plane of the Milky Way, since we expect to detect isolated neutron stars only within our own galaxy. Crucially, the prior is intrinsically subjective. Two different experimenters may well know different pieces of information, and so will naturally adopt two different priors. While this subjectivity is sometimes viewed as a weakness of Bayesian statistics, I believe it should instead be viewed as a strength – the direct appearance of the prior in Eq. (5.41) forces us to be transparent, making explicit the assumptions that, stated or unstated, invariably underlie *all* statistical analyses.

The evidence $p(d)$, in contrast, is much less important for parameter estimation. In Eq. (5.41), the evidence is simply a constant normalization factor, obtained by integrating the numerator over all possible source parameters:

$$p(d) = \int d\vec{\theta}\, p(d|\vec{\theta})p(\vec{\theta}). \tag{5.43}$$

The evidence is sometimes called the **fully-marginalized likelihood**, as we have integrated out the dependence on all parameters $\vec{\theta}$.

The evidence plays a much more central role in the second-most common application of Bayes' theorem: **hypothesis testing** (also called model selection). Given data d, say we now wish to determine which of two hypotheses is better supported by this data? For instance, an extremely common question is whether data better supports the presence of a true gravitational-wave signal (hypothesis $\mathcal{S}$) or whether it shows only evidence for noise (hypothesis $\mathcal{N}$). In general, we can compute the probability $P(\mathcal{A}|d)$ that any one hypothesis $\mathcal{A}$ is correct using Bayes' theorem:

$$p(\mathcal{A}|d) = \frac{p(d|\mathcal{A})p(\mathcal{A})}{p(d)}. \tag{5.44}$$

Note that, in Eq. (5.44), it is the *evidence* $p(d|\mathcal{A})$ (abbreviated just as $p(d)$ in Eqs. (5.41)-(5.43) above) that appears in place of the likelihood. There also appears a (necessarily subjective) prior probability $p(\mathcal{A})$ for hypothesis $\mathcal{A}$. The meaning of the factor $p(d)$ in the denominator of Eq. (5.44), however, is somewhat less clear. Recall that, when performing parameter estimation, the denominator $p(d|\mathcal{S})$ of Eq. (5.41) represented a normalizing integral taken over all possible source parameters $\vec{\theta}$. Analogously, the denominator $p(d)$ of Eq. (5.44) can only be given by a sum over *all possible hypotheses*:

$$p(d) = \sum_{\mathcal{A}} p(d|\mathcal{A})p(\mathcal{A}). \tag{5.45}$$

Obviously, short of omniscient knowledge of all the infinite possible hypotheses, this sum is impossible to compute.[3]

Fortunately, we don't actually need to compute $p(d)$. In any real-world scenario, we are concerned not with $p(\mathcal{A}|d)$ alone but with the *relative* probabilities between two competing hypotheses $\mathcal{A}$ and $\mathcal{B}$. This relative probability is

[3] I sometimes hear $p(d)$ referred to as "God's evidence."

expressed via an **odds ratio**:

$$\begin{aligned}\mathcal{O}^{\mathcal{A}}_{\mathcal{B}} &= \frac{p(\mathcal{A}|d)}{p(\mathcal{B}|d)} \\ &= \boxed{\frac{p(d|\mathcal{A})}{p(d|\mathcal{B})}\frac{p(\mathcal{A})}{p(\mathcal{B})}.}\end{aligned} \tag{5.46}$$

By virtue of taking a ratio, all incalculable factors of $p(d)$ have vanished! The odds ratio quantifies our relative belief in hypotheses $\mathcal{A}$ and $\mathcal{B}$, following the measurement of our new data d. In fact, $\mathcal{O}^{\mathcal{A}}_{\mathcal{B}}$ can be interpreted as literal betting odds between each hypothesis. Odds ratios $\mathcal{O}^{\mathcal{A}}_{\mathcal{B}} \gg 1$ (or $\ln \mathcal{O}^{\mathcal{A}}_{\mathcal{B}} \gg 0$) indicate strong belief in $\mathcal{A}$ over $\mathcal{B}$, whereas odds $\mathcal{O}^{\mathcal{A}}_{\mathcal{B}} \ll 1$ ($\ln \mathcal{O}^{\mathcal{A}}_{\mathcal{B}} \ll 0$) imply the opposite. Within Eq. (5.46), the factor $p(\mathcal{A})/p(\mathcal{B})$ is called the **prior odds** between hypotheses. As with all priors, this factor is again subjective, quantifying how much evidence we require to tip our beliefs in favor of one hypothesis or the other. The ratio $p(d|\mathcal{A})/p(d|\mathcal{B})$, meanwhile, is known as the **Bayes factor**. Many studies work purely in terms of Bayes factors, rather than odds ratios.[4]

In Ch. 4 above, we used maximum likelihood ratios to determine which of two model energy-density spectra (e.g. power-law vs. astrophysical) better described the gravitational-wave background visible to Advanced LIGO. There exists a conceptually simple relationship between maximum likelihood ratios and the admittedly more involved odds ratios described in this chapter.

Consider the one-dimensional likelihood sketched in Fig. 5.14 (our argument will generalize straightforwardly to multi-dimensional likelihoods). Assume that we're adopting a flat prior over a parameter range $\Delta\theta$, and that the likelihood is reasonably well-peaked about an interval of width $\delta\theta \ll \Delta\theta$, with maximum-likelihood value $\mathcal{L}^{\text{ML}}$. Under these assumptions, we can approximate the evidence $p(d|\mathcal{A})$ as

$$\begin{aligned}p(d|\mathcal{A}) &= \int d\theta\, p(d|\theta,\mathcal{A})p(\theta,\mathcal{A}) \\ &\approx \int d\theta\, p(d|\theta,\mathcal{A})\frac{1}{\Delta\theta} \\ &\approx \mathcal{L}^{\text{ML}}_{\mathcal{A}}\frac{\delta\theta}{\Delta\theta},\end{aligned} \tag{5.47}$$

[4] I find this practice misleading. The use of Bayes factors alone, rather than odds ratios, purports to sidestep the subjective choice of prior odds. But whenever Bayes factors are used as measures of statistical significance, there exists an *implicit* choice of equal prior odds $p(\mathcal{A})/p(\mathcal{B}) = 1$ between the relevant hypotheses.

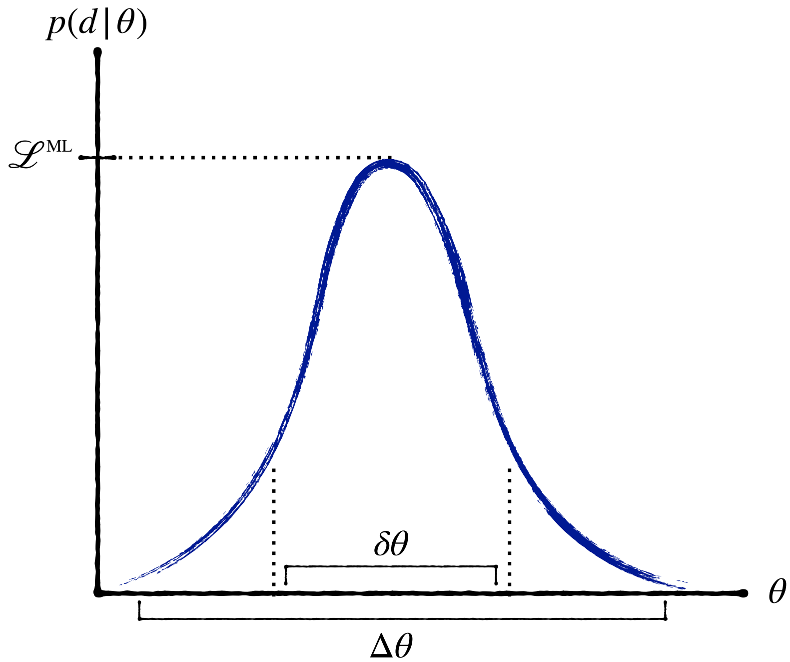

Figure 5.14: Sketch of a toy likelihood distribution, used in Eq. (5.47) to illustrate the relationship between a Bayesian odds ratio and a maximum likelihood ratio. In this toy example, we assume that the likelihood is strongly peaked about a range $\delta\theta$ and that our prior is flat across a range $\Delta\theta$.

where in the second line we used $p(\theta, \mathcal{A}) \approx 1/\Delta\theta$ and in the third line we assumed that the integral over the likelihood can be approximated by the likelihood's height times its width.

Now consider a second hypothesis $\mathcal{B}$ with model parameters ϕ. Again using Eq. (5.47) to approximate the corresponding evidence $p(d|\mathcal{B})$, the odds ratio between $\mathcal{A}$ and $\mathcal{B}$ is approximated by

$$\begin{aligned} \mathcal{O}^{\mathcal{A}}_{\mathcal{B}} &= \frac{p(d|\mathcal{A})}{p(d|\mathcal{B})}\frac{p(\mathcal{A})}{p(\mathcal{B})} \\ &\approx \boxed{\frac{\mathcal{L}^{\mathrm{ML}}_{\mathcal{A}}}{\mathcal{L}^{\mathrm{ML}}_{\mathcal{B}}}\frac{\delta\theta/\Delta\theta}{\delta\phi/\Delta\phi}\frac{p(\mathcal{A})}{p(\mathcal{B})}.} \end{aligned} \tag{5.48}$$

Thus the odds ratio is approximated as a standard likelihood ratio, multiplied by an extra term $\frac{\delta\theta/\Delta\theta}{\delta\phi/\Delta\phi}$ (as well as our prior odds). The factor $\frac{\delta\theta/\Delta\theta}{\delta\phi/\Delta\phi}$ is known as the **Occam's factor**. The fractions $\delta\theta/\Delta\theta$ and $\delta\phi/\Delta\phi$ can be thought of as the *fraction of available parameter space* supported by the data under each model. A simple hypothesis with few free variables, for example, has a relatively small parameter space $\Delta\theta$, and thus the fraction $\delta\theta/\Delta\theta$ is likely to be large. In contrast, an extremely complex model with many free parameters may fit the data quite well, but is likely to do so only in a very small corner $\delta\theta/\Delta\theta$ of the overall available parameter space. The Occam's factor is therefore a formal realization of Occam's razor, imposing an extra penalty on the more complex of the two hypotheses and safe-guarding against the overfitting of data.

5.6.2 Application to Polarization Measurements

With this Bayesian formalism in hand, we can now define two odds ratios to characterize a generically-polarized gravitational-wave background. First, we will define an odds ratio $\mathcal{O}^{\rm SIG}_{\rm N}$ between signal (SIG) and noise (N) hypotheses to determine if a stochastic background (of any polarization) has been observed in our data. Once a background is detected, we can then construct a second odds ratio $\mathcal{O}^{\rm NGR}_{\rm GR}$ to determine if the background contains only tensor polarization (GR hypothesis) or if there is evidence of alternative polarizations (the NGR hypothesis). Unlike previously existing detection methods that assume a pure tensor background, this scheme will allow for the detection of generically-polarized stochastic backgrounds. It encapsulates the optimal detection of tensor, vector, and scalar polarizations as described in Ch. 5.5, and moreover enables the detection of more complex backgrounds of mixed polarization.

To construct odds ratios we need three ingredients: a likelihood, priors on the parameters of each hypothesis considered, and prior odds between hypotheses. As in Ch. 4, we will adopt a Gaussian likelihood for measuring cross-correlation $\hat{C}(f)$ within a single frequency bin of width df [83, 185, 226]:

$$\mathcal{L}\left(\hat{C}(f)|\theta,\mathcal{A}\right) \propto \exp\left(-\frac{\left[\hat{C}(f)-\gamma_A(f)\Omega^A_{\mathcal{A}}(\theta;f)\right]^2}{2\sigma^2(f)}df\right). \tag{5.49}$$

Here, $\gamma_A(f)\Omega^A_{\mathcal{A}}(\theta;f)$ is our model energy-density spectrum under hypothesis $\mathcal{A}$ with parameters θ (recall the implicit summation over polarizations A) and the variance $\sigma^2(f)$ is given by Eq. (5.32). The full likelihood $\mathcal{L}(\{\hat{C}\}|\theta,\mathcal{A})$ for a spectrum of cross-correlation measurements is the product of the individual likelihoods in each frequency bin:

$$\begin{aligned}\mathcal{L}(\{\hat{C}\}|\theta,\mathcal{A}) &\propto \prod_f \mathcal{L}(\hat{C}(f)|\theta,\mathcal{A}) \\ &= \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\hat{C}-\gamma_A\Omega^A_{\mathcal{A}} \,|\, \hat{C}-\gamma_B\Omega^B_{\mathcal{A}}\right)\right],\end{aligned} \tag{5.50}$$

using the inner product defined in Eq. (5.34).

Under the noise hypothesis (N), we assume that no signal is present at all, such that $\Omega^A_{\mathrm{N}}(f) = 0$. Then the corresponding likelihood is simply

$$\mathcal{L}(\{\hat{C}\}|\mathrm{N}) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\hat{C} \,|\, \hat{C}\right)\right]. \tag{5.51}$$

The signal hypothesis (SIG) is somewhat more complex. First, explicitly define a "TVS" hypothesis that allows for the simultaneous presence of tensor, vector, and scalar gravitational-wave polarizations. In this case, we would model the stochastic energy-density spectrum as the sum of three power laws,

$$\Omega_{\mathrm{TVS}}(f) = \Omega^T_0 \left(\frac{f}{f_0}\right)^{\alpha_T} + \Omega^V_0 \left(\frac{f}{f_0}\right)^{\alpha_V} + \Omega^S_0 \left(\frac{f}{f_0}\right)^{\alpha_S}, \tag{5.52}$$

with free parameters Ω^A_0 and α_A setting the amplitude and spectral index of each polarization sector. In defining the TVS hypothesis, though, we have explicitly assumed that tensor, vector, and scalar radiation are all indeed present. This is not the only possibility, of course. A second distinct possibility, for instance, is that only tensor and vector polarizations exist; call this our "TV" hypothesis with the corresponding energy-density spectrum

$$\Omega_{\mathrm{TV}}(f) = \Omega^T_0 \left(\frac{f}{f_0}\right)^{\alpha_T} + \Omega^V_0 \left(\frac{f}{f_0}\right)^{\alpha_V}. \tag{5.53}$$

In a similar fashion, we must ultimately define seven such hypotheses, denoted TVS, TV, TS, VS, T, V, and S, to encompass all combinations of tensor, vector, and scalar gravitational-wave backgrounds. Our complete signal hypothesis is given by the union of these seven sub-hypotheses [190, 218]. Each of the signal sub-hypotheses are logically independent, and so the net odds ratio $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ between signal and noise hypotheses is given by the sum of odds ratios between the noise hypothesis and each of the seven signal sub-hypotheses [190, 218]:

$$\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}} = \sum_{\mathcal{A}\in\{\mathrm{T,V,S,TV},\dots\}} \mathcal{O}^{\mathcal{A}}_{\mathrm{N}}. \tag{5.54}$$

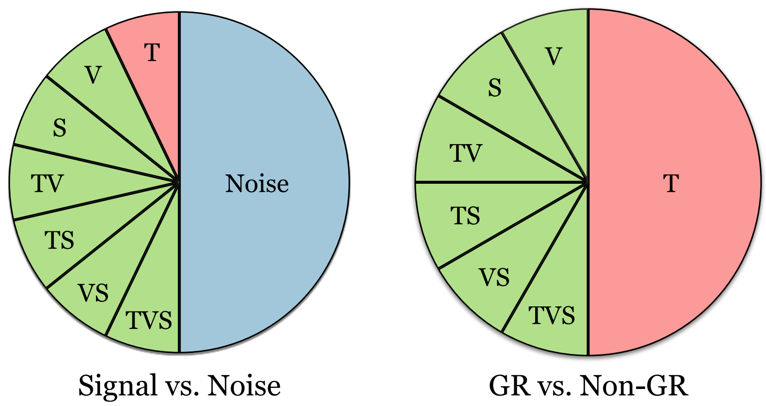

Figure 5.15: Illustration of the prior odds assigned to models and sub-hypotheses when searching for a stochastic background of alternative gravitational-wave polarizations. When constructing $\mathcal{O}^{\text{SIG}}_{\text{N}}$, we assign equal prior probability to the noise and signal hypotheses; within the signal hypothesis, equal probability is given to the seven signal sub-hypotheses $\{\text{T}, ..., \text{TVS}\}$. Similarly, when constructing $\mathcal{O}^{\text{NGR}}_{\text{GR}}$, we give equal probability to the NGR and GR hypotheses and identically weight the six non-GR sub-hypotheses $\{\text{V}, ..., \text{TVS}\}$.

As illustrated in Fig. 5.15, we assign equal prior probability to the signal and noise hypotheses. Within the signal hypothesis, we weight each of the signal sub-hypotheses equally, such that the prior odds between e.g. the T and N hypotheses is $p(\text{T})/p(\text{N}) = 1/7$. This choice of prior probabilities is not unique; other choices may be equally defensible.

The odds ratio $\mathcal{O}^{\text{NGR}}_{\text{GR}}$ is constructed similarly. In this case, we are selecting between the hypothesis that the stochastic background is purely tensor-polarized (GR), or the hypothesis that additional polarization modes are present (NGR). The GR hypothesis is identical to our tensor-only hypothesis T from above, with energy-density

$$\Omega_{\text{GR}}(f) = \Omega_0^T \left(\frac{f}{f_0}\right)^{\alpha_T}. \tag{5.55}$$

The NGR hypothesis, on the other hand, will be the union of the six signal sub-hypotheses that are inconsistent with general relativity: V, S, TV, TS, VS, and TVS. The complete odds ratio between NGR and GR hypothesis is then

$$\mathcal{O}^{\text{NGR}}_{\text{GR}} = \sum_{\mathcal{A}\in\{\text{V},\text{S},\text{TV},...\}} \mathcal{O}^{\mathcal{A}}_{\text{T}}. \tag{5.56}$$

As shown in Fig. 5.15, we have assigned equal priors to the GR and NGR hypotheses as well as identical priors to the six NGR sub-hypotheses.

We still need priors for the various parameters governing each model for the stochastic background. There are two classes of parameters in the various energy-density models presented above: amplitudes Ω_0^A and spectral indices α_A of the background's different polarization components. For each amplitude parameter, we will use the prior

$$p(\Omega_0) \propto \begin{cases} 1/\Omega_0 & (\Omega_{\mathrm{Min}} \leq \Omega_0 \leq \Omega_{\mathrm{Max}}) \\ 0 & (\mathrm{Otherwise}) \end{cases}. \tag{5.57}$$

This corresponds to a uniform prior in the log-amplitudes between $\log \Omega_{\mathrm{Min}}$ and $\log \Omega_{\mathrm{Max}}$. In order for this prior to be normalizable, we cannot let it extend all the way to $\Omega_{\mathrm{Min}} = 0$. Instead, we must choose a finite lower bound. While this lower bound is somewhat arbitrary, our results depend only weakly on the specific choice of bound [218]. Here, we take $\Omega_{\mathrm{Min}} = 10^{-13}$, an amplitude that is indistinguishable from noise with Advanced LIGO. Our upper bound, meanwhile, is $\Omega_{\mathrm{Max}} = 10^{-6}$, consistent with upper limits placed by Initial LIGO and Virgo [229].

We adopt a *triangular* prior on α, centered at zero:

$$p(\alpha) = \begin{cases} \frac{1}{\alpha_{\mathrm{Max}}} \left(1 - \frac{|\alpha|}{\alpha_{\mathrm{Max}}}\right) & (|\alpha| \leq \alpha_{\mathrm{Max}}) \\ 0 & (\mathrm{Otherwise}) \end{cases}. \tag{5.58}$$

This prior has several desirable properties. First, it captures a natural tendency for spectral index posteriors to peak symmetrically about $\alpha = 0$ given uninformative data. As a result, our α posteriors reliably recover this prior in the absence of a stochastic detection (see Fig. 5.26, for example). Second, this prior preferentially weights shallower energy-density spectra. This quantifies our expectation that the stochastic background's energy density is likely to be distributed somewhat uniformly across logarithmic frequency intervals (at least in the LIGO band), rather than entirely at very high or very low frequencies.

Alternatively, Eq. (5.58) can be viewed as corresponding to identical log-uniform priors on the background strength at two different frequencies. A spectral index α may be written as a function of background amplitudes Ω_0 and Ω_1 at two frequencies f_0 and f_1 (see Fig. 5.16):

$$\alpha(\Omega_0, \Omega_1) = \frac{\log\left(\Omega_1/\Omega_0\right)}{\log\left(f_1/f_0\right)}. \tag{5.59}$$

The prior probability of a particular slope α is equal to the probability of drawing any two amplitudes Ω_0 and Ω_1, both with random log-amplitudes between $\log \Omega_{\mathrm{Min}}$ and $\log \Omega_{\mathrm{Max}}$, that together satisfy $\log(\Omega_0/\Omega_1) = \alpha \log(f_0/f_1)$:

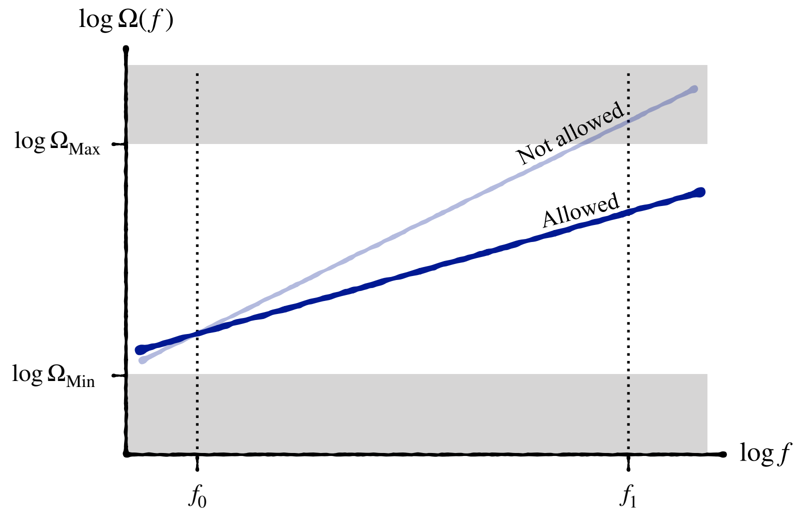

Figure 5.16: Spectral indices allowed by identical log-uniform priors on the gravitational-wave background's amplitude at frequencies f_0 and f_1.

$$p(\alpha) = \int p(\Omega_1) p(\Omega_0)\, \delta\left[\log(\Omega_1/\Omega_0) - \alpha \log(f_1/f_0)\right]\, d\Omega_1. \tag{5.60}$$

For simplicity, we will choose $f_1 = 10 f_0$, so that $\log(f_1/f_0) = 1$. Then Eq. (5.60) reduces to Eq. (5.58).

Although so far we've focused only on the Advanced LIGO Hanford-Livingston baseline, below we will additionally explore how prospects are improved by the three-detector Advanced LIGO-Virgo network. The Bayesian framework discussed here is easily extended to accommodate multiple detector pairs. Together, the three LIGO and Virgo detectors measure three cross-correlation spectra: $\hat{C}^{\rm HL}(f)$, $\hat{C}^{\rm HV}(f)$, and $\hat{C}^{\rm LV}(f)$. In the small signal limit ($\Omega^A(f) \ll 1$), the covariances between these measurements vanish at leading order and so the three baselines can be treated as statistically independent [83]. We can therefore factorize the joint likelihood for the three baselines:

$$
\begin{aligned}
\mathcal{L}(\{\hat{C}^{\text{HL}},&\hat{C}^{\text{HV}},\hat{C}^{\text{LV}}\}|\theta,\mathcal{A}) \\
&= \mathcal{L}(\{\hat{C}^{\text{HL}}\}|\theta,\mathcal{A})\,\mathcal{L}(\{\hat{C}^{\text{HV}}\}|\theta,\mathcal{A})\,\mathcal{L}(\{\hat{C}^{\text{LV}}\}|\theta,\mathcal{A}) \\
&= \frac{1}{(2\pi)^{3/2}}\exp\Bigg\{-\frac{1}{2}\Big[\left(\hat{C}^{\text{HL}}-\gamma_A^{\text{HL}}\Omega_{\mathcal{A}}^A\,|\,\hat{C}^{\text{HL}}-\gamma_B^{\text{HL}}\Omega_{\mathcal{A}}^B\right) \\
&\qquad\qquad + \left(\hat{C}^{\text{HV}}-\gamma_A^{\text{HV}}\Omega_{\mathcal{A}}^A\,|\,\hat{C}^{\text{HV}}-\gamma_B^{\text{HV}}\Omega_{\mathcal{A}}^B\right) \\
&\qquad\qquad + \left(\hat{C}^{\text{LV}}-\gamma_A^{\text{LV}}\Omega_{\mathcal{A}}^A\,|\,\hat{C}^{\text{LV}}-\gamma_B^{\text{LV}}\Omega_{\mathcal{A}}^B\right)\Big]\Bigg\},
\end{aligned}
\tag{5.61}
$$

substituting likelihoods of the form (5.50) for each pair of detectors. Note that we have explicitly distinguished between the overlap reduction functions for each baseline. Other than the above change to the likelihood, all other details of the odds ratio construction are unchanged when including three detectors.

Our Bayesian approach differs in important ways from previous proposed methods with which to search for gravitational-wave backgrounds of alternative polarizations [214–216]. These previous methods endeavor to separate and measure the background's tensor, vector, and scalar content within *each frequency bin.* To solve for these three unknowns, three pairs of gravitational-wave detectors are required to break the degeneracy between polarizations. A nice feature of these methods is that they allow for the separation of polarization modes without the need for a parametrized model of the background's energy-density spectrum. However, they have several drawbacks. First, as already stated, such component separation schemes require at least three detectors. Although we now in fact have three detectors, with the addition of Advanced Virgo, component separation remains not very sensitive; large covariances between polarization modes mean that only very loud backgrounds can be separated and independently detected with reasonable confidence. Finally, component separation methods are largely concerned with the *detection* of a background, not the characterization of its spectral shape. Ref. [216] does discuss parameter estimation on the stochastic background using a Fisher matrix formalism, but Fisher matrices suffer from well-known problems [230].

Instead of attempting to resolve the relative polarization content separately within each frequency bin, we assume a broadband model for the energy-density spectrum in each polarization mode. While our approach is potentially susceptible to bias if our model poorly fits the true background, it is a reasonable model for astrophysically plausible scenarios. Even if the true background differs significantly from this model, we will see in Ch. 5.8 below that potential bias is negligible. Another advantage of our method is that it can be used with only two detectors and hence can be applied *today*, rather than waiting for the construction of future gravitational-wave detectors. Additionally, in Ch. 5.7 we will be able to perform full Bayesian parameter estimation on the stochastic background, properly taking into account the full degeneracies between background parameters (something a Fisher matrix analysis cannot do).

To compute odds ratios $\mathcal{O}^{\rm SIG}_{\rm N}$ and $\mathcal{O}^{\rm NGR}_{\rm GR}$, we use the PyMultiNest package [231], which implements a Python wrapper for the nested sampling software MultiNest [232–234]. MultiNest, an implementation of the nested sampling algorithm [235, 236], is designed to efficiently evaluate Bayesian evidences in high-dimensional parameter spaces, even in the case of large and possibly-curving parameter degeneracies. At little additional computational cost, MultiNest also returns posterior probabilities for each model parameter, allowing for parameter estimation in addition to model selection. Details associated with running MultiNest are given in Ch. 5.B.

5.6.3 Backgrounds of Single Polarizations

As a first demonstration of our Bayesian machinery, we explore the simple cases of purely tensor, vector, or scalar-polarized stochastic backgrounds. Shown in Fig. 5.17 are the distributions of odds ratios $\mathcal{O}^{\rm SIG}_{\rm N}$ obtained for simulated observations of pure tensor and scalar backgrounds, each of slope $\alpha = 2/3$ (the characteristic slope of a tensor binary black hole background). For each polarization, we consider two choices of amplitude, corresponding to $\langle \mathrm{SNR}_{\rm OPT}\rangle = 5$ and 10 after three years of observation with design-sensitivity Advanced LIGO. For comparison, the hatched grey distribution show odds ratios obtained in the presence of pure Gaussian noise. Gaussian noise yields a narrow odds ratio distribution centered at $\ln \mathcal{O}^{\rm SIG}_{\rm N} \approx -1.0$. In contrast, the simulated observations of tensor and scalar backgrounds yield large, positive odds ratios, well-separated from Gaussian noise. Note that the tensor and scalar distributions lie nearly on top of one another, as $\mathcal{O}^{\rm SIG}_{\rm N}$ depends primarily on the optimal SNR of a background and not its polarization content.

Figure 5.18, in turn, shows the odds ratios $\mathcal{O}^{\rm NGR}_{\rm GR}$ quantifying the evidence

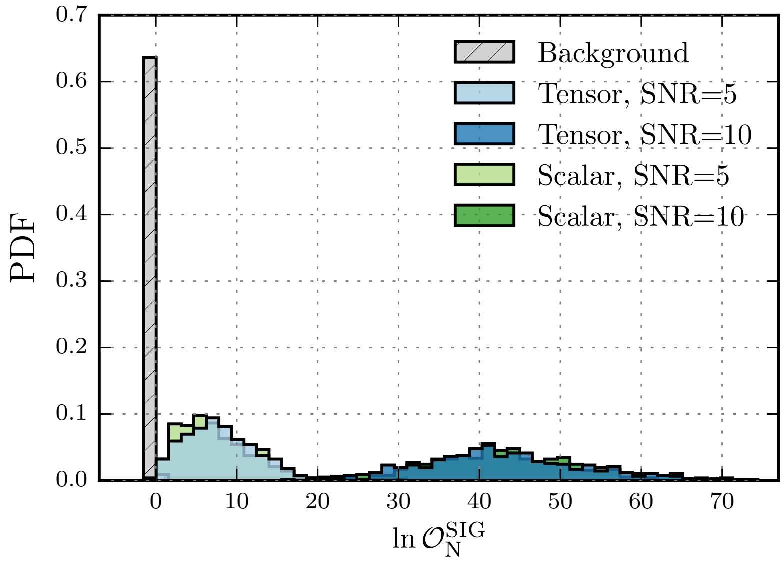

Figure 5.17: Distributions of odd ratios $\mathcal{O}^{\rm SIG}_{\rm N}$ between signal and noise hypotheses for simulated observations of tensor (blue) and scalar (green) stochastic backgrounds of slope $\alpha = 2/3$, assuming three years of observation with design-sensitivity Advanced LIGO. We consider two different strengths for each polarization, corresponding to $\langle \mathrm{SNR}_{\rm OPT} \rangle = 5$ and 10. For each background strength, the tensor and scalar odds ratios lie nearly on top of one another. Also shown is the background distribution of odds ratios obtained when observing pure Gaussian noise (hatched grey). In the presence of a stochastic background, the recovered odds ratios quadratically grow as $\ln \mathcal{O}^{\rm SIG}_{\rm N} \propto \mathrm{SNR}^2_{\rm OPT}$, showing increasingly large preference for the signal hypothesis.

for alternative polarization modes. In the case of pure Gaussian noise, we again see a narrow distribution of odds ratios, centered at $\ln \mathcal{O}^{\rm NGR}_{\rm GR} \approx -0.4$. In the absence of informative data, our analysis thus slightly favors the GR hypothesis. This can be understood as a consequence of the implicit Bayesian "Occam's factor," which penalizes the more complex NGR hypothesis over the simpler GR hypothesis. Simulated observations of scalar backgrounds, in turn, yield large positive values for $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$, correctly preferring the NGR hypothesis. In contrast, pure tensor backgrounds yield negative $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$. Interestingly, the recovered odds ratios do not grow increasingly negative with larger tensor amplitudes, but instead saturate at $\ln \mathcal{O}^{\rm NGR}_{\rm GR} \approx -1.4$. This reflects the fact that a non-detection of vector or scalar polarizations can never strictly rule out their presence, but only place an upper limit on their amplitudes. In other words, a strong detection of a pure tensor stochastic background cannot provide evidence *for* the GR hypothesis, but at best only offers no evidence *against* it. This behavior is in part due to our choice of amplitude priors, which

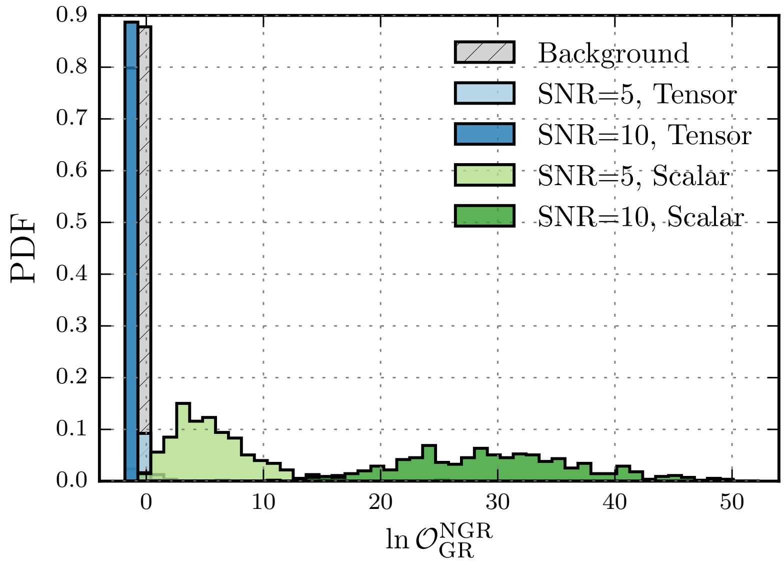

Figure 5.18: Odds ratios $\mathcal{O}^{\rm NGR}_{\rm GR}$ between NGR and GR hypotheses obtained for the same simulated Advanced LIGO observations considered in Fig. 5.17. In the presence of a tensor-polarized background, we recover narrow distributions of odds ratios centered at $\ln \mathcal{O}^{\rm NGR}_{\rm GR} \approx -1.4$, reflecting consistency with the GR hypothesis. A scalar background, on the other hand, yields large positive odds ratios, correctly showing a strong preference for our NGR hypothesis.

allow for finite but immeasurably small vector and scalar energy densities.

Figures 5.19-5.21 illustrate more generally how $\mathcal{O}^{\rm SIG}_{\rm N}$ and $\mathcal{O}^{\rm NGR}_{\rm GR}$ scale with the amplitudes of purely tensor, vector, and scalar-polarized stochastic backgrounds. Black points mark odds ratios computed from individual realizations of simulated data, while the solid curves and shaded regions trace their smoothed mean and standard deviation. We again see $\ln \mathcal{O}^{\rm SIG}_{\rm N}$ increasing monotonically with injected amplitude for all three polarizations. Specifically, $\mathcal{O}^{\rm SIG}_{\rm N}$ depends inversely on the noise-hypothesis likelihood [defined by Eq. (5.51)] and therefore scales as

$$\ln \mathcal{O}^{\rm SIG}_{\rm N} \propto {\rm SNR}^2_{\rm OPT}. \tag{5.62}$$

As we saw earlier in Fig. 5.18, $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ saturates at -1.4 for loud tensor backgrounds. In the case of vector and scalar backgrounds, on the other hand, $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ grows quadratically with increasing amplitude. In particular, $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ is proportional to the squared SNR of the *residuals* between the observed $\hat{C}(f)$ and the best-fit tensor model. We begin to see a strong preference for the NGR hypothesis when these residuals become statistically significant.

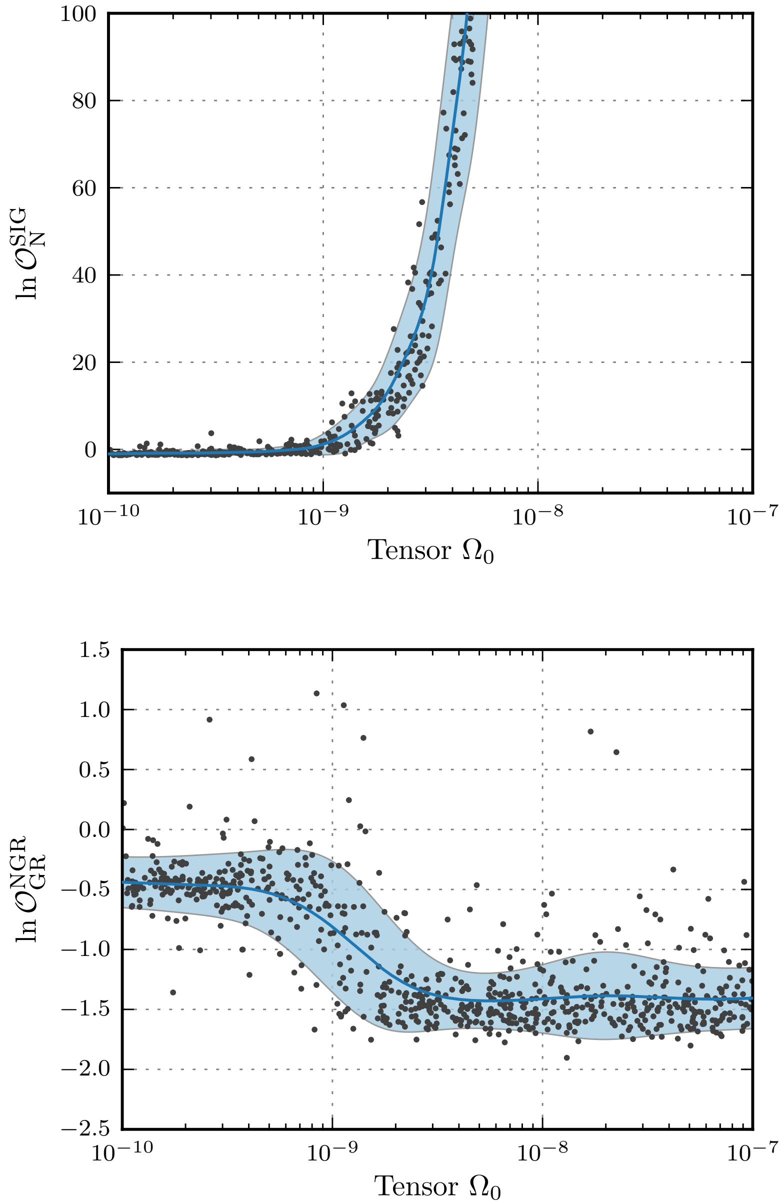

Figure 5.19: Odds ratios $\mathcal{O}_{\mathrm{N}}^{\mathrm{SIG}}$ (top) and $\mathcal{O}_{\mathrm{GR}}^{\mathrm{NGR}}$ (bottom) for simulated Advanced LIGO observations of purely tensor-polarized stochastic backgrounds. Within each plot, we show 750 simulated observations, with random log-amplitudes chosen uniformly over the range $-10 < \log \Omega_0^T < -7$. Black points mark the results from individual realizations, while the solid curves and shaded regions show the moving mean and standard deviations (smoothed with a Gaussian kernel) of these realizations. For each polarization, $\ln \mathcal{O}_{\mathrm{N}}^{\mathrm{SIG}}$ scales quadratically with the amplitude of the stochastic background. The values of $\ln \mathcal{O}_{\mathrm{GR}}^{\mathrm{NGR}}$, meanwhile, saturate at approximately -1.4.

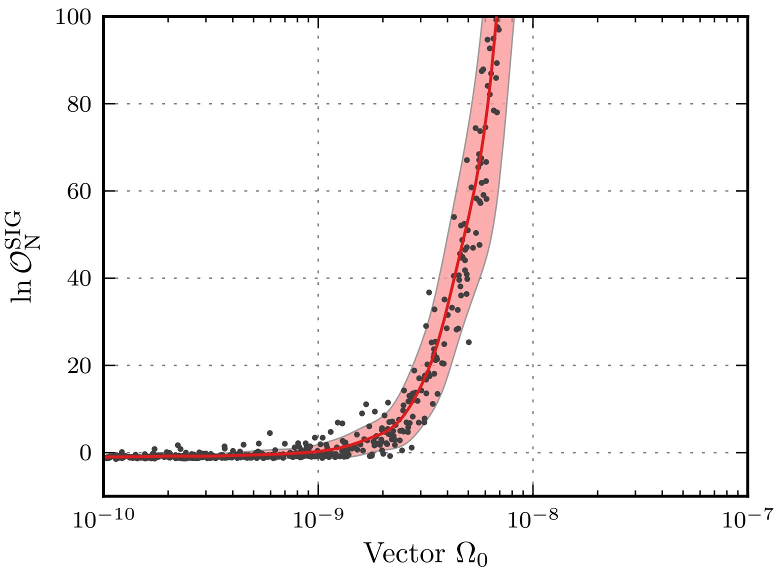

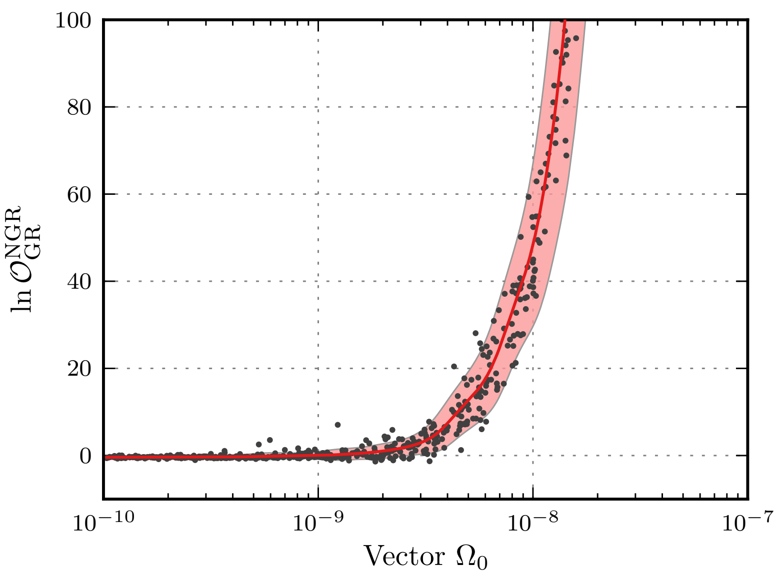

Figure 5.20: As in Fig. 5.19, but for simulated Advanced LIGO observations of purely vector-polarized stochastic backgrounds. In contrast to Fig. 5.19, both $\ln \mathcal{O}_{\mathrm{N}}^{\mathrm{SIG}}$ and $\ln \mathcal{O}_{\mathrm{GR}}^{\mathrm{NGR}}$ increase quadratically with the strength of our injected signals.

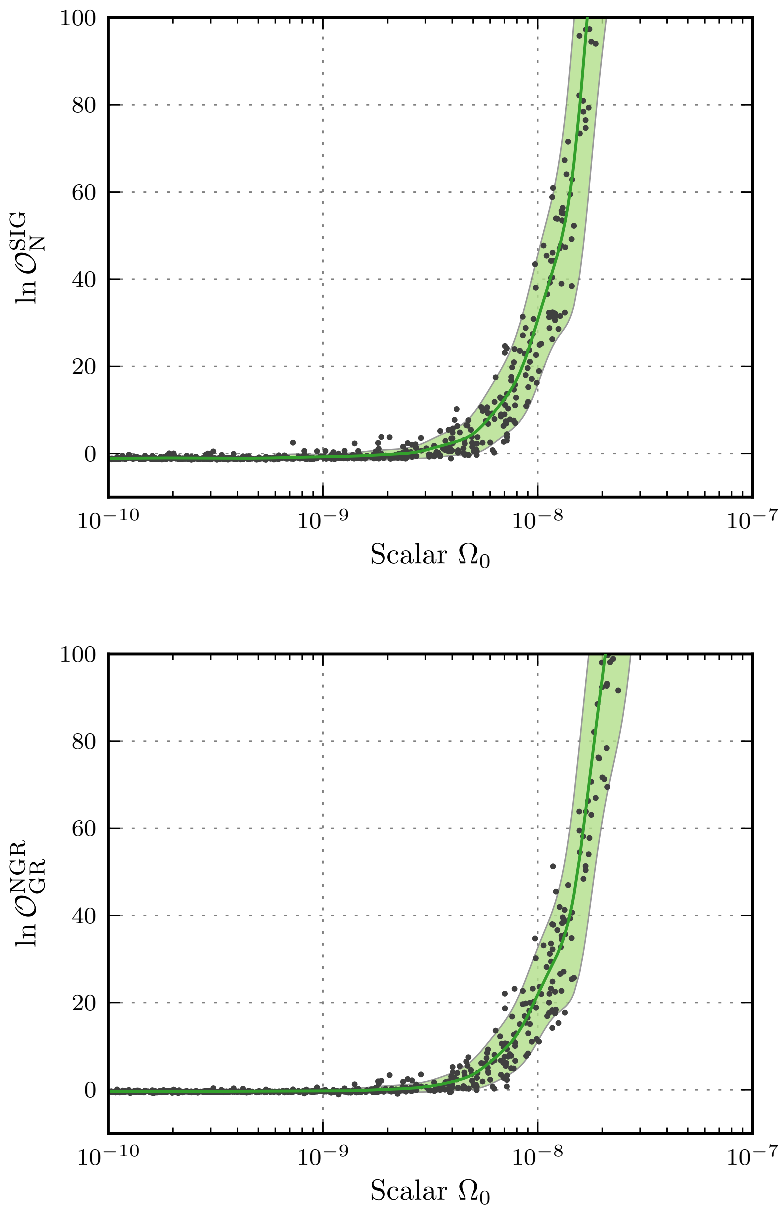

Figure 5.21: As in Fig. 5.19, but for simulated Advanced LIGO observations of purely scalar-polarized stochastic backgrounds.

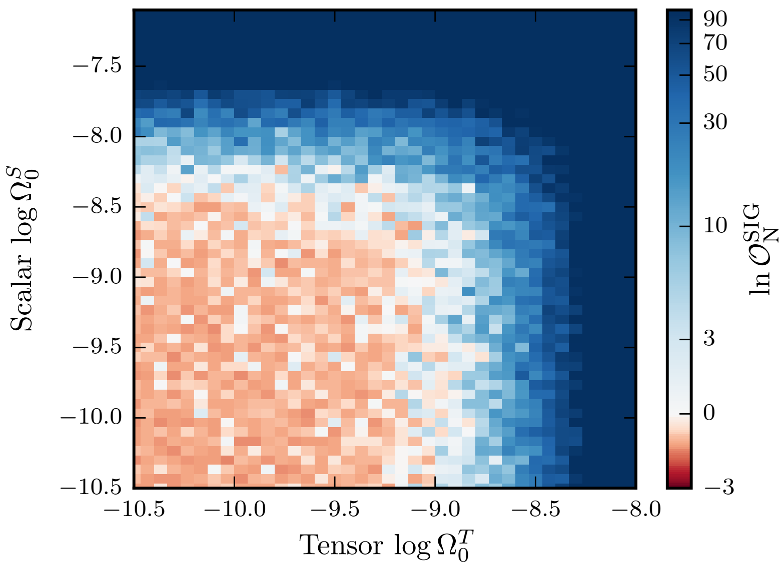

Figure 5.22: Odds ratios $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ for simulated Advanced LIGO measurements of stochastic backgrounds containing both tensor and scalar polarizations, assuming three years of observation at design sensitivity. The tensor and scalar components have slopes $\alpha_T = 2/3$ and $\alpha_S = 0$, respectively. The observed values of $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ effectively trace contours in total background energy. Thus the detectability of a background depends largely on its total power, not its polarization content.

5.6.4 Backgrounds of Mixed Polarization

So far we have considered only cases of *pure* tensor, vector, or scalar polarization. Plausible alternative theories of gravity, however, would typically predict a mixed background of multiple polarization modes. How does our Bayesian machinery handle a background of mixed polarization? To answer this question, we will investigate backgrounds of mixed tensor and scalar polarization. Figures 5.22 and 5.23 illustrate values of $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ and $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$, respectively, as a joint function of the amplitude of each polarization. While we vary the amplitudes of our injected signals, we fix the tensor and scalar slopes to $\alpha_T = 2/3$ (as predicted for binary black hole backgrounds) and $\alpha_S = 0$.

Within Fig. 5.22, the recovered values of $\ln\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ simply trace contours of total energy. Thus the *detectability* of a mixed background depends only on its total measured energy, rather than its polarization content. Meanwhile, Fig. 5.23 exhibits three distinct regions. First, for small tensor and small scalar amplitudes ($\log\Omega^T_0 \lesssim -9.0$ and $\log\Omega^S_0 \lesssim -8.5$), we obtain $\ln\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}} \approx -0.4$. In this region, the mixed background cannot be detected and so simply we recover

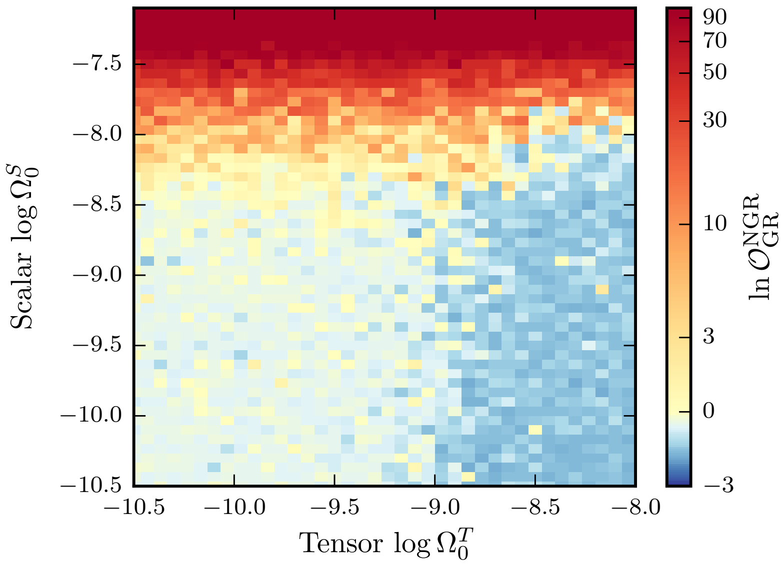

Figure 5.23: As in Fig. 5.22, but now showing odds $\mathcal{O}^{\rm NGR}_{\rm GR}$ between NGR and GR hypotheses. Advanced LIGO would confidently identify the presence of the scalar background component when $\log \Omega_0^S \gtrsim -7.9$. LIGO's sensitivity to the scalar component is nearly independent of the strength of the tensor component; the minimum identifiable scalar amplitude Ω_0^S rises only slightly with increasing Ω_0^T.

the slight Occam's bias towards the GR hypothesis as noted above. Second, for small scalar amplitudes but large tensor amplitudes ($\log \Omega_0^T \gtrsim -9.0$), the recovered odds ratios decrease to $\ln \mathcal{O}^{\rm NGR}_{\rm GR} \approx -1.4$. This corresponds to the detection of the tensor component alone; the decrease in odds ratios is the same behavior previously seen in Figs. 5.18 and 5.19. Finally, when Ω_0^S is large, the scalar component is detectable and so $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ increases rapidly to large, positive values. The threshold value of Ω_0^S at which $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ becomes positive shows only little dependence on the amplitude of any tensor background which might also be present. When Ω_0^T is small, for instance, scalar amplitudes of size $\log \Omega_0^S \gtrsim -7.9$ are required to preference the NGR model. When Ω_0^T is large, this requirement increases only slightly to $\log \Omega_0^S \gtrsim -7.8$. Thus, we should expect Advanced LIGO to be able to both detect and *identify as non-tensorial* a flat scalar background of amplitude $\log \Omega_0^S \gtrsim -8$, regardless of the presence of an additional tensor component.

It should be pointed out that positive values of $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ indicates only that there exists evidence for alternative polarizations. From this odds ratio alone we cannot infer which specific polarizations – vector and/or scalar – are present in the background. Although Advanced LIGO might succeed in identifying mixed tensor-scalar backgrounds as non-tensorial when $\log \Omega_0^S \gtrsim -8$, this *does not* necessarily imply that we can successfully identify the scalar component as such, only that our measurements are not consistent with tensor polarization alone; this qualification will be discussed more in Ch. 5.7 below.

The future addition of new gravitational wave detectors will extend the reach of stochastic searches and help to break degeneracies between backgrounds of different polarizations. This expansion has begun with the completion of Advanced Virgo, which joined Advanced LIGO during its O2 observing run in August 2017 [34, 44]. It is therefore interesting to investigate how the introduction of Advanced Virgo improves the above results. Given detectors indexed by $i \in \{1, 2, ...\}$, the total SNR of a stochastic background is the quadrature sum of SNRs from each detector pair [83]:

$$\mathrm{SNR}^2 = \sum_i \sum_{j>i} \mathrm{SNR}_{ij}^2, \tag{5.63}$$

where each SNR_{ij} is computed according to Eq. (5.33). Naively, the SNR with which a background is observed is expected to increase as $\mathrm{SNR} \propto \sqrt{N}$, where N is the total number of available detector pairs (three in the case of the Advanced LIGO-Virgo network). However, both the Hanford-Virgo and Livingston-Virgo pairs exhibit reduced sensitivity to the stochastic background due to their large physical separations. This fact is reflected in their respective overlap reduction functions, which are a factor of several smaller in magnitude than the Hanford-Livingston overlap reduction functions; see Figs. 5.7 and 5.8.

Given three independent detector pairs (and hence three independent measurements at each frequency), one can in principle directly solve for the unknown tensor, vector, and scalar contributions to the background in each frequency bin [214–216, 219]. This component separation scheme can be performed without resorting to a model for the stochastic energy-density spectrum. However, frequency-by-frequency component separation is unlikely to be successful using the LIGO-Virgo network, due to the large uncertainties in the measured background at each frequency. Instead, when considering joint Advanced LIGO-Virgo observations we will again apply the Bayesian framework introduced above, leveraging measurements made at many frequencies in order to constrain the power-law amplitude and slope of each polarization mode.

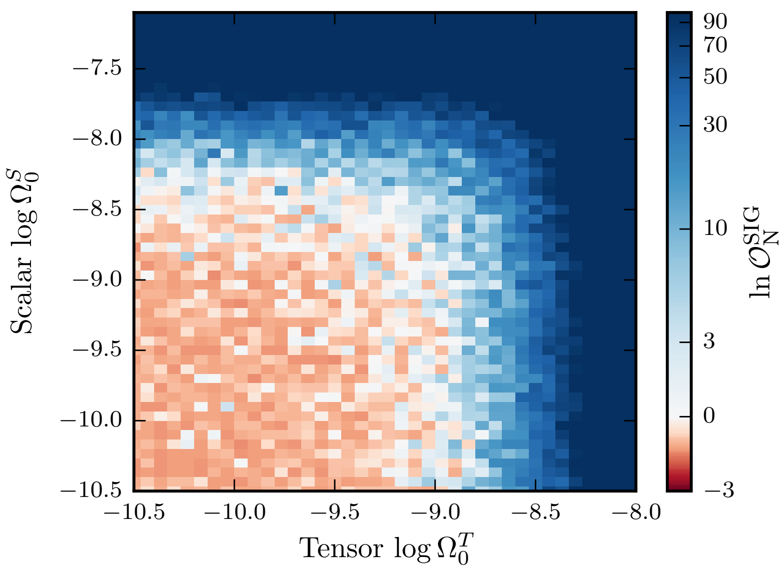

Figure 5.24: As in Fig. 5.22, but for simulated three-year observations with the joint Advanced LIGO-Virgo network at design sensitivity. Despite the inclusion of Advanced Virgo, the sensitivity of this three-detector network is nearly identical to that of Advanced LIGO alone.

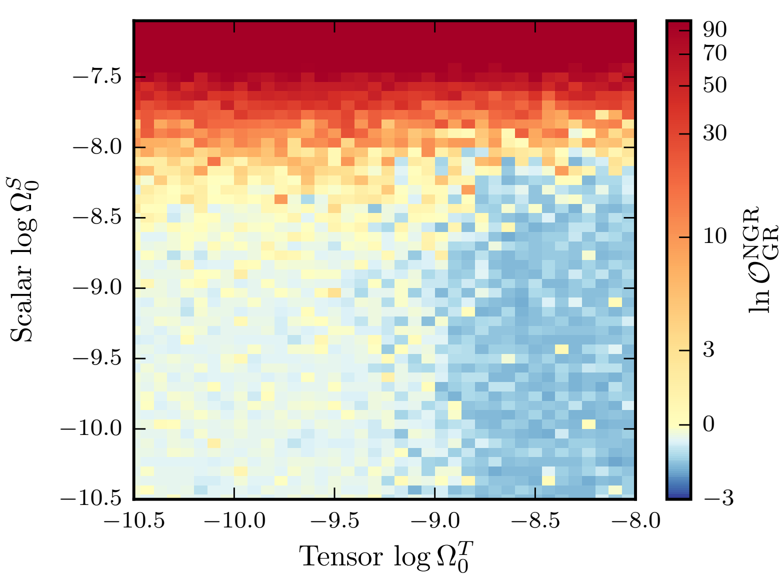

Figure 5.25: As in Fig. 5.23, but assuming three-years of observation with the joint Advanced LIGO-Virgo network at design sensitivity. Once again, the inclusion of Advanced Virgo yields negligible improvement in network sensitivity.

Table 5.1: Stochastic background parameters used for each case study presented. For each case, the vector amplitude is set to zero. Also shown are the odds ratios computed for each simulated observation.

Case	$\log \Omega_0^T$	α_T	$\log \Omega_0^S$	α_S	H1-L1		H1-L1-V1	
					$\ln \mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$	$\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$	$\ln \mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$	$\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$
1. Noise	-	-	-	-	-1.1	-0.4	-1.1	-0.4
2. Tensor	-8.78	0.67	-	-	8.4	-1.4	8.8	-1.4
3. Tensor+Scalar	-8.48	0.67	-7.83	0.0	193.5	16.1	197.3	19.3

To quantify the extent to which Advanced Virgo aids in the detection of the stochastic background, we again consider simulated observations of a mixed tensor (slope $\alpha_T = 2/3$) and scalar (slope $\alpha_S = 0$) background, this time with a three-detector Advanced LIGO-Virgo network. The odds ratios obtained from our simulated Advanced LIGO-Virgo observations are shown in Figs. 5.24 and 5.25 for various tensor and scalar amplitudes. The inclusion of Advanced Virgo yields no clear improvement over the results shown in Figs. 5.22 and 5.23 using Advanced LIGO alone. Due to its large distance from either LIGO detector, Advanced Virgo does not contribute more than a small fraction of the total observed SNR. As a result, the combined Hanford-Livingston-Virgo network both detects (as indicated with $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$) and identifies (via $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$) the scalar background component with virtually the same sensitivity as the Hanford-Livingston network alone.

5.7 Parameter Estimation on Mixed Backgrounds

Parameter estimation will be the final step in a search for a stochastic background of generic polarization. If a gravitational-wave background is detected (as inferred from $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$), how well can Advanced LIGO constrain the properties of the background? Alternatively, if no detection is made, what upper limits can Advanced LIGO place on the background amplitudes of each polarization mode? We investigate these questions through three case studies: an observation of pure Gaussian noise, a standard tensor stochastic background, and a background of mixed tensor and scalar polarizations. The simulated background parameters used for each case are listed in Table 5.1.

When performing model selection above, the odds ratios $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ and $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$ were constructed by independently allowing for each combination of tensor, vector, and scalar modes. Parameter estimation, on the other hand, must necessarily be performed in the context of one *specific* background model. For the case studies below, we will adopt the broadest possible hypothesis, allowing for all three polarization modes – we previously termed this the "TVS" hypothesis, with an energy-density spectrum given by Eq. (5.52). This choice will allow us to place simultaneous constraints on the presence of tensor, vector, and scalar polarizations in the stochastic background. Parameter estimation is achieved using MultiNest, which returns samples drawn from the measured posterior distributions.

There are several key subtleties that must be understood when interpreting the parameter estimation results presented below. First, whereas standard tensor upper limits have conventionally been defined by assuming a single, *fixed* slope [176, 229], we will quote amplitude limits obtained *after* marginalization over all possible spectral indices. This approach concisely combines information from the entire posterior parameter space to offer a single limit on each polarization considered. As a result, however, our simulated upper limits presented here should not be directly compared to those from past searches for tensor backgrounds. Secondly, parameter estimation results are contingent upon the choice of a specific model. While we will demonstrate parameter estimation results under our TVS hypothesis, other hypotheses may be better suited to answering other experimental questions. For example, if we were specifically interested in constraining scalar-tensor theories (which *a priori* do not allow vector polarizations), we would instead perform parameter estimation under the TS hypothesis. And if our goal was to perform a standard stochastic search for a purely tensor-polarized background, we would restrict to the T hypothesis. Although these various hypotheses all contain an analogous parameter Ω_0^T, the resulting upper limits on Ω_0^T will generically be different in each case. In short, different experimental questions will yield different answers.

5.7.1 Case 1: Gaussian Noise

First, we consider the case of pure noise, producing a simulated three-year observation of Gaussian noise at Advanced LIGO's design sensitivity. The resulting TVS posteriors are shown in Fig. 5.26. The colored histograms along the diagonal show the marginalized 1D posteriors for the amplitudes and slopes of the tensor, vector, and scalar components (blue, red, and green, respectively). The priors placed on each parameter are indicated with a dashed grey curve.

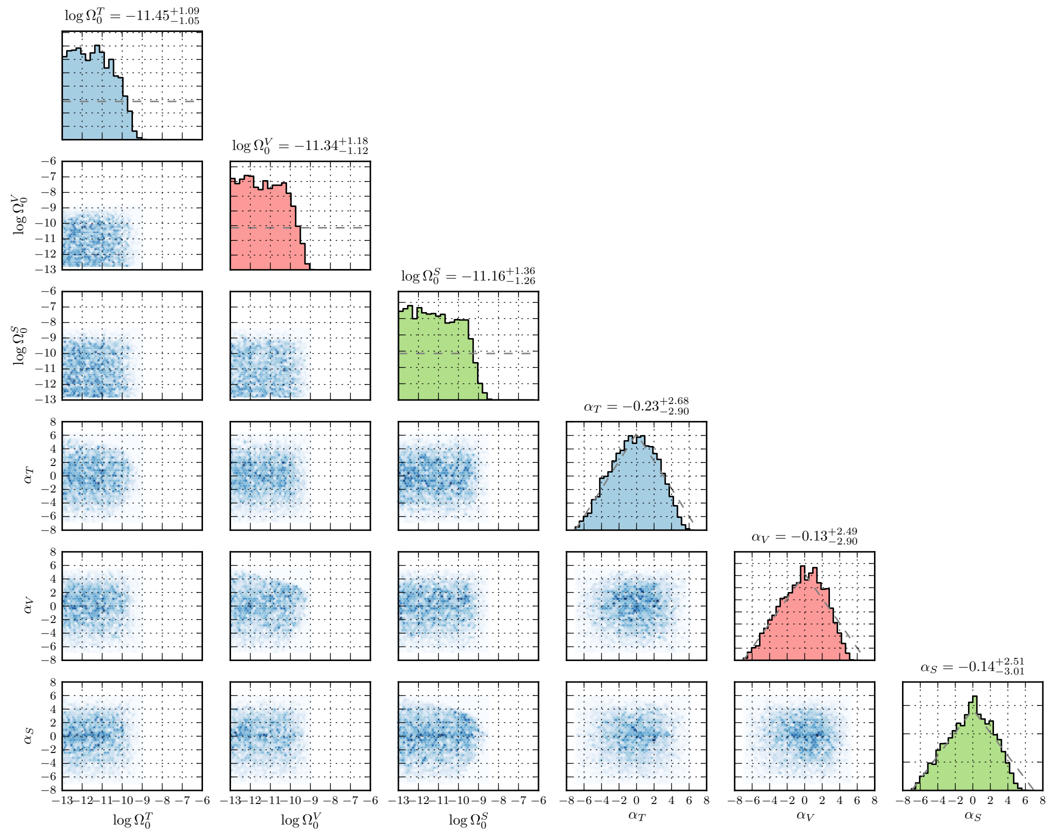

Figure 5.26: Posteriors obtained for a simulated Advanced LIGO observation of pure Gaussian noise (Case 1 in Table 5.1), under the TVS hypothesis. The subplots along the diagonal show marginalized posteriors for the amplitudes and slopes of the tensor, vector, and scalar backgrounds (blue, red, and green, respectively), while the remaining subplots show the 2D posterior between each pair of parameters. Each amplitude posterior is consistent with our lower prior bound, reflecting the non-detection of a stochastic background.

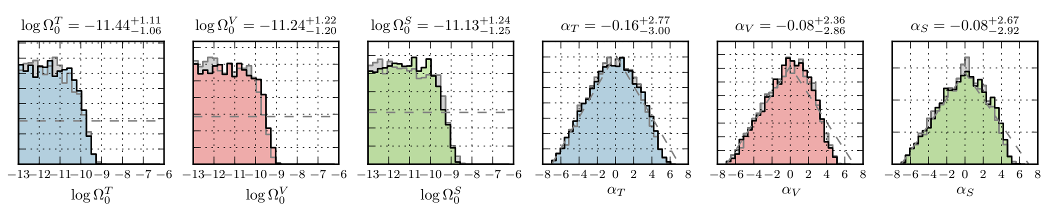

Figure 5.27: Marginalized amplitude and slope posteriors for the Gaussian noise observation in Fig. 5.26, after the additional inclusion of design-sensitivity Advanced Virgo. The light grey histograms show the Advanced LIGO-only results from Fig. 5.26. As above, dashed grey lines show the priors placed on each parameter. We see that the inclusion of Advanced Virgo does not significantly affect the parameter estimation results.

Above each posterior we quote the median posterior value as well as ±34% credible limits. The remaining subplots illustrate the joint 2D posteriors between each pair of parameters.

For this simulated Advanced LIGO observation, we obtain $\ln \mathcal{O}^{\text{SIG}}_{\text{N}} = -1.1$, consistent with a null detection. Accordingly, the posteriors on Ω_0^T, Ω_0^V, and Ω_0^S are each consistent with the lower bound of our amplitude prior (at $\log \Omega_{\text{Min}} = -13$). Meanwhile, the posteriors on spectral indices α_T, α_V, and α_S simply recover our chosen prior. The 95% credible upper limits on each amplitude are $\log \Omega_0^T < -9.8$, $\log \Omega_0^V < -9.7$, and $\log \Omega_0^S < -9.3$.

In Fig. 5.27 we show the posteriors obtained if we additionally include design-sensitivity Advanced Virgo (incorporating simulated measurements for the HV and LV detector pairs). For reference, the grey histograms show the posteriors from Fig. 5.26 obtained by Advanced LIGO alone. The Advanced LIGO-Virgo posteriors are virtually identical to those obtained from Advanced LIGO alone, with 95% credible upper limits of $\log \Omega_0^T < -9.9$, $\log \Omega_0^V < -9.6$, and $\log \Omega_0^S < -9.4$. In the case of a null-detection, then, the inclusion of Advanced Virgo does not notably improve the upper limits placed on the amplitudes of tensor, vector, and scalar backgrounds.

5.7.2 Case 2: Tensor Background

Next, we produce a simulated observation of a pure tensor background with amplitude $\log \Omega_0^T = -8.78$ and spectral index $\alpha_T = 2/3$. The amplitude is chosen such that the background would be detected by Advanced LIGO with expected $\langle \text{SNR}_{\text{OPT}} \rangle = 5$ after three years of observation at design-sensitivity. The odds ratios obtained for this simulated observation are $\ln \mathcal{O}^{\text{SIG}}_{\text{N}} = 8.4$ and $\ln \mathcal{O}^{\text{NGR}}_{\text{GR}} = -1.4$, indicating a strong detection consistent with general relativity.

The corresponding parameter posteriors are shown in Fig. 5.28. In this case, the injected parameter values are shown via dot-dashed black lines. The $\log \Omega_0^T$ posterior is strongly peaked near the true value, with a central 68% credible interval of $-9.0 \leq \log \Omega_0^T \leq -8.7$ and a median value of $\log \Omega_0^T = -8.8$. The vector and scalar amplitudes, in turn, are consistent with the lower bound on our prior, with 95% credible upper limits of $\log \Omega_0^V < -9.2$ and $\log \Omega_0^S < -9.0$.

The parameter estimation results when additionally including Advanced Virgo are given in Fig. 5.29. Once again, the grey histograms show parameter esti-

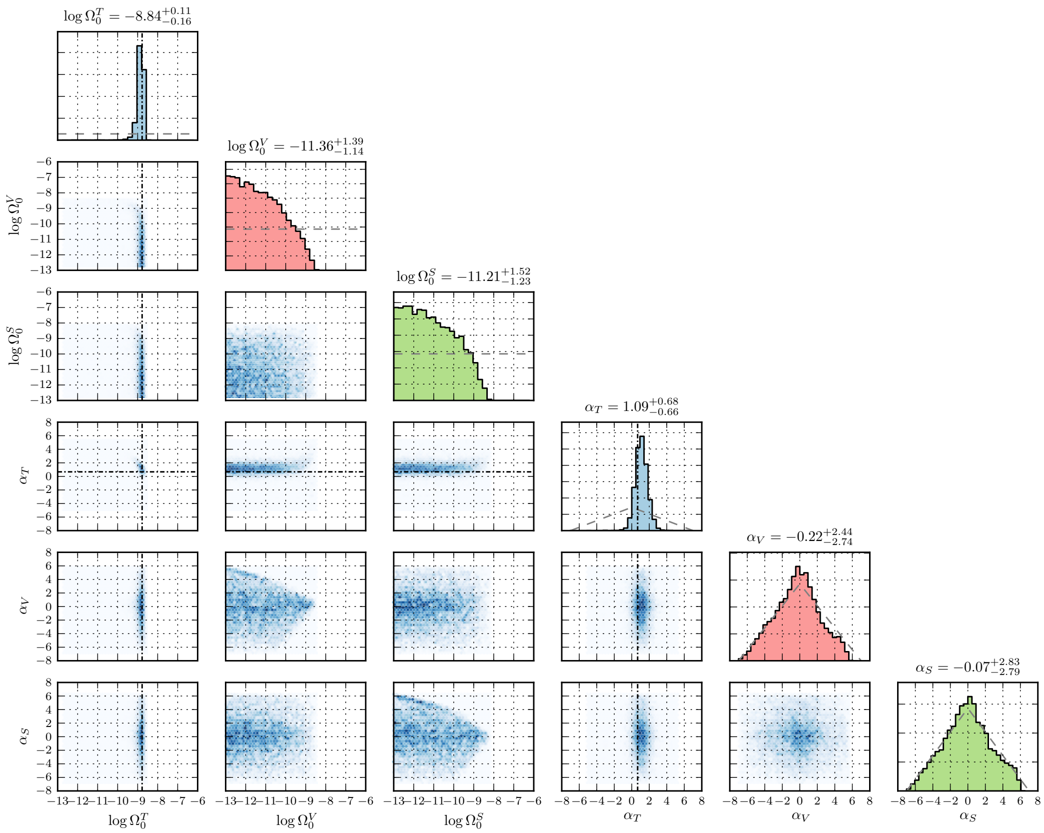

Figure 5.28: As in Fig. 5.26, for a simulated observation of a pure tensor background (Case 2 in Table 5.1). The injected tensor amplitude and slope are indicated by dot-dashed black lines. The tensor amplitude and slope posteriors are peaked about their true values. The vector and scalar amplitude posteriors, meanwhile, are consistent with our lower prior bound.

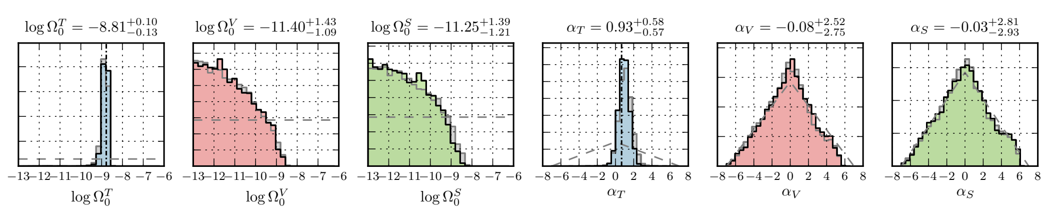

Figure 5.29: Marginalized amplitude and slope posteriors for the tensor background observation in Fig. 5.28, after the additional inclusion of design-sensitivity Advanced Virgo. For reference, the light grey histograms show the Advanced LIGO-only results from Fig. 5.28. The joint LIGO-Virgo parameter estimation yields a slightly tighter measurement of Ω_0^T, as well as marginally improved upper limits on Ω_0^V and Ω_0^S.

mation results from Advanced LIGO alone. Although Virgo does not improve our confidence in the detection, it *can* serve to break degeneracies present between different polarization modes. We begin to see this behavior in Fig. 5.29, in which the vector and scalar log-amplitude posteriors are pushed to smaller values in the joint LIGO-Virgo analysis. When including Advanced Virgo, we obtain a marginally tighter 68% credible interval of $-8.9 \leq \log \Omega_0^T \leq -8.7$ on the tensor amplitude, and slightly improved upper limits of $\log \Omega_0^V < -9.3$ and $\log \Omega_0^S < -9.2$ on vector and scalar amplitudes.

5.7.3 Case 3: Tensor and Scalar Backgrounds

As discussed above, most alternative theories of gravity would predict a stochastic background of mixed polarization. For our final case study, we therefore consider a mixed background with both tensor ($\log \Omega_0^T = -8.48$ and $\alpha_T = 2/3$) and scalar ($\log \Omega_0^S = -7.83$ and $\alpha_S = 0$) components. The amplitudes are chosen such that each component is individually observable with $\langle \mathrm{SNR}_{\mathrm{OPT}} \rangle = 10$ after three years of observation. Analysis with MultiNest yields odds ratios $\ln \mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}} = 193.5$ and $\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}} = 16.1$, representing an extremely loud detection with very strong evidence for the presence alternative polarizations.

The posteriors obtained for this data are shown in Fig. 5.30. Despite the strength of the simulated stochastic signal, we see that parameter estimation results are dominated by degeneracies between the different polarization modes. Although the tensor and scalar amplitude posteriors are locally peaked about their true values, much of the background's energy is misattributed to vector modes, illustrating that potentially severe degeneracies persist even at high SNRs. These degeneracies are exacerbated for backgrounds with small or negative spectral indices, as in the present case. Such backgrounds preferentially weight low frequencies where the Advanced LIGO overlap reduction functions are all similar (see Fig. 5.6). This example serves to illustrate that, while Advanced LIGO can likely identify the *presence* of alternative polarizations through the odds ratio $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$, Advanced LIGO alone is unable to determine which modes (vector or scalar) have been detected.

In contrast, the degeneracies in Fig. 5.30 are completely broken with the inclusion of Advanced Virgo. Whereas the Ω_0^V posterior is strongly peaked in Fig. 5.30, we see in Fig. 5.31 that the posterior is instead entirely consistent with our lower prior bound when including Advanced Virgo. The tensor and

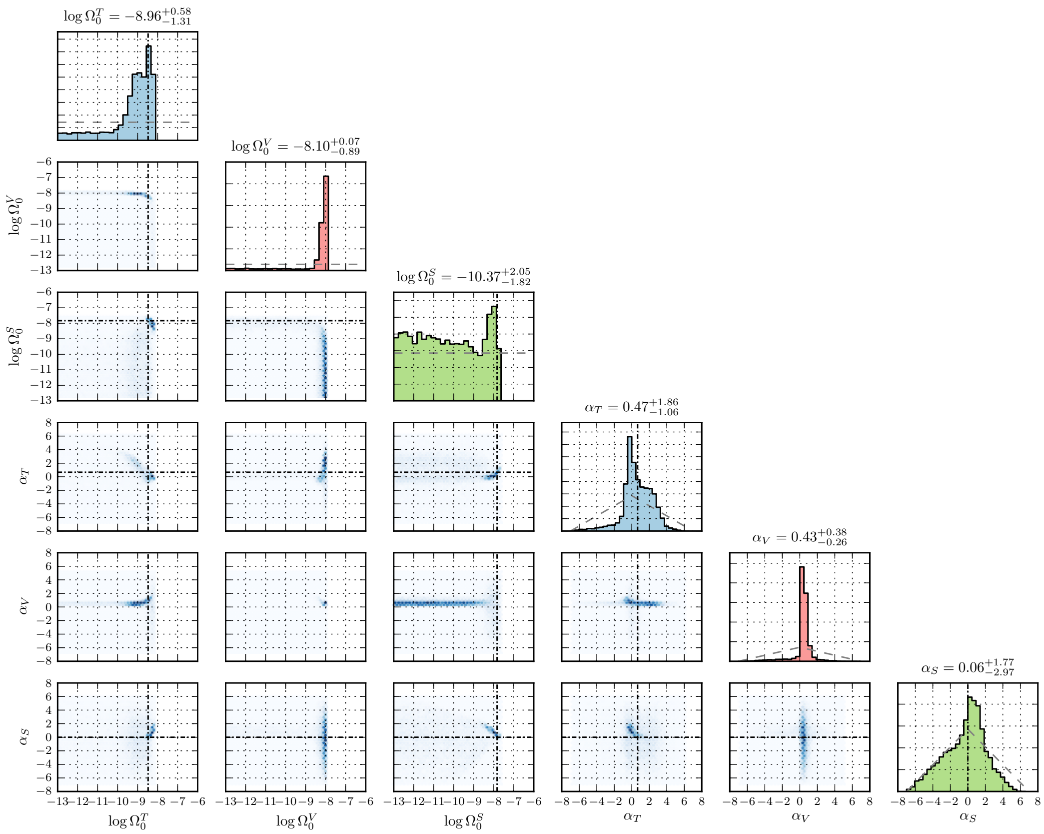

Figure 5.30: As in Figs. 5.26 and 5.28, for a simulated observation of a mixed tensor and scalar background (Case 3 in Table 5.1). While the Ω_0^T and Ω_0^S posteriors are locally peaked about the true values, much of the observed energy is mistaken for vector polarizations. Thus Advanced LIGO alone is unable to break the degeneracy between tensor, vector, and scalar amplitudes.

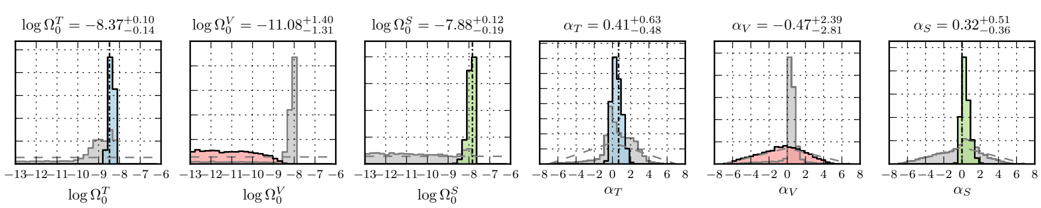

Figure 5.31: Marginalized amplitude and slope posteriors for the mixed tensor and scalar background observation in Fig. 5.30, after the additional inclusion of design-sensitivity Advanced Virgo. For reference, the light grey histograms show the Advanced LIGO-only results from Fig. 5.30. In contrast to the results in Fig. 5.30, the degeneracy between polarization modes is completely broken when including Advanced Virgo. Thus, while Advanced Virgo does not particularly improve prospects for the detection of a mixed background, it can significantly improve our ability to perform parameter estimation on multiple modes simultaneously.

scalar amplitude posteriors, meanwhile, are each more strongly-peaked about their correct values and are now inconsistent with the lower amplitude bound. Thus, while Advanced Virgo generally does not improve our ability to *detect* a stochastic background, we see that it can significantly improve prospects for simultaneous parameter estimation of multiple polarizations.

5.8 Safeguarding against Mismodeling of the Gravitational-Wave Background

The stochastic search method presented here offers a means to search for alternative gravitational-wave polarizations in a nearly model-independent way. Unlike direct searches for compact binary coalescences, our search makes minimal assumptions about the source and nature of the stochastic background. We do, however, make one notable assumption: that the energy density spectra $\Omega^A(f)$ are well-described by power laws in the Advanced LIGO frequency band. This is expected to be a reasonable approximation for most predicted astrophysical sources of gravitational waves. The backgrounds expected from stellar-mass binary black holes [226], core-collapse supernovae [147], and rotating neutron stars [112, 115, 128], for instance, are all well-modelled by power laws in the Advanced LIGO band. It may be, however, that the stochastic background is in fact *not* well-described by a single power law. This may be the case if, for instance, the background is dominated by high-mass binary black holes, an excess of systems at high redshift, or previously-unexpected sources of gravitational waves.

Given that our search allows only for power-law background models, how would we interpret a non-power-law background? In particular, if the stochastic background is purely tensorial (obeying general relativity) but is not well-described by a power-law, would our search mistakenly claim evidence for alternative polarizations?

To investigate this question, we consider simulated Advanced LIGO observations of pure tensor backgrounds described by broken power laws:

$$\Omega^T(f) = \begin{cases} \Omega_0 \left(\frac{f}{f_k}\right)^{\alpha_1} & (f < f_k) \\ \Omega_0 \left(\frac{f}{f_k}\right)^{\alpha_2} & (f \geq f_k). \end{cases} \tag{5.64}$$

Here, Ω_0 is the background's amplitude at the "knee frequency" f_k, while α_1 and α_2 are the slopes below and above the knee frequency, respectively. We

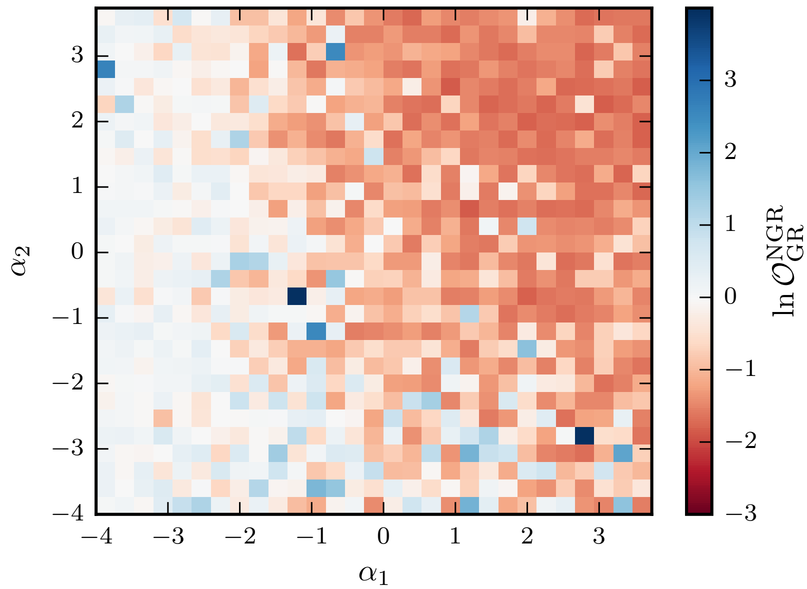

Figure 5.32: Odds ratios $\mathcal{O}^{\rm NGR}_{\rm GR}$ obtained for simulated Advanced LIGO observations of tensor-polarized broken power-law backgrounds with energy-density spectra given by Eq. (5.64), with knee frequencies $f_k = 30\,\mathrm{Hz}$ in the center of the stochastic sensitivity band. We scale the amplitude Ω_0 of each injected background such that it is optimally detectable with $\langle \mathrm{SNR}_{\rm OPT}\rangle = 5$ after the simulated observation period. By design, these backgrounds are not well-described by single power laws, the form explicitly assumed in our search. Despite this fact, we find that these backgrounds are *not* systematically misclassified as containing vector or scalar polarizations.

will set the knee frequency to $f_k = 30\,\mathrm{Hz}$, placing the backgrounds' knees in the most sensitive band of the LIGO-Virgo stochastic search. We will deliberately analyze simulated observations of (purely-tensorial) broken power laws with an *incorrect* model, assuming an ordinary power-law form for $\Omega^A(f)$.

The resulting odds ratios $\mathcal{O}^{\rm NGR}_{\rm GR}$ are shown in Fig. 5.32 as a function of the two slopes α_1 and α_2. Each simulation assumes three years of observation at design-sensitivity, and the injected amplitudes Ω_0 are scaled such that each simulated stochastic background has expected $\langle \mathrm{SNR}_{\rm OPT}\rangle = 5$ after this time. Any trends in Fig. 5.32 are therefore due to the backgrounds' spectral shapes rather than their amplitudes.

If tensor broken power laws are indeed misclassified by our search, we should expect large, positive $\ln \mathcal{O}^{\rm NGR}_{\rm GR}$ values in Fig. 5.32. Instead, we see that broken power laws are *not* systematically misclassified. When α_1 and α_2 are each positive, we recover $\ln \mathcal{O}^{\rm NGR}_{\rm GR} \approx -1.5$, correctly classifying backgrounds as tensorial despite the fact that they are not described by power laws. When $\alpha_1 < 0$, meanwhile, we recover odds ratios scattered about $\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}} \approx 0$. This simply reflects the fact that when α_1 is negative the majority of a background's SNR is collected at low frequencies where Advanced LIGO's tensor, vector, and scalar overlap reduction functions are degenerate. In such a case we do not show preference for either model over the other. Note that we find $\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}} \approx 0$ even along the line $\alpha_2 = \alpha_1$ (for $\alpha_1 < 0$), where the background *is* described by a single power law.

We expect broken power laws to be most problematic when $\alpha_1 > 0$ and $\alpha_2 < 0$; in this case a background's SNR is dominated by a small frequency band around the knee itself. This would be the case if, for instance, the stochastic background were dominated by unexpectedly massive binary black hole mergers [226]. Figure 5.32 does suggest a larger scatter in $\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$ for such backgrounds. Even in this region, however, there is not a systematic bias towards larger values of $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$, and the largest recovered odds ratios have $\ln \mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}} \lesssim 2.5$, well below the level required to confidently claim evidence for the presence of alternative polarizations.

Despite the fact that we assume purely power-law models for the stochastic energy-density spectra, our search appears reasonably robust against broken power-law spectra that are otherwise purely tensor-polarized. In particular, in order to be mistakenly classified by our search, a tensor stochastic background would have to emulate the pattern of positive and negative cross-power associated with the vector and/or scalar overlap reduction functions (see, for instance, Fig. 5.12). This is simply not easy to do without a pathological background. While we have demonstrated this only for Advanced LIGO, we find similarly robust results for three-detector Advanced LIGO-Virgo observations.

Nevertheless, when interpreting odds ratios $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$ it should be kept in mind that the true stochastic background may deviate from a power law. Even if a broken tensor background is not misclassified in our analysis, the *parameter estimation* results we obtain would likely be incorrect (another example of so-called "stealth bias"). It should be pointed out, though, that our analysis is not fundamentally restricted to power-law models. While we adopt power-law models here for computational simplicity, our analysis can be straightforwardly expanded in the future to include more complex models for the stochastic energy-density spectrum.

5.9 Discussion

The direct detection of gravitational waves by Advanced LIGO and Virgo has opened up new and unique prospects for testing general relativity. One such avenue is the search for vector and scalar gravitational-wave polarizations, predicted by some alternative theories of gravity but prohibited by general relativity. Observation of vector or scalar polarizations in the stochastic background would therefore represent a clear violation of general relativity. While the first preliminary measurements have recently been made of the polarization of GW170814 and GW170817, our ability to study the polarizations of transient gravitational-wave signals is generally limited by the number and orientation of current-generation detectors. In contrast, searches for long-duration sources like the stochastic background offer a promising means of directly measuring gravitational-wave polarizations with existing detectors.

In this chapter, we explored a procedure by which Advanced LIGO can detect or constrain the presence of vector and scalar polarizations in the stochastic background. In Ch. 5.5, we found that a stochastic background dominated by alternative polarization modes may be missed by current searches optimized only for tensor polarizations. In particular, backgrounds of vector and scalar polarizations with large, positive slopes may take up to ten times as long to detect with current methods, relative to a search optimized for alternative polarizations. In Ch. 5.6, we therefore proposed a Bayesian method with which to detect a generically-polarized stochastic background. This method relies on the construction of two odds ratios. The first serves to determine if a stochastic background has been detected, while the second quantifies evidence for the presence of alternative polarizations in the background. This search has the advantage of being entirely generic; it is capable of detecting and identifying stochastic backgrounds containing any combination of gravitational-wave polarizations. With this method, we demonstrated that a flat scalar-polarized background of amplitude $\Omega_0^S \approx 2 \times 10^{-8}$ can be confidently identified as non-tensorial with Advanced LIGO.

In Ch. 5.7, we then considered the ability of Advanced LIGO to perform simultaneous parameter estimation on tensor, vector, and scalar components of the stochastic background. After three years of observation at design sensitivity, Advanced LIGO will be able to limit the amplitudes of tensor, vector, and scalar polarizations to $\Omega_0^T < 1.6 \times 10^{-10}$, $\Omega_0^V < 2.0 \times 10^{-10}$, and

$\Omega_0^S < 5.0 \times 10^{-10}$, respectively, at 95% credibility. If, however, a stochastic background of mixed polarization is detected, Advanced LIGO alone cannot precisely determine the parameters of the tensor, vector, and/or scalar components simultaneously due to large degeneracies between modes.

We also considered how the addition of Advanced Virgo to the Hanford-Livingston network affects the search for alternative polarizations. In Ch. 5.6, we found that addition of Advanced Virgo does not particularly increase our ability to detect or identify backgrounds of alternative polarizations. However, in Ch. 5.7 we saw Advanced Virgo *does* significantly improve our ability to perform parameter estimation on power-law backgrounds, breaking the degeneracies that plagued the Hanford-Livingston analysis.

Relative to other modeled searches for gravitational waves, the stochastic search described here has the advantage of being nearly model-independent. We have, however, made one large assumption: that the tensor, vector, and scalar energy-density spectra are well-described by power laws in the Advanced LIGO band. In Ch. 5.8 we explored the implications of this assumption, asking the question: would tensor backgrounds that are *not* described by power laws be mistaken for alternative polarizations in our search? We found that our proposed Bayesian method is reasonably robust against this possibility. In particular, even pure tensor backgrounds with sharply-broken power law spectra are not systematically misidentified by our search.

The non-detection of alternative polarizations in the stochastic background may yield interesting experimental constraints on extended theories of gravity. Meanwhile, any experimental evidence *for* alternative polarizations in the stochastic background would be a remarkable step forward for experimental tests of gravity. Of course, if future stochastic searches do yield evidence for alternative polarizations, careful study would be required to verify that this result is not due to unmodeled effects like non-Gaussianity [26, 100, 103, 104, 106] or anisotropy [237–239] in the stochastic background. Comparison to polarization measurements of other long-lived sources like rotating neutron stars [217, 218] will additionally aid in the interpretation of stochastic search results.

∾

Appendix 5.A Comparison to Nishizawa *et al.* (2009)

In this chapter we defined the overlap reduction functions for tensor, vector, and scalar polarizations as

$$\begin{aligned}
\gamma_T(f) &= \frac{5}{8\pi}\sum_{A\in\{+,\times\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
\gamma_V(f) &= \frac{5}{8\pi}\sum_{A\in\{x,y\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
\gamma_S(f) &= \frac{5}{8\pi}\sum_{A\in\{b,l\}}\int d\hat{\mathbf{n}}\, F_1^A(\hat{\mathbf{n}})\, F_2^A(\hat{\mathbf{n}})\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c},
\end{aligned} \tag{5.65}$$

such that the cross-correlation statistic $\hat{C}(f)$ has expectation value

$$\langle \hat{C}(f)\rangle = \gamma_A(f)\Omega^A(f), \tag{5.66}$$

in terms of the canonical energy densities $\Omega^A(f)$ [Eq. (5.30)]. The scalar overlap reduction function $\gamma_S(f)$ differs from that given in Eq. (24) of Ref. [214]:

$$\tilde{\gamma}_S(f) = \frac{15}{4\pi}\frac{1}{1+2\kappa}\int d\hat{\mathbf{n}}\,\left[\tilde{F}_1^b(\hat{\mathbf{n}})\,\tilde{F}_2^b(\hat{\mathbf{n}}) + \kappa\tilde{F}_1^l(\hat{\mathbf{n}})\,\tilde{F}_2^l(\hat{\mathbf{n}})\right]\, e^{-2\pi i f\Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c}, \tag{5.67}$$

where $\kappa = \tilde{\Omega}^l/\tilde{\Omega}^b$ is defined to be the ratio between energy densities in longitudinal and breathing modes. Note that we will use tildes to denote quantities defined according to the conventions of Ref. [214].

The difference between Eqs. (5.65) and (5.67) is due to different definitions of the longitudinal polarization's basis tensor. We adopt

$$\hat{\mathbf{e}}^l_{ij} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{5.68}$$

while Ref. [214] uses

$$\tilde{\mathbf{e}}^l_{ij} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sqrt{2} \end{pmatrix}. \tag{5.69}$$

As a consequence, Ref. [214] obtains a longitudinal antenna pattern

$$\tilde{F}^l(\hat{\mathbf{n}}) = \frac{1}{\sqrt{2}}\sin^2\theta\cos 2\phi, \tag{5.70}$$

which differs by a factor of $\sqrt{2}$ from the conventional form

$$F^l(\hat{\mathbf{n}}) = \frac{1}{2}\sin^2\theta\cos 2\phi. \tag{5.71}$$

Thus the LIGO/Virgo antenna patterns for longitudinal and breathing modes are, in the convention of Ref. [214], related via $|\tilde{F}^l(\hat{\mathbf{n}})| = \sqrt{2}|\tilde{F}^b(\hat{\mathbf{n}})|$. Given this relation, Eq. (5.67) can be rewritten

$$\begin{aligned}
\tilde{\gamma}_S(f) &= \frac{15}{4\pi}\int d\hat{\mathbf{n}}\, \tilde{F}^b_1(\hat{\mathbf{n}})\, \tilde{F}^b_2(\hat{\mathbf{n}})\, e^{-2\pi i f \Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&= \frac{15}{4\pi}\int d\hat{\mathbf{n}}\, F^b_1(\hat{\mathbf{n}})\, F^b_2(\hat{\mathbf{n}})\, e^{-2\pi i f \Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&= \frac{15}{8\pi}\int d\hat{\mathbf{n}}\, \left[F^b_1(\hat{\mathbf{n}})\, F^b_2(\hat{\mathbf{n}}) + F^l_1(\hat{\mathbf{n}})\, F^l_2(\hat{\mathbf{n}})\right]\, e^{-2\pi i f \Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} \\
&= 3\gamma_S(f);
\end{aligned} \tag{5.72}$$

note the conversion from $\tilde{F}^b_1$ to F^b_1 in the second line.

If $\tilde{F}^l = \sqrt{2}F^l$, then the strain amplitudes as defined in Ref. [214] must in turn be reduced by $\tilde{h}_l = h_l/\sqrt{2}$, so that the signals $\tilde{F}^l\tilde{h}_l$ and $F^l h_l$ registered by a given detector are equal. Hence the quantity $\tilde{\Omega}^l(f)$ appearing in Ref. [214], which is quadratic in $\tilde{h}_l$, is *half* of what we define as the canonical energy density $\Omega^l(f)$ of longitudinal gravitational waves:

$$\tilde{\Omega}^l(f) = \frac{1}{2}\Omega^l(f). \tag{5.73}$$

A related algebraic result that will be useful is the relation

$$\begin{aligned}
\Omega^S(f) &= \Omega^b(f) + \Omega^l(f) \\
&= \tilde{\Omega}^b(f) + 2\tilde{\Omega}^l(f) \\
&= (1+2\kappa)\,\tilde{\Omega}^b(f) \\
&= \frac{1+2\kappa}{1+\kappa}\left(\tilde{\Omega}^b(f) + \tilde{\Omega}^l(f)\right) \\
&= \frac{1+2\kappa}{1+\kappa}\tilde{\Omega}^S(f)
\end{aligned} \tag{5.74}$$

between $\Omega^S(f)$ and $\tilde{\Omega}^S(f)$.

To verify that both approaches are equivalent, note that in the convention of Ref. [214], the contribution of scalar modes to the cross-correlation statistic is $\langle\hat{C}(f)\rangle = \xi\tilde{\gamma}_S(f)\tilde{\Omega}^S(f)$, where $\xi = \frac{1}{3}\left(\frac{1+2\kappa}{1+\kappa}\right)$. Converting to our own notation

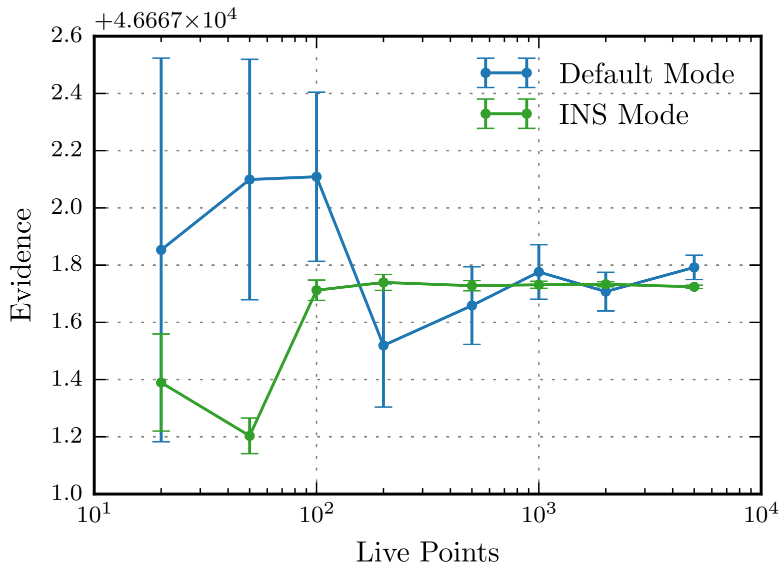

Figure 5.33: MultiNest Bayesian evidences for a single simulated stochastic background observation as a function of the number of live points chosen. For the simulated data, we assume a tensor-polarized background (with $\Omega_0^T = 2 \times 10^{-8}$ and $\alpha_T = 2/3$) observed for one year with design-sensitivity Advanced LIGO, and compute evidence using the T model. Results are shown for both MultiNest's Default and INS modes; also shown are the error estimates provided by each mode. To compute the results presented in this chapter, we used $n = 2000$ live points.

using Eqs. (5.72) and (5.74),

$$\begin{aligned}\langle \hat{C}(f) \rangle &= \xi \tilde{\gamma}_S(f) \tilde{\Omega}^S(f) \\ &= \frac{1}{3}\left(\frac{1+2\kappa}{1+\kappa}\right)\left(3\gamma_S(f)\right)\left(\frac{1+\kappa}{1+2\kappa}\Omega^S(f)\right) \\ &= \gamma_S(f)\Omega^S(f),\end{aligned} \tag{5.75}$$

consistent with Eq. (5.66).

Appendix 5.B Evaluating Evidences with MultiNest

Here I discuss several details associated with using MultiNest to evaluate Bayesian evidences for various models of the stochastic background. The MultiNest algorithm allows for several user-defined parameters, including the number n of live points used to sample the prior volume and the sampling efficiency ϵ, which governs acceptance rate of new proposed live points (see e.g. Ref. [233] for details). MultiNest also provides the option to run in De-

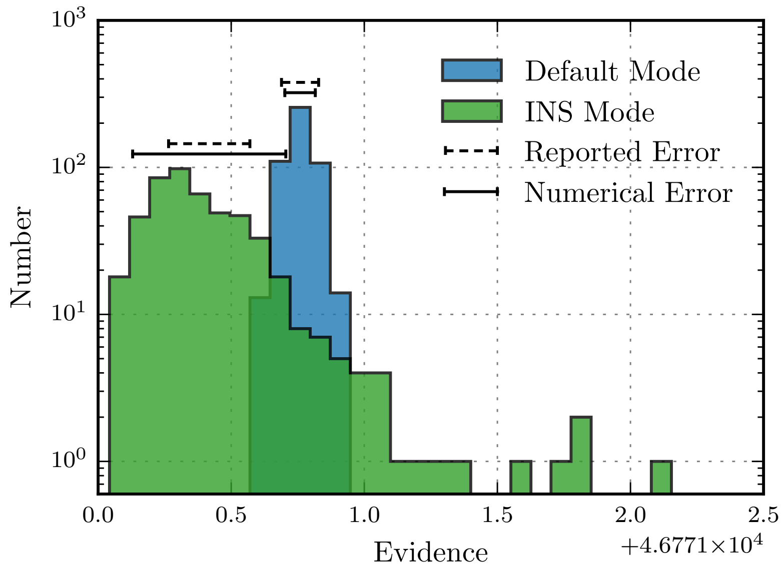

Figure 5.34: Histograms of MultiNest evidences (for the TVS model) obtained by evaluating a single simulated data set 500 times in both the Default and INS modes. To generate the simulated data, we assume a one-year observation of a tensor background ($\Omega_0^T = 2 \times 10^{-8}$ and $\alpha_T = 2/3$) with design-sensitivity Advanced LIGO. The dashed error bars show the mean 68% confidence interval reported by each method, while the solid error bars show the true 68% confidence interval computed from the evidence distributions.

fault or Importance Nested Sampling (INS) modes, each of which use different methods to evaluate evidences [234].

To set the number of live points, we investigated the convergence of MultiNest's evidence estimates with increasing values of n. For a single simulated observation of a tensorial background (with amplitude $\Omega_0^T = 2 \times 10^{-8}$ and slope $\alpha_T = 2/3$), for instance, Fig. 5.33 shows the recovered evidence for the tensor-only T hypothesis as a function of n, using both the Default (blue) and INS modes (green). The results are reasonably stable for $n \gtrsim 1000$; we choose $n = 2000$ live points. Meanwhile, our recovered evidence estimates do not exhibit noticeable dependence on the sampling efficiency; we choose the recommended values $\epsilon = 0.3$ for evidence evaluation and $\epsilon = 0.8$ for parameter estimation [233].

In addition to computing Bayesian evidences, MultiNest also returns an estimate of the numerical error associated with each evidence calculation. See, for instance, the error bars in Fig. 5.33. To gauge the accuracy of these error estimates, we construct a single simulated Advanced LIGO observation of a purely-tensorial stochastic background (again with $\Omega_0^T = 2 \times 10^{-8}$ and $\alpha_T = 2/3$). We then use MultiNest to compute the corresponding TVS evidence 500 times, in both Default and INS modes. The resulting distributions of evidences are shown in Fig. 5.34. The dashed error bars show the averaged $\pm 1\sigma$ intervals reported by MultiNest, while the solid bars show the true $\pm 1\sigma$ scatter observed in the ensemble of runs. We see that the errors reported by MultiNest's Default mode appear to accurately reflect the numerical error in the evidence calculation, while the errors reported by the INS mode are underestimated by a factor of ~ 2.

Additionally, Fig. 5.34 illustrates several systematic differences between the Default and INS results. First, Default mode appears significantly more precise than INS mode, giving rise to a much narrower distribution of evidences. Not only is the INS evidence distribution wider, but it exhibits a large tail extending several units in evidence above the mean. We find that similarly long tails also appear for other pairs of injected signals and recovered models. For this reason, we choose to use MultiNest's Default mode in all evidence calculations. Typical numerical errors in Default mode are of order $\delta(\text{evidence}) \sim 0.1$, and so the uncertainty associated with a log-odds ratio is $\delta(\ln \mathcal{O}) \sim \sqrt{2}\delta(\text{evidence})$, again of order 0.1. Additionally, we see that the peaks of the Default and INS distributions do not coincide. In general, the peaks of evidence distributions from the Default and INS modes lie ~ 0.3 units apart. Thus there may be additional systematic uncertainties in a given evidence calculation. However, as long as we consistently use one mode or the other (in our case, Default mode), any uniform systematic offset in the evidences will simply cancel when we ultimately compute a log-odds ratio.

Chapter 6
Bayesian Constraints on the Gravitational-Wave Background from the O1 and O2 Observing Runs

∾

This chapter contains work published in:

> P. B. Covas *et al.*, *Identification and Mitigation of Narrow Spectral Artifacts that Degrade Searches for Persistent Gravitational Waves in the First Two Observing Runs of Advanced LIGO*, Phys. Rev. D **97**, 082002 (2018).
>
> LIGO Scientific Collaboration & Virgo Collaboration, *Search for Tensor, Vector, and Scalar Polarizations in the Stochastic Gravitational-Wave Background*, Phys. Rev. Letters **120**, 201102 (2018).
>
> LIGO Scientific Collaboration & Virgo Collaboration, *A Search for the Isotropic Stochastic Background Using Data from Advanced LIGO's Second Observing Run*, Phys. Rev. D (in press)

I developed and ran the software to identify narrowband Hanford-Livingston correlations in the O1 and O2 observing run; these results were used in LIGO & Virgo (2018) and (2019), and documented in Covas *et al.* (2018). I conceived of and led the study published within LIGO & Virgo (2018). Finally, I produced the parameter estimation results and most figures appearing in LIGO & Virgo (2018); this manuscript was co-written by Andrew Matas, Rich Ormisten, and myself. The content of Ch. 6.A has not been previously published.

∾

In the previous chapter, we developed a Bayesian methodology with which to search for stochastic backgrounds of alternative gravitational-wave polarizations. This same Bayesian machinery, though, can be perfectly well applied to searches for a standard tensorial gravitational-wave background. Here we will do exactly this. In this chapter I will present Bayesian analyses of O1 and O2 Advanced LIGO data. The results are the best upper limits to date on the amplitude and spectral shape of the stochastic gravitational-wave background. We will also have a chance to apply the full machinery developed in Ch. 5, additionally computing constraints on the presence of alternative polarizations in the gravitational-wave background.

6.1 The Advanced LIGO O1 Observing Run

6.1.1 Data

Here we will analyze Advanced LIGO data recording during the O1 Observing Run between September 18, 2015 15:00 UTC and January 12, 2016 16:00 UTC. We will deliberately exclude times that contain the binary black hole signals GW150914, GW151226, and GW151012.

The initial data processing proceeds as in previous LIGO/Virgo stochastic analyses [176, 229]. Time-domain strain measurements from the LIGO-Hanford and LIGO-Livingston detectors are down-sampled from 16384 Hz to 4096 Hz and divided into half-overlapping 192 s segments. Each time segment is then Hann-windowed, Fourier transformed, and high-pass filtered using a 16th order Butterworth filter with a knee frequency of 11 Hz. Finally, the strain data are coarse-grained to a frequency resolution of 0.03125 Hz and restricted to a frequency band from 20–1726 Hz. Within each segment, we compute the LIGO-Hanford and LIGO-Livingston strain auto-power spectral densities using Welch's method [240].

Standard data quality cuts are performed in both the time and frequency domains to mitigate the effects of non-Gaussian instrumental and environmental noise [176, 239, 241]. In the time domain, we discard 35% of data due to nonstationary detector noise, leaving 29.85 days of coincident observing time. In the frequency domain, an additional 21% of data is discarded to remove correlated narrow-band features between LIGO-Hanford and LIGO-Livingston [176, 239, 241]. These narrow-band correlations are due to a variety of sources, including injected calibration signals, power mains, and GPS timing systems. The final cross-correlation spectrum between Hanford and Livingston is shown

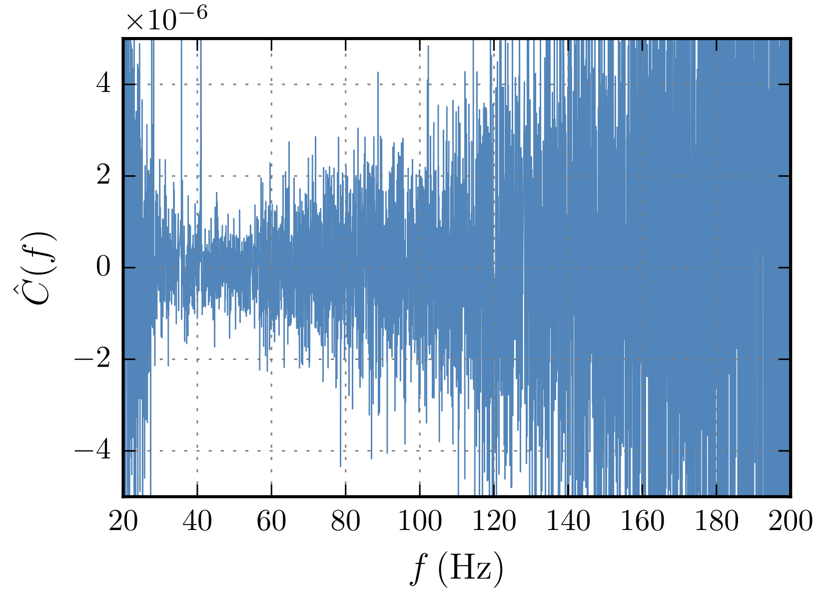

Figure 6.1: The cross-correlation spectrum $\hat{C}(f)$ measured between Advanced LIGO's Hanford and Livingston detectors during the O1 observing run. The estimator is normalized so that $\langle \hat{C}(f) \rangle = \gamma(f)\Omega(f)$. The black traces mark $\pm 1\sigma$ uncertainties on the measured cross-correlations. Coherent lines that were identified to have an instrumental cause have been removed from the spectrum.

in Fig. 6.1.

To estimate possible contamination due to terrestrial Schumann resonances [242–244], we additionally monitored coherences between magnetometers installed at both detectors. Schumann resonances were found to contribute negligibly to the O1 stochastic measurement [176, 241].

We assume conservative 4.8% and 5.4% calibration uncertainties on the strain amplitude measured by LIGO-Hanford and LIGO-Livingston, respectively [245]. Phase calibration is a much smaller source of uncertainty and is therefore neglected [176, 246]. All results below are obtained after marginalization over amplitude uncertainties.

6.1.2 Methods

We will analyze the O1 Advanced LIGO data using the method described in Ch. 5 above to search for a generically-polarized background. As in Ch. 5, we assume that the gravitational-wave background is stationary, isotropic, and Gaussian. We also assume that the background is uncorrelated between polarization modes, and that the tensor and vector contributions to the background are individually unpolarized (with equal contributions, for instance, from the tensor plus and cross modes). Recall that certain theories may violate one or more of these assumptions. For example, the stochastic background is unlikely to remain strictly unpolarized in the presence of gravitational-wave birefringence, as in Chern-Simons gravity [86, 88, 89], while theories violating Lorentz invariance may yield a departure from isotropy [208, 209]. The violation of one or more of our assumptions would likely reduce our search's sensitivity to the stochastic background.

Given the above assumptions, the expected cross-correlation between two detectors in the presence of a stochastic background is of the form [83, 214–216]

$$\langle \tilde{s}_1(f)\tilde{s}_2^*(f')\rangle = \frac{1}{16\pi}\delta(f-f')\sum_A \Gamma_A(f)\mathcal{H}^A(f) \tag{6.1}$$

(see Ch. 5.4). $\mathcal{H}^A(f)$ is the (two-sided) gravitational-wave strain power spectral density of the net tensor ($A=T$), vector (V), and scalar (S) contributions to the stochastic background and $\Gamma_A(f)$ are the *unnormalized* overlap reduction functions, defined [83, 186, 214, 247]

$$\Gamma_A(f) = \frac{1}{8\pi}\sum_{a\in A}\int d\hat{n}\, F_1^a(\hat{n})F_2^a(\hat{n})\, e^{2\pi i f \hat{n}\cdot\Delta x/c}. \tag{6.2}$$

$F_I^a(\hat{n})$ is the antenna response function of detector I to signals of polarization a, Δx is the separation vector between detectors, and c is the speed of light. The integral is taken over all sky directions $\hat{n}$.

As in previous chapters, we will work not with $\Gamma_A(f)$, but rather with the *normalized* overlap reduction functions $\gamma_A(f) = \lambda\Gamma_A(f)$, where the constant λ is chosen such that $\gamma_T(f)=1$ for co-located and co-oriented detectors. For Advanced LIGO, $\lambda = 5/8\pi$, but in general its value will vary for other experiments like LISA and pulsar timing arrays [219]. The normalized overlap reduction functions for LIGO's Hanford-Livingston baseline are shown above in Fig. 5.6. Because tensor, vector, and scalar modes each have distinct overlap reduction functions, the shape of a measured cross-correlation spectrum [Eq. (6.1)] will reflect the polarization content of the stochastic background [214, 247]. Of the three Hanford-Livingston overlap reduction functions, the scalar overlap reduction function is smallest in magnitude. This reflects the fact that the Advanced LIGO detectors have weaker geometrical responses to scalar-polarized gravitational waves than to tensor- and vector-polarized signals.

Also as in previous chapters, we will adopt the standard convention of working in terms of the energy-density spectra $\Omega^A(f)$ in different polarization modes, not the strain power $\mathcal{H}^A(f)$. It should be reiterated, though, that the precise relationship between $\Omega^A(f)$ and $\mathcal{H}^A(f)$ is theory dependent. Under any theory obeying Isaacson's formula for the stress-energy of gravitational waves [225], the energy-density spectrum is related to $\mathcal{H}(f)$ by [83, 175, 247]

$$\Omega^A(f) = \frac{4\pi^2}{3H_0^2} f^3 \mathcal{H}^A(f). \tag{6.3}$$

Although Eq. (6.3) does not necessarily hold in general [175], for ease of comparison with previous studies we will take Eq. (6.3) as the *definition* of the canonical energy-density spectra $\Omega^A(f)$. The canonical energy-density spectra can be directly identified with true energy densities under any theory obeying Isaacson's formula. For other theories, $\Omega^A(f)$ can instead be understood simply as a function of the detector-frame observable $\mathcal{H}^A(f)$.

Within each 192 s time segment (indexed by i), we form an estimator of the visible cross power between LIGO-Hanford and LIGO-Livingston:

$$\hat{C}_i(f) = \frac{1}{\Delta T} \frac{20\pi^2}{3H_0^2} f^3 \tilde{s}_{1,i}(f) \tilde{s}^*_{2,i}(f), \tag{6.4}$$

normalized such that the estimator's mean is

$$\langle \hat{C}_i(f) \rangle = \sum_A \gamma_A(f) \Omega^A(f) \tag{6.5}$$

and its variance is $\langle C(f) C(f') \rangle = \delta(f - f') \sigma^2(f)$, with

$$\sigma_i^2(f) = \frac{1}{\Delta T} \left(\frac{10\pi^2}{3H_0^2} \right)^2 f^6 P_{1,i}(f) P_{2,i}(f), \tag{6.6}$$

respectively. Within Eqs. (6.4) and (6.6), ΔT is the segment duration and $P_{I,i}(f)$ is the one-sided auto-power spectral density of detector I in time segment i. Finally, the cross-power estimators from each segment are optimally combined via a weighted sum to form a single cross-power spectrum for the O1 observing run,

$$\hat{C}(f) = \frac{\sum_i \hat{C}_i(f) \sigma_i^{-2}(f)}{\sum_i \sigma_i^{-2}(f)}, \tag{6.7}$$

with the corresponding variance

$$\sigma^{-2}(f) = \sum_i \sigma_i^{-2}(f). \tag{6.8}$$

Given the final measured cross-power spectrum $\hat{C}(f)$, we compute Bayesian evidence for the various hypotheses defined in Ch. 5 describing the presence and polarization of a possible stochastic signal within our data:

- Gaussian noise (N): No stochastic signal is present in our data, and the measured cross power is due entirely to Gaussian noise [see Eq. (5.51)].
- Signal (SIG): A stochastic background of any polarization(s) is present [Eq. (5.52) and surrounding text].
- Tensor-polarized (GR): The data contains a purely tensor-polarized stochastic signal, consistent with general relativity [Eq. (5.55)].
- Non-standard polarizations (NGR): The data contains a stochastic signal with vector and/or scalar contributions [Eq. (5.56)].

We calculate Bayesian evidences for each hypothesis using PyMultiNest [231], a Python interface to the nested sampling code MultiNest [232–236]. These evidences are combined to form two Bayesian odds [247]:

- Odds $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ for the presence of a stochastic signal relative to pure noise
- Odds $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$ for the presence of nonstandard polarizations versus ordinary tensor modes

As demonstrated in Ch. 5, $\mathcal{O}^{\mathrm{SIG}}_{\mathrm{N}}$ quantifies evidence for the *detection* of a generically polarized stochastic background, and generally depends only on a background's total power, not its polarization content. $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$, meanwhile, indicates if the background's polarization is inconsistent with general relativity. In particular, the sensitivity of $\mathcal{O}^{\mathrm{NGR}}_{\mathrm{GR}}$ to nonstandard polarizations is not significantly affected by the strength of any tensor polarization which may also be present [247].

We parameterize all energy densities as power laws

$$\Omega^A(f) = \Omega^A_0 \left(\frac{f}{f_0}\right)^{\alpha_A} \tag{6.9}$$

using a reference frequency $f_0 = 25\,\text{Hz}$. For purposes of computing Bayesian evidences, we adopt the same priors as in Ch. 5, with a log-uniform amplitude prior between $10^{-13} \leq \Omega_0^A \leq 10^{-5}$. The upper amplitude bound (10^{-5}) is consistent with limits placed by Initial LIGO and Virgo [229]. In order to be normalizable, the log-uniform prior requires a nonzero lower bound; although parameter estimation results will depend on the specific choice of lower bound, in general this dependence is weak [218]. Our lower bound (10^{-13}) is chosen to encompass small energy densities well below the reach of LIGO and Virgo at design sensitivity [48, 176]. Meanwhile, we use a triangular prior $p(\alpha_A) \propto 1 - |\alpha_A|/\alpha_{\text{MAX}}$ for $|\alpha_A| \leq \alpha_{\text{MAX}}$ on spectral indices [247]. This prior preferentially weights flat energy-density spectra, penalizing spectra which are more steeply positively or negatively sloped in the Advanced LIGO band; see Ch. 5.6. We conservatively choose $\alpha_{\text{MAX}} = 8$, allowing for energy-density spectra significantly steeper than is predicted for known astrophysical sources (like compact binary mergers). Finally, we choose the same prior odds between hypotheses as shown in Fig. 5.15.

When performing parameter estimation, we will present results obtained with two different amplitude priors – the same log-uniform prior between $10^{-13} \leq \Omega_0^A \leq 10^{-5}$ as well as a uniform prior between $0 \leq \Omega_0^A \leq 10^{-5}$. Although the log-uniform prior more faithfully represents our true uncertainty in the gravitational-wave background amplitude, the uniform prior implicitly reproduces the maximum likelihood analysis used in previous studies, and therefore allows more direct comparison to previous stochastic search results [176, 229].

In Ch. 5 above, we analyzed only simulated cross-correlation measurements between detectors. As we are analyzing real data here, we have to deal with an additional layer of complication – calibration uncertainty. The strain measured by Hanford and Livingston is not known perfectly, but is subject to non-zero calibration uncertainty. For imperfectly calibrated data, the expectation value of cross-power measurements $\hat{C}(f)$ is *not* $\sum_A \gamma_A(f)\Omega^A(f)$, but is instead given by

$$\langle \hat{C}(f) \rangle = \kappa \sum_A \gamma_A(f) \Omega^A(f), \tag{6.10}$$

where κ is some multiplicative (and possibly frequency-dependent) factor [246]. Perfect calibration would yield $\kappa = 1$, but in general κ is unknown.

To account for calibration uncertainty, we include the calibration factor κ as an additional parameter in MultiNest, so that the likelihood function becomes

$$\mathcal{L}\left(\hat{C}(f)|\kappa, \Omega_M^A(f)\right) \propto \prod_f \exp\left[-\frac{\left(\hat{C}(f) - \kappa \sum_A \gamma_A(f)\Omega_M^A(f)\right)^2}{2\sigma^2(f)} df\right], \quad (6.11)$$

where $\Omega_M^A(f)$ is our model energy-density spectrum and $\hat{C}(f)$ and $\sigma^2(f)$ are given by Eqs. (6.7) and (6.8), respectively. We place a Gaussian prior on κ, centered at $\kappa = 1$:

$$p(\kappa) \propto \exp\left(-\frac{(\kappa-1)^2}{2\epsilon^2}\right). \quad (6.12)$$

The standard deviation ϵ encapsulates the amplitude calibration uncertainty. During O1, Hanford and Livingston had maximum estimated amplitude uncertainties of 4.8% and 5.4%, respectively, within the 20-1726 Hz frequency band [245]. These uncertainty estimates are retrospectively improved relative to the the uncertainties previously adopted in Refs. [176, 239, 241]. For our analysis, we take $\epsilon = 0.072$, the quadrature sum of the Hanford and Livingston uncertainties [176]. All results are given after marginalization over κ.

In this prescription for calibration uncertainty we have made two simplifying assumptions. First, we have neglected phase calibration uncertainty, which is expected to be a sub-dominant source of uncertainty in the stochastic analysis [245, 246]. Secondly, although calibration uncertainties are frequency-dependent, for simplicity we've assumed uniform amplitude uncertainties across all frequencies. Our quoted amplitude uncertainties are conservative, encompassing the largest calibration uncertainties in the stochastic sensitivity band [176, 245].

6.1.3 Results

Using the cross power measured between Hanford and Livingston during Advanced LIGO's O1 observing run, we obtain the Bayes factors given in Table 6.1 between our various signal hypotheses and Gaussian noise. If we adopt equal prior odds between Gaussian noise and a standard tensor-polarized signal, then we $\ln \mathcal{O}_{\mathrm{N}}^{\mathrm{T}} = -0.33$, indicating no evidence for the detection of a gravitational-wave background.

Given this non-detection, we can place upper limits on the amplitude of the stochastic gravitational-wave background. We perform parameter estimation using posterior samples obtained by PyMultiNest. Figures 6.2 and 6.3 show our posteriors on the amplitude (at $f_0 = 25$ Hz) and spectral index of the stochastic energy-density spectrum, under both log-uniform and uniform amplitude priors and following marginalization over calibration uncertainty. The diagonal subplots show marginalized posteriors on the amplitude Ω_0 and slope α of the stochastic background, while the interior subplots show the joint posteriors between Ω_0 and α.

In Fig. 6.2, we recover spectral index posteriors that are largely consistent with our choice of prior, but indicate a slight bias against large positive spectral indices. The posterior preference towards small or negative spectral indices is far more pronounced in Fig. 6.3. This preference reflects the fact that Advanced LIGO is most sensitive to backgrounds of large, positive slopes [247]. Hence the non-detection of a stochastic background therefore constrains larger amplitudes to have small and/or negative spectral indices; see the joint Ω_0– α posterior in Fig. 6.3. All together, after marginalizing over calibration uncertainty and spectral index, we limit the amplitude of the gravitational-wave background to

$$\log \Omega_0 \leq -7.21 \qquad \text{(Log-uniform prior)} \tag{6.13}$$

or

$$\log \Omega_0 \leq -6.60 \qquad \text{(Uniform prior)}, \tag{6.14}$$

depending on one's choice of amplitude prior.

Table 6.1: Bayes factors between each signal sub-hypothesis and the Gaussian noise hypothesis, as computed by MultiNest. These Bayes factors are combined following Eqs. (5.54) and (5.56) to obtain odds $\mathcal{O}^{\text{SIG}}_{\text{N}}$ between Signal and Gaussian noise hypotheses, and odds $\mathcal{O}^{\text{NGR}}_{\text{GR}}$ between NGR and GR hypotheses.

Hypothesis	$\ln \mathcal{B}^{\mathcal{A}}_{\text{N}}$
T	-0.33
V	-0.33
S	-0.31
TV	-0.66
TS	-0.65
VS	-0.65
TVS	-0.99

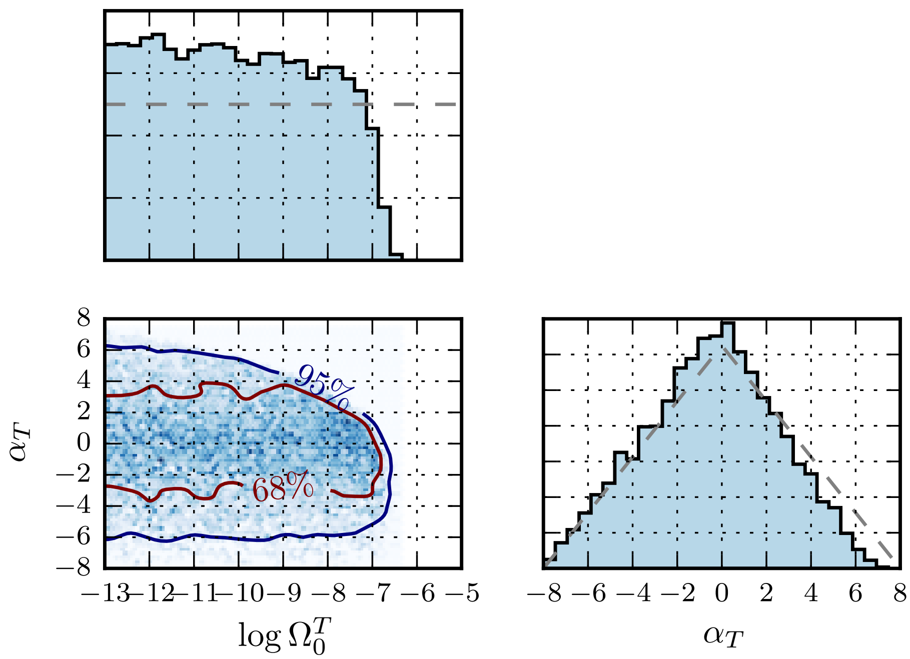

Figure 6.2: Posterior on the amplitude and spectral index of the stochastic gravitational-wave background, following Advanced LIGO's O1 observing run. Here we assume a log-uniform prior on the background's amplitude; the priors on both $\log \Omega_0$ and α are shown as dashed gray lines.

For reference, Fig. 6.3 additionally contains a dashed black curve in the joint $\Omega_0 - \alpha$ posterior; this represents the exclusion curve previously published in Ref. [176] using O1 data. Care should be taken when comparing previously published upper limits like these to our new limits obtained here. Three important distinctions should be kept in mind. First, our amplitude limits are obtained after marginalization over spectral index. Previous analysis, on the other hand, typically assume specific fixed spectral indices (e.g. Table I of Ref. [176]) or present exclusion curves in the $\Omega_0 - \alpha$ plane [176]. Second, Bayesian upper limits may be strongly influenced by one's adopted prior. Uniform amplitude priors, for instance, preferentially weight larger signals and hence yield larger upper limits, while log-uniform priors support smaller signal amplitudes, giving tighter limits. The exclusion curve of Ref. [176], meanwhile, implicitly assumes uniform spectral index priors, unlike our peaked priors here. Finally, these results are obtained under a specific signal hypothesis allowing only for tensor polarizations. These limits will not be identical, say, to the limits on the tensor amplitude Ω_0^T under alternative hypotheses that allow for the simultaneous presence of additional gravitational-wave polarizations

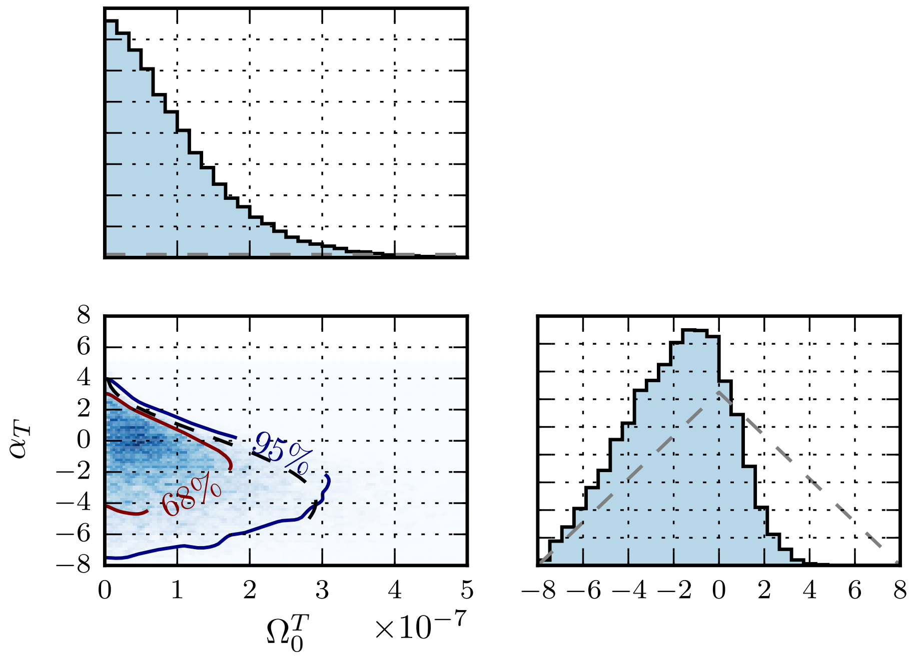

Figure 6.3: As in Fig. 6.2, but here assuming a uniform prior on Ω_0. For reference, the dashed black curve in the 2D subplot shows the exclusion line previously published in Ref. [176]. This previously published limit does not account for calibration uncertainty, and it additionally implicitly assumed a uniform prior for α.

Now let's relax our restriction to purely tensor polarized backgrounds. If we allow for a generically-polarized gravitational-wave background, we obtain odds $\ln \mathcal{O}^{\textsc{sig}}_{\textsc{n}} = -0.53$ between our generic Signal and Gaussian noise hypotheses, again indicating a non-detection of the stochastic gravitational-wave background. Additionally, we find $\ln \mathcal{O}^{\textsc{ngr}}_{\textsc{gr}} = -0.25$, consistent with values expected in the presence of Gaussian noise [247]. These odds ratios are obtained using the Bayes factors in Table 6.1, combined with the prior odds shown in Fig. 5.15 between our various sub-hypotheses. These prior odds are necessarily somewhat arbitrary; different choices will invariably yield different values of $\mathcal{O}^{\textsc{sig}}_{\textsc{n}}$ and $\mathcal{O}^{\textsc{ngr}}_{\textsc{gr}}$. Using the results in Table 6.1, one can easily recompute these odds under different prior choices.

We can now place generic upper limits on the presence of tensor, vector, and scalar contributions to the stochastic background. To simultaneously constrain the properties of each polarization, we will restrict to the "TVS" hypotheses that assumes the existence of tensor, vector, and scalar-polarized signals. Under this hypothesis, the total canonical energy density of the stochastic

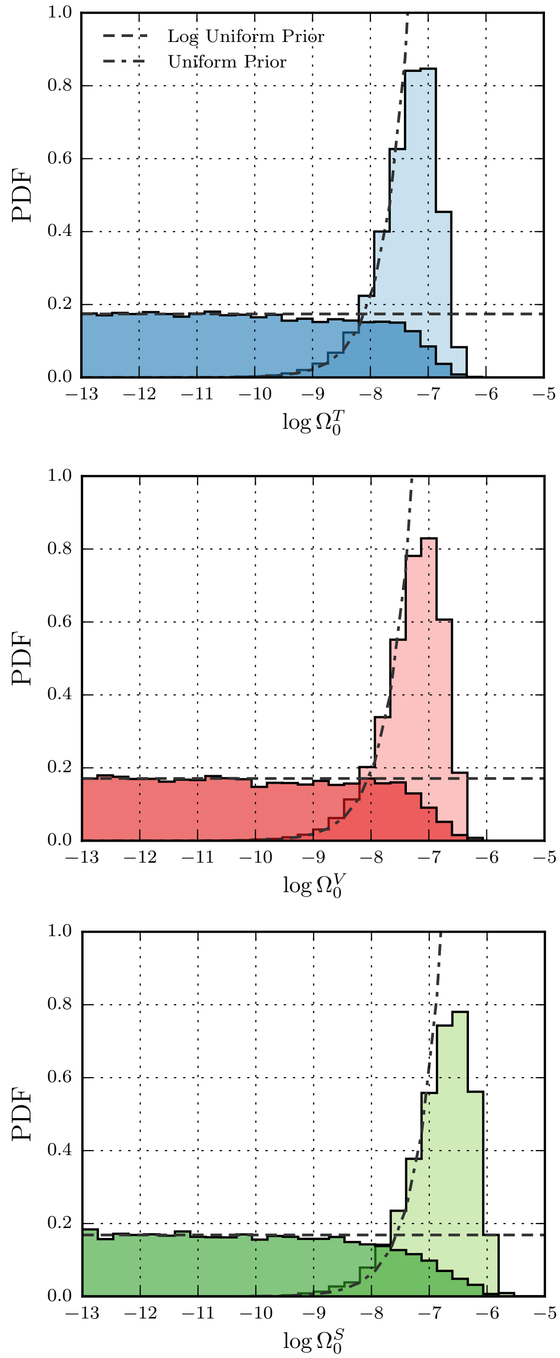

Figure 6.4: Posteriors on the tensor (top), vector (center), and scalar (bottom) stochastic background amplitudes at reference frequency $f_0 = 25$ Hz. Within each subplot, dark posteriors show results obtained assuming log-uniform priors (dashed curves) on Ω_0^A, while light posteriors show results corresponding to uniform amplitude priors (dot-dashed curves). The prior curves shown here have been renormalized by constant factors to illustrate consistency with the posteriors below our measured upper limits. These posteriors correspond to the 95% credible upper limits listed in Table 6.2. Relative to the log-uniform priors, the uniform amplitude priors preferentially weight loud stochastic signals and therefore yield more conservative upper limits.

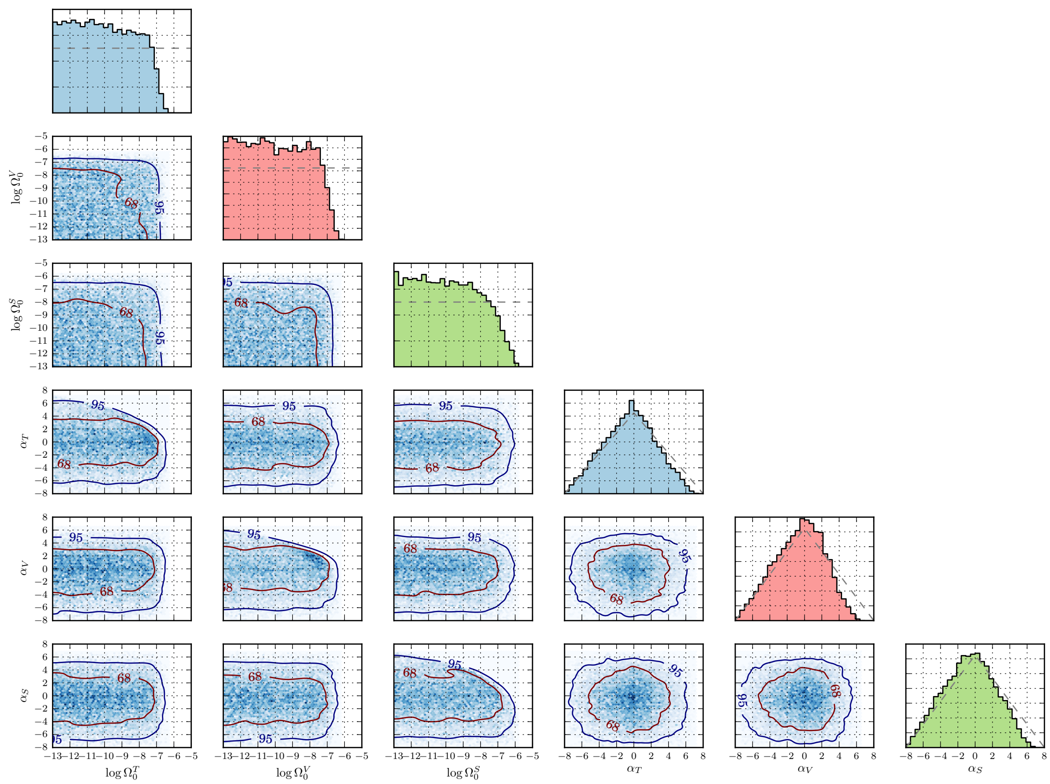

Figure 6.5: Full posteriors for the power-law amplitudes and slopes of tensor, vector, and scalar contributions to the stochastic background, assuming the TVS hypothesis and log-uniform amplitude priors. The corresponding 95% credible amplitude and spectral index limits are listed in Table 6.2.

background is modeled as a sum of power laws:

$$\Omega(f) = \Omega_0^T \left(\frac{f}{f_0}\right)^{\alpha_T} + \Omega_0^V \left(\frac{f}{f_0}\right)^{\alpha_V} + \Omega_0^S \left(\frac{f}{f_0}\right)^{\alpha_S}. \tag{6.15}$$

Figure 6.4 shows posteriors on the tensor, vector, and scalar background amplitudes, under each choice of amplitude prior. The dashed and dot-dashed curves are proportional to the log-uniform and uniform amplitude priors, respectively; each prior curve has been renormalized by a constant factor to illustrate consistency between our priors and posteriors at small Ω_0^A. When allowing for the presence of all three polarization sectors, we obtain upper limits

$$\begin{aligned} \log \Omega_0^T &\leq -7.25 \\ \log \Omega_0^V &\leq -7.20 \qquad \text{(Log-uniform priors)} \\ \log \Omega_0^S &\leq -6.96 \end{aligned} \tag{6.16}$$

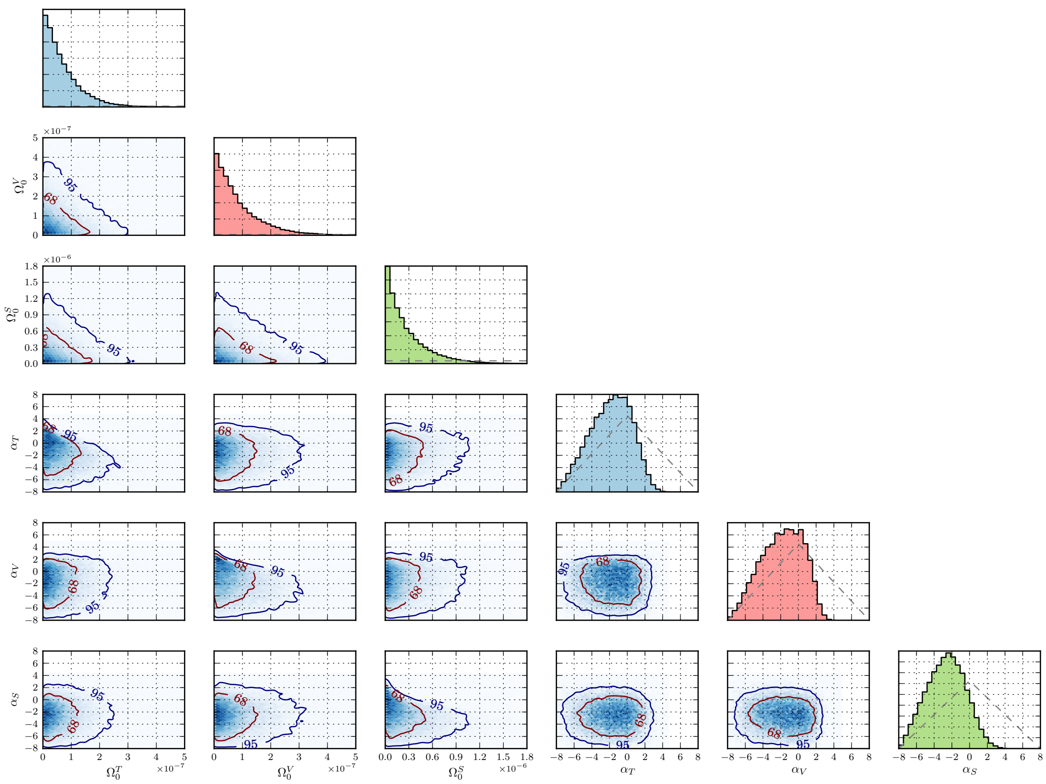

Figure 6.6: As in Fig. 6.5, but assuming a uniform prior on background amplitudes. The resulting 95% credible limits are listed in Table 6.3.

and

$$\begin{aligned} \log \Omega_0^T &\leq -6.70 \\ \log \Omega_0^V &\leq -6.59 \qquad \text{(Uniform priors)} \\ \log \Omega_0^S &\leq -6.07. \end{aligned} \tag{6.17}$$

For completeness, in Figs. 6.5 and 6.6 we also show the full six-dimensional parameter estimation results for the amplitude and slopes of the TVS model for each choice of prior. Tables 6.2 and 6.3, meanwhile, list these upper limits as well as central 95% credible intervals on the spectral index of each polarization sector.

As mentioned above, upper limits obtained under one hypothesis are not, in general, equal to those obtained under some different hypothesis. The upper limit $\log \Omega_0^T \leq -7.25$ given by the TVS model, for instance, is not the same as the upper limit $\log \Omega_0 \leq -7.21$ we found earlier when considering a standard tensor polarizations alone. Therefore, in addition to upper limits under the T and TVS hypotheses, Tables 6.2 and 6.3 additionally list 95% credible limits on the parameters of *every* signal sub-hypothesis.

Table 6.2: Parameter estimation results for each signal sub-hypothesis, obtained with log-uniform priors on the amplitude of each polarization component. Columns 2-4 give 95% credible upper limits on $\log\Omega_0^A$, and columns 5-7 show the corresponding limits on Ω_0^A for convenience. Columns 8-10, meanwhile, show the central 95% credible bounds on spectral indices α_A. The second (T) and eighth (TVS) rows correspond to the posteriors in Figs. 6.2 and 6.5.

Hypothesis	$\log\Omega_0^{T,95\%}$	$\log\Omega_0^{V,95\%}$	$\log\Omega_0^{S,95\%}$	$\Omega_0^{T,95\%}$	$\Omega_0^{V,95\%}$	$\Omega_0^{S,95\%}$	α_T	α_V	α_S
T	-7.21	-	-	6.23×10^{-8}	-	-	$-0.4^{+5.5}_{-5.9}$	-	-
V	-	-7.16	-	-	6.96×10^{-8}	-	-	$-0.4^{+4.9}_{-5.8}$	-
S	-	-	-6.86	-	-	1.38×10^{-7}	-	-	$-0.5^{+5.4}_{-5.8}$
TV	-7.25	-7.16	-	5.56×10^{-8}	6.95×10^{-8}	-	$-0.4^{+5.6}_{-5.9}$	$-0.4^{+4.8}_{-5.9}$	-
TS	-7.25	-	-6.93	5.68×10^{-8}	-	1.18×10^{-7}	$-0.4^{+5.5}_{-5.9}$	-	$-0.5^{+5.4}_{-5.8}$
VS	-	-7.17	-6.90	-	6.74×10^{-8}	1.25×10^{-7}	-	$-0.4^{+5.0}_{-5.9}$	$-0.5^{+5.4}_{-5.8}$
TVS	-7.25	-7.20	-6.96	5.58×10^{-8}	6.35×10^{-8}	1.08×10^{-7}	$-0.4^{+5.5}_{-5.9}$	$-0.4^{+4.8}_{-5.9}$	$-0.6^{+5.4}_{-5.8}$

Table 6.3: As in Table 6.2, but now assuming uniform amplitude priors.

Hypothesis	$\log\Omega_0^{T,95\%}$	$\log\Omega_0^{V,95\%}$	$\log\Omega_0^{S,95\%}$	$\Omega_0^{T,95\%}$	$\Omega_0^{V,95\%}$	$\Omega_0^{S,95\%}$	α_T	α_V	α_S
T	-6.60	-	-	2.53×10^{-7}	-	-	$-0.4^{+3.8}_{-4.7}$	-	-
V	-	-6.48	-	-	3.29×10^{-7}	-	-	$-1.5^{+3.3}_{-4.9}$	-
S	-	-	-5.97	-	-	1.08×10^{-6}	-	-	$-2.6^{+4.1}_{-4.0}$
TV	-6.66	-6.55	-	2.18×10^{-7}	2.85×10^{-7}	-	$-1.8^{+3.9}_{-4.7}$	$-1.7^{+3.6}_{-4.9}$	-
TS	-6.65	-	-6.03	2.23×10^{-7}	-	9.28×10^{-7}	$-1.7^{+3.8}_{-4.7}$	-	$-2.6^{+4.0}_{-4.1}$
VS	-	-6.53	-6.04	-	2.96×10^{-7}	9.21×10^{-7}	-	$-1.7^{+3.6}_{-4.7}$	$-2.8^{+4.1}_{-4.0}$
TVS	-6.70	-6.59	-6.07	2.02×10^{-7}	2.54×10^{-7}	8.44×10^{-7}	$-1.7^{+3.8}_{-4.8}$	$-1.7^{+3.7}_{-4.7}$	$-2.6^{+4.0}_{-4.1}$

6.2 The Advanced LIGO O2 Observing Run

The O2 observing run began in November 2016 and ran until August 2017. Much of O2 consisted of the Advanced LIGO Hanford and Livingston detectors alone. However, in June 2017, Advanced Virgo became operational and joined LIGO for the final few months of O2, yielding the first triple-coincidence gravitational-wave observations [44, 46].

In this section, we will incorporate Hanford and Livingston's O2 data to update the O1 stochastic search results presented above. Although O2 data from Advanced Virgo is available, the combination of Virgo's relative insensitivity and the the poor Hanford-Virgo and Livingston-Virgo overlap reduction functions means that Virgo would contribute little to the sensitivity of the stochastic search. We will therefore not include Advanced Virgo in this analysis.

6.2.1 Data

We will incorporate data recorded between 16:00:00 UTC on 30 November, 2016 and 22:00:00 UTC on 25 August, 2017. We make use of "cleaned" Hanford and Livingston strain time series [248, 249], from which linearly coupled noise has been removed via Wiener filtering [250]. By comparing Hanford-Livingston coherence spectra with and without this Wiener filtering, we have verified that this noise subtraction doesn't introduce any spurious correlated artifacts into the strain time series.

Data analysis proceeds exactly as described in Ch. 6.1 above, with strain data downsampled, high-pass filtered, Fourier transformed, and finally coarse-grained to a resolution of 1/32 Hz with 192 s segment lengths. We remove times when detector noise is sufficiently non-stationary, as well as times containing known gravitational-wave signals [45]; these cuts remove 16% of available coincident time, leaving a total of 99 days available to analyze. In the frequency domain, we remove frequencies containing narrowband coherent features that have known instrumental or environmental causes [81], excluding 15% of the total frequency band (but only 4% of the band below 300 Hz, where the stochastic search is most sensitive).

Figure 6.7 shows the final Hanford-Livingston cross-correlation spectrum obtained after these data quality cuts. Note that, since I'm following the normalization convention $\langle \hat{C}(f) \rangle = \gamma(f)\Omega(f)$, Fig. 6.7 differs by a factor of $\gamma(f)$

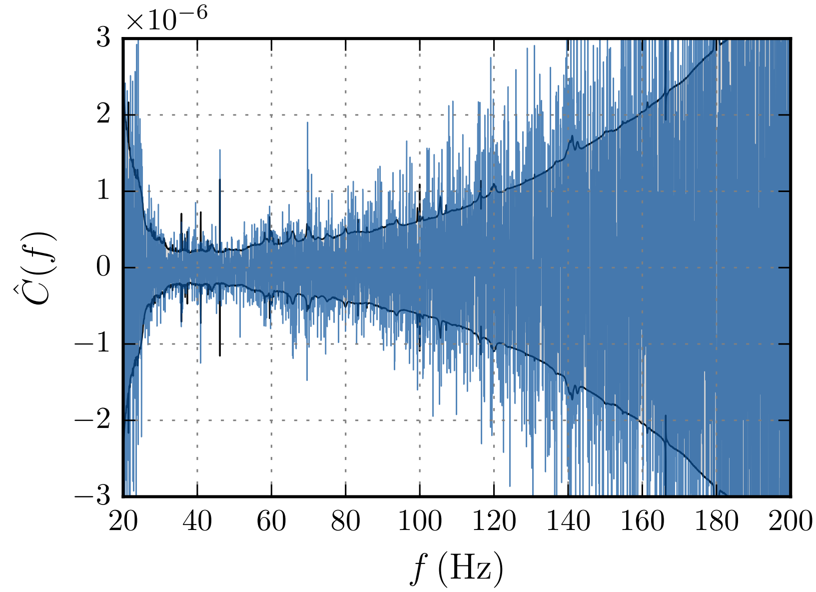

Figure 6.7: The cross-correlation spectrum $\hat{C}(f)$ measured between Advanced LIGO's Hanford and Livingston detectors during the O2 observing run. The estimator is normalized so that $\langle \hat{C}(f) \rangle = \gamma(f)\Omega(f)$. The black traces mark the $\pm 1\sigma$ uncertainties on the measured cross-correlations. Coherent lines that were identified to have an instrumental cause have been removed from the spectrum.

from Fig. 1 of Ref. [84].

6.2.2 Joint analysis and marginalization over calibration uncertainty

In order to jointly analyze O1 and O2 data (with measured cross-correlation spectra $\hat{C}_{\mathrm{O1}}$ and $\hat{C}_{\mathrm{O2}}$) in search of a gravitational-wave background, we will take advantage of the fact that $\hat{C}_{\mathrm{O1}}$ and $\hat{C}_{\mathrm{O2}}$ represent independent trials and hence have a factorizable likelihood:

$$\mathcal{L}(\hat{C}_{\mathrm{O1}}, \hat{C}_{\mathrm{O2}} | \Omega_M) = \mathcal{L}(\hat{C}_{\mathrm{O1}} | \Omega_M) \mathcal{L}(\hat{C}_{\mathrm{O2}} | \Omega_M). \tag{6.18}$$

We will handle calibration uncertainty slightly differently than when we analyzed O1 alone. As in Ch. 6.1 above, we incorporate imperfect calibration by adopting a likelihood

$$\begin{aligned}\mathcal{L}\left(\hat{C}|\kappa,\Omega_M\right) &\propto \prod_f \exp\left[-\frac{\left(\hat{C}(f)-\kappa\gamma(f)\Omega_M(f)\right)^2}{2\sigma^2(f)}df\right] \\ &= \mathcal{N}\exp\left[-\sum_i \frac{\left(\hat{C}(f_i)-\kappa\gamma(f_i)\Omega_M(f_i)\right)^2}{2\sigma^2(f_i)}df\right]\end{aligned} \tag{6.19}$$

where $\Omega_M(f)$ is our predicted energy-density spectrum based on a model for the gravitational-wave background and κ is an unknown multiplicative factor that encapsulates potential calibration inaccuracy. Although in the past we've written similar expressions as an integral over continuous frequencies f, here we'll explicitly retain the summation over discrete frequencies f_i. $\mathcal{N}$, meanwhile, is a normalization factor that we'll carry through this calculation. As above, we will treat κ as a random variable variable drawn from a normal distribution centered at 1 (perfect calibration) but with a variance ϵ^2:

$$p(\kappa) \propto \exp\left[-\frac{1}{2\epsilon^2}\left(\kappa-1\right)^2\right]. \tag{6.20}$$

When analyzing O1 data above, we incorporated κ as an additional parameter that we sampled with MultiNest, numerically achieving the marginalization over calibration uncertainty. This strategy worked well when we had only a single calibration factor to worry about. However, this strategy isn't sustainable when jointly analyzing data from multiple runs and/or multiple detector pairs. In this case, we must include a separate independent calibration parameter for every detector pair and every observing run (e.g. $\kappa^{\mathrm{HL}}_{\mathrm{O1}}$, $\kappa^{\mathrm{HL}}_{\mathrm{O2}}$, etc.). These additional parameters rapidly increase the dimensionality of our parameter space, making numerical marginalization over calibration uncertainties infeasible. Instead, we will shift strategies and analytically marginalize over calibration uncertainties, prior to parameter estimation on the gravitational-wave background with MultiNest.

First, we can impose the additional constraint that κ be positive – that is, we might not know the amplitude of the measured strain correct, but we are confident in its sign. In this case, the probability distribution for κ becomes

$$p(\kappa) = \sqrt{\frac{2}{\pi}}\frac{1}{\epsilon\left[1+\mathrm{Erf}(\frac{1}{\sqrt{2\epsilon^2}})\right]}\exp\left[-\frac{1}{2\epsilon^2}\left(\kappa-1\right)^2\right], \tag{6.21}$$

normalized to unity on the interval $\kappa \in (0, \infty)$. Equation (6.21) will function as our prior on κ. Next, with Eq. (6.21) in hand, we can marginalize Eq. (6.19) over κ:

$$\begin{aligned}
&\mathcal{L}_{\text{Marg}}(\hat{C}|\Omega_M) \\
&\quad = \int \mathcal{L}(\hat{C}|\kappa, \Omega_M) p(\kappa) d\kappa \\
&\quad = \mathcal{N}\sqrt{\frac{2}{\pi}} \frac{1}{\epsilon\left[1 + \text{Erf}(\frac{1}{\sqrt{2\epsilon^2}})\right]} \\
&\qquad \times \int_0^\infty \exp\left[-\frac{1}{2}\sum_i \frac{\left(\hat{C}(f_i) - \kappa\gamma(f_i)\Omega_M(f_i)\right)^2}{\sigma^2(f_i)} - \frac{1}{2}\frac{(\kappa-1)^2}{\epsilon^2}\right] d\kappa.
\end{aligned} \tag{6.22}$$

Expanding the exponential and grouping powers of κ:

$$\begin{aligned}
&\mathcal{L}_{\text{Marg}}(\hat{C}|\Omega_M) \\
&\quad = \mathcal{N}\sqrt{\frac{2}{\pi}} \frac{1}{\epsilon\left[1 + \text{Erf}(\frac{1}{\sqrt{2\epsilon^2}})\right]} \\
&\qquad \times \int_0^\infty \exp\Bigg[-\frac{1}{2}\Bigg\{ \kappa^2 \left(\frac{1}{\epsilon^2} + \sum_i \frac{[\gamma(f_i)\Omega_M(f_i)]^2}{\sigma^2(f_i)}\right) \\
&\qquad\qquad - 2\kappa\left(\frac{1}{\epsilon^2} + \sum_i \frac{\hat{C}(f_i)\gamma(f_i)\Omega_M(f_i)}{\sigma^2(f_i)}\right) \\
&\qquad\qquad + \left(\frac{1}{\epsilon^2} + \sum_i \frac{\hat{C}^2(f_i)}{\sigma^2(f_i)}\right)\Bigg\}\Bigg] d\kappa \\
&\quad \equiv \mathcal{N}\sqrt{\frac{2}{\pi}} \frac{1}{\epsilon\left[1 + \text{Erf}(\frac{1}{\sqrt{2\epsilon^2}})\right]} \int_0^\infty \exp\left[-\frac{1}{2}\left(A\kappa^2 - 2B\kappa + C\right)\right] d\kappa,
\end{aligned} \tag{6.23}$$

where for convenience we've defined

$$A = \frac{1}{\epsilon^2} + \sum_i \frac{[\gamma(f_i)\Omega_M(f_i)]^2}{\sigma^2(f_i)}, \tag{6.24}$$

$$B = \frac{1}{\epsilon^2} + \sum_i \frac{\hat{C}(f_i)\gamma(f_i)\Omega_M(f_i)}{\sigma^2(f_i)}, \tag{6.25}$$

and

$$C = \frac{1}{\epsilon^2} + \sum_i \frac{\hat{C}^2(f_i)}{\sigma_i^2}. \tag{6.26}$$

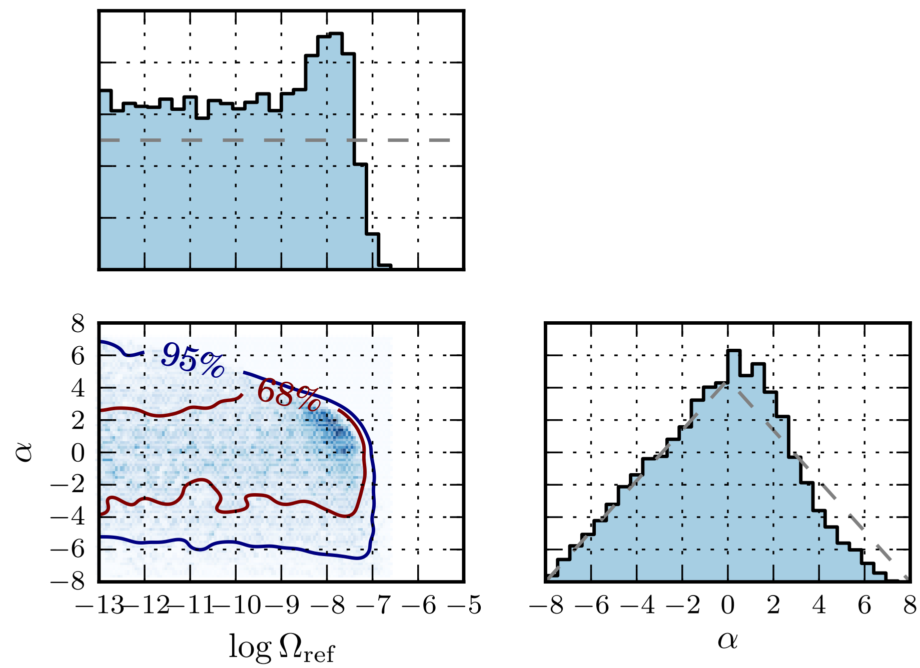

Figure 6.8: Posterior on the amplitude and spectral index of the gravitational-wave background following Advanced LIGO's O1 and O2 observing runs. Here we assume a log-uniform prior on the background's amplitude. Although the posterior does show a marginal peak near $\log \Omega_0 \approx -8$ and $\alpha \approx 1$, this peak is not statistically significant.

The final line of Eq. (6.23) is now in a form that we can integrate analytically by hand, giving

$$\boxed{\mathcal{L}_{\rm Marg}(\hat{C}|\Omega_M) = \mathcal{N}\frac{1}{\epsilon\sqrt{A}}\left[\frac{1+\mathrm{Erf}(\frac{B}{\sqrt{2A}})}{1+\mathrm{Erf}(\frac{1}{\sqrt{2\epsilon^2}})}\right]\exp\left[-\frac{1}{2}\left(C-\frac{B^2}{A}\right)\right].} \tag{6.27}$$

The final likelihood that we implement in MultiNest is

$$\mathcal{L}(\hat{C}_{\rm O1},\hat{C}_{\rm O2}|\Omega_M) = \mathcal{L}_{\rm Marg}(\hat{C}_{\rm O1}|\Omega_M)\mathcal{L}_{\rm Marg}(\hat{C}_{\rm O2}|\Omega_M). \tag{6.28}$$

Once again, we take $\epsilon = 0.072$ for O1, the root-mean-square of the 4.8% and 5.4% amplitude uncertainties in Hanford and Livingston [245]. In O2, amplitude uncertainties were estimated to be less than 2.6% and 3.85% in Hanford and Livingston across our 20-1726 Hz frequency band [245, 251]; we use the root-mean-square $\epsilon = 0.046$ for our O2 uncertainty.

6.2.3 O1 and O2 limits on the gravitational-wave background

Given the O1 and O2 cross-correlation spectra between Hanford and Livingston, we find a Bayes factor $\mathcal{B} = -0.33$ between the presence of a (tensor-

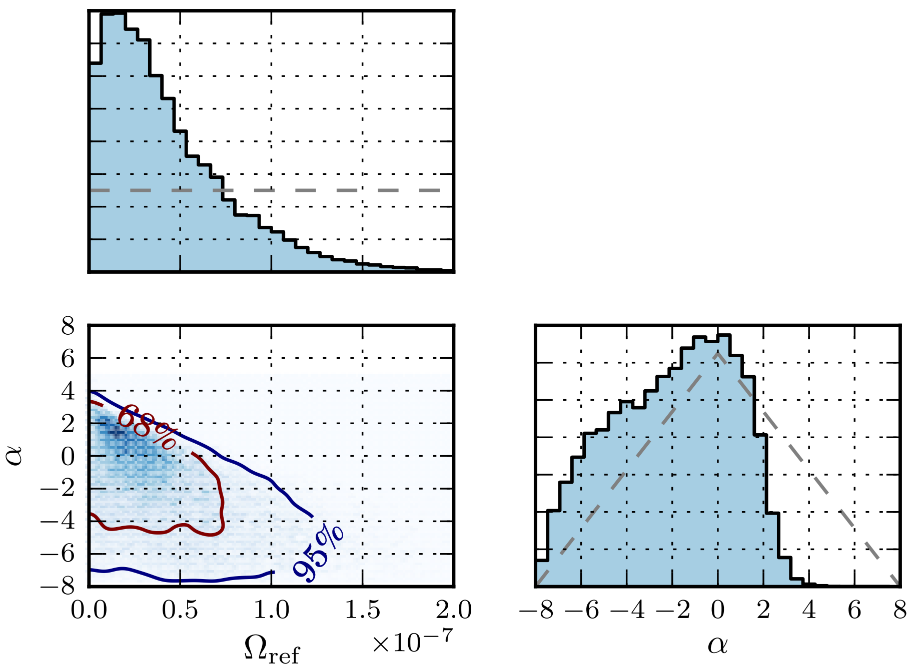

Figure 6.9: As in Fig. 6.8, but assuming a uniform amplitude prior.

polarized) gravitational-wave background and Gaussian noise. Hence we have again made no detection of a stochastic gravitational-wave signal.

We can, however, noticeably improve our upper limits on the possible strength of the gravitational-wave background. Figure 6.8 shows the joint posterior between the background amplitude Ω_0 and spectral index α following O1 and O2, assuming a log-uniform prior on Ω_0. For comparison to past frequentist results, Fig. 6.9 additionally shows the joint posterior when assuming a uniform amplitude prior. After marginalizing over spectral index, we obtain 95% credible upper limits

$$
\begin{aligned}
&\log \Omega_0 \leq -7.47 && \text{(Log-uniform prior)} \\
&\log \Omega_0 \leq -6.95 && \text{(Uniform prior)}
\end{aligned}
\tag{6.29}
$$

Comparing to the O1-only limits in Ch. 6.1 above, these represent an improvement by roughly a factor of two. These upper limits, along with 95% credible bounds on α, are listed in Tables 6.5 and 6.6.

We can additionally update our constraints on the presence of non-standard polarizations in the gravitational-wave background. After jointly analyzing O1 and O2 data, we obtain the Bayes factors shown in Table 6.4 between each

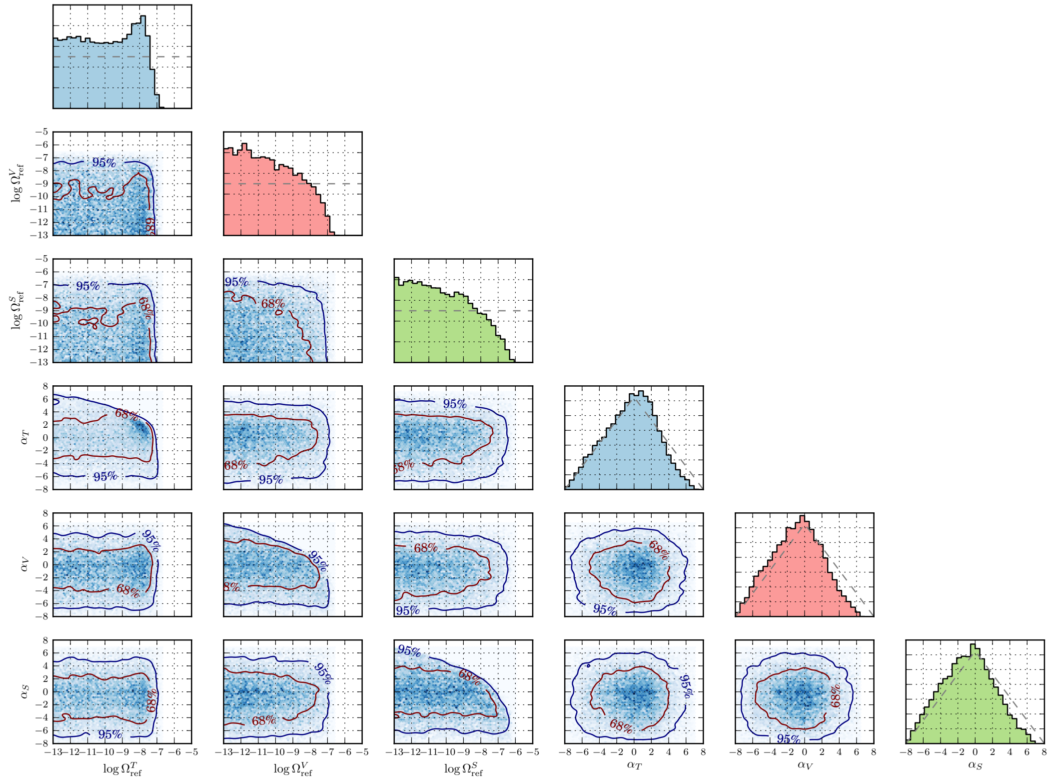

Figure 6.10: Posterior on possible tensor, vector, and scalar-polarized contributions to the gravitational-wave background, following the O1 and O2 observing runs. These results assume a log-uniform prior on the amplitude of each background contribution.

signal sub-hypothesis and Gaussian noise. If we adopt the same prior odds as in Chs. 5 and 6.1 above (see Fig. 5.15), we are left with odds $\mathcal{O}^{\rm SIG}_{\rm N} = -0.53$ between a generically-polarized stochastic background, and odds $\mathcal{O}^{\rm NGR}_{\rm GR} = -0.25$ for the presence of alternative gravitational-wave polarizations. Both are, unsurprisingly, consistent with Gaussian noise. Figures 6.10 and 6.11 show the full posteriors for the TVS model allowing the existence of tensor, vector, and scalar modes, and 95% credible limits are listed in Tables 6.5 and 6.6 for all unique combinations of polarizations.

It is worth mentioning that, unlike our O1-only results above, the our joint O1+O2 posteriors are not featureless. In fact, Fig. 6.8 contains a rather noticeable peak,[1] centered at roughly $\Omega_0 \approx 10^{-8}$ and $\alpha \approx 1$, suggestively compatible with the expected slope $\alpha = 2/3$ due to compact binary mergers. Even more notably, this peak appears in Fig. 6.10 showing joint posteriors for tensor, vector, and scalar-polarized contributions to the stochastic background. Within

[1]Infamously known to the Stochastic Group as "The Bump."

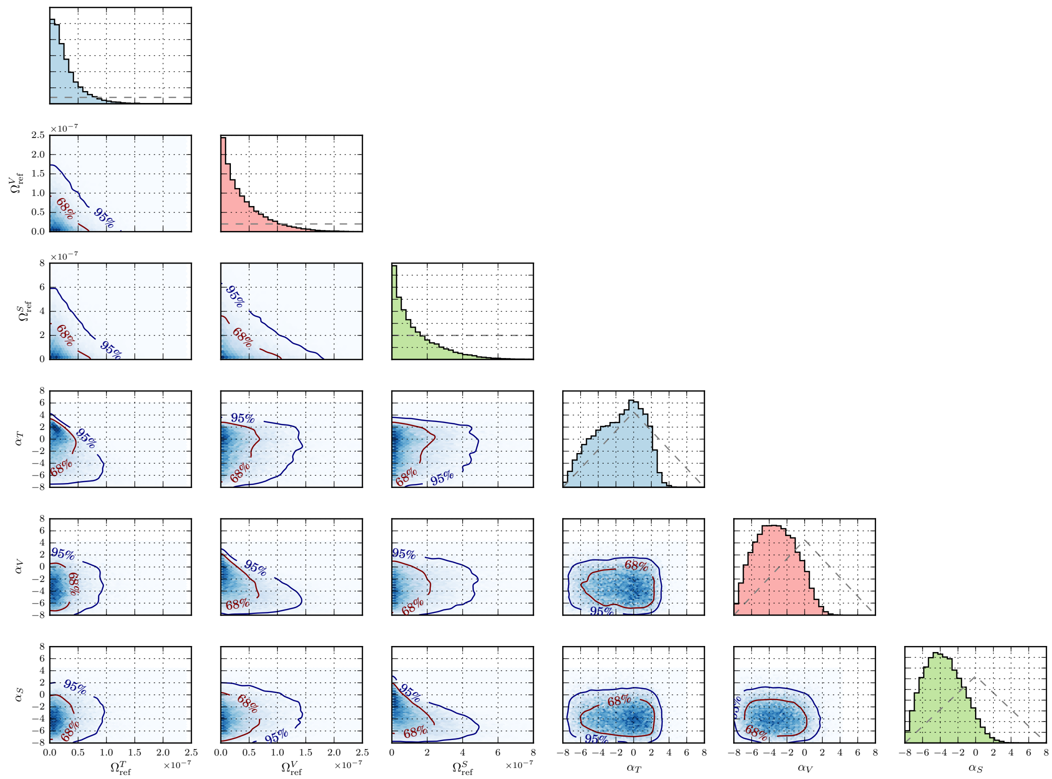

Figure 6.11: As in Fig. 6.10, but assuming a uniform amplitude prior.

Hypothesis	$\ln \mathcal{B}$
T	-0.33
V	-0.33
S	-0.31
TV	-0.66
TS	-0.65
VS	-0.65
TVS	-0.99

Table 6.4: Log Bayes factors between each signal sub-hypothesis considered and the Gaussian noise hypothesis.

Fig. 6.10, the peak appears *only in the tensor sector*. Recall that different gravitational-wave polarizations are modulated by different overlap reduction functions in our cross-correlation measurements $\hat{C}(f)$. Gaussian noise (or even non-Gaussian terrestrial sources of coherence) does not "know" about the overlap reduction functions; naively, we might therefore expect a loud noise instantiation to yield a peak in *all* polarization sectors. In contrast, Fig. 6.10 shows exactly the behavior we might expect of a real gravitational-wave background obeying general relativity.

As illustrated below in Ch. 6.A, this intuition is misleading – loud noise realizations *can*, in fact, lead to marginal posterior peaks that appear asymmetrically between the tensor, vector, and scalar sectors. This behavior can be traced back to the overlap reduction functions. Because the overlap reduction functions for the various gravitational-wave polarizations have different magnitudes, in any given frequency band certain polarizations will appear more strongly in $\hat{C}(f)$ than others. Thus elevated noise in $\hat{C}(f)$ can, depending on the particular frequency, be more readily identified with certain polarizations than than others.

6.3 Looking Ahead

So far we have detected no signs of the gravitational-wave background. Even so, our new upper limits represent a remarkable improvement over past results. The best upper limit set by Initial LIGO was, restricting to a spectral index $\alpha = 0$, $\Omega_0 \leq 5.6 \times 10^{-6}$ [229]. With our combined analysis of O1 and O2 data from Advanced LIGO, we have achieved an upper limit of $\Omega_0 \leq 3.4 \times 10^{-8}$ (marginalized across all α), or 3.5×10^{-8} if we set a delta-function prior at $\alpha = 0$. Despite our non-detection, we have improved on previous upper limits by more than two orders of magnitude!

As I write, the next O3 observing run is currently underway, with both the Advanced LIGO Hanford and Livingston detectors as well as Advanced Virgo joining from the start. Current projections estimate that the astrophysical background due to compact binary mergers *may* be detectable with the Advanced LIGO-Virgo network in the next several years [84]. The coming observing runs will also see the addition of KAGRA [35, 37]. Despite the large separation between KAGRA, Virgo, and LIGO, this increasing number of baselines will further accelerate our progress towards detection of the gravitational-wave

Table 6.5: Marginalized parameter estimation results from O1 and O2, for each unique combination of tensor, vector, and scalar polarizations and assuming log-uniform amplitude priors. Specifically, we list the 95% credible upper limits on each amplitude parameter, and the median recovered spectral indices with uncertainties encompassing the central 95% credible region.

Hypothesis	$\log\Omega_0^{T,95\%}$	$\log\Omega_0^{V,95\%}$	$\log\Omega_0^{S,95\%}$	$\Omega_0^{T,95\%}$	$\Omega_0^{V,95\%}$	$\Omega_0^{S,95\%}$	α_T	α_V	α_S
T	-7.47	-	-	3.38×10^{-8}	-	-	$-0.2^{+5.3}_{-6.1}$	-	-
V	-	-7.48	-	-	3.32×10^{-8}	-	-	$-0.8^{+5.4}_{-5.7}$	-
S	-	-	-7.16	-	-	6.98×10^{-8}	-	-	$-0.7^{+5.8}_{-5.8}$
TV	-7.48	-7.52	-	3.30×10^{-8}	3.01×10^{-8}	-	$-0.2^{+5.3}_{-6.1}$	$-0.9^{+5.4}_{-5.6}$	-
TS	-7.47	-	-7.21	3.36×10^{-8}	-	6.19×10^{-8}	$-0.2^{+5.3}_{-6.2}$	-	$-0.8^{+5.8}_{-5.7}$
VS	-	-7.50	-7.19	-	3.17×10^{-8}	6.39×10^{-8}	-	$-0.9^{+5.5}_{-5.6}$	$-0.8^{+5.9}_{-5.7}$
TVS	-7.50	-7.55	-7.22	3.20×10^{-8}	2.85×10^{-8}	6.07×10^{-8}	$-0.2^{+5.3}_{-6.1}$	$-0.8^{+5.4}_{-5.7}$	$-0.8^{+5.8}_{-5.7}$

Table 6.6: As in Table 6.5, but assuming uniform amplitude priors.

Hypothesis	$\log \Omega_0^{T,95\%}$	$\log \Omega_0^{V,95\%}$	$\log \Omega_0^{S,95\%}$	$\Omega_0^{T,95\%}$	$\Omega_0^{V,95\%}$	$\Omega_0^{S,95\%}$	α_T	α_V	α_S
T	-6.95	-	-	1.13×10^{-7}	-	-	$-1.7^{+4.0}_{-5.3}$	-	-
V	-	-6.77	-	-	1.68×10^{-7}	-	-	$-3.6^{+4.3}_{-3.6}$	-
S	-	-	-6.24	-	-	5.78×10^{-7}	-	-	$-4.2^{+4.5}_{-3.1}$
TV	-7.02	-6.85	-	9.48×10^{-8}	1.42×10^{-7}	-	$-1.4^{+4.0}_{-5.4}$	$-3.4^{+4.4}_{-3.8}$	-
TS	-7.03	-	-6.32	9.39×10^{-8}	-	4.81×10^{-7}	$-1.2^{+3.8}_{-5.5}$	-	$-3.9^{+4.6}_{-3.4}$
VS	-	-6.84	-6.32	-	1.44×10^{-7}	4.80×10^{-7}	-	$-3.3^{+4.3}_{-3.9}$	$-3.9^{+4.6}_{-3.4}$
TVS	-7.09	-6.92	-6.38	8.20×10^{-8}	1.22×10^{-7}	4.16×10^{-7}	$-1.4^{+3.9}_{-5.4}$	$-3.3^{+4.4}_{-3.9}$	$-3.7^{+4.6}_{-3.5}$

background.

With this future in mind, it is crucial to begin preparing *now* for the day that a gravitational-wave background is apparently detected. The tentative detection of the stochastic background is, in many ways, the *easiest* part of our endeavor. Once we have identified statistically-significant correlated signal across our detector network, we will then have to address an even more challenging question: how can we be sure that this signal is due to the gravitational-wave background, rather than instrumental and/or terrestrial noise? We'll explore this question next in Ch. 7.

$\sim$

Appendix 6.A The polarization footprint of loud noise

In Ch. 6.2 above we saw that, although consistent with the non-detection of the gravitational-wave background, the joint O1+O2 stochastic analyses exhibited a marginal peak in our Ω_0 posterior. Moreover, when performing parameter estimation using our TVS model that encompasses alternative polarizations, we found that the peak appeared only in the tensor sector, just as we would expect for a true marginal detection of a gravitational-wave background obeying general relativity. Random noise, the argument went, does not "know" about gravitational-wave polarizations and their corresponding overlap reduction functions; if the "bump" were due to noise alone, we might therefore expect it to appear in all three polarization sectors, rather than standard tensor polarizations alone.

To semi-quantitatively test this intuition, we can purposefully inject mock instantiations of elevated noise and inspect the resulting parameter estimation on the gravitational-wave background. Figure 6.12, for instance, shows simulated measurements of Gaussian noise consistent with the variance of our true O2 cross-correlation spectrum. To this Gaussian noise we have added a frequency comb with a (negative) amplitude of 3×10^{-7}. The resulting amplitude posteriors obtained when analyzing this data under the TVS hypothesis are shown in Fig. 6.13. Contrary to intuition, the injected comb doesn't contribute apparent power symmetrically across polarizations, but instead yields a peak confined entirely to the tensor sector, qualitatively quite similar to the posterior peak that appeared in our O2 analysis above.

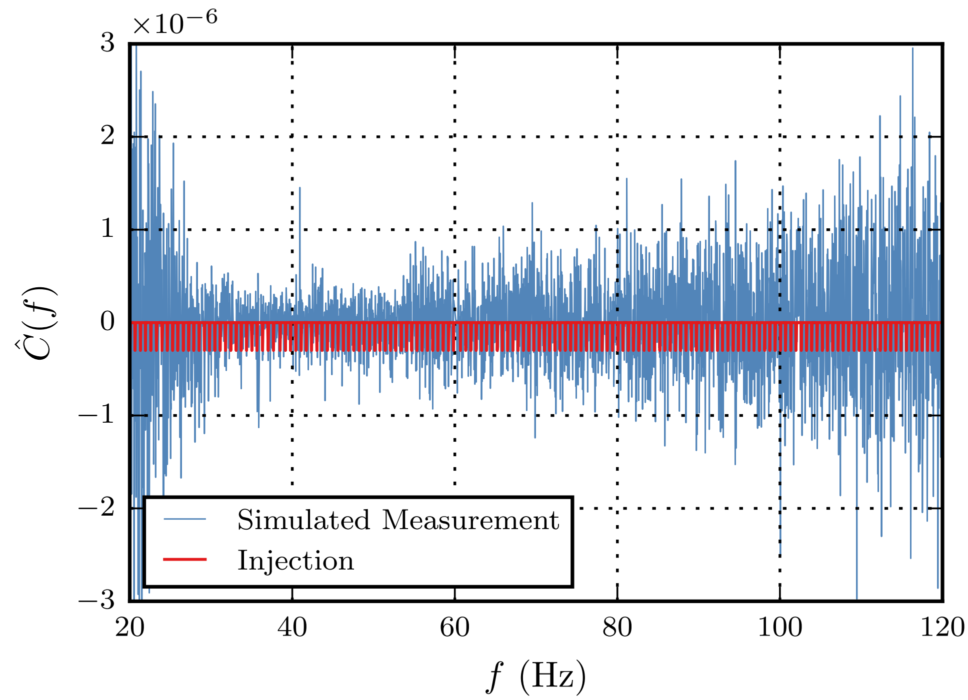

Figure 6.12: Simulated Hanford-Livingston cross-correlation spectrum, consisting of a frequency comb plus Gaussian noise. The Gaussian noise is consistent with the $\sigma^2(f)$ uncertainty spectrum given by the O2 stochastic analysis. The filled grey region spans $\pm 1\,\sigma(f)$. The corresponding parameter estimation results are shown in Fig. 6.13.

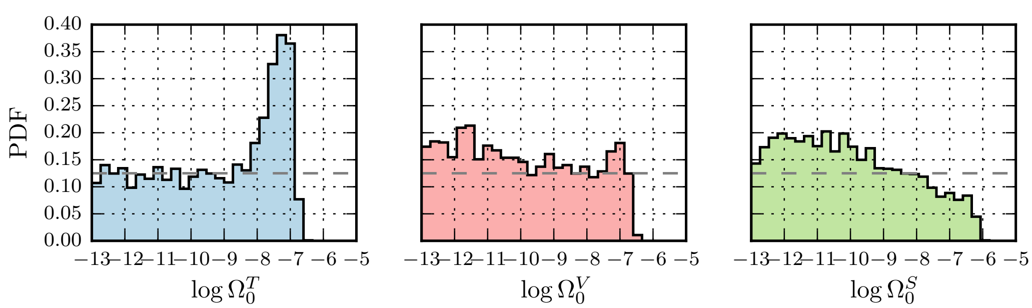

Figure 6.13: Tensor, vector, and scalar amplitude posteriors (under the TVS model) given by analysis of the simulated comb in Fig. 6.12. These results are qualitatively similar to that of our O2 stochastic search (see Figs. 6.8 and 6.10), with a peak that is confined entirely to Ω_0^T.

Besides non-Gaussian features like correlated combs, we can alternatively obtain apparent excess of noise if we *underestimate* the uncertainties $\sigma(f)$ on the measured cross-correlation spectrum $\hat{C}(f)$. Such underestimates can occur, for example, if estimates of the detectors' PSDs are biased by non-stationary data. Figure 6.14 shows a simulated observation of random Gaussian noise, with variance consistent *not* with the $\sigma^2(f)$ measured in O2 (black curves), but instead with a deliberately inflated variance $1.7\sigma^2(f)$ (red curves). Analysis of this inflated noise yields the amplitude posteriors shown in Fig. 6.15. We again see a peak confined to one polarization mode alone, but this time the vector sector. The exact behavior of elevated noise varies significantly from realization to realization; the polarization sector in which a peak appears varies between trials, and often peaks appear in two (but very rarely three) sectors.

Our initial intuition – that terrestrial artifacts or elevated noise do not "know" about polarizations and hence should project equally onto tensor, vector, and scalar modes – appears incorrect. Thus the kinds of outlier that appeared in the O2 analysis, confined entirely to tensor polarizations, are in fact consistent with noise and should not necessarily be taken as evidence of a marginal astrophysical signal.

The reason that elevated noise can contribute so asymmetrically to the different polarization sectors can be traced back to the overlap reduction functions. The fact that the overlap reduction functions for different polarizations have distinct spectral shapes means that, at a given frequency, elevated noise can masquerade as certain polarization modes much more readily than others. Recall that the signal-to-noise ratio of a measured gravitational-wave signal is

$$\mathrm{SNR} = \frac{\left(\hat{C}|\gamma_A \Omega_M^A\right)}{\sqrt{\left(\gamma_B \Omega_M^B|\gamma_C \Omega_M^C\right)}}. \tag{6.30}$$

Let's restrict to the cases of a pure-tensor (hypothesis "T"), pure-vector (V), and pure-scalar (S) models for the gravitational-wave background. Furthermore, assume flat energy-density spectra, with spectral indices $\alpha_A = 0$. Then

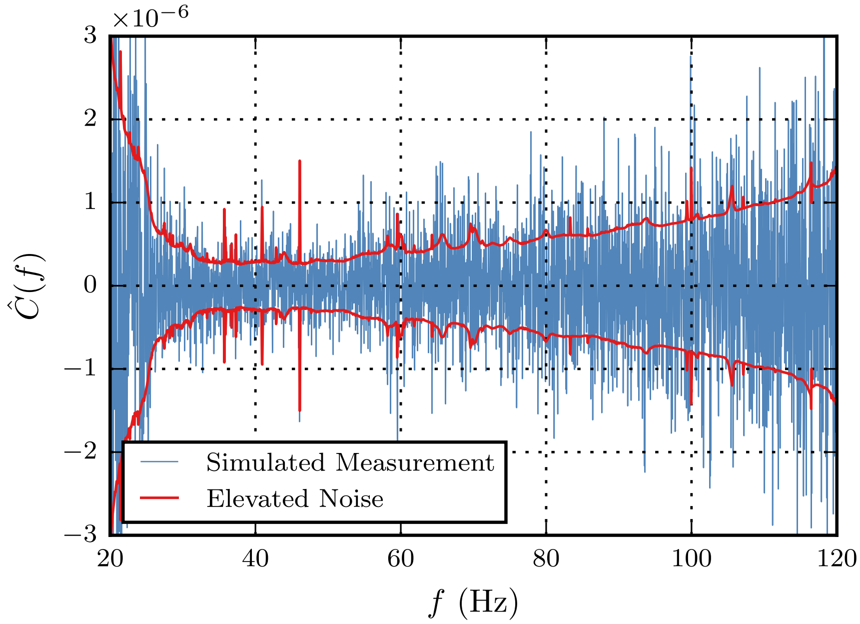

Figure 6.14: Simulated Hanford-Livingston cross-correlation spectrum, consisting of elevated Gaussian noise. The shaded grey region spans $\pm 1\,\sigma(f)$, as given by our O2 stochastic search results. The simulated measurements, however, are randomly drawn from a deliberately broadened Gaussian distribution with variance $1.7\sigma^2(f)$; the inflated ± 1 standard deviation band is shown in red. The corresponding parameter estimation results are shown in Fig. 6.15.

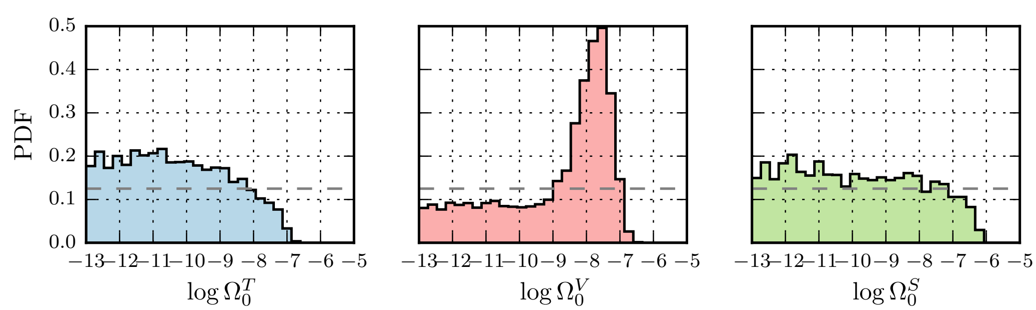

Figure 6.15: Tensor, vector, and scalar amplitude posteriors (under the TVS model) given by analysis of the simulated elevated noise in Fig. 6.14. Like the simulated comb, here elevated noise contributes asymmetrically to the inferred polarization content of the gravitational-wave background, this time yielding a peak in the vector amplitude Ω_0^V.

the signal-to-noise ratio inferred under our T hypothesis, for example, is

$$\begin{aligned}\mathrm{SNR}_T &= \frac{\left(\hat{C}|\gamma_T\right)}{\sqrt{(\gamma_T|\gamma_T)}} \\ &= \left(\int \frac{\gamma_T^2(f)}{\sigma^2(f)}df\right)^{-1/2} \int \frac{\hat{C}(f)\gamma_T(f)}{\sigma^2(f)}df.\end{aligned} \tag{6.31}$$

Similarly, the signal-to-noise ratios obtained under the V and S hypotheses are

$$\mathrm{SNR}_V = \left(\int \frac{\gamma_V^2(f)}{\sigma^2(f)}df\right)^{-1/2} \int \frac{\hat{C}(f)\gamma_V(f)}{\sigma^2(f)}df \tag{6.32}$$

and

$$\mathrm{SNR}_S = \left(\int \frac{\gamma_S^2(f)}{\sigma^2(f)}df\right)^{-1/2} \int \frac{\hat{C}(f)\gamma_S(f)}{\sigma^2(f)}df. \tag{6.33}$$

Given these three models and associated broadband SNRs, define the **SNR response** $\mathrm{SNR}_A(f)$ as the recovered signal-to-noise ratio due to a delta-function peak at frequency f when assuming polarization A:

$$\begin{aligned}\mathrm{SNR}_T(f) &= \left(\int \frac{\gamma_T^2(f)}{\sigma^2(f)}df\right)^{-1/2} \frac{\gamma_T(f)}{\sigma^2(f)} \\ \mathrm{SNR}_V(f) &= \left(\int \frac{\gamma_V^2(f)}{\sigma^2(f)}df\right)^{-1/2} \frac{\gamma_V(f)}{\sigma^2(f)} \\ \mathrm{SNR}_S(f) &= \left(\int \frac{\gamma_S^2(f)}{\sigma^2(f)}df\right)^{-1/2} \frac{\gamma_S(f)}{\sigma^2(f)}.\end{aligned} \tag{6.34}$$

These SNR responses are plotted in Fig. 6.16. The relative signs and amplitudes indicate how elevated noise at different frequencies projects onto different gravitational-wave polarizations. In region "A," for instance, all three SNR responses are negative, and $\mathrm{SNR}_T(f)$ and $\mathrm{SNR}_V(f)$ have roughly equal amplitudes. Thus elevated *negative* coherence in the 20-40 Hz range is expected to project onto all three polarization sectors, but to be primarily identified with tensor or vector-polarized gravitational waves. In region "B," meanwhile, the scalar response is positive while the tensor response has the largest negative amplitude. Parameter estimation of excess positive coherence in this 40-45 Hz range would therefore prefer a scalar-polarized signal, whereas excess negative coherence would be most readily identified as a tensorial gravitational-wave background.

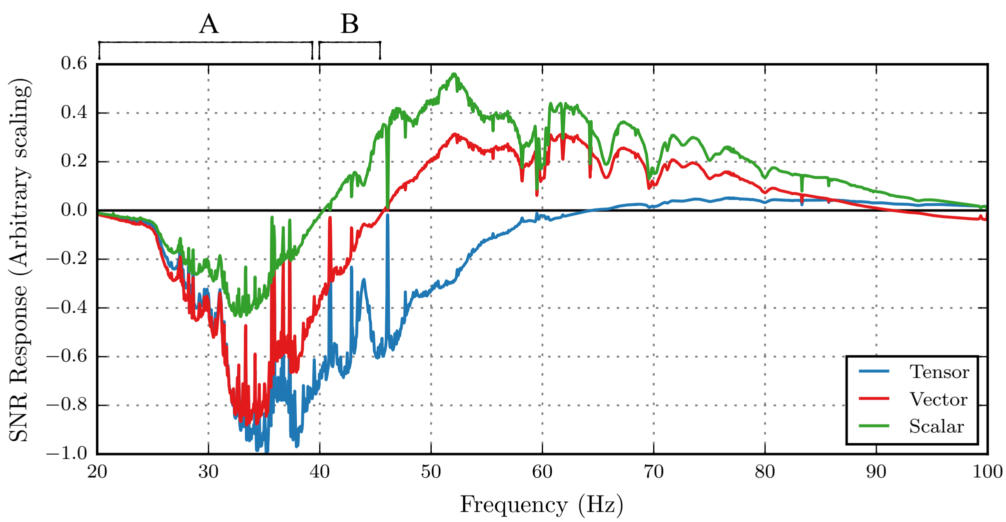

Figure 6.16: Relative SNRs inferred from delta-function cross-correlations $\hat{C}$ at different frequencies, when assuming tensor, vector, and scalar polarized gravitational-wave signals.

Chapter 7
Validating Detection of the Stochastic Gravitational-Wave Background

∾

This chapter contains work published in:

> T. A. Callister, M. W. Coughlin, J. B. Kanner, *Gravitational-wave geodesy: A new tool for validating detection of the stochastic gravitational-wave background*, Astrophys. J. Letters **869**, L28 (2018).

I conceived of this project, produced all results, and wrote the majority of the published manuscript.

∾

7.1 The Danger of Correlated Noise

So far no observation has been made of the stochastic gravitational-wave background. In Ch. 6, we saw that analysis of Advanced LIGO's O1 and O2 data showed no evidence for the detection of a stochastic signal; instead we limited its amplitude to $\Omega_0 \leq 3.4 \times 10^{-8}$ at 25 Hz, improving on previous limits by two orders of magnitude. Our situation may soon change, though. The latest binary black hole and binary neutron star merger rates inferred by Advanced LIGO and Virgo [45, 46, 111] indicate that the astrophysical stochastic gravitational-wave background may in fact be detectable in the next several years [26, 48, 84, 176].

The cross-correlation search for the stochastic background relies on the assumption that, in the absence of a gravitational-wave signal, the outputs of different gravitational-wave detectors are fundamentally uncorrelated. The LIGO-Hanford and LIGO-Livingston detectors, for instance, are separated by

3000 km, with a light travel time of ≈ 0.01 s between sites. One might therefore reasonably expect them to be safely uncorrelated at frequencies greater than roughly $(0.01\,\mathrm{s})^{-1} \sim 100$ Hz, in the frequency band of interest for ground-based detectors. All our work performed so far in Chs. 4-6 has been predicated on this fundamental assumption.

In reality, however, terrestrial gravitational-wave detectors are *not* truly uncorrelated. Hanford-Livingston coherence spectra consistently show correlated features that, if not properly identified and removed, can severely contaminate searches for the stochastic gravitational-wave background [81]. Schumann resonances are one expected source of terrestrial correlation [252, 253]. Global electromagnetic excitations in the cavity formed by the Earth and its ionosphere, Schumann resonances may magnetically couple to Advanced LIGO and Advanced Virgo's test mass suspensions and induce a correlated signal between detectors [84, 186, 242, 243, 254, 255]. Another expected source of correlation is the joint synchronization of electronics at each detector to Global Positioning System (GPS) time. In Advanced LIGO's O1 observing run, for instance, a strongly correlated 1 Hz comb was traced to blinking LED indicators on timing systems independently synchronized to GPS [81].

Any undiagnosed terrestrial correlations may yield a false-positive detection of the stochastic gravitational-wave background. We even just saw in Ch. 6.A above that a comb can yield spuriously sensible parameter estimation results, such as energy-density confined entirely to tensor modes as expected from general relativity. More concerningly, while Schumann resonances and frequency combs represent two known classes of correlation, other currently unknown classes may also exist. The validation of any apparent observation of the stochastic background will therefore require us to answer the following question:

How likely is an observed correlated signal to be of astrophysical origin, rather than a yet-unidentified source of terrestrial correlation?

The gravitational-wave community lacks the tools to quantitatively answer this question. Searches for gravitational-wave transients can address this issue through the use of time-slides: the artificial time-shifting of data from one detector relative to another's. This process eliminates any coherent gravitational-wave signals while preserving all other properties of the data, allowing for accurate estimation of the false-positive detection rate. In cross-correlation searches for the stochastic background, however, time-slides would not only remove a gravitational-wave signal but also any correlated terrestrial contamination as well. Time-slides are therefore of limited use in searches for the gravitational-wave background.

In this chapter I will explore a novel technique with which to evaluate the astrophysical significance of an apparent correlated stochastic signal. Our method is inspired by the field of radio geodesy, in which interferometric observations of the radio sky serve to precisely localize radio telescopes on the Earth. In the same way, we will see that measurements of the gravitational-wave background can be similarly reverse-engineered to infer the separations and relative orientations of gravitational-wave detectors. By demanding that a true gravitational-wave background yield results consistent with the *known* geometry of our detectors, we can separate true gravitational-wave signals from spurious terrestrial correlations.

7.2 Gravitational-Wave Geodesy

Recall that searches for the stochastic background seek to measure the gravitational-wave energy density $\Omega(f)$ by computing the cross-correlation spectrum $\hat{C}(f)$ between pairs of gravitational-wave detectors:

$$\hat{C}(f) = \frac{1}{\Delta T}\frac{20\pi^2}{3H_0^2} f^3\,\mathrm{Re}\left[\tilde{s}_1^*(f)\tilde{s}_2(f)\right], \tag{7.1}$$

where ΔT is the time duration of data analyzed, and $\tilde{s}_I(f)$ is the (Fourier domain) strain measured by detector I. Equation (7.1) is normalized such that, for Advanced LIGO, the expectation value of $\hat{C}(f)$ is [83]

$$\langle \hat{C}(f) \rangle = \gamma(f)\Omega(f). \tag{7.2}$$

In the weak signal limit, the variance of $\hat{C}(f)$ is given by $\langle \hat{C}(f)\hat{C}(f')\rangle = \delta(f-f')\sigma^2(f)$, with

$$\sigma^2(f) = \frac{1}{\Delta T}\left(\frac{10\pi^2}{3H_0^2}\right)^2 f^6 P_1(f)P_2(f), \tag{7.3}$$

where $P_I(f)$ is the one-sided noise power spectral density of detector I. The overlap reduction function $\gamma(f)$ encodes the dependence of the measured correlations on the detector baseline geometry – the detectors' locations and relative

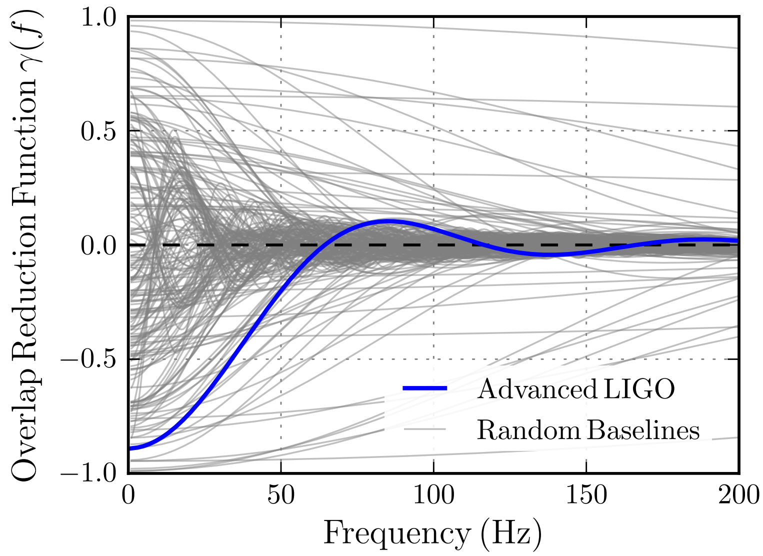

Figure 7.1: Overlap reduction function $\gamma(f)$ (blue) for the Advanced LIGO's Hanford-Livingston detector baseline. Alternative baseline geometries have different overlap reduction functions as illustrated by the collection of grey curves, which show overlap reduction functions between hypothetical detectors randomly positioned on Earth's surface.

orientations [186]. Advanced LIGO's normalized overlap reduction function is given by

$$\gamma(f) = \frac{5}{8\pi} \sum_A \int_{\text{Sky}} F_1^A(\hat{\mathbf{n}}) F_2^A(\hat{\mathbf{n}}) e^{2\pi i f \Delta\mathbf{x}\cdot\hat{\mathbf{n}}/c} d\hat{\mathbf{n}}. \tag{7.4}$$

Here, $F_I^A(\hat{\mathbf{n}})$ is the antenna response of detector I to gravitational waves of polarization A and $\Delta\mathbf{x}$ is the separation vector between detectors. The integral is performed over all sky directions $\hat{\mathbf{n}}$ and a sum is taken over both the "plus" and "cross" gravitational-wave polarizations. The leading factor of $5/8\pi$ normalizes the overlap reduction function such that identical, coincident, and co-aligned detectors would have $\gamma(f) = 1$.

Overlap reduction functions are strongly dependent upon baseline geometry – different pairs of gravitational-wave detectors generically have very different overlap reduction functions. To illustrate this, the overlap reduction function for the Hanford and Livingston baseline is shown in blue in Fig. 7.1. The collection of grey curves, meanwhile, illustrates alternative overlap reduction functions for hypothetical pairs of detectors placed randomly on the surface of the Earth.

The strong dependence of $\gamma(f)$ on baseline geometry raises an interesting possibility. Given cross-correlation measurements $\hat{C}(f)$ between two detectors, we

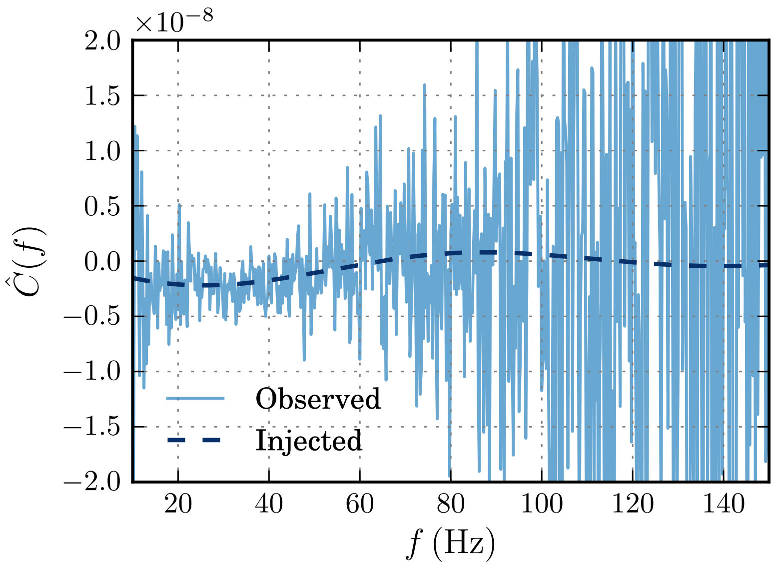

Figure 7.2: Simulated Advanced LIGO cross-correlation measurements (blue) following a three-year observation of an isotropic stochastic gravitational-wave background. The injected background has energy-density $\Omega(f) = 3.33 \times 10^{-9}\,(f/25\,\mathrm{Hz})^{2/3}$, corresponding to an expected signal-to-noise ratio of 10 after three years of observation. The dashed curve shows the expected cross-correlation in the absence of measurement noise, and the grey band indicates $\pm 1\sigma$ uncertainties. The distance between Hanford and Livingston as inferred from this observation is shown in Fig. 7.3.

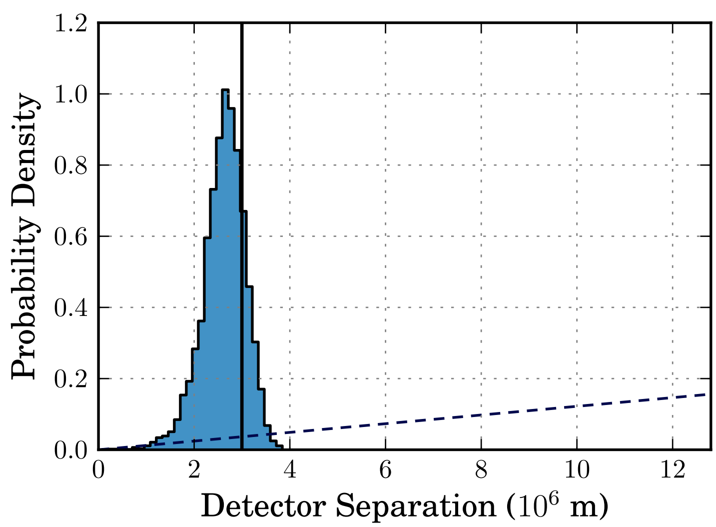

Figure 7.3: The posterior on the distance between the LIGO Hanford and Livingston detectors, obtained using the simulated cross-correlation measurements shown in Fig. 7.2. The dashed line indicates the distance prior used and the vertical black line marks the true Hanford-Livingston separation. Using the gravitational-wave sky, we self-consistently recover a posterior compatible with the true distance between detectors. Details regarding parameter estimation are explained in Ch. 7.3 below.

can use the measurements themselves to infer the baseline's geometry. In the electromagnetic domain, a very similar technique has long been used in the field of **geodesy**: the experimental study of Earth's geometry. While most commonly used to study the radio sky, very-long baseline interferometry can instead be used to precisely localize radio telescopes on the Earth, allowing for measurements of tectonic motion to better than $\sim 0.1\,\mathrm{mm\,yr^{-1}}$ [256, 257]. Similarly, here we will use the *gravitational-wave sky* to determine our detectors' relative positions and orientations.

As an initial demonstration, Fig. 7.2 illustrates a simulated observation of the stochastic gravitational-wave background with design-sensitivity Advanced LIGO. We assume a stochastic energy-density spectrum given by $\Omega(f) = 3.3 \times 10^{-9}\,(f/25\,\mathrm{Hz})^{2/3}$, chosen to yield SNR $= 10$ after three years of observation. The dashed curve indicates the mean correlation spectrum $\langle \hat{C}(f) \rangle$ corresponding to this injection, while the solid trace shows a simulated cross-correlation spectrum $\hat{C}(f)$ after three years of observation. By fitting to $\hat{C}(f)$ (as will be described below in Ch. 7.3), we can attempt to estimate the geometry of the Hanford-Livingston baseline. The resulting posterior on the separation between the Hanford and Livingston detectors is shown in Fig. 7.3. This posterior is consistent with the true separation between detectors ($\approx 3000\,\mathrm{km}$).

7.3 Differentiating Astrophysical and Terrestrial Sources of Correlation

Of course, the physical separations between current gravitational-wave detectors are already known to far better precision than can be obtained through gravitational-wave geodesy. Nevertheless, the ability to measure baseline geometry with the gravitational-wave sky suggests a powerful consistency test for any possible detection of the gravitational-wave background.

In the presence of an isotropic, astrophysical stochastic background, the measured cross-correlation spectrum $\hat{C}(f)$ *must* exhibit amplitude modulations and zero-crossings consistent with the baseline's overlap reduction function. Thus, when using the data $\hat{C}(f)$ to infer the baseline's geometry, we must obtain results that are consistent with the known separations and orientations of the detectors. In contrast, spurious sources of terrestrial correlation are *not* bound to trace the overlap reduction function. Hence, there is no *a priori* reason that a correlated terrestrial signal should prefer the true baseline geometry over any other possible detector configuration.

We can more rigorously define this consistency check within the framework of Bayesian hypothesis testing. Given a measured cross-correlation spectrum $\hat{C}(f)$, we can ask which of the following hypotheses better describes the data:

- Hypothesis $\mathcal{H}_\gamma$: The measured cross-correlation is consistent with the true baseline geometry (and hence the baseline's true overlap reduction function).

- Hypothesis $\mathcal{H}_{\text{Free}}$: The cross-correlation spectrum is consistent with a model in which we *do not* impose the baseline's known geometry, instead (unphysically) treating the detectors' positions and orientations as free variables to be determined by the data.

An isotropic, astrophysical stochastic signal will be consistent with both $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ (assuming that the true baseline geometry is among the possible configurations supported in $\mathcal{H}_{\text{Free}}$). As the simpler hypothesis, however, $\mathcal{H}_\gamma$ will be favored by the Bayesian "Occam's factor" that penalizes the more complex model. So a true isotropic astrophysical stochastic background will favor $\mathcal{H}_\gamma$. A generic terrestrial signal, on the other hand, is unlikely to follow the baseline's true overlap reduction function, and so will be better fit by the additional degrees of freedom allowed in $\mathcal{H}_{\text{Free}}$. Terrestrial sources of correlation are therefore likely to favor $\mathcal{H}_{\text{Free}}$.

This procedure is similar to the "sky scramble" technique used in pulsar timing searches for very low-frequency gravitational waves [258, 259]. In pulsar timing experiments, the analogue to the overlap reduction function is the Hellings and Downs curve, which quantifies the expected correlations between pulsars as a function of their angular separation on the sky [260].[1] By artificially shifting pulsar positions on the sky, one can seek to disrupt this spatial correlation and produce null data devoid of gravitational-wave signal but that retains other (possibly correlated) noise features.

Given a tentative detection of the stochastic background, we can compute an odds ratio $\mathcal{O}$ between hypotheses $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ to determine which is favored

[1] Whereas the overlap reduction function quantifies correlations between detectors as a function of their distance, in units of wavelengths.

by the data. We assume Gaussian likelihoods, such that the probability of measuring $\hat{C}(f)$ given a model spectrum $C_{\mathcal{H}}(\Theta; f)$ with parameters Θ is

$$p(\{\hat{C}\}|\Theta, \mathcal{H}) \propto \exp\left[-\frac{1}{2}\left(\hat{C} - C_{\mathcal{H}}(\Theta)|\hat{C} - C_{\mathcal{H}}(\Theta)\right)\right], \tag{7.5}$$

in terms of the inner product

$$(A|B) = 2\int_0^\infty \frac{A^*(f)B(f)}{\sigma^2(f)}df. \tag{7.6}$$

For both hypotheses, we adopt a standard power-law form for the background's energy-density spectrum, defined by a reference amplitude Ω_0 and a spectral index α:

$$\Omega(f) = \Omega_0\left(\frac{f}{25\,\mathrm{Hz}}\right)^\alpha. \tag{7.7}$$

Our model for the cross-correlation spectrum under $\mathcal{H}_\gamma$ is therefore

$$C_\gamma(\Omega_0, \alpha; f) = \gamma_{\mathrm{True}}(f)\,\Omega_0\,(f/25\,\mathrm{Hz})^\alpha, \tag{7.8}$$

where $\gamma_{\mathrm{True}}(f)$ is the true overlap reduction function for the given baseline.

For our unphysical hypothesis $\mathcal{H}_{\mathrm{Free}}$, we additionally need a parametrized model encompassing the various possible baseline geometries. We use the scheme illustrated in Fig. 7.4. Given two detectors on the surface of the Earth (which we approximate as a sphere of radius $R_\oplus = 6.4\times10^6$ m), we will choose coordinates such that the first detector lies at the pole and the second along the meridian, in the $x-z$ plane. We then have three remaining degrees of freedom that specify the baseline geometry: the polar angle θ between detectors, and the rotation angles ϕ_1 and ϕ_2 of each detector about its local zenith. Specifically, if $\hat{u}_i$ and $\hat{v}_i$ are unit vectors aligned with the arms of detector i, then we define ϕ_i as the angles between $\hat{v}_i$ and the y-axis:

$$\phi_i = \cos^{-1}\left(\hat{v}_i\cdot\hat{y}\right). \tag{7.9}$$

For convenience, we will actually work in terms of the distance $\Delta x = 2R_\oplus\sin\theta/2$ between detectors, rather than the polar angle. All together, our model cross-correlation spectrum under hypothesis $\mathcal{H}_{\mathrm{Free}}$ is

$$C_{\mathrm{Free}}(\Omega_0, \alpha, \Delta x, \phi_1, \phi_2; f) = \gamma(\Delta x, \phi_1, \phi_2; f)\,\Omega_0\,(f/25\,\mathrm{Hz})^\alpha, \tag{7.10}$$

where $\gamma(\Delta x, \phi_1, \phi_2; f)$ is the overlap reduction function corresponding to detectors separated by Δx with local rotation angles ϕ_1 and ϕ_2.

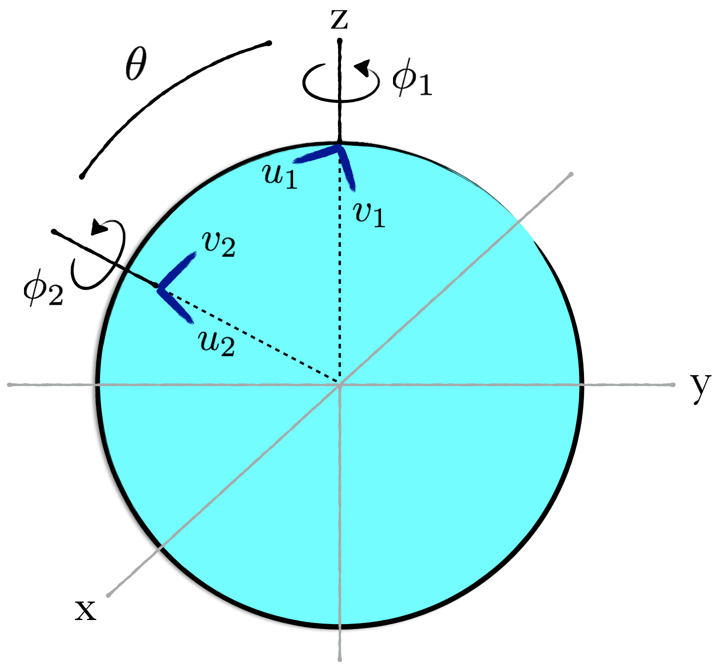

Figure 7.4: Parametrized geometry of an arbitrary detector baseline on the Earth's surface. We initially choose coordinates such that the detectors lie in the $x - z$ plane, with one detector at the pole. The remaining degrees of freedom are the polar angle θ between detectors, and the rotation angles ϕ_1 and ϕ_2 specifying the orientation of each detector.

We set a log-uniform prior on Ω_0 between $(10^{-12}, 10^{-6})$, extending well above and well below Advanced LIGO's sensitivity, and uniform priors on ϕ_1 and ϕ_2 on $(0, 2\pi)$. Similarly, we use a uniform prior on $\cos\theta$ between $(-1, 1)$, corresponding to a prior $p(\Delta x) \propto \Delta x$ on the distance between detectors. We adopt a triangular prior on the background's spectral index: $p(\alpha) \propto 1 - |\alpha|/\alpha_{\text{Max}}$, with $\alpha_{\text{Max}} = 6$. This prior penalizes very steeply sloped backgrounds, while still accommodating backgrounds that are much steeper than those predicted from known sources. Finally, we will choose equal prior odds between both hypotheses.

7.4 A Demonstration

To explore our ability to differentiate terrestrial correlation from an astrophysical background, we will simulate Advanced LIGO measurements of three different sources of correlation: an isotropic stochastic background, a correlated frequency comb, and magnetic Schumann resonances. These latter two sources are terrestrial, and hence should disfavor $\mathcal{H}_\gamma$ over $\mathcal{H}_{\text{Free}}$. We should

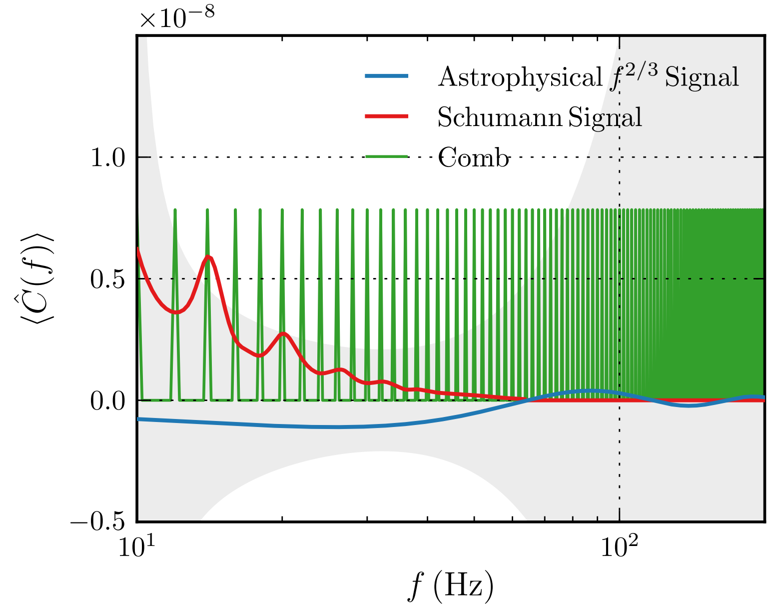

Figure 7.5: Mean cross-correlation spectra used to simulate stochastic search measurements with the Advanced LIGO Hanford and Livingston detectors. We consider an isotropic astrophysical stochastic background, with energy density $\Omega(f) \propto f^{2/3}$ [blue; Eq. (7.11)]. We additionally consider two sources of terrestrial, non-astrophysical correlation: a signal due to magnetic Schumann resonances [red; Eq. (7.13)] and a correlated frequency comb with $\Delta f = 2$ Hz spacing [green; Eq. (7.12)]. The amplitudes of the spectra have been scaled such that each is expected to be detected with SNR $= 10$ after three years of observation with design-sensitivity Advanced LIGO. For comparison, the grey band illustrates the $\pm 1\sigma$ uncertainties of a cross-correlation search after three years of integration.

point out that there exist dedicated strategies for identifying and mitigating combs and Schumann resonances [81, 243]. Here, we use combs and Schumann resonances simply as proxies for any as-of-yet *unknown* sources of terrestrial correlation that could contaminate stochastic search efforts and lead to a false detection claim.

We adopt the following models for the cross-correlation spectra expected in each test case:

1. Isotropic stochastic gravitational-wave background: We assume that the stochastic gravitational-wave background is well described by a power law with spectral index $\alpha = 2/3$, as predicted for compact binary mergers. The corresponding expected cross-correlation spectrum is

$$\langle C(f) \rangle_{\mathrm{Stoch}} = \gamma_{\mathrm{LIGO}}(f)\, \Omega_0 \left(\frac{f}{25\,\mathrm{Hz}} \right)^{2/3}, \tag{7.11}$$

where $\gamma_{\mathrm{LIGO}}(f)$ is the overlap reduction function for the Hanford-Livingston baseline.

2. Frequency comb: We consider a correlated comb of uniformly spaced lines, separated in frequency by Δf and with heights set by C_0:

$$\langle C(f)\rangle_{\rm Comb} = C_0 \Delta f \sum_{n=0}^{\infty} \delta(f - n\Delta f). \tag{7.12}$$

Note that the leading factor of Δf in Eq. (7.12) ensures that C_0 is dimensionless. In the examples below, we use a comb spacing of $\Delta f = 2$ Hz.

3. Magnetic Schumann resonances: Given an environmental magnetic field $\tilde{m}(f)$, the strain induced in a gravitational-wave detector is $\tilde{s}(f) = T(f)\tilde{m}(f)$, where $T(f)$ is a transfer function with units [strain/Tesla]. If there exists a correlated magnetic power spectrum $M(f) = \langle \tilde{m}_1^*(f)\tilde{m}_2(f)\rangle$ between the sites of two gravitational-wave detectors, then from Eq. (7.1) the resulting strain correlation will be of the form $\hat{C}(f) \propto f^3 |T(f)|^2 \,\mathrm{Re}\, M(f)$. We take $M(f)$ to be the median Schumann auto-power spectrum measured at the Hylaty station in Poland, as reported in Ref. [255]. This may not exactly match the magnetic cross-power spectrum between Hanford and Livingston. Most notably, we take $\mathrm{Re}\, M(f)$ to be everywhere positive, as the (potentially frequency-dependent) sign of the Schumann cross-power between the LIGO detectors is not well known. Nevertheless, this model captures the qualitative features expected of a Schumann signal. The magnetic transfer functions for the LIGO detectors are expected to be power laws, but their spectral indices are also not well known; we somewhat arbitrarily choose $T(f) \propto f^{-2}$. Our Schumann signal model is therefore

$$\langle C(f)\rangle_{\rm Schumann} = S_0 \left(\frac{f}{25\,\mathrm{Hz}}\right)^{-1} \frac{\mathrm{Re}\, M(f)}{\mathrm{Re}\, M(25\,\mathrm{Hz})}, \tag{7.13}$$

normalized so that S_0 is the cross-correlation measured at the reference frequency 25 Hz.

The mean cross-correlation spectra for the astrophysical, Schumann, and comb models are shown in Fig. 7.5. For each source of correlation, we simulate Advanced LIGO measurements of 300 signal and noise realizations, with expected signal-to-noise ratios ranging from 0.1 to 100. To produce each realization, we scale the relevant amplitude parameter (Ω_0, C_0, or S_0) to obtain the desired SNR and add random Gaussian measurement noise $\delta C(f)$ with variance given

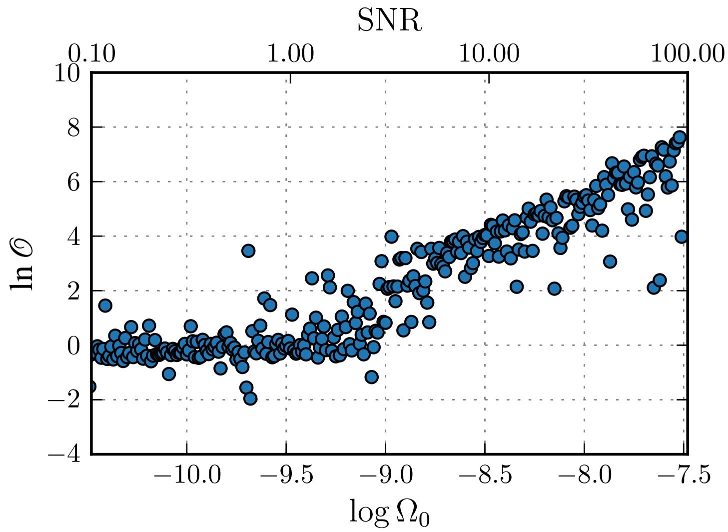

Figure 7.6: Log-odds between the physical and unphysical hypotheses $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ as a function of the amplitude Ω_0 of an injected astrophysical stochastic background [Eq. (7.11)]. To enable a direct comparison between injection types, the upper horizontal axes shows the signal-to-noise ratios of these injections. The recovered values of $\ln \mathcal{O}$ increase linearly with the strength of the astrophysical injections, indicating consistency with the correct (known) detector geometry.

by Eq. (7.3). For each simulated measurement, we then compute an odds ratio between $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ to determine whether the data physically favors the correct detector geometry, or unphysically favors some alternate geometry. We compute Bayesian evidences using MultiNest [232, 233], an implementation of the nested sampling algorithm [235, 236]. We make use of PyMultiNest, which provides a Python interface to MultiNest [231].

The resulting odds plotted in Figs. 7.6-7.8 as a function of injected signal amplitude. As physically distinct parameters, the power-law, Schumann, and comb amplitudes should not be directly compared to one another. Instead, we show the injections' expected signal-to-noise ratios (which *can* be directly compared) on the upper horizontal axes. To compute these SNRs, we assume recovery with a power-law model of slope $\alpha = 2/3$. Thus the SNRs of the power-law injections are optimal. Although SNRs for the comb and Schumann injections are *not* optimal (as the recovery model and injections are not identical), they do represent the signal-to-noise ratios at which such signals would contaminate searches for the stochastic background.

At signal-to-noise ratios SNR $\lesssim 1$, the log-odds for all three sources of corre-

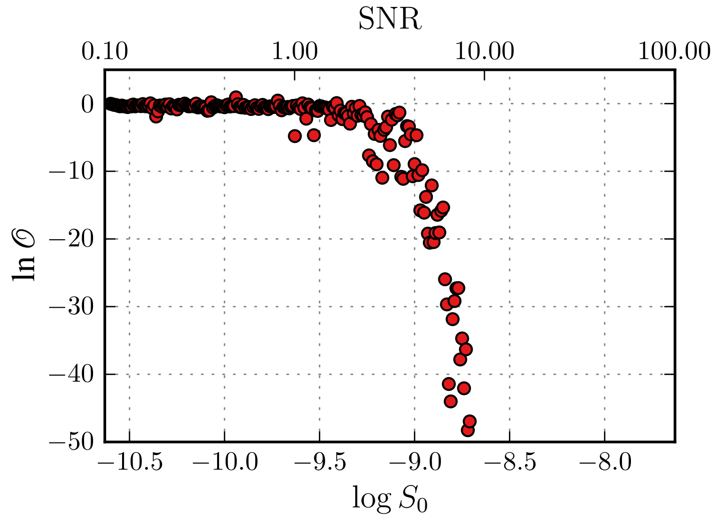

Figure 7.7: As in Fig. 7.6, but for injections of magnetic Schumann resonances [Eq. (7.13)]. The recovered log-odds decrease exponentially with the strength of the Schumann signal, *disfavoring* the correct geometry. Thus $\ln\mathcal{O}$ therefore successfully discriminates between an astrophysical gravitational-wave signal and terrestrial Schumann contamination.

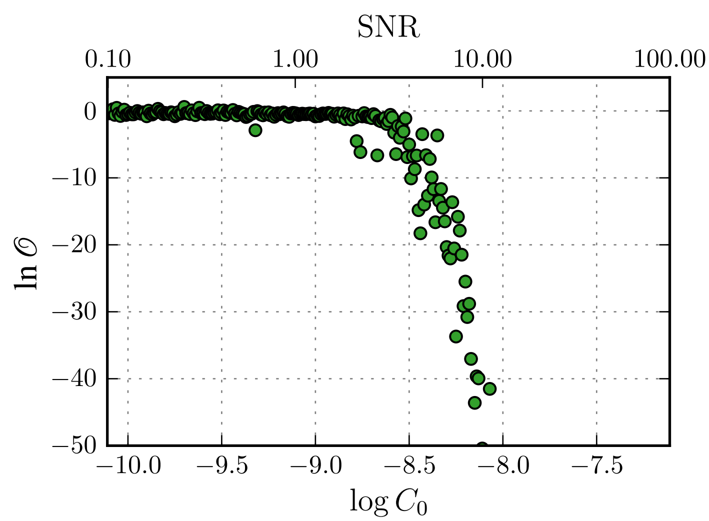

Figure 7.8: As in Fig. 7.6, but for injections of a correlated frequency comb [Eq. (7.12)]. The recovered log-odds again decrease exponentially, correctly rejecting the comb as terrestrial.

lation cluster near $\ln \mathcal{O} \sim 0$. For an astrophysical signal above SNR ~ 1, $\ln \mathcal{O}$ becomes positive, growing approximately linearly with $\log \Omega_0$. In contrast, $\ln \mathcal{O}$ falls exponentially to large negative values as we increase the amplitude of Schumann and comb injections.

It is instructive to look at parameter estimation results for specific astrophysical, comb, and Schumann injections. Figures 7.10-7.12 show posteriors on the parameters of $\mathcal{H}_{\text{Free}}$ given simulated observations of an astrophysical background, a frequency comb, and a Schumann signal, each injected with SNR $= 10$. Figure 7.9 shows these three injections as well as the posteriors obtained on each cross-correlation spectrum. With the five free parameters afforded by $\mathcal{H}_{\text{Free}}$, we succeed in reasonably fitting each of the three spectra. Note that, although we appear to poorly recover the correlated comb injection, the posterior on $C(f)$ closely matches the constant *frequency-averaged* correlation spectrum.

Although the gravitational-wave background, comb, and Schumann injections are all reasonably well fit under $\mathcal{H}_{\text{Free}}$, they yield very different posteriors on Advanced LIGO's baseline length Δx and detector orientation angles ϕ_1 and ϕ_2. Figure 7.10 shows the parameter posteriors given by the simulated astrophysical gravitational-wave background. The diagonal subplots show marginalized one-dimensional posteriors on each parameter, while the central subplots show joint posteriors between each pair of parameters. The solid black lines indicate true parameter values and dashed curves show the priors placed on each parameter. We recover posteriors consistent with the amplitude and spectral index of the injected stochastic signal. More importantly, though, we also obtain well-behaved posteriors on Advanced LIGO's geometry, with a distance posterior[2] consistent with the true separation between detectors. Interestingly, although neither ϕ_1 nor ϕ_2 are well constrained, their *difference* is well measured. This can be seen in the joint posterior between both angles, which strongly supports diagonal bands of constant $\phi_1 - \phi_2$, including the true rotation angles of Hanford and Livingston. We therefore have strong support for the correct detector geometry, in this case yielding $\ln \mathcal{O} = 3.6$ ($\mathcal{O} = 36.6$) in favor of $\mathcal{H}_\gamma$.

Figure 7.11, meanwhile, shows parameter estimation results obtained for the comb injection. As seen in Fig. 7.9, we have enough freedom to fit the (average)

[2]This is the same posterior shown in Fig. 7.3

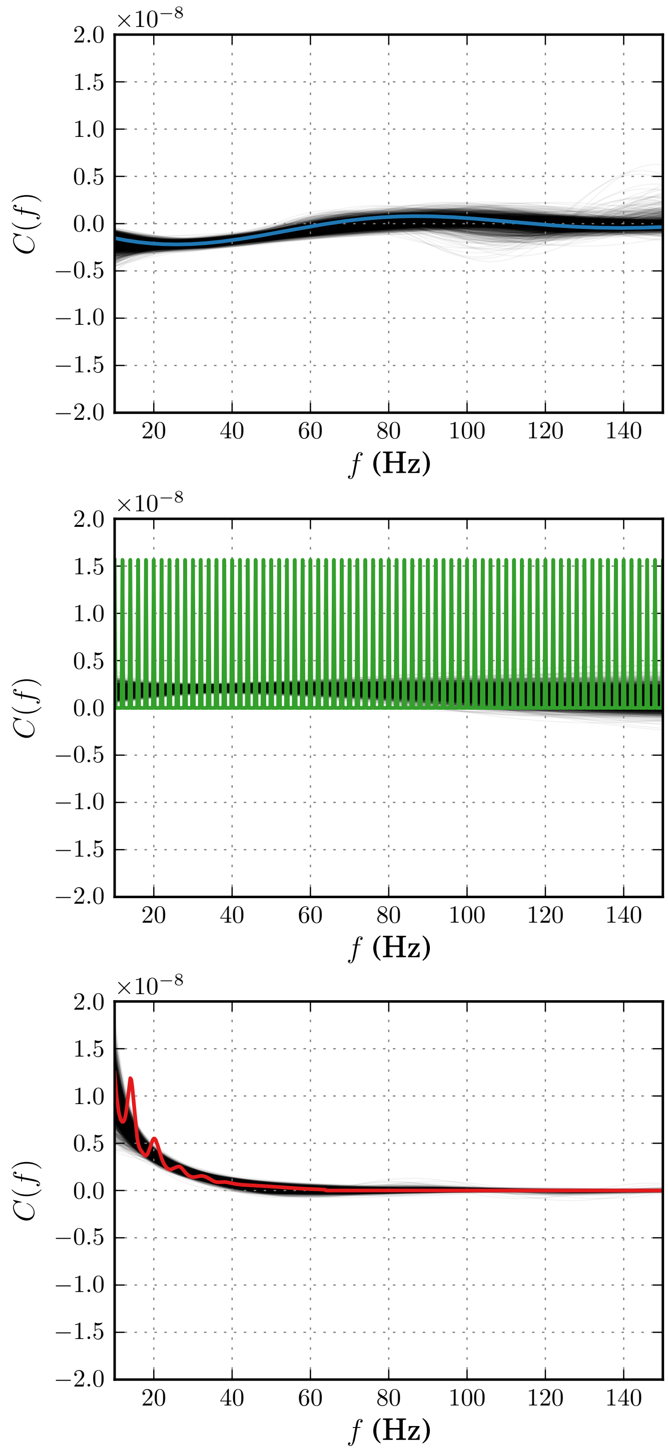

Figure 7.9: Reconstructed cross-correlation spectra using simulated Advanced LIGO observations of an isotropic gravitational-wave background (*top*), a correlated frequency comb (*middle*), and Schumann resonances (*bottom*). The blue, green, and red curves show the injected gravitational-wave, comb, and Schumann spectra, respectively, while the shaded bands indicate the $\pm 1\sigma$ uncertainty region on the simulated measurements. The collections of grey curves show the resulting posteriors on the injected cross-correlation spectrum under the unphysical $\mathcal{H}_{\mathrm{Free}}$ hypothesis in which baseline geometry is allowed to vary.

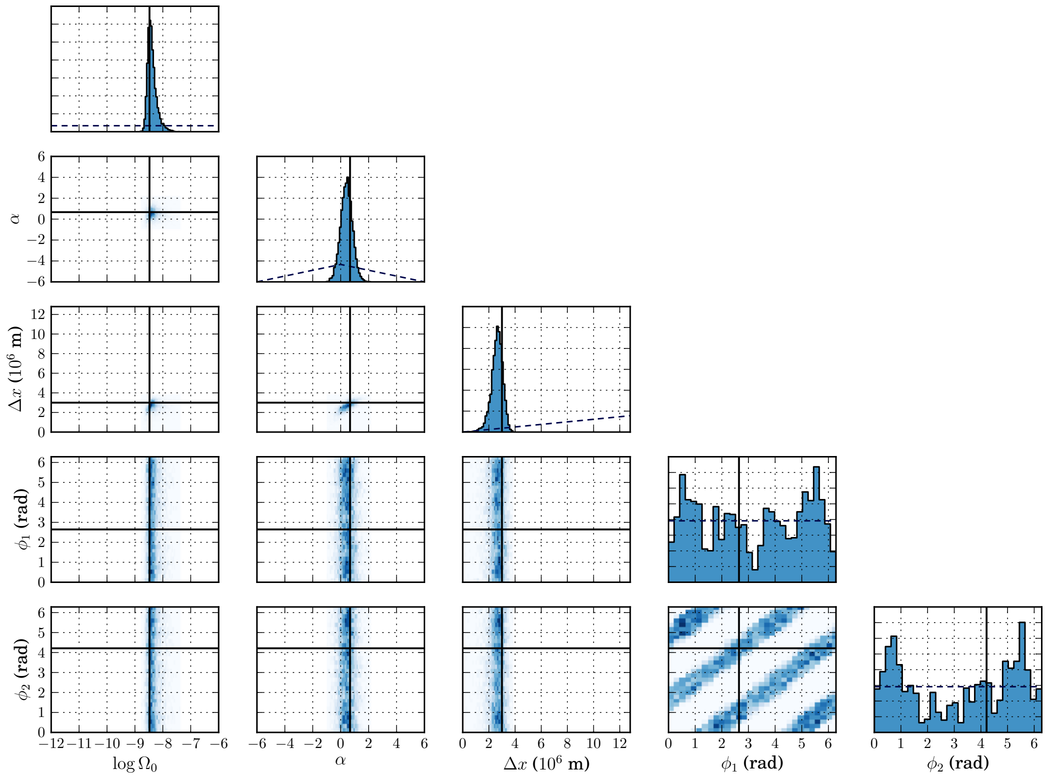

Figure 7.10: Posterior on the stochastic background amplitude Ω_0 and spectral index α, as well the separation Δx and rotation angles ϕ_1 and ϕ_2 of the Advanced LIGO detectors, given a simulated three-year observation of an isotropic astrophysical background with an expected signal-to-noise ratio of 10. Dashed lines in the one-dimensional marginalized posteriors show the prior adopted for each parameter, while solid black lines mark the injected background parameters and the true Advanced LIGO geometry. In addition to recovering the amplitude and spectral index of the injected stochastic signal, we obtain posteriors consistent with the true separation and rotation angles of the Advanced LIGO detectors.

cross-correlation spectrum, yielding reasonably peaked posteriors in Fig. 7.11. However, the posteriors on detector separation and orientation are unphysical, excluding the known Hanford-Livingston geometry. We therefore obtain $\ln \mathcal{O} = -58.5$ ($\mathcal{O} = 3.9 \times 10^{-26}$). Similarly, Fig. 7.12 gives parameter estimation results for the Schumann injection. Interestingly, the distance posterior for this injection *is* consistent with the true Hanford-Livingston separation. The rotation angle posteriors, though, exclude the true detector orientations, yielding $\ln \mathcal{O} = -62.7$ ($\mathcal{O} = 5.9 \times 10^{-28}$).

The observed dependence of $\mathcal{O}$ on injected astrophysical and terrestrial signal strengths can be analytically understood, at least approximately, using the Laplace approximation. As discussed in Ch. 5.6.1, the Laplace approximation involves two assumptions. First, our prior $p(\Theta|\mathcal{H})$ on the parameters of hy-

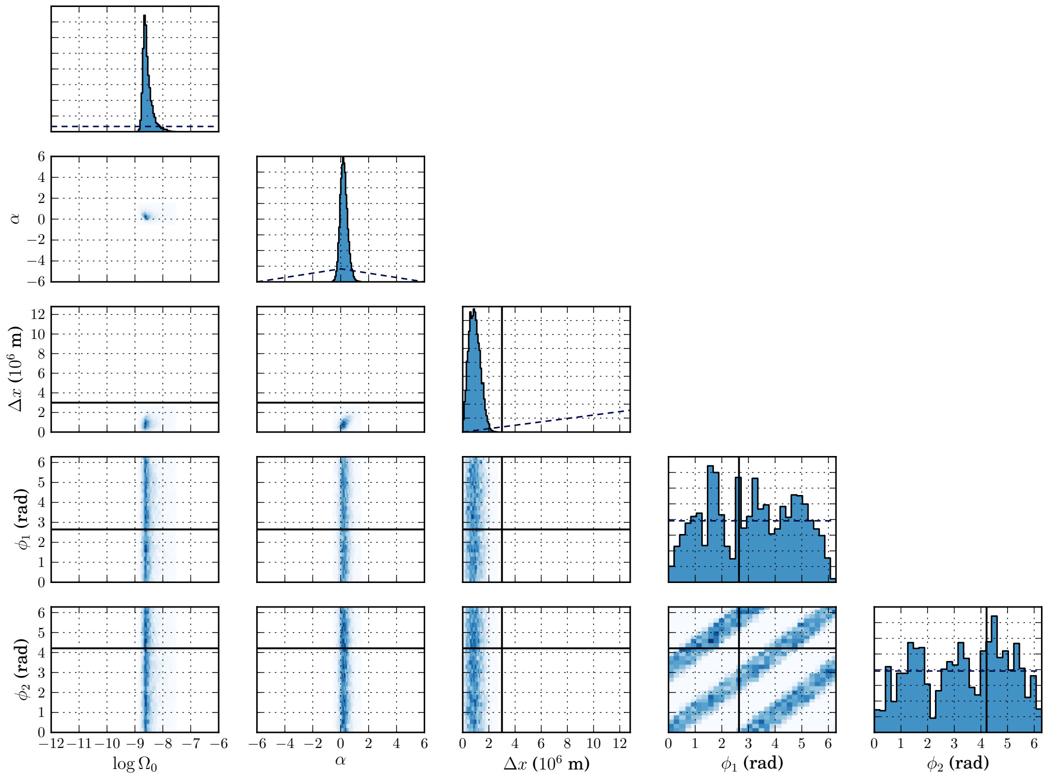

Figure 7.11: As in Fig. 7.10 above, but for a simulated measurement of a correlated frequency comb with SNR = 10. The correlated comb is not well fit by the Advanced LIGO overlap reduction function, and so our recovered posteriors on Hanford and Livingston's separation and rotation angles are inconsistent with their known values (solid black lines).

pothesis $\mathcal{H}$ is assumed flat over a range $\Delta\Theta$, so that $p(\Theta|\mathcal{H}) = 1/\Delta\Theta$. Second, the likelihood $p(\hat{C}|\Theta, \mathcal{H})$ is assumed to be strongly peaked about maximum-likelihood parameter values $\overline{\Theta}$ and a peak value $\overline{\mathcal{L}}$. The width of the peak is $\delta\Theta$ (see Fig. 5.14). Under these assumptions, a Bayesian evidence may be approximated as

$$\begin{aligned} p(\hat{C}|\mathcal{H}) &= \int p(\hat{C}|\Theta, \mathcal{H}) p(\Theta|\mathcal{H}) d\Theta \\ &\approx \frac{\delta\Theta}{\Delta\Theta}\overline{\mathcal{L}}. \end{aligned} \tag{7.14}$$

Given two hypotheses $\mathcal{H}_A$ and $\mathcal{H}_B$, the odds ratio between them becomes

$$\begin{aligned} \mathcal{O}^A_B &= \frac{p(\hat{C}|\mathcal{H}_A)}{p(\hat{C}|\mathcal{H}_B)} \frac{p(\mathcal{H}_A)}{p(\mathcal{H}_B)} \\ &\approx \frac{\delta\Theta_A/\Delta\Theta_A}{\delta\Theta_B/\Delta\Theta_B} \frac{\overline{\mathcal{L}}_A}{\overline{\mathcal{L}}_B} \frac{p(\mathcal{H}_A)}{p(\mathcal{H}_B)}. \end{aligned} \tag{7.15}$$

The ratio $\overline{\mathcal{L}}_A/\overline{\mathcal{L}}_B$ is the standard maximum likelihood ratio between $\mathcal{H}_A$ and $\mathcal{H}_B$, while $p(\mathcal{H}_A)/p(\mathcal{H}_B)$ is our prior odds between hypotheses. The leading

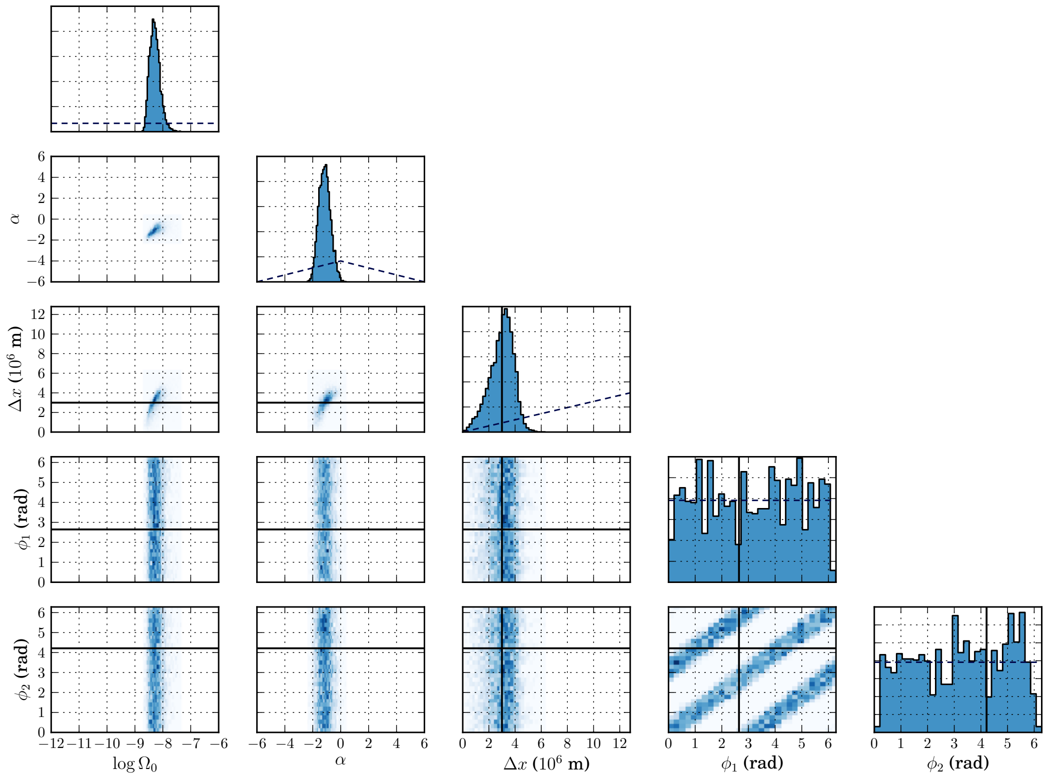

Figure 7.12: As in Figs. 7.10 and 7.11 above, for a simulated observation of a correlated Schumann signal with SNR = 10. While the posterior does encompass the correct Hanford-Livingston separation, it is incompatible with the detectors' true rotation angles.

term, known as the "Occam's factor," penalizes the more complex hypothesis with the larger available parameter space.

Using the Laplace approximation, our odds ratio between hypotheses $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ may be written

$$\begin{aligned}\mathcal{O} &= \frac{p(\hat{C}\,|\mathcal{H}_\gamma)}{p(\hat{C}|\mathcal{H}_{\text{Free}})} \\ &\approx \left[\frac{\delta\Omega_0}{\Delta\Omega_0}\frac{\delta\alpha}{\Delta\alpha}\right]_\gamma \left[\frac{\delta\Omega_0}{\Delta\Omega_0}\frac{\delta\alpha}{\Delta\alpha}\frac{\delta\theta}{\Delta\theta}\frac{\delta\phi_1}{\Delta\phi_1}\frac{\delta\phi_2}{\Delta\phi_2}\right]^{-1}_{\text{Free}} \\ &\quad\times \frac{\exp\left[-\frac{1}{2}\left(\hat{C}-\overline{C}_\gamma|\hat{C}-\overline{C}_\gamma\right)\right]}{\exp\left[-\frac{1}{2}\left(\hat{C}-\overline{C}_{\text{Free}}|\hat{C}-\overline{C}_{\text{Free}}\right)\right]},\end{aligned} \tag{7.16}$$

where $\overline{C}_\gamma$, for instance, is the maximum-likelihood fit to the data under the $\mathcal{H}_\gamma$ hypothesis and we have set our prior odds to unity.

First, consider the case of an isotropic astrophysical background of amplitude Ω_0. In this case, both hypotheses $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ can successfully fit the resulting cross-correlation spectrum. Then our residuals are $\hat{C}-\overline{C}_\gamma \approx \hat{C}-\overline{C}_{\text{Free}} \approx 0$ and the likelihood ratio in Eq. (7.16) is approximately unity. Because both models can fit the data, posteriors on each parameter (of each hypothesis) are peaked, with fractional widths (e.g. $\delta\theta/\Delta\theta$) that scale as $\text{SNR}^{-1} \propto \Omega_0^{-1}$. Then, in the presence of an astrophysical background, we might expect Eq. (7.16) to scale as $\mathcal{O} \propto \Omega_0^3$, or

$$\ln \mathcal{O} \sim 3 \log \Omega_0, \tag{7.17}$$

up to additive constants. This linear slope is, in fact, a reasonably good approximation to the linear trend in Fig. 7.6.

Next, consider a correlated signal of terrestrial origin, characterized by some amplitude C_0. In this case, it is likely that $\mathcal{H}_\gamma$ is unable to accommodate the measured correlations, but that $\mathcal{H}_{\text{Free}}$, with a greater number of free parameters, *can* successfully fit the data to some extent. Then $\hat{C} - \overline{C}_{\text{Free}} \approx 0$ but $\hat{C} - \overline{C}_\gamma \neq 0$. So the resulting likelihood ratio in Eq. (7.16) is not constant, but will depend exponentially on C_0. Ignoring the leading Occam's factors (which can scale at most as a power law in C_0), our odds ratio becomes

$$\begin{aligned} \mathcal{O} &\propto \exp\left[-\frac{1}{2}\left(\hat{C}-\overline{C}_\gamma|\hat{C}-\overline{C}_\gamma\right)\right] \\ &\propto \exp\left[-\frac{1}{2}\left(\hat{C}|\hat{C}\right)+\left(\hat{C}|\overline{C}_\gamma\right)-\frac{1}{2}\left(\overline{C}_\gamma|\overline{C}_\gamma\right)\right], \end{aligned} \tag{7.18}$$

giving

$$\ln \mathcal{O} \sim -\frac{1}{2}\left(\hat{C}|\hat{C}\right)+\left(\hat{C}|\overline{C}_\gamma\right)-\frac{1}{2}\left(\overline{C}_\gamma|\overline{C}_\gamma\right). \tag{7.19}$$

The maximum likelihood value of Ω_0 [the amplitude of our model spectrum $C_\gamma(f)$] is given by [226]

$$\overline{\Omega}_0 = \frac{\left(f^{2/3}|\hat{C}\right)}{\left(f^{2/3}|f^{2/3}\right)}. \tag{7.20}$$

Although this does scale proportionally with C_0, the inner product $(f^{2/3}|\hat{C})$ may be small if the measured correlation spectrum has a very different shape from an astrophysical power law. In this case the cross term $(\hat{C}|\overline{C}_\gamma)$ in Eq. (7.19) may be neglected, and we recover $\ln \mathcal{O} \propto -C_0^2$, or

$$\ln \mathcal{O} \propto -10^{2 \log C_0}. \tag{7.21}$$

This again provides a reasonably good match to the scaling seen in Figs. 7.7 and 7.8.

We've now seen that the geodesy technique can discriminate between an astrophysical stochastic background and spurious, terrestrial sources of correlation. One possibility that we've ignored, however, is that of false positives: non-astrophysical correlation spectra that, purely by chance, yield posteriors consistent with Advanced LIGO's geometry. To carefully calculate the probability of a false positive at a particular value of $\mathcal{O}$, one could analyze a set of random cross-correlation spectra (e.g. drawn from the space of spectra supported by $\mathcal{H}_{\text{Free}}$) and construct a null distribution of the resulting Bayes factors. Alternatively, we can quickly estimate the probability of false positives at a given $\ln \mathcal{O}$ using the results in Fig. 7.6. At SNR $= 10$, injected isotropic signals yield $\ln \mathcal{O} \approx 4$, indicating $e^4 : 1$ odds that these signals are drawn from $\mathcal{H}_\gamma$ rather than $\mathcal{H}_{\text{Free}}$. Taken at face value, this implies that we would need to simulate $e^4 + 1 \approx 56$ random correlation spectra from $\mathcal{H}_{\text{Free}}$ with SNR $= 10$ before finding one that yields $\ln \mathcal{B} \gtrsim 4$ by chance. Interpreted in this way, the geodesy formalism not only offers a means of rejecting non-astrophysical correlations, but can bolster the statistical significance of a real stochastic signal.

7.5 Complications

Our analysis in this chapter has relied on several important assumptions. Most notably, we have assumed that our model energy-density spectrum (a power law) is a good descriptor of the true stochastic background. This assumption was guaranteed by design, as our injected stochastic energy-density spectrum *was* a power law. While most gravitational-wave sources are predicted to yield power-law energy-density spectra in the Advanced LIGO and Virgo band, there do exist speculative sources like superradiant axion clouds [261–263] that may instead yield more complex spectra.

It is worthwhile to investigate what might happen if we mistakenly adopt an *incorrect* model – one that is a poor descriptor of the background's energy-density spectrum. In this case, would we risk rejecting a real stochastic background as a terrestrial signal? To test this, we simulate observations of a broken-power law background with energy density

$$\Omega(f) = \begin{cases} \Omega_0 \left(f/f_0\right)^{\alpha_1} & f \leq f_0 \\ \Omega_0 \left(f/f_0\right)^{\alpha_2} & f > f_0 \end{cases}. \tag{7.22}$$

We will analyze these observations, meanwhile, using an ordinary power-law

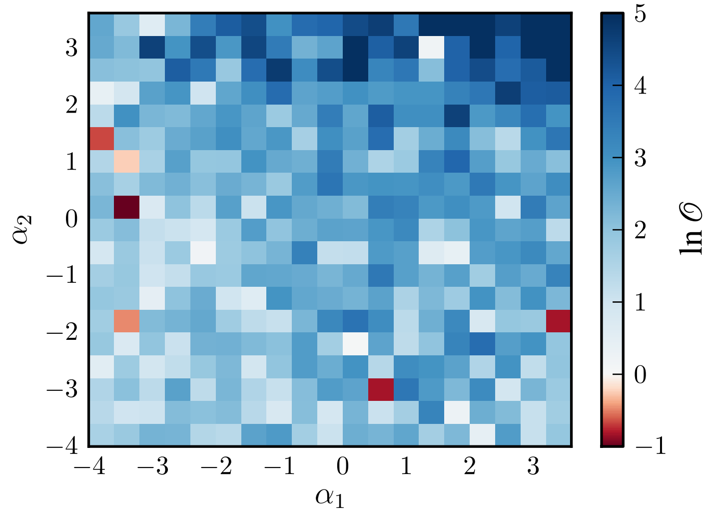

Figure 7.13: Log-odds between $\mathcal{H}_\gamma$ and $\mathcal{H}_{\text{Free}}$ when deliberately analyzing astrophysical broken power-law signals with an *incorrect* power-law model. Each injected signal has a knee frequency of $f_0 = 30$ Hz and an amplitude Ω_0 scaled such that the signal has SNR $= 5$ after three years of observation with design-sensitivity Advanced LIGO. Despite the signal-model mismatch, we correctly classify the majority of the simulated signals, with no evidence of increased false-dismissals due to the mismatch.

model, deliberately choosing an incorrect description of the injected stochastic signal.

Figure 7.13 illustrates the resulting odds ratios for simulated observations with α_1 and α_2 each ranging between -4 and 4. For each injection we chose $f_0 = 30$ Hz, placing the broken power-law's "knee" in the center of the stochastic sensitivity band, and scaled the amplitudes Ω_0 such that each observation has SNR $= 5$ when naively recovered with an ordinary power law. The vast majority of these simulations yield positive log-odds factors, correctly classifying these signals despite our poor choice of model. Note that the injections falling along the line $\alpha_1 = \alpha_2$ *are* power laws. If the signal-model mismatch significantly degraded our ability to classify stochastic signals, then Fig. 7.13 would exhibit a color gradient as we move perpendicularly off the $\alpha_1 = \alpha_2$ line, away from power laws and toward increasingly sharp signal spectra. Instead, Fig. 7.13 shows no such gradient, and our method remains robust even in the case of poorly fitting models.

We attribute this robustness to the fact that the isotropic energy-density spectrum and baseline geometry have very different effects on the expected cross-correlation spectrum $\langle C(f) \rangle = \gamma(f)\Omega(f)$. The energy density spectrum $\Omega(f)$ is everywhere positive, and so different energy-density spectra can change only the *amplitude* of $C(f)$, not its sign. The sign of $C(f)$ is set by the overlap reduction function, which alternates between positive and negative values with zero-crossings fixed by the baseline geometry. Even if our model for $C(f)$ assumes an incorrect energy-density spectrum (as above), our $\mathcal{H}_\gamma$ hypothesis nevertheless predicts the correct zero-crossings of the observed cross-correlation spectrum. This offers some robustness against false-dismissal of a true stochastic signal, even if our model energy-density spectrum is imperfect. At the same time, it prevents us from *over-fitting* spurious terrestrial correlations [whose sign is unrelated to the sign of $\gamma(f)$], mitigating the risk of false positives.

We made a number of other assumptions about the character of the gravitational-wave background – that it is isotropic, unpolarized, and free of the non-standard "vector" and "scalar" gravitational-wave polarizations predicted by modified theories of gravity. These assumptions are unlikely to all be strictly true. The stochastic background may be polarized by a variety of early universe effects [85], as well as the scattering of gravitational waves by massive objects during propagation [264]. Meanwhile, the solar system's motion with respect to the cosmic microwave background will likely impart a small apparent dipole moment to the stochastic gravitational-wave background. Additional anisotropies might arise from structure in the local universe [96, 265], together with the fact that, over a finite integration time, we observe only a discrete set of gravitational-wave events [100].

A stochastic background containing anisotropies, polarization asymmetries, or non-standard polarizations would yield correlations that are *not* consistent with the standard overlap reduction function, but that instead obey some different effective overlap reduction function. If we naively analyzed such a signal with the method presented in the main text, we would likely find a preference for the (unphysical) hypothesis $\mathcal{H}_{\mathrm{Free}}$ over $\mathcal{H}_\gamma$ and risk rejecting the signal as terrestrial.

In practice, deviations from our ideal stochastic background model are expected to be small, and so these complicating factors are unlikely to significantly affect our analysis. For example, the solar system moves with speed $v_\oplus \approx 370$ km/s with respect to the cosmic microwave background [91], and so the stochastic background's apparent dipole moment is expected to be a factor of $v_\oplus/c \sim 10^{-3}$ weaker than the isotropic monopole moment. The intrinsic anisotropy and polarization of the astrophysical background are also predicted to be small. Considering multipole moments up to $l = 20$ (the approximate angular resolution limit of the LIGO Hanford-Livingston baseline [238]), the observed energy density is expected to vary by no more than $\sim 10\%$ with direction [96, 265]. Any net polarization arising from scattering is predicted to be further suppressed by many orders of magnitude in the frequency band of Advanced LIGO and Virgo [264].

If any of these complications were a significant concern, however, the formalism of Ch. 7.3 can be straightforwardly extended to accommodate these effects. As an example, here we demonstrate how to extend our formalism to the case of an anisotropic stochastic background.

When allowing for anisotropy, the observed energy-density of the stochastic background will generically have directional dependence on our viewing angle $\hat{\mathbf{n}}$. It is generally assumed that an anisotropic energy-density spectrum can be factored via $\Omega(\hat{\mathbf{n}}, f) = \mathcal{F}(f)\mathcal{P}(\hat{\mathbf{n}})$, where $\mathcal{F}(f)$ and $\mathcal{P}(\hat{\mathbf{n}})$ encode the frequency and directional dependence of $\Omega(\hat{\mathbf{n}}, f)$, respectively. We can further decompose $\mathcal{P}(\hat{\mathbf{n}})$ into a sum of spherical harmonics $Y_{lm}(\hat{\mathbf{n}})$, giving [237–239]

$$\Omega(\hat{\mathbf{n}}, f) = \mathcal{F}(f) \sum_{l,m} \mathcal{P}^{lm} Y_{lm}(\hat{\mathbf{n}}) \tag{7.23}$$

for some set of coefficients $\mathcal{P}^{lm}$. We use a normalization convention in which $\int |Y_{lm}(\hat{\mathbf{n}})|^2 \mathrm{d}\hat{\mathbf{n}} = 1$.

Over the course of a sidereal day, gravitational-wave detectors have varying sensitivities to different sky directions $\hat{\mathbf{n}}$. In the presence of an anisotropic background, the expected cross-correlation between detectors is therefore time-dependent:

$$\langle C(f,t)\rangle = \mathcal{F}(f) \sum_{l,m} \mathcal{P}^{lm} \gamma_{lm}(t, f), \tag{7.24}$$

where t is periodic over a sidereal day. This expression is similar in form to our standard expectation in the presence of an isotropic signal, but contains a sum over spherical harmonics and distinct (time-dependent) overlap reduction functions for each spherical harmonic [237, 238]:

$$\gamma_{lm}(t, f) = \frac{5}{2\sqrt{4\pi}} \sum_A \int_{\mathrm{Sky}} Y_{lm}(\hat{\mathbf{n}}) F_1^A(\hat{\mathbf{n}}, t) F_2^A(\hat{\mathbf{n}}, t) \mathrm{e}^{2\pi i f \Delta\mathbf{x}(t)\cdot\hat{\mathbf{n}}/c} \mathrm{d}\hat{\mathbf{n}}. \tag{7.25}$$

In Eq. (7.25), the detectors' antenna patterns $F_i^A(\hat{\mathbf{n}}, t)$ and separation vector $\Delta\mathbf{x}(t)$ are time-dependent, rotating with the Earth over the course of a sidereal day. The normalization of Eq. (7.25) is chosen such that monopole overlap reduction function $\gamma_{00}(t, f)$ reduces to Eq. (7.4) for our standard isotropic overlap reduction function. The time-dependence of Eq. (7.25) can be conveniently factored out via [237, 238]

$$\gamma_{lm}(t, f) = \gamma_{lm}(0, f)\mathrm{e}^{2\pi i m(t/T)}, \tag{7.26}$$

where T is the length of one sidereal day.

If we incorrectly assumed an isotropic background and averaged our cross-correlation measurements over a sidereal day, we would measure cross-correlation

$$\begin{aligned} \langle C(f)\rangle &= \frac{1}{T}\int_0^T \langle C(f,t)\rangle dt \\ &= \mathcal{F}(f)\sum_{l,m}\mathcal{P}^{lm}\gamma_{lm}(0,f)\frac{1}{T}\int_0^T \mathrm{e}^{2\pi i m(t/T)}dt \\ &= \mathcal{F}(f)\sum_{l}\mathcal{P}^{l0}\gamma_{l0}(0,f), \end{aligned} \tag{7.27}$$

where the integral vanishes for all $m \neq 0$. Equation (7.27) does not trace the isotropic overlap reduction function, but instead follows a linear combination of the anisotropic $\gamma_{l0}(f)$'s. Thus, if the background were significantly anisotropic (with some $\mathcal{P}^{l0}$ comparable in magnitude to the monopole amplitude $\mathcal{P}^{00}$), we would incorrectly conclude that the resulting correlated signal is incompatible with our detector geometry and dismiss it as terrestrial.

To safeguard against this possibility, we could extend hypothesis $\mathcal{H}_\gamma$ to include anisotropic terms, adopting a model cross-correlation spectrum of the form

$$C_\gamma(\Theta, \mathcal{P}^{lm}; f) = \mathcal{F}(\Theta; f)\sum_{l,m}\mathcal{P}^{lm}\gamma_{lm}^{\mathrm{True}}(t, f), \tag{7.28}$$

where $\gamma_{lm}^{\mathrm{True}}(t, f)$ is the baseline's known overlap reduction function for spherical harmonic (l, m) and Θ represents the variables parametrizing $\mathcal{F}(f)$. Analogously, the unphysical hypothesis $\mathcal{H}_{\mathrm{Free}}$ would become

$$C_{\mathrm{Free}}(\Theta, \mathcal{P}^{lm}, \theta, \phi_1, \phi_2; f) = \mathcal{F}(\Theta; f)\sum_{l,m}\mathcal{P}^{lm}\gamma_{lm}(\theta, \phi_1, \phi_2; t, f). \tag{7.29}$$

7.6 Discussion

As searches for the stochastic gravitational-wave background grow increasingly sensitive, we may be nearing the first detection of the background. This prospect, though, comes with significant risk, namely the high cost of a false positive detection. To minimize this risk, it will be important to develop methods to validate tentative detections of the gravitational-wave background. Specifically, when assessing any apparent detection, it will be necessary to argue not just that an observed correlation is statistically significant, but that it is *astrophysical* – that it is due to gravitational waves and not some terrestrial process. While well-developed methods exist to quantify the statistical significance of measured correlations, until now no generic method has existed to gauge whether or not a statistically significant cross-correlation is indeed astrophysical.

In this chapter, we explored how gravitational-wave geodesy – the use of the stochastic gravitational-wave background itself to determine the positions and orientations of gravitational-wave detectors – can form the basis for a novel consistency check on an apparent detection of the background. If the measured correlation between detectors truly represents a gravitational-wave signal, then the reconstructed detector orientations and positions must be compatible with their true, known values. Correlations due to any terrestrial source, on the other hand, have no reason to prefer the baseline's true geometry over any other possible arrangement. By demanding that gravitational-wave geodesy yield results consistent with the true baseline geometry, we can discriminate between astrophysical and terrestrial sources of correlation. Used in this fashion, gravitational-wave geodesy provides a second independent measure of detection significance besides a standard signal-to-noise ratio.

7.7 Looking Ahead

This chapter concludes our exploration of the stochastic gravitational-wave background. For the remainder of this thesis, I will shift gears and focus on an altogether different aspect of gravitational-wave astronomy: the search for electromagnetic counterparts to gravitational-wave events.

Chapter 8
Prompt Radio Emission from Compact Binary Mergers

So far we've largely focused on study of the stochastic gravitational-wave background – how we might detect it and what we might learn from it regarding distant compact binaries and even the nature of gravity itself. In this chapter we will switch gears entirely, leaving the stochastic background and instead investigating the electromagnetic signatures that might accompany binary neutron star and black hole mergers.

8.1 Electromagnetic Emission from Compact Binary Mergers

It should not come as a surprise that compact binary mergers yield observable electromagnetic counterparts. In the gravitational-wave driven merger of a binary with total mass M, a considerable fraction of the binary's energy $E \sim Mc^2$ is radiated in time $t \sim GM/c^3$, giving a luminosity roughly of order [266]

$$L \sim \frac{c^5}{G} \sim 10^{59}\,\mathrm{erg\,s^{-1}}. \tag{8.1}$$

For comparison, the brightest supernova ever witnessed had an inferred bolometric luminosity of $10^{45}\,\mathrm{erg\,s^{-1}}$ [267]. Less than *one-trillionth* of the energy lost by a merging compact binary need be converted to electromagnetic energy to yield an overwhelmingly luminous electromagnetic signal. Any mechanism that might extract a fraction of the binary's energy, however small, is therefore a good candidate for powering an electromagnetic counterpart.

The binary neutron star merger GW170817 confirmed these expectations. Alongside its gravitational-wave signal [46], GW170817 was observed in virtually every electromagnetic band. Figure 8.1 presents a timeline of these observations, spanning gamma rays [53, 59], ultraviolet through infrared [60, 61, 268–271], x-rays [272], and, beginning roughly two weeks after the event, radio [62]. Together, these observations yielded an extraordinary amount of information, including measurements of nucleosynthesis in kilonova [63, 273–

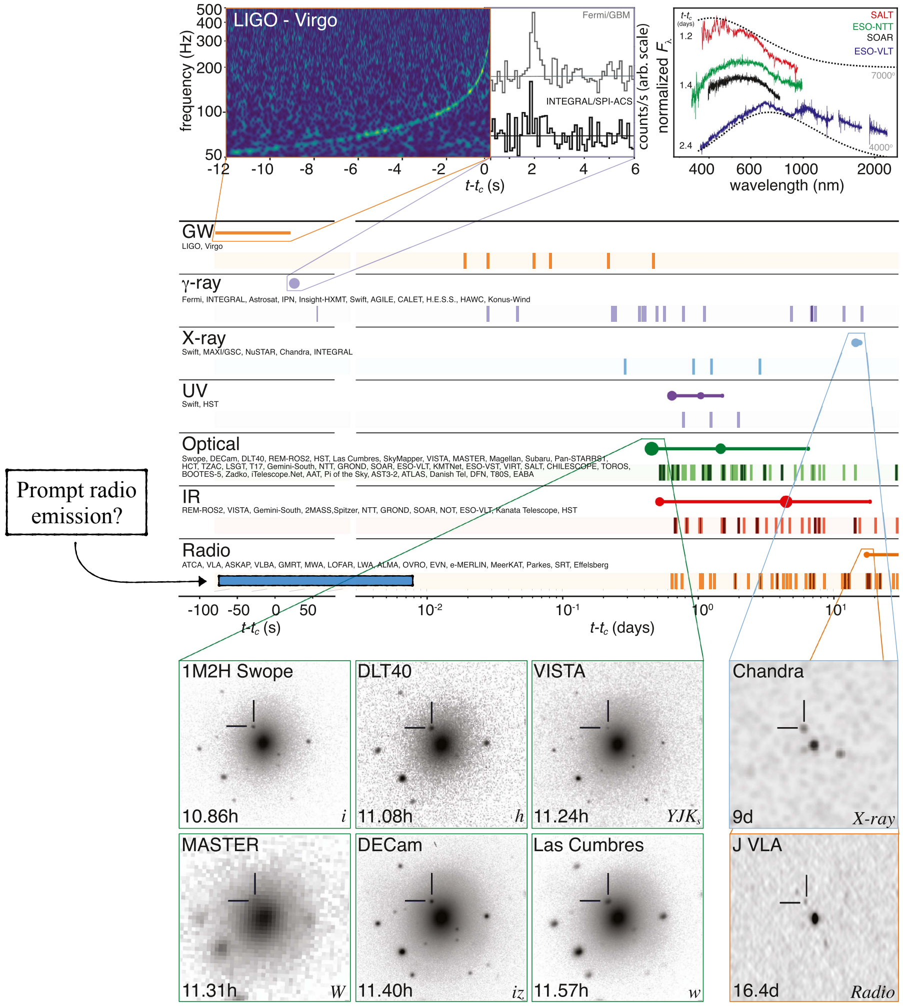

Figure 8.1: Timeline of the multi-messenger observations of the binary neutron star merger GW170817, adapted from Ref. [54]. Beyond the electromagnetic counterparts associated with GW170817, binary neutron stars are additionally predicted to give rise to prompt radio emission radiated in coincidence with the gravitational-wave signal, as shown in the annotation.

275], new insights into gamma-ray burst mechanisms [276, 277], constraints on the neutron star equation of state [49, 278, 279], an independent measurement of the Hubble constant [64, 280], and powerful new constraints on alternative theories of gravity [52, 53, 281–283].

8.2 Prompt Radio Emission

Beyond this already-diverse range of counterparts, binary neutron stars are predicted to yield yet another electromagnetic signature: **prompt radio**

emission. Unlike the late-time radio emission from GW170817, which was due to synchrotron radiation from a relativistic jet, this theorized prompt radio emission is generated at or near the time of merger by a distinct process (or processes) operating in the *immediate* vicinity of the merging neutron stars.

A diverse range of processes have been theorized to generate prompt radio emission. As detailed in the following subsections, possible emission mechanisms can be loosely sorted into three categories based on the time at which they predict prompt radio emission to be generated. Each of these mechanisms is individually quite speculative; it is not at all certain which (if any) are at play in real binary neutron star mergers. Conveniently, though, all have been suggested to give rise to phenomenologically similar signals – sub-second, coherent radio pulses generated simultaneously with the final gravitational-wave "chirp."

8.2.1 Before Merger

A host of processes have been proposed to convert the orbital and/or magnetic field energy of a compact binary into electromagnetic radiation in the final moments just before merger. Figure 8.2 depicts a late-stage binary neutron star system. The leftmost neutron star is assumed to be highly magnetized and slowly spinning, while the rightmost neutron star has a negligible intrinsic magnetic field. Paired high- and low-field neutron stars like this are predicted to arise from binary stellar evolution. The older, firstborn neutron star accretes matter from its (then) stellar companion, increasing its rotation rate while decreasing its magnetic field, much like the recycling of millisecond pulsars. The second, younger neutron star is later born with a comparatively stronger magnetic field and slower rotation rate.

The system's net magnetic field offers several means of energy extraction [266]. The rapid orbit of the primary's dipole field generates a time-varying quadrupole, giving rise to magnetic quadrupole radiation. Meanwhile, the primary induces a dipole magnetic field around the unmagnetized companion; the rotation of this induced dipole with the binary's orbit yields dipole radiation. Lense-Thirring precession can also source dipole emission via the modulation of the primary's magnetic field orientation.

The relativistic motion of the unmagnetized companion through the primary's magnetic field also induces a strong voltage across the companion. This volt-

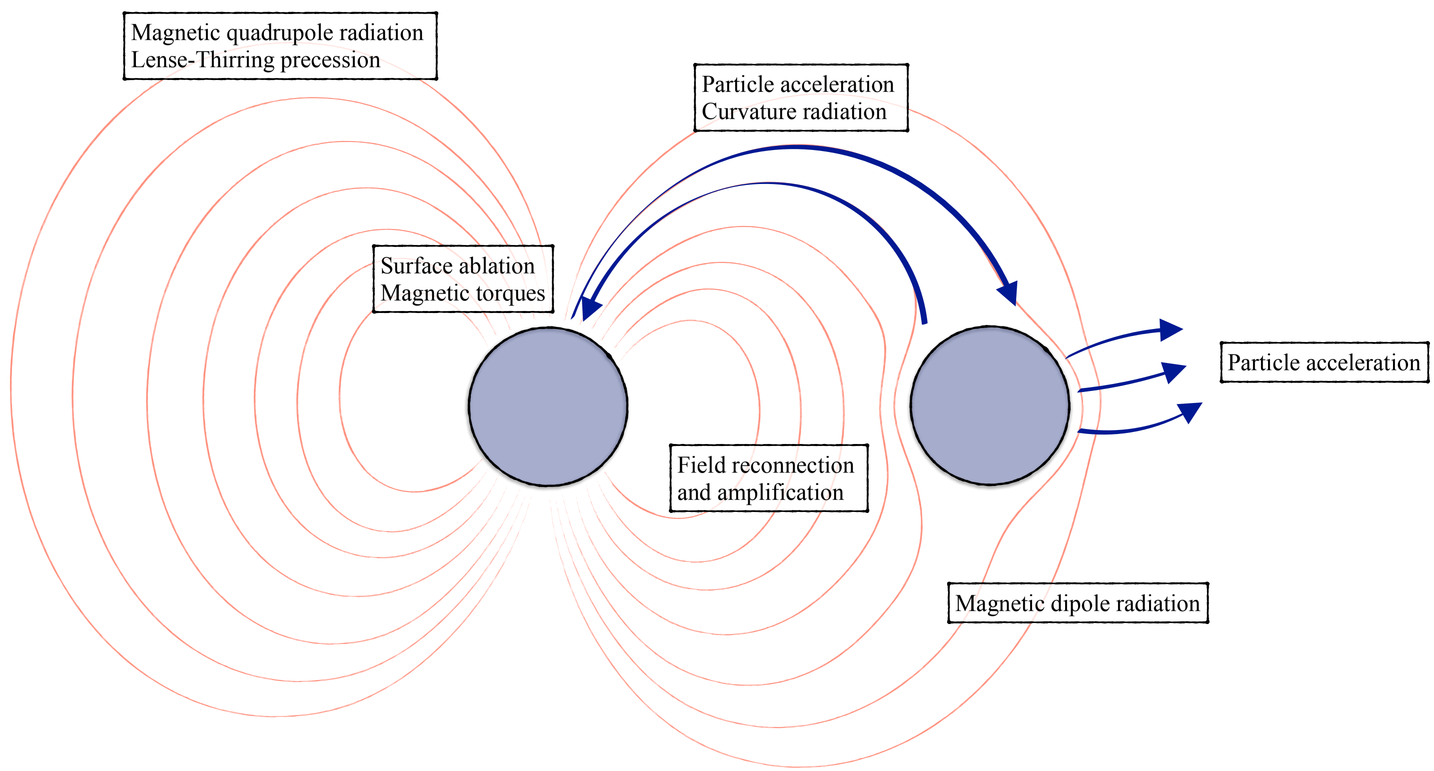

Figure 8.2: Cartoon depiction of a late-time binary neutron star merger. Labeled are various mechanisms that have been theorized to power prompt radio emission. These include magnetic dipole [266, 284] and quadrupole [266] radiation, amplification or reconnection of interacting magnetic fields [285, 286], particle acceleration by induced electromotive forces [286–289], or the establishment of a circuit between neutron stars. The latter, in turn, can yield curvature radiation [290], further particle acceleration within the magnetosphere [291–293], magnetic torques on the neutron stars [291, 292], or even ablation of the neutron stars's surfaces [292].

age might accelerate particles away from the binary, driving a relativistic wind of charged particles [287–289]. Alternatively, it might drive particles inward toward the primary, establishing a current loop – effectively a DC circuit – guided by magnetic field lines [287–294]. This is a so-called "unipolar inductor," the same mechanism that powers radio emission from the Jupiter-Io system [295]. The unipolar inductor mediates several additional channels for orbital energy extraction. Curvature radiation might arise from charged particles moving relativistically along the magnetic field lines [290]. Interaction of the current with the primary's magnetic field may torque the binary and spin up the primary [291, 292]. With the circuit analogy in mind, resistive losses might deposit energy in the magnetosphere, driving additional particle acceleration and eventual synchrotron radiation [291–293]. Resistive losses on the neutron stars themselves, meanwhile, will heat and potentially ablate portions of their outer surfaces [292].

Yet more channels are possible if *both* stars have non-negligible magnetic fields. If the neutron stars' dipole fields are misaligned, magnetic reconnection can occur in the space between stars [285, 286]. Alternatively, aligned dipoles can sequester energy via the amplification of the stars' parallel magnetic fields [285, 286]. Finally, at very late times the neutron stars may become tidally locked, with rotational periods matching their orbital period. Corotation may be accompanied by a burst of energy from reconnecting field lines [286] or by coherent dipole radiation due to the synchronized rotation of both stars' magnetic fields [284].

An altogether different pre-merger energy source is the outgoing gravitational radiation itself. Magnetized plasma is unstable to perturbations from gravitational waves. The intense gravitational radiation generated in the moments before merger can trigger these instabilities in the plasma surrounding the binary, driving the growth of magnetoacoustic and Alfven waves that propagate parallel to the gravitational waves [296–300]. The subsequent interaction of these magnetohydrodynamic waves with relativistic outflowing winds may be sources of detectable synchrotron and/or inverse Compton radiation [298].

8.2.2 During Merger

Prompt emission may alternatively arise *during* the binary's final merger, due to processes associated with possible short-lived merger remnants.

It is generally expected that neutron stars can sustain maximum masses of roughly $2.3\,M_\odot$ [301]. Merger remnants heavier than this maximum will ultimately collapse into black holes. This collapse is not necessarily instant, though. Rotating "supramassive" neutron star remnants with $\sim 2.5\,M_\odot$ masses can centrifugally support themselves (temporarily) against gravitational collapse. Highly-magnetized remnants may have lifetimes between 10 s and 10^4 s after merger [302], while weakly-magnetized remnants might survive for thousands of years [303]. When these remnants inevitably collapse, the newborn black hole must shed the magnetic "hair" of the supramassive remnant. The result is the rapid destruction of the system's magnetosphere via reconnection, accelerating particles to relativistic speeds and potentially producing coherent radio emission through curvature radiation [302–304].

Even heavier "hypermassive" remnants might exist, supported by differential (rather than solid-body) rotation for roughly 10 ms before collapsing to a black hole. Differential rotation can strongly amplify the remnants' initial magnetic field to magnetar field strengths ($\sim 10^{16}$ G). Subsequent magnetic braking of the hypermassive remnant might then power coherent pulsar-like emission for the duration of the remnant's lifetime [305]. Following gravitational collapse, there presumably may also be an electromagnetic transient associated with the destruction of the hypermassive remnant's amplified magnetic field, analogous to the case of supramassive remnants above.

Not all authors agree. In Refs. [288, 289, 306] it is argued that, in the presence of a plasma-filled magnetosphere, the "no-hair" theorem is invalid and the final black hole retains the remnant's magnetic field. In this case there is no rapid destruction of the magnetosphere. Instead, the newborn black hole is magnetically braked just like a neutron star, generating a potentially observable electromagnetic signal [288, 306].

8.2.3 Post-Merger

Finally, prompt radio emission may even arise in the moments *after* merger, due to a highly-magnetized relativistic wind launched by the binary merger. As the relativistic wind interacts with ambient interstellar gas, it is decelerated and its magnetic field weakened. When, in the wind's rest frame, the energy densities of the oncoming gas and the wind's magnetic field become comparable, the wind-gas interaction becomes turbulent, exciting non-stationary currents at the wind front and generating low-frequency electromagnetic waves [307–309].

8.3 Why Radio?

In Ch. 8.2 we've seen many means of energy extraction from binary neutron stars. We've seen conspicuously fewer explicit means of *radio wave* production. Instead, most authors cited above posit some microphysical process, generally parametrized by an unknown efficiency factor, that converts the extracted energy into coherent radio emission.

Why should we trust that these mechanisms do in fact yield radio emission? The answer is two-fold. First, many of these ideas describe conditions resembling the environments of pulsars, with relativistic charges moving along the field lines of a highly magnetized and rapidly rotating central engine. And, although we don't understand how, we *know* that pulsars ultimately produce coherent radio emission (see e.g. Ref. [310]). It may not be too large a leap, therefore, to suspect that whichever mechanism(s) are responsible for pulsar radio emission might also be at play here.

Second, in many ways these authors' focus on radio signatures anticipates observational selection effects. We are most likely to detect the brightest electromagnetic signals, which in turn are most efficiently produced by coherent emission. And coherent emission is most readily obtained at low frequencies and long wavelengths. There are also *temporal* selection effects to consider. As will be discussed more below, the speed at which electromagnetic waves travel through the intergalactic and interstellar media is frequency-dependent. Prompt emission at high frequencies propagates at the speed of light, and so will arrive at Earth instantaneously with the gravitational-wave signal. Barring independent detection by wide-field survey instruments, this high-frequency emission will have come and gone by the time the gravitational-wave signal is registered and announced to electromagnetic observers. Prompt emission at radio frequencies, on the other hand, propagates below the (vacuum) speed of light; the resulting delay between the arrival of gravitational-wave and radio signals may offer just enough time to detect a gravitational wave, alert and point radio telescopes, and catch the slightly slower prompt radio counterpart.

8.4 Looking Ahead

In the next two chapters we'll explore what we can presently say about the nature of prompt radio emission. First, in Ch. 9 I'll ask if prompt radio emission has *already* been detected in the form of **fast radio bursts** (FRBs). Powerful radio transients of extragalactic origin, the source of FRBs remains unknown. Many authors speculate that FRBs may constitute prompt radio emission from distant compact binaries. We'll investigate whether this association is plausible by comparing the inferred rates of FRBs and compact binary mergers.

In Ch. 10 I will then present the first search for prompt radio emission associated with a gravitational-wave detection. This search target the binary black hole event GW170104 [42]. While binary black holes are generally *not* expected to yield electromagnetic emission, this study will develop and demonstrate the tools needed for prompt radio follow-up of compact binaries, tools that are currently being used to follow-up gravitational-wave candidates in the present O3 observing run.

Chapter 9
Fast Radio Bursts as Prompt Emission from Compact Binaries

∾

This chapter contains work published in:

> T. Callister, J. Kanner, and A. Weinstein, *Gravitational-Wave Constraints on the Progenitors of Fast Radio Bursts*, Astrophys. J. Letters **825**, L12 (2016).

I led this study, producing all results and writing the published manuscript.

Today the detection of fast radio bursts (FRBs) seems nearly commonplace – high time-resolution radio surveys are routinely detecting dozens of FRB signals. In 2016, though, when this work was published, there were only 17 confirmed fast radio bursts. While I have updated Ch. 9.1 with the latest FRB count and appropriate references, the remainder of this chapter is written from the perspective of 2016, assuming only those FRBs publicly announced when Callister *et al.* (2016) was in preparation. This also means that the rates of compact binary mergers quoted in Ch. 9.2 will be rather dated. In 2016, our knowledge of the binary black hole merger rate was based on only a single event – GW150914. And we had only *upper limits* on the rate of binary neutron star mergers. Nevertheless, the conclusions drawn in this chapter are extremely robust, and do not change when using the most up-to-date FRB and compact binary merger rates. Although not reflected in the text, I *have* annotated Fig. 9.2 to include present-day constraints on compact binary merger rates.

∾

9.1 Fast Radio Bursts

The discussion of prompt radio emission in Ch. 8 above may well remind you of another class of radio transient – fast radio bursts. First discovered in 2007 [311], fast radio bursts (FRBs) are radio transients characterized by millisecond durations, ∼Jansky flux densities, and dispersion measures (DMs) consistent with sources at Gpc distances. Although first observed only with the Parkes Radio Telescope (and partially conflated with a misbehaving microwave), FRBs are now observed with a growing number of instruments, including Parkes [311–318], Arecibo [319–321], Green Bank [322], Molonglo [323–325], the SKA Pathfinder [326–330], and CHIME [331, 332]. FRBs also seem to be quite numerous. To date, 73 confirmed FRBs have been reported [333].[1] After correcting for sky coverage and observing cadence, this implies that between 10^3 and 10^4 occur on the sky per day [313, 334]. That is, a hypothetical telescope array observing continuously with complete sky coverage would observe between 1000 and 10,000 FRBs per day!

A large number of theories have been put forward as to the possible source(s) of FRBs. Theorized sources include (but are certainly not limited to) supergiant neutron star pulses [335, 336], pulsar-planet systems [337], bremmstrahlung from gamma-ray bursts or active galactic nuclei [338], and galactic flare stars [339, 340]. More exotic sources include the explosions of white holes [341], primordial black hole evaporation [342], cosmic string decay [343], or even lasers used by advanced alien civilizations to propel light sails [344].

As tempting as alien light sails may be, there is another attractive possibility: compact binary mergers. The observed properties of FRBs are tantalizingly similar to those expected of prompt radio emission. Many of the mechanisms discussed in Ch. 8, in fact, were first proposed as possible explanations for the origin of FRBs [284, 290, 291, 302–304]. Mechanisms like the unipolar inductor [294] may also be at play in neutron star-black hole [288, 345–347] and even binary black hole [288, 348] mergers.

The possibility that binary coalescences are FRB progenitors is particularly appealing. This association would further bolster the case for prompt radio monitoring of compact binary mergers [349–356] Although the recent discovery of two repeating fast radio bursts [320, 332] points to a non-cataclysmic origin for at least some fraction of FRBs, FRBs may not constitute a single population [320, 346]; there may instead exist multiple FRB populations, each arising from a different class of progenitor.

[1] See the FRBCAT: http://www.astronomy.swin.edu.au/pulsar/frbcat/

If binary coalescences are to be considered plausible models for one such progenitor population, then their astrophysical rates must be less than or equal to the inferred rate of FRBs. In this chapter, we will explore this consistency. We will find that existing and future gravitational-wave measurements of the rates of binary coalescences can be leveraged to place novel constraints on the nature of FRB progenitors. In some cases, we can confidently rule out certain classes of binary coalescences as dominant FRB progenitors.

9.2 Rates of Compact Binary Coalescences

The Advanced LIGO and Virgo detection of the BBH merger GW150914 [25] produced the first direct measurement of the binary black hole merger rate per comoving volume (the so-called "rate density") in the nearby Universe. From this event, it was inferred that the BBH merger rate density lies between 2 and $400\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ [27].

While the rate densities of BNS and NSBH mergers remain unknown, binary pulsar observations and population synthesis models place rough bounds on the expected BNS and NSBH rates, respectively.[2] BNS and NSBH merger rate densities are predicted to plausibly fall between $R_{\mathrm{BNS}} = 10 - 10^4\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ and $R_{\mathrm{NSBH}} = 0.6 - 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ [357–360]. Note, however, that Ref. [361] predicts NSBH rate densities as low as $0.04\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$. Gravitational-wave experiments have not yet begun to probe these predicted ranges; the best experimental results, placed by jointly by Initial LIGO and Initial Virgo, limit BNS and NSBH merger rate densities to $R_{\mathrm{BNS}} < 1.3 \times 10^5\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ and $R_{\mathrm{NSBH}} < 3.1 \times 10^4\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, respectively [362].

Although these Initial LIGO/Virgo limits are well above the most optimistic predictions from population synthesis and binary pulsars, Advanced LIGO's recently concluded first observing run (O1) is expected to measure rate densities down to $R_{\mathrm{BNS}} \approx 3 \times 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ and $R_{\mathrm{NSBH}} \approx 750\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, experimentally probing for the first time the range of astrophysically plausible merger rates [363]. In 2017-18, Advanced LIGO's second O2 observing run is projected to be sensitive[3] to rate densities as low as $R_{\mathrm{BNS}} \approx 450\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$

[2] We now know the binary neutron star merger rate to be $R_{\mathrm{BNS}} = 0.1 - 3.8 \times 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ [45]. The binary black hole merger rate has been refined to $R_{\mathrm{BBH}} = 0.1 - 1.1 \times 10^2\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, and an improved upper limit of $R_{\mathrm{NSBH}} < 600\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ has been placed on the neutron star-black hole merger rate [45, 111].

[3] Clearly this projection was optimistic; see the previous footnote.

and $R_{\rm NSBH} \approx 100\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$, while O3 may further push Advanced LIGO's sensitivity to $R_{\rm BNS} \approx 100\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$ and $R_{\rm NSBH} \approx 20\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$ [363].

9.3 Rates of FRBs

The predicted and measured rates of binary coalescences allow for direct constraints on the nature of FRB progenitors by comparison to the inferred FRB rate per comoving volume. Other authors have considered the physical rate of FRBs, but these calculations are typically not shown in detail and significant disagreement exists in the literature, e.g. Ref. [284] vs. Refs. [364, 365]. Our goal in this section is therefore a careful accounting of the FRB rate density. As we will show below, the FRB rate per comoving volume is potentially far higher than the corresponding rate densities of binary coalescences. Thus, it is unlikely that the coalescence of stellar-mass compact binaries represents more than a small fraction of FRB progenitors. Because of this rate discrepancy, the *lowest* FRB rate estimates are *most compatible* with CBC progenitors. In the following, we will therefore deliberately seek a lower limit on the FRB rate density in order to most generously assess the plausibility of CBC progenitors of FRBs.

The inferred FRB rate per comoving volume is approximately $(3r_{\rm obs})/(4\pi D^3)$. Here, D is the comoving distance containing the observed FRB population and $r_{\rm obs}$ is the observed rate at which FRBs occur on the sky. For simplicity, we will assume this rate density is constant and neglect evolution with redshift. If FRB emission is *beamed*, then the rate $r_{\rm obs}$ is undercounted due to selection effects – beamed FRBs, like pulsars or GRBs, are only observed if the Earth lies within the path of the beam. In general, the FRB rate per comoving volume is

$$R_{\rm FRB} \approx \frac{3r_{\rm obs}}{\Omega D^3}, \tag{9.1}$$

where Ω is a typical solid angle over which emission is beamed.

Although few FRBs have been observed, their inferred rate on the sky is large. With four FRB detections at high Galactic latitudes using Parkes, Ref. [313] inferred that $r_{\rm obs} = 1.0^{+0.6}_{-0.5} \times 10^4$ FRBs occur on the sky per day. However, there remains considerable disagreement as to the true value of $r_{\rm obs}$, with subsequent radio surveys producing differing rate estimates, often defined with respect to different fluence limits and different assumptions about search systematics. Ref. [334], for instance, points out that FRB detection is subject to significant selection effects, such as survey incompleteness below a fluence of $\sim 2\,\mathrm{Jy\,ms}$, suboptimal recovery of broad radio pulses, and potential obscuration of FRBs in the galactic plane. They estimate a fluence-limited detectable FRB rate of $2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$ above $\sim 2\,\mathrm{Jy\,ms}$. Ref. [366] also arrives at $r_{\mathrm{obs}} \approx 2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$ but by different means, suggesting that the apparent FRB rate at high latitudes is enhanced by interstellar scintillation. Ref. [367], meanwhile, adopts a Bayesian approach, combining several published rate estimates to obtain $r_{\mathrm{obs}} = 4.4^{+5.2}_{-3.1} \times 10^{3}\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$ above $4.0\,\mathrm{Jy\,ms}$. On the other hand, Ref. [368] argues that previously published single-dish rate estimates are biased below their true values and that, once potential biases are corrected, previous estimates are consistent with $r_{\mathrm{obs}} = 1.2 \times 10^{4}\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$ above $1.7\,\mathrm{Jy\,ms}$.

It is not obvious which value to select for r_{obs} (or even which range of uncertainties to adopt). In order to place a robust lower limit on the FRB rate density, however, we will take $r_{\mathrm{obs}} = 2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$, consistent with the lowest of the above estimates. To additionally allow for various search selection effects, we will define η as the FRB detection efficiency, the fraction of otherwise detectable FRBs (e.g. with intrinsic signal-to-noise ratios above some threshold detection value) which are actually recovered in a radio transient search. The physical rate of FRBs on the sky is then $r_{\mathrm{obs}}/(\eta\Omega)$.

Distances to FRB sources may be estimated using their reported DMs, which we obtained from the FRBCAT [333]. Assuming that the intergalactic medium (IGM) is homogeneous and fully ionized, the dispersion measure $\mathrm{DM}_{\mathrm{IGM}}$ due to propagation through the IGM is related to source redshift via [369, 370]

$$\mathrm{DM}_{\mathrm{IGM}}(z) = \frac{\overline{n}_e c}{H_0} \int_0^z \frac{(1+z')\,dz'}{\sqrt{\Omega_{\mathrm{M}}(1+z')^3 + \Omega_\Lambda}}, \tag{9.2}$$

where $\overline{n}_e = \rho_c \Omega_{\mathrm{B}}/m_p = 2.5 \times 10^{-7}\,\mathrm{cm}^{-3}$ is the local free electron density in a fully ionized Universe. Here, Ω_{B}, Ω_{M}, and Ω_Λ are the energy densities of baryons, matter, and dark energy, respectively, m_p is the proton mass, and $\rho_c = 3H_0^2/8\pi G$ is the critical energy density required to close the Universe. G is Newton's constant, c the speed of light, and H_0 the Hubble constant; we use $H_0 = 67.7\,\mathrm{km\,s^{-1}Mpc^{-1}}$, $\Omega_{\mathrm{B}} = 0.049$, $\Omega_{\mathrm{M}} = 0.31$, and $\Omega_\Lambda = 0.69$ [371]. The comoving distance corresponding to redshift z is given by

$$D(z) = \frac{c}{H_0} \int_0^z \frac{dz'}{\sqrt{\Omega_{\mathrm{M}}(1+z')^3 + \Omega_\Lambda}}. \tag{9.3}$$

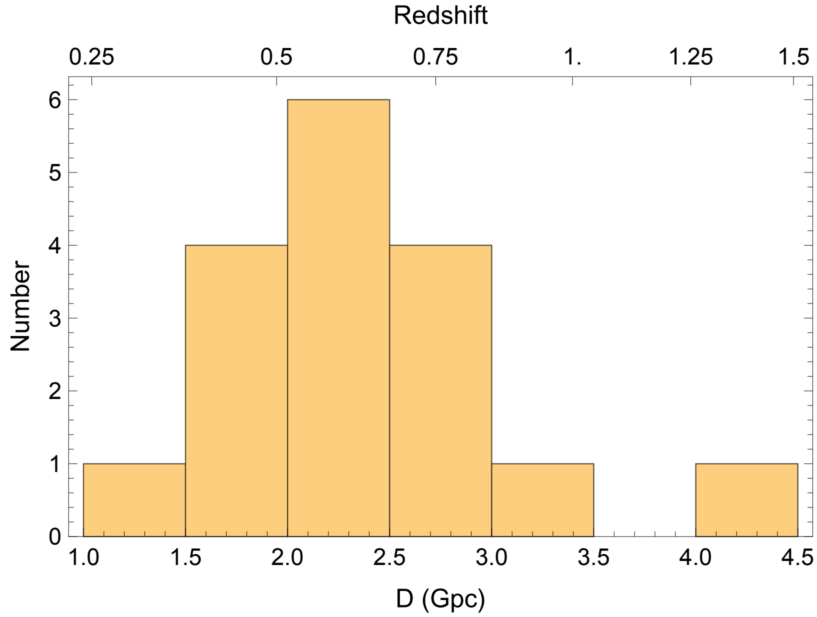

Figure 9.1: Distribution of inferred distances to known FRBs, assuming a homogeneous, fully ionized intergalactic medium and neglecting dispersion measure contributions from both the Milky Way and FRB host galaxies. We take 3 Gpc as a fiducial distance bounding the observed FRB population.

In the small redshift limit this reduces to $D \approx cz/H_0$, and Eq. (9.2) becomes $\mathrm{DM}_{\mathrm{IGM}}(z) \approx \overline{n}_e D$. Using Eqs. (9.2) and (9.3), the inferred comoving distances to the 17 known FRBs are shown in Fig. 9.1. Based on this sample, we will take $D = 3\,\mathrm{Gpc}$ as a fiducial distance encompassing the observed FRB population.

We have made several assumptions in computing the distances shown in Fig. 9.1. Since a factor of 2 error in the fiducial distance will result in a factor of 2^3 error in the FRB rate density, it is important to highlight these assumptions and understand how they affect our result.

First, we have assumed that the observed radio dispersions are entirely due to propagation through the IGM. In reality, the Milky Way may contribute up to $\sim 20\%$ of the observed DM [333]. The distances in Fig. 9.1 may therefore be overestimated by a factor of ~ 1.25. If we also allow for a comparable DM contribution by the FRB's host galaxy (as well contributions from any matter over-densities along the line of sight to the FRB), then the distances may be overestimated by at least a factor of 1.7. This implies that our FRB rate density is *underestimated* by a factor between 2 and 5.

Second, we assumed a fully ionized Universe. While valid for hydrogen, this is not necessarily true for helium, which may be either singly or fully ionized.

Helium makes up approximately 24% of the IGM by mass [370]; if this helium is only singly ionized, then the free electron density $\overline{n}_e$ will be reduced by roughly 10%. Finally, Ω_{B} is an overestimate of the baryon density in the IGM, since $\sim$ 10% of baryons are sequestered in galaxies [372]. Together, these two approximations cause $\overline{n}_e$ to be (at most) 20% larger than the true free electron density in the IGM. Hence the fiducial distance D is underestimated by a factor of 1.25, and the FRB rate density is correspondingly *overestimated* by a factor of 2.

Of the two assumptions described above, the first (uncertainty in the galactic and intergalactic DM) will cause the fiducial distance to be underestimated, while the second and third (uncertainty in $\overline{n}_e$) cause the distance to be overestimated. Of these uncertainties, the potentially large overestimate of the intergalactic DM is expected to dominate. Thus our choice of $D = 3\,\mathrm{Gpc}$ is likely an *upper bound* on the fiducial FRB distance. This is in line with our goal – to obtain a robust lower limit on the rate density of FRBs. Because $R_{\mathrm{FRB}} \propto D^{-3}$, any decrease in the fiducial distance D will only increase our estimated FRB rate density, further increasing the tension identified below between the rates of FRBs and binary coalescences.

All together, the FRB rate per comoving volume is

$$\begin{aligned} R_{\mathrm{FRB}} = 8.1 \times 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1} \left(\frac{r_{\mathrm{obs}}}{2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}} \right) \left(\frac{1}{\eta} \right) \\ \times \left(\frac{3\,\mathrm{Gpc}}{D} \right)^3 \left(\frac{4\pi\,\mathrm{Sr}}{\Omega} \right). \end{aligned} \tag{9.4}$$

The characteristic parameters in Eq. (9.4), $r_{\mathrm{obs}} = 2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$, a perfect detection efficiency $\eta = 1$, and isotropic radio emission $\Omega = 4\pi\,\mathrm{Sr}$, have been chosen to yield the lowest possible FRB rate density consistent with observations:

$$R_{\mathrm{FRB}}^{\mathrm{Low}} = 8.1 \times 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}. \tag{9.5}$$

A more plausible rate density, on the other hand, is obtained by assuming $r_{\mathrm{obs}} = 5000\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$, an imperfect detection efficiency of $\eta = 0.5$, and moderately beamed emission with a 30° half-opening angle. These values give

$$R_{\mathrm{FRB}}^{\mathrm{Realistic}} = 4.8 \times 10^5\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}, \tag{9.6}$$

nearly two orders of magnitude larger than Eq. (9.5).

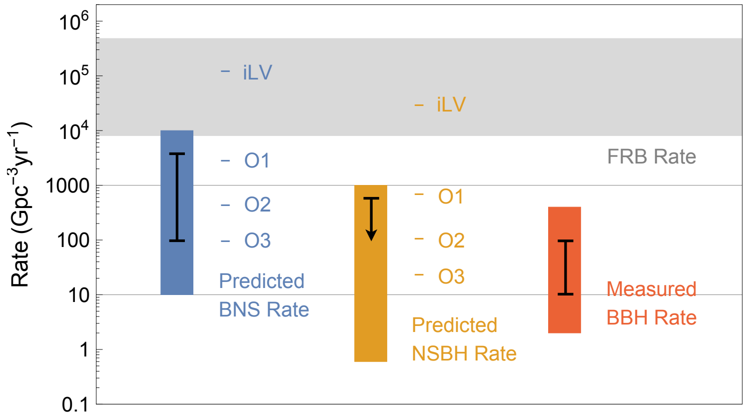

Figure 9.2: Binary coalescence rates compared to the inferred rate of FRBs. Solid bars indicate the range of BNS (blue) and NSBH (orange) merger rates predicted by binary pulsar observations and population synthesis models [359], as well as the *measured* LIGO/Virgo rate of BBH mergers (red) [27]. Also shown are existing Initial LIGO/Virgo (iLV) limits and projected O1, O2, and O3 sensitivities [362, 363]. The gray band indicates a range of potential FRB rate densities, from the lowest plausible value in Eq. (9.5) to a more realistic estimate in Eq. (9.6). (**Note**: Shown in black are the latest estimates of compact binary merger rates [45, 111], following the O2 observing run.)

Our lower limit on the FRB rate density agrees well with the rate previously estimated by Ref. [284]. It is, however, more than an order of magnitude higher than the more recent results computed in Refs. [364, 365]. The discrepancy lies in the fact that Refs. [364, 365] mistakenly use the luminosity distance $D_L = D(1+z)$ rather than the comoving distance to calculate the FRB rate density. Refs. [364, 365] choose $z = 1$ as a fiducial redshift; the corresponding comoving and luminosity distances are $D = 3.4\,\mathrm{Gpc}$ and $D_L = 6.8\,\mathrm{Gpc}$, respectively. This factor of 2 error in distance leads to a factor of 8 error in the FRB rate density. Plugging in $D = 3.4\,\mathrm{Gpc}$ in Eq. (8) of Ref. [365] gives an FRB rate density of $5.8 \times 10^3\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, in reasonably good agreement with our lower limit.

9.4 Compact Binaries as FRB Progenitors?

By comparing R_{FRB} from Ch. 9.3 to the binary coalescence rates in Ch. 9.2, we can constrain the fraction of FRBs that can be explained via compact binary coalescences. Fig. 9.2 shows a range of potential FRB rate densities, from the lowest plausible estimate given in Eq. (9.5) (assuming $r_{\rm obs} = 2500\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$, efficiency $\eta = 1$, and isotropic FRB emission) to a more realistic value in Eq. (9.6) (which assumes $r_{\rm obs} = 5000\,\mathrm{sky}^{-1}\,\mathrm{day}^{-1}$, efficiency $\eta = 0.5$, and FRB beaming with a half-opening angle of 30°). Solid bars indicate the range of BNS and NSBH merger rate densities predicted by binary pulsar observations and population synthesis models, as well as the measured BBH rate density. Also shown are existing Initial LIGO/Virgo limits, as well as the projected sensitivities of the O1, O2, and O3 observing runs.

Binary black holes: The measured rate of binary black holes mergers is *at most* $\sim 5\%$ of the inferred FRB rate. Thus, BBHs cannot explain more than a small fraction of the observed FRB population. Previous claims that the rates of FRBs and BBH mergers are consistent [364, 365] are based on an erroneous calculation of the FRB rate density, as discussed in Ch. 9.3.

Neutron star-black hole binaries: Population synthesis predictions are highly inconsistent with the theory that NSBH mergers are FRB progenitors, with predicted NSBH merger rates equal to at most $\sim 12\%$ of the FRB rate. This fraction assumes isotropic radio emission, and hence should be taken as a highly optimistic upper limit on the FRB fraction compatible with NSBH binaries. Even moderate beaming, with a half-opening angle of e.g. 30°, reduces the predicted FRB fraction to $\sim 0.8\%$. Assuming the realistic FRB rate density in Eq. (9.6) further lowers this fraction by a factor of four.

Although Initial LIGO/Virgo upper limits are uninformative (limiting the most optimistic NSBH fraction of FRB progenitors to $R_{\rm NSBH}/R_{\rm FRB}^{\rm Low} \lesssim 4$), Advanced LIGO is capable of measuring significantly smaller NSBH merger rates. A non-detection during the O1 and O2 observing runs, for instance, would limit the NSBH FRB fraction to $\lesssim 9\%$ and $\lesssim 1\%$, respectively (assuming isotropic emission).

Binary neutron stars: There exist competing claims as to whether the rates of FRBs and binary neutron star mergers are [284, 290] or are not [313] compatible. Note that Ref. [290] adopts the FRB rate estimate from Ref. [365], and hence erroneously finds strong consistency between the rates of BNS mergers and FRBs. We find that the most optimistic BNS rate density predictions are consistent with the lowest possible FRB rate density, with $R_{\rm BNS}/R_{\rm FRB}^{\rm low} \approx 1.2$. Therefore, BNS mergers could constitute a subpopulation of FRB progenitor if multiple FRB subclasses do indeed exist. This compatibility is quite tenuous, however, simultaneously requiring the highest possible BNS rates and the lowest possible FRB rates (with, e.g., perfect FRB detection efficiency and isotropic radio emission). FRB models that predict even moderately beamed emission are incompatible with BNS progenitors.

If BNS mergers are indeed FRB progenitors, then it is likely that Advanced LIGO will observe a large number of BNS sources in the O1 observing run. If no such detections are made, then the resulting rate limits will increasingly cast doubt on the role of BNSs as FRB progenitors. An Advanced LIGO non-detection during O1 and O2 would limit the most optimistic fraction of FRBs compatible with BNS mergers to $\lesssim 40\%$ and $\lesssim 6\%$, respectively. If we assume moderate FRB beaming (again with a half-opening angle of 30°), then O1 and O2 non-detections imply even more stringent FRB fractions of $\lesssim 2\%$ and $\lesssim 0.4\%$, respectively. Note that these limits also apply equally well to short-lived products of BNS mergers, such as hypermassive neutron stars.

9.5 Conclusions

A diverse range of FRB progenitor models have been proposed, including the binary coalescences of neutron stars and/or black holes. Existing or future limits from gravitational-wave observations can serve to severely constrain such models. The recent Advanced LIGO/Virgo measurement of the local BBH merger rate density largely rules out stellar-mass binary black holes as progenitors of the observed FRB population. Meanwhile, predictions of NSBH merger rate densities from population synthesis are in strong tension with the inferred rate density of FRBs; upcoming observations by Advanced LIGO and Virgo could rule out NSBHs as FRB progenitors.

Under highly generous assumptions (broadly beamed radio emission, large FRB distances, and low underlying FRB rates), the rate of BNS mergers may be consistent with a subpopulation of FRB progenitors. In order for this subpopulation to be significant, however, the BNS merger rate density must be on the order of $\sim 10^4\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, comparable to the most optimistic predictions from population synthesis. Additionally, FRB emission must be largely isotropic; models that predict even moderately beamed emission are inconsistent with BNS rates. If BNS mergers are indeed FRB progenitors, then Advanced LIGO and Virgo will soon begin to observe a large number of such systems. If no such observations are made, then the resulting rate limits will increasingly constrain the ability of BNSs to explain even a subclass of the FRB population.

Chapter 10
A First Search for Prompt Radio Emission Accompanying a Gravitational-Wave Event

This chapter contains work published in:

> T. A. Callister, M. M. Anderson, G. Hallinan, *et al.*, *A First Search for Prompt Radio Emission from a Gravitational-Wave Event*, Astrophys. J. Letters **877**, L39 (2019).

I produced all results (with a great deal of help and guidance from Marin Anderson) and led the preparation of the published manuscript.

10.1 Introduction

Beyond the gamma-ray burst, kilonova, and radio jet of GW170817, we argued in Ch. 8 that binary neutron stars may additionally be accompanied by prompt radio emission. Unlike the late-time radio afterglow associated with GW170817, due to the interaction of relativistic ejecta with the ambient medium, prompt radio emission is theorized to be generated by an altogether different process (or processes) in the immediate vicinity of the merging objects. In particular, prompt emission may take the form of a short (likely sub-second) coherent radio pulse generated near the instant of merger.

The detection of prompt radio emission from a binary neutron star would yield an immense amount of information. Prompt radio emission would serve as a probe of the binary's immediate magnetic environment near the time of merger. The observed dispersion and/or scattering of the emission acquired during propagation might enable study of the intergalactic medium and Milky Way halo. The detection of prompt radio emission would also incidentally provide precise $\sim$arcminute constraints on the *location* of its compact binary progenitor.

The successful observation of prompt emission is made difficult by several factors, however [352, 353, 373]. First, gravitational-wave detectors provide only poor localization of gravitational-wave sources. Even for the three-detector Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo network, the median binary neutron star localization is expected to be $120 - 180\,\mathrm{deg}^2$ during the upcoming O3 observing run [38]. Second, low-frequency prompt emission released at time of merger may arrive at Earth as little as one minute after the gravitational-wave signal, slowed only by free electrons encountered during propagation. Searches for prompt radio emission are therefore typically limited by the latency with which gravitational-wave candidates are announced – notices released more than minutes after a gravitational wave's arrival may well come too late.

All previous searches for prompt radio emission have targeted short gamma-ray bursts [350, 351, 354, 356, 374] or were carried out too late to detect any prompt emission that may have been present [375, 376]. Here, I will describe the first search for prompt radio emission coincident with a gravitational-wave signal, using the Owens Valley Radio Observatory Long Wavelength Array (OVRO-LWA).

The OVRO-LWA is a low-frequency interferometry radio array in Owens Valley, California. The array comprises 288 dual-polarization antenna spanning a 1.5 km maximum baseline, and observes between 27 and 84 MHz. Cross-correlation of 256 of these antenna using the LEDA (Large-Aperture Experiment to Detect the Dark Age) correlator allows for all-sky imaging with 24 kHz frequency resolution and ~ 10 arcminute spatial resolution [354, 377, 378].

The OVRO-LWA is uniquely suited to the challenge of detecting prompt radio emission. Its nearly hemispherical field of view can capture much of the LIGO-Virgo localization region within a single image. Additionally, the OVRO-LWA operates in a continuous buffered mode, temporarily saving all visibilities to disk for up to 24 hours. This alleviates (although does not eliminate; see Ch. 10.5) the need for rapid LIGO-Virgo notices. As long as LIGO and Virgo release a notice within one day of the gravitational-wave event, the relevant on-source data can be retrieved from the buffer and written permanently to disk.

In previous discussion of prompt radio emission, we have focused almost entirely on the case of binary neutron stars. The OVRO-LWA was indeed observing at the time of GW170817. Unfortunately, however, the binary neutron star merger occurred below the OVRO-LWA's horizon [46], and so we were unable to search for prompt radio emission associated with GW170817. Instead, here I will describe the results of an analogous search targeting the binary black hole merger GW170104 [42, 379].

According to conventional wisdom, stellar-mass binary black hole mergers are generally not expected to yield electromagnetic transients. However, the Fermi-GBM detection of a marginally-significant gamma-ray transient coincident with GW150914 [380] (and also a Fermi-LAT outlier at the time of the binary black hole GW170608 [43, 381, 382]) has sparked new interest in possible counterparts to LIGO/Virgo's binary black hole events. In particular, binary black holes might conceivably generate electromagnetic transients if one or more of the black holes is charged [364, 383–385] or if the system is subjected to an ambient magnetic field [288, 348]. Electromagnetic transients may also occur in the presence of a circumstellar or circumbinary disk [386, 387], or in the case of black hole "twins" born from the collapse of a single massive star [388, 389]. Although the statistical significance of the Fermi-GBM candidate remains under debate [390–393], the plethora of models predicting electromagnetic counterparts makes binary black hole mergers an interesting (if speculative) observational target.

While valuable in its own right, the search for prompt radio emission from GW170104 will additionally serve as a powerful proof of principle. GW170104 exemplifies the challenges facing detection of prompt radio emission. First, its accompanying localization is poor, spanning a significant fraction of the sky. Second, the LIGO/Virgo alert announcing the detection of GW170104 was released hours after the gravitational-wave event, long after the expected arrival of any prompt radio emission. Despite these challenges, we place stringent upper limits on the prompt radio luminosity of GW170104, demonstrating the capability of the OVRO-LWA to follow up future compact binary mergers.

10.2 GW170104 and OVRO-LWA Observations

The gravitational-wave signal GW170104 was measured on 2017 January 4 at 10:11:58.6 UTC by the Advanced LIGO experiment [42, 379]. Arising from a $31 + 19\,M_\odot$ binary black hole merger, the signal was initially localized to a ~ 1600 square degree band on the sky (see Fig. 10.1) and its redshift estimated

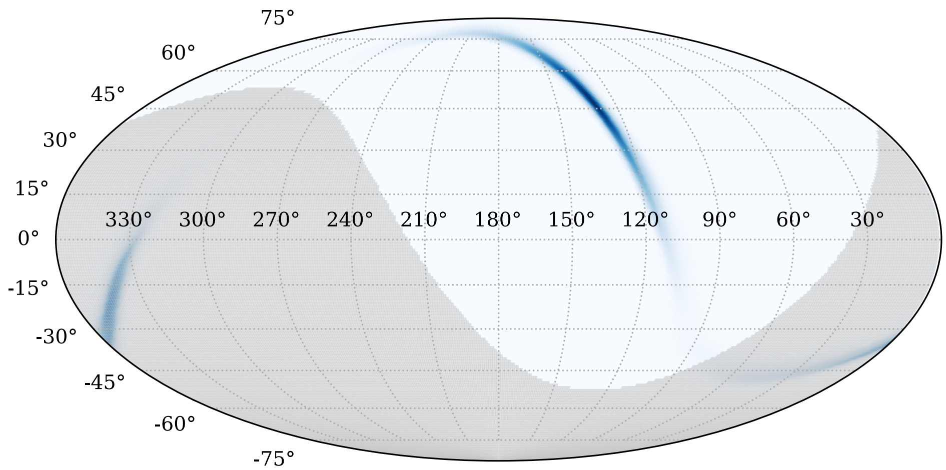

Figure 10.1: Posterior probability distribution (in blue) on the sky position of the binary black hole merger GW170104. Also shown is the OVRO-LWA's field of view at GW170104's time of arrival; areas below the OVRO-LWA horizon at this time are shaded in grey. The total localization region provided by Advanced LIGO spans approximately 1600 square degrees, with a 54% probability of GW170104 occurring within the OVRO-LWA field of view.

to be $z = 0.18^{+0.08}_{-0.07}$. Following a delay related to the calibration of Advanced LIGO's Hanford detector, an alert with preliminary event localization was released at 16:49:56 UTC, six hours after the gravitational wave's arrival [394]. At this time, the OVRO-LWA was under continuous operation, temporarily storing 13 s integrations in a continuously-overwritten 24-hour buffer. Upon receiving the gravitational-wave event notice, buffered data spanning 09:00:03 to 14:11:11 UTC were copied to long-term storage.

Data are flagged (i.e. vetoed) on a per antenna, baseline, and channel basis. We flag antennas showing anomalous autopower spectra, cutting an average of 54 antennas ($\sim$ 38% of visibilities). An additional 398 baselines are flagged to mitigate cross-talk between adjacent signal paths and eliminate other spurious excess power. Finally, loud individual channels are automatically flagged to reduce radio frequency interference (RFI), removing $\sim$ 12% of the 2398 frequency channels.

Cassiopeia (Cas) A and Cygnus (Cyg) A are the brightest sources in the low-frequency radio sky and therefore make opportune calibration sources. We calibrate our visibility data using a single integration recorded roughly ten hours earlier, at 22:44:04 January 03 UTC (21:49:47 local sidereal time), when both Cas A and Cyg A are close to zenith. The per-channel complex gains of each antenna are determined using a simplified sky model comprising three point sources – Cas A, Cyg A, and the Sun [395, 396].

Following this initial calibration, there persist residual errors due to unmodeled directional variations in antenna gains. To combat sidelobe contamination arising from these errors, we "peel" bright sources, performing an additional direction-dependent calibration and subtraction of these sources [397]. At the time of GW170104, Cas A and Taurus (Tau) A are the brightest sources in the OVRO-LWA field of view (Cyg A had since set below the horizon). We peel both Cas A and Tau A, as well as a generic near-field source to remove a stationary noise pattern likely caused by cross-talk between electronics [378]. Because Cas A is nearly on the OVRO-LWA's horizon, this peeling procedure fails for a small number of integrations; these integrations are manually flagged.

Figure 10.2 shows a peeled and deconvolved 13 s image of the OVRO-LWA sky at the time of GW170104 with 0.125 deg resolution. Deconvolution is performed using the WSCLEAN algorithm with a Briggs weighting of 0 and a multiscale bias of 0.6[398]. The blue contours show the 68% and 95% credible bounds on the sky location of GW170104's progenitor, restricted to the OVRO-LWA's field of view. The 95% credible contour contains 72,556 pixels, each of which we search for a dispersed radio signal.

10.3 Search for a Dispersed Signal

Radio waves of frequency ν propagating through the interstellar and/or intergalactic media experience a dispersion delay [399, 400]

$$\begin{aligned} t &= \left(\frac{e^2}{2\pi m_e c}\right)\frac{\mathrm{DM}}{\nu^2} \\ &= \left(4.149\times 10^3\,\mathrm{s}\right)\left(\frac{\mathrm{DM}}{\mathrm{pc\,cm^{-3}}}\right)\left(\frac{\nu}{\mathrm{MHz}}\right)^{-2} \end{aligned} \tag{10.1}$$

relative to signals of infinite frequency. Here, e is the fundamental charge, m_e is the electron mass, c is the speed of light, and the dispersion measure DM is the integrated column density of free electrons along the wave's path. The dispersion measure may contain contributions from the immediate environment and/or host galaxy of GW170104, the intergalactic medium, and the interstellar medium of the Milky Way.

Parameter estimation on the gravitational-wave signal GW170104 constrains its redshift to $z = 0.173^{+0.072}_{-0.071}$ with an effective-precession waveform model and

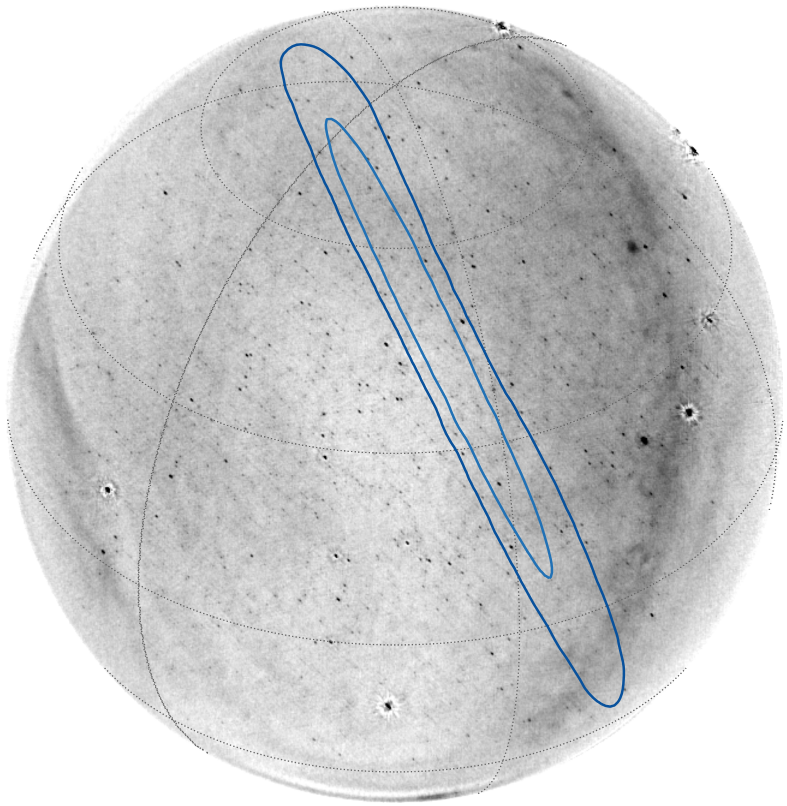

Figure 10.2: Total intensity image of the OVRO sky from 27 to 84 MHz in a 13 s interval centered at 10:11:54.1 UTC, containing GW170104's time of arrival. The dark and light blue contours show the 95% and 68% credible bounds on the location of GW170104's progenitor, respectively, conditioned on the OVRO-LWA's field of view.

$z = 0.182^{+0.081}_{-0.078}$ using a model capturing full spin-precession effects [42, 379]. We conservatively assume that GW170104's progenitor lies between $0.1 \leq z \leq 0.3$. When including the effects of cosmology, the dispersion measure due to propagation through the intergalactic medium is [369, 370]

$$\mathrm{DM}_{\mathrm{IGM}} = \overline{n}_e c \int_0^z \frac{(1+z')}{H(z')} dz', \tag{10.2}$$

where $H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$, H_0 is the Hubble constant, and Ω_m and Ω_Λ are the dimensionless energy-densities of matter and dark energy, respectively. We take $H_0 = 67.7\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, $\Omega_m = 0.31$, and $\Omega_\Lambda = 0.69$. As in Ch. 9.3, we obtain an upper limit on the mean electron density $\overline{n}_e$ in the intergalactic medium by assuming the Universe's baryonic density $\Omega_B = 0.049$

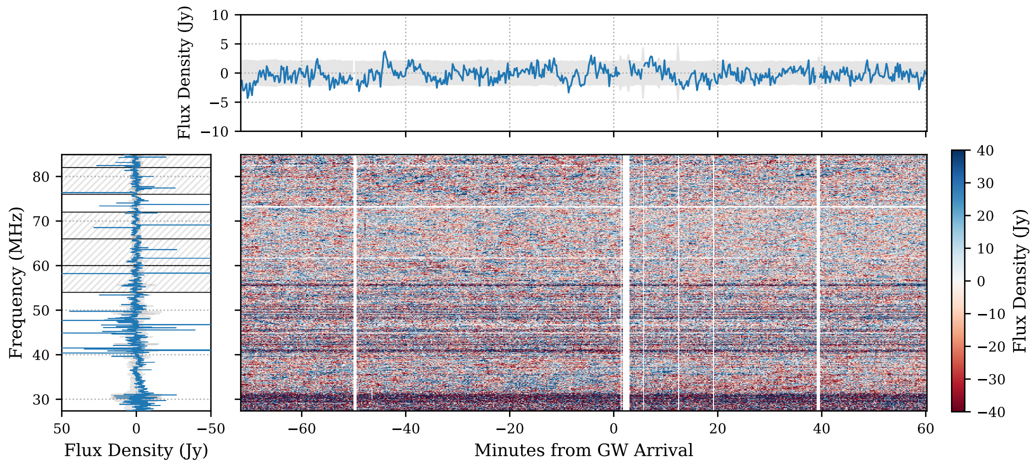

Figure 10.3: Dynamic spectrum of a randomly-chosen sky location within the GW170104 localization region (Fig. 10.2), after subtraction of the median flux measured in an annulus surrounding the target location. White vertical and horizontal bands correspond to times and frequency channels that have been flagged due to excess antenna power or RFI. The left and upper subplots show the time- and frequency-averaged flux densities, respectively. The filled grey region within each subplot marks the $\pm 3\sigma$ band as measured in the background annulus. Broadcast television channels are denoted by hatched regions in the time-averaged spectrum; these channels represent common sources of RFI due to meteor reflection events.

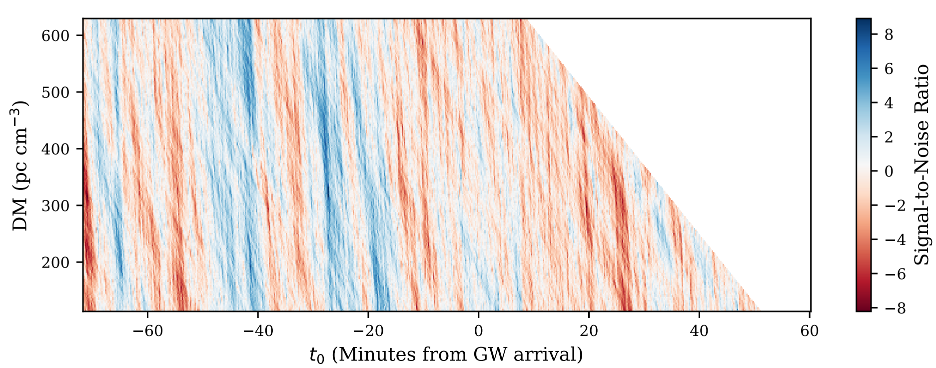

Figure 10.4: Signal-to-noise ratios as a function of dispersion measure DM and the initial time t_0 at which a signal is presumed to enter the OVRO-LWA band, targeting the same sky location as Fig.!10.3. We search for signals with dispersion measures $113\,\mathrm{pc\,cm^{-3}} \leq \mathrm{DM} \leq 630\,\mathrm{pc\,cm^{-3}}$ in a roughly one-hour window around the GW170104's time of arrival. The blank region on the figure's right-hand side corresponds to time-frequency tracks that extend beyond the duration of our data set.

is composed entirely of ionized hydrogen [369]. Then the mean electron number density is $\overline{n}_e = \Omega_B \rho_c / m_p$, where $\rho_c = 3H_0^2/8\pi G$ is the closure density of the Universe, G is Newton's constant, and m_p is the proton mass. As the Universe is neither fully ionized nor composed purely of hydrogen, this approximation yields an overestimate of $\overline{n}_e$ and hence a conservative overestimate of the intergalactic dispersion measure; see the discussion in Ch. 9.3. Assuming that GW170104's progenitor lies within $0.1 \leq z \leq 0.3$, we estimate $113\,\mathrm{pc\,cm^{-3}} \leq \mathrm{DM}_{\mathrm{IGM}} \leq 350\,\mathrm{pc\,cm^{-3}}$.

The GW170104 localization region, meanwhile, spans a broad range of Galactic latitudes, corresponding to a wide range of possible galactic dispersion measures. A lower bound on the dispersion measure to due the Milky Way is simply zero. An upper bound is given by assuming a line of sight directly through the Galactic disk, yielding $0 \leq \mathrm{DM}_{\mathrm{MW}} \leq 180\,\mathrm{pc\,cm^{-3}}$ [401, 402].

Finally, we need to account for the progenitor's host galaxy and/or immediate environment. We have no knowledge about either of these, and so the best we can do is add an additional and somewhat arbitrary term to our total dispersion measure budget. We naively assume $\mathrm{DM}_{\mathrm{Host}} + \mathrm{DM}_{\mathrm{Env}} \leq 100\,\mathrm{pc\,cm^{-3}}$.

Combining the individual contributions from the Milky Way, the intergalactic medium, and GW170104's environment and host galaxy, we bound the dispersion measure of radio transients associated with GW170104 to $113\,\mathrm{pc\,cm^{-3}} \leq \mathrm{DM}_{\mathrm{IGM}} \leq 630\,\mathrm{pc\,cm^{-3}}$. This corresponds to time delays ranging from 640 to 3600 s at the bottom of the OVRO-LWA band, relative to electromagnetic signals of infinite frequency. We therefore analyze data up to one hour after the gravitational-wave event. Recall that many models for prompt radio emission predict a *precursor* signal released before binary merger [285–287, 289, 293], and so we additionally analyze the 70 minutes of buffered data recorded before the event. Our final data set comprises 610 integrations spanning 09:00:03 to 11:12:00 UTC, each 13 s in duration.

Because of our finite 24 kHz frequency resolution, radio transients will also be dispersed *within* each frequency channel. This intra-channel dispersion is strongest in the lowest channel, in which signals separated by 24 kHz are delayed by a maximum of 6.4 s with respect to one another. Although this delay is smaller than our 13 s integration time, it is sufficiently long that a randomly placed transient might conceivably be split across adjacent integrations, potentially degrading our search sensitivity at low frequencies.

Scatter broadening is unlikely to affect our search. Assuming a ν^{-4} frequency dependence [400], the estimated Milky Way scattering timescale of $0.06\,\mu$s at 1 GHz corresponds only to 0.1 s at 28 MHz, much less than our 13 s integration time [401]. We might expect similarly negligible contributions from GW170104's host galaxy. Additionally, fast radio bursts show minimal scattering due to the intergalactic medium [403].

Our search window spans approximately 130 minutes. In this time the sky rotates considerably, and so we must track the movement of a given source across the OVRO-LWA's field of view. Just as a sufficiently broadened pulse could span multiple time integrations at a given frequency, it is possible for the Earth's rotation to smear emission across multiple image pixels within a single 13 s integration. Since the array's synthesized beam (with 0.50 deg and 0.24 deg major and minor axes at 56 MHz) is larger than our 0.125 deg pixel size, any astrophysical emission will manifest in multiple neighboring pixels. We are therefore unlikely to miss a significant fraction of a source's emission as we follow it from one image pixel to the next.

As an example, Fig. 10.3 shows the dynamic spectrum obtained by tracking a randomly chosen location within the GW170104 localization region. To account for slow temporal variations and sidelobes from bright, nearby sources, we have subtracted away the median flux measured in an annulus extending five to seven beamwidths around the target location (see Fig. 10.6 below). We search all such dynamic spectra for significant dispersed transients, stepping through dispersion measures and times t_0 at which a proposed signal enters the OVRO-LWA band. The spacing δDM between our dispersion measure trials is set by our $t_{\mathrm{int}} = 13\,\mathrm{s}$ integration time and the bounds $\nu_1 = 27.384\,\mathrm{MHz}$ and $\nu_2 = 84.912\,\mathrm{MHz}$ on the OVRO-LWA band:

$$\frac{\delta\mathrm{DM}}{\mathrm{pc\,cm^{-3}}} = \frac{t_{\mathrm{int}}}{4.149\times10^3\,\mathrm{s}}\left[\left(\frac{\nu_1}{\mathrm{MHz}}\right)^{-2} - \left(\frac{\nu_2}{\mathrm{MHz}}\right)^{-2}\right]^{-1}, \tag{10.3}$$

giving $\delta\mathrm{DM} = 2.62\,\mathrm{pc\,cm^{-3}}$. For each dispersed track, we estimate the corresponding flux density with the weighted average

$$\hat{F} = \frac{\sum_i \hat{F}_i/\sigma_i^2}{\sum_j 1/\sigma_j^2}, \tag{10.4}$$

where $\hat{F}_i$ is the measured flux density in the track's ith time-frequency pixel and σ_i^2 is the corresponding variance, estimated using the background annulus.

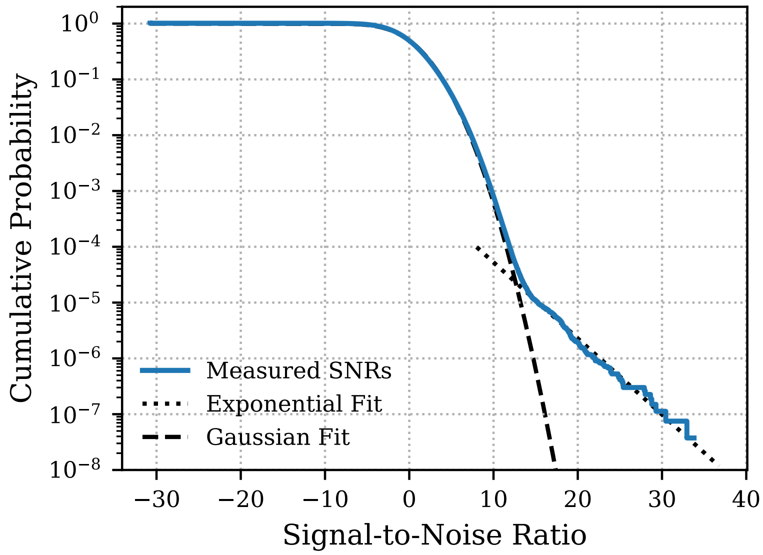

Figure 10.5: Cumulative background distribution of signal-to-noise ratios from a subset of sky directions and dispersion trials. The distribution is well fit by a central Gaussian and a exponential tail dominated by meteor reflection events. Based on this distribution, we manually follow up any dispersion trial giving S/N > 20.

In the presence of a true radio transient of flux density F, the expectation value and variance of $\hat{F}$ are

$$\langle \hat{F} \rangle = F \tag{10.5}$$

and

$$\sigma^2 = \frac{1}{\sum_i 1/\sigma_i^2}, \tag{10.6}$$

respectively. The signal-to-noise ratio (S/N) of each dispersion trial is defined by combining Eqs. (10.4) and (10.6):

$$\mathrm{S/N} = \frac{\hat{F}}{\sigma}. \tag{10.7}$$

Figure 10.4, for example, shows the signal-to-noise ratios obtained from dedispersing the dynamic spectrum in Fig. 10.3.

With 72,556 sky pixels and 92,763 DM and t_0 trials per pixel, a dedispersion search over the entire GW170104 localization region yields 6.7×10^9 total trials. To determine a suitable S/N threshold for manual follow-up, in Fig. 10.5 we plot the cumulative distribution of signal-to-noise ratios obtained from a random subset of sky locations, dispersion measures, and initial times t_0. We find our signal-to-noise ratios to be fairly Gaussian distributed. The bulk of the distribution is well fit by a somewhat broadened Gaussian centered at zero with a

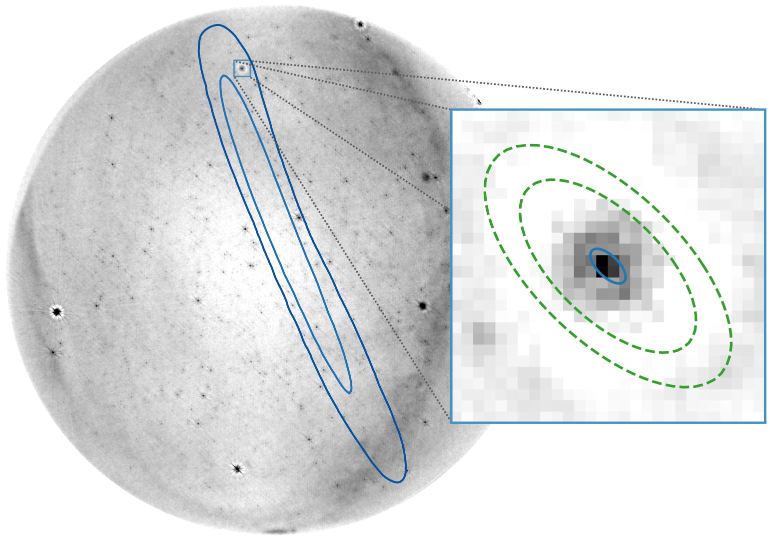

Figure 10.6: Full-band dirty image of an example meteor reflection event. Within the inset, the blue contour gives the OVRO-LWA's synthesized beam and the dashed green contours mark the annulus used for background estimation. Meteor reflections occur in-atmosphere and so appear as resolved sources.

variance of 1.44. At high significances, however, Fig. 10.5 shows the emergence of a non-Gaussian tail. This tail is dominated by meteor reflection events, in which patches of atmosphere temporarily ionized by passing meteors act as reflective surfaces, redirecting RFI from beyond the horizon into the OVRO-LWA (see more below). The tail is well-fit by $\log_{10} P = a\,(\mathrm{S/N})+b$, with $a = -0.136$ and $b = -2.913$. Using this fit, we choose our threshold for manual inspection to be $\mathrm{S/N} = 20$, above which we expect $(6.7 \times 10^9)\,10^{\,a(\mathrm{S/N})+b} \approx 1.5 \times 10^4$ outliers.

After searching across the entire GW170104 localization region, we find 6,828 outliers exceeding our threshold. This suggests that extrapolation of the subset of data shown in Fig. 10.5 overestimates the rate of high significance events by a factor of two. All candidates warranting manual follow-up are identified as meteor reflection events [404, 405]. Figures 10.6 and 10.7 illustrate the properties of a typical reflection event. First, meteor reflections occur within the atmosphere (well inside the array's $2D^2/\lambda \sim 1000\,\mathrm{km}$ far-field limit, where D is the array size and λ a characteristic wavelength) and hence appear as resolved sources. Second, their spectra show emission confined to one or more broadcast television channels. The reflection event in Figs. 10.6 and 10.7, for

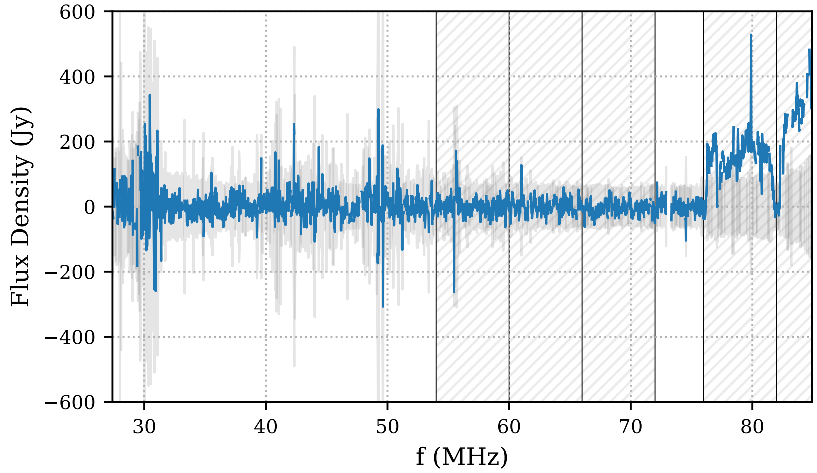

Figure 10.7: The spectrum of the meteor reflection event in Fig. 10.6. As is typical, the observed emission from this event is confined exactly to broadcast television channels 5 (76–82 MHz) and 6 (82–88 MHz), and so is readily identifiable as terrestrial in origin.

instance, is confined to channels 5 (76–82 MHz) and 6 (82–88 MHz). As meteor reflections currently dominate our search background, the automated identification and rejection of meteor reflections will be a crucial step in improving the sensitivity of future searches.

10.4 Radio Luminosity Limits

Having rejected all outliers as reflection events, we place upper limits on the prompt radio emission associated with GW170104. Let's symbolically represent our radio data via $\mathbf{d}$. Our first goal is to compute the posterior $p(F|\mathbf{d},\hat{\Omega})$ on the radio flux of GW170104 for every possible progenitor sky location $\hat{\Omega}$. This posterior is obtained by marginalizing over the parameters of our signal model, namely dispersion measure DM and initial signal time t_0:

$$p(F|\mathbf{d},\hat{\Omega}) = \int d\mathrm{DM} \int dt_0\, p(F|\mathbf{d},\hat{\Omega},\mathrm{DM},t_0)p(\mathrm{DM})p(t_0), \tag{10.8}$$

where $p(\mathrm{DM})$ and $p(t_0)$ are our prior probabilities on a signal's dispersion measure and initial time. Next, using Bayes' theorem, we can relate the posterior $p(F|\mathbf{d},\hat{\Omega},\mathrm{DM},t_0)$ to the likelihood $p(\mathbf{d}|F,\hat{\Omega},\mathrm{DM},t_0)$ of having measured $\mathbf{d}$:

$$p(F|\mathbf{d},\hat{\Omega}) \propto \int d\mathrm{DM} \int dt_0\, p(\mathbf{d}|F,\hat{\Omega},\mathrm{DM},t_0)p(F)p(\mathrm{DM})p(t_0), \tag{10.9}$$

where $p(F)$ is our flux density prior.

We will assume Gaussian likelihoods, centered at $\hat{F}(\hat{\Omega}, \mathrm{DM}, t_0)$ with a variance $\sigma^2(\hat{\Omega}, \mathrm{DM}, t_0)$ as defined by Eqs. (10.4) and (10.6). Meanwhile, for simplicity we assume flat priors over the ranges $\Delta\mathrm{DM}$ and Δt_0 considered: $p(\mathrm{DM}) = 1/\Delta\mathrm{DM}$ and $p(t_0) = 1/\Delta t_0$. We similarly assume a uniform (improper) prior over all positive F. All together,

$$p(F|\mathbf{d}, \hat{\Omega}) = \int \frac{d\mathrm{DM}}{\Delta\mathrm{DM}} \int \frac{dt_0}{\Delta t_0} \frac{1}{\mathcal{N}(\hat{\Omega}, \mathrm{DM}, t_0)} \exp\left(-\frac{\left[\hat{F}(\hat{\Omega}, \mathrm{DM}, t_0) - F\right]^2}{2\sigma^2(\hat{\Omega}, \mathrm{DM}, t_0)}\right), \tag{10.10}$$

with normalization factor

$$\mathcal{N}(\hat{\Omega}, \mathrm{DM}, t_0) = \int_0^\infty dF \exp\left(-\frac{\left[\hat{F}(\hat{\Omega}, \mathrm{DM}, t_0) - F\right]^2}{2\sigma^2(\hat{\Omega}, \mathrm{DM}, t_0)}\right). \tag{10.11}$$

With the flux posterior $p(F|\mathbf{d}, \hat{\Omega})$ in hand, the 95% credible flux upper limit in direction $\hat{\Omega}$ corresponds to the flux F_{95} satisfying

$$0.95 = \int_0^{F_{95}} dF p(F|\mathbf{d}, \hat{\Omega}). \tag{10.12}$$

Figure 10.8 shows these 95% credible flux upper limits for each image pixel within the 95% credible gravitational-wave localization region. We exclude pixels containing persistent point sources detected at 5σ prior to dedispersion, yielding the "holes" seen in Fig. 10.8. We additionally trim the southernmost points that set below the OVRO-LWA's horizon during the observation. All together, we cover 94% of the localization region contained within the OVRO-LWA's field of view, and 54% of GW170104's global probability map. We achieve a median upper limit of 2.4 Jy. Our sensitivity is degraded at low elevations due to the $(\sin\theta)^{1.6}$ scaling of the antennas' primary beam with elevation angle θ [406]. Flux upper limits are also impacted by sidelobes in the vicinity of particularly bright point sources.

With the sky and distance localization provided by Advanced LIGO, we can re-express our flux limits as constraints on the equivalent isotropic radio luminosity of GW170104. To do this, we need to compute the posterior $p(L|\mathbf{d})$ on the equivalent isotropic luminosity of GW170104, marginalized over all possible progenitor sky locations $\hat{\Omega}$ and distances D:

$$p(L|\mathbf{d}) = \int dD \int d\hat{\Omega}\, p(L|\mathbf{d}, D, \hat{\Omega}) p(D, \hat{\Omega}). \tag{10.13}$$

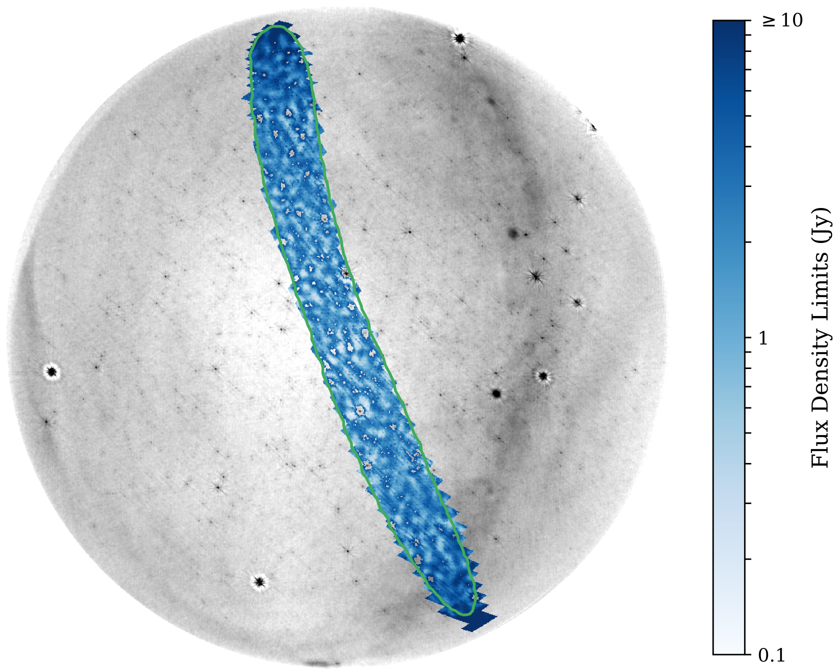

Figure 10.8: 95% credible upper limits on the flux density of prompt radio emission from GW170104, as a function of its presumed sky location. The "holes" mark locations of persistent point sources excluded from our analysis. For reference, the contour traces the 95% credible localization of GW170104 within the OVRO-LWA's field of view. Our median upper limit across the sky is 2.4 Jy. Marginalizing over the sky location and distance constraints due to the gravitational-wave signal, we limit GW170104's equivalent isotropic luminosity between 27 and 84 MHz to $L \leq 2.5 \times 10^{41}\,\mathrm{erg\,s^{-1}}$ at 95% credibility.

Here, $p(D, \hat{\Omega})$ is the probability distribution on the progenitor location of GW170104; we take this to be the localization provided by Advanced LIGO. As in Eq. (10.12) above, the 95% credible upper limit is given by the luminosity L_{95} satisfying $0.95 = \int_0^{L_{95}} p(L|\mathbf{d})dL$, or

$$0.95 = \int dD \int d\hat{\Omega} \int_0^{L_{95}} dL\, p(L|\mathbf{d}, D, \hat{\Omega}) p(D, \hat{\Omega}). \tag{10.14}$$

As currently written, this equation requires the probabilities $p(L|\mathbf{d}, D, \hat{\Omega})$ on luminosity as a function of direction and distance. We can recast Eq. (10.14) in terms of our known flux posteriors $p(F|\mathbf{d}, \hat{\Omega})$ [Eq. (10.10)] by substituting $p(L|\mathbf{d}) = p(F|\mathbf{d})\, dF/dL = p(F|\mathbf{d})/4\pi D^2$ and $dL = 4\pi D^2 dF$, giving

$$0.95 = \int dD \int d\hat{\Omega} \int_0^{F(D, L_{95})} dF\, p(F|\mathbf{d}, \hat{\Omega}) p(D, \hat{\Omega}). \tag{10.15}$$

In practice, Eq. (10.15) is somewhat easier to evaluate when rearranged as

$$0.95 = \int d\hat{\Omega}\, p(\hat{\Omega}) \int dD\, p(D|\hat{\Omega}) \int_0^{F(D,L_{95})} dF\, p(F|\mathbf{d},\hat{\Omega}), \tag{10.16}$$

evaluating the flux and distance integrals for each sky location searched.

After performing this marginalization over all sky locations and distances, we obtain a limit on GW170104's equivalent isotropic luminosity between 27 and 84 MHz if $L_{\rm radio} \leq 2.5\times10^{41}\,{\rm erg\,s^{-1}}$ at 95% credibility, assuming the source lies within the OVRO-LWA's field of view. The total energy radiated by GW170104 was $E_{\rm GW} = 2.0^{+0.6}_{-0.7} M_\odot c^2$ [42]. We therefore limit the fraction of the total energy converted to prompt radio emission to $L_{\rm radio} t_{\rm int}/E_{\rm GW} \leq 1.4\times10^{-12}$, using the lower bound on $E_{\rm GW}$.

For reference, the equivalent isotropic luminosity of the *Fermi-GBM* outlier associated with GW150914 was $1.8^{+1.5}_{-1.0}\times10^{49}\,{\rm erg\,s^{-1}}$ [380]. Note that GW150914 and GW170104 occurred at luminosity distances of approximately 410 Mpc and 880 Mpc, respectively [25, 42]. If the OVRO-LWA had been operating at the time of GW150914, we would therefore have been sensitive to any associated radio transient with luminosity $2.5\times10^{41}\,{\rm erg\,s^{-1}}\,(410\,{\rm Mpc}/880\,{\rm Mpc})^2 \approx 5.5\times10^{40}\,{\rm erg\,s^{-1}}$. Hence in the future, if additional gamma-ray outliers are identified in coincidence with gravitational-wave events, simultaneous observations with the OVRO-LWA will limit the ratio of radio and gamma-ray luminosities to $\lesssim 3\times10^{-9}$.

Similar limits will be possible for future binary neutron star mergers. GW170817 occurred at a distance of 40.7 Mpc [407]. Rescaling our flux limit to 40.7 Mpc, OVRO-LWA follow-up of binary neutron stars at this distance will yield luminosity limits of approximately $5\times10^{38}\,{\rm erg\,s^{-1}}$. The equivalent isotropic luminosity of GRB 170817A was estimated to be $1.6\times10^{47}\,{\rm erg\,s^{-1}}$ (with total energy 3.1×10^{46} erg) in the 1 keV-10 MeV band [53]. The limits attainable with the OVRO-LWA would therefore limit the ratio of radio and gamma-ray luminosities to $\lesssim 3\times10^{-9}$, and the ratio of total radiated energies to $\lesssim 2\times10^{-7}$.

10.5 The Third LIGO/Virgo Observing Run and Beyond

Advanced LIGO & Virgo's third observing run (O3) began in April 2019 and is scheduled to run for one calendar year. During this time, between 1-50 binary neutron star detections are expected [38]. The OVRO-LWA will operate in continuous buffering mode during O3, searching for prompt radio transients associated with compact binary mergers.

The sensitivity of this study to sub-second radio transients is limited by the 13 s resolution of buffered visibilities [356]. The buffering of future data with higher time resolution will increase the signal-to-noise ratio of temporally unresolved transients. We are additionally exploring options to buffer the incoherent sum of antenna powers at their raw 197 MHz sampling rate. The incoherent sum will provide no directional information, but the vastly increased time resolution and temporal coincidence with gravitational-wave events will enable sensitive measurements of prompt radio transients.

A more ambitious goal is the buffering and coherent dedispersion of all 512 signal paths at 197 MHz. This endeavor has previously required prohibitively large buffer disk space due to significant latency in the release of LIGO/Virgo alerts. In their upcoming O3 observing run, however, LIGO & Virgo will transition to automated alerts released within 1-10 minutes of a gravitational-wave candidate [408]. If successful, this reduced latency may make the buffering of raw antenna voltages computationally feasible.

Finally, the OVRO-LWA will soon be undergoing upgrades towards its "Stage 3" design, consisting of 352 correlated antennas over an extended 2.5 km maximum baseline. Also included in this design is the buffering of raw antenna voltages, allowing high time-resolution searches triggered by automated LIGO & Virgo alerts. With these improvements, Stage 3 OVRO-LWA promises to enable even more sensitive detection and precise localization of prompt radio emission from compact binary mergers.

Chapter 11
The Next Steps

I consider myself unreasonably fortunate in the timing of my graduate studies. These early days of practical gravitational-wave astronomy have been truly exhilarating, and even more inspiring to witness as a graduate student in the field. It is an incredibly dynamic time – we learn something new seemingly every day, and facts are revised almost as soon as they can be written down. As of August 13, 2019, while I write this final chapter, LIGO and Virgo's current O3 observing run has yielded 22 public gravitational-wave candidates; the coming days (or even the coming hours) will surely bring even more. And, as I wrote in Ch. 1, gravitational-wave astronomy is in the enviable stage where our open questions far outnumber our available answers.

In this thesis I explored just two of these many questions:

First, I asked what present-day gravitational-wave experiments might learn from observation of the stochastic gravitational-wave background. In Ch. 4, I quantified which redshifts dominate the astrophysical gravitational-wave background accessible to Advanced LIGO, and identified which spectral features of the background we might hope to constrain. Chapter 5, meanwhile, demonstrated that the gravitational-wave background may be utilized as a tool with which to study fundamental physics, enabling direct measurements of the polarizations of gravitational waves. Having studied the *interpretation* of a stochastic gravitational-wave signal, in Ch. 6 we turned to the actual problem of *detection*. In this chapter I presented a search for the stochastic gravitational-wave background with Advanced LIGO. Although we did not uncover a clear signal, we were able to improve upon previous upper limits by two orders of magnitude. Then, in Ch. 7, I asked how one might *confirm* an apparent detection of the gravitational-wave background, developing a method to differentiate a true astrophysical signal from spurious terrestrial noise.

Second, I investigated whether the compact binary mergers detected with Advanced LIGO might also give rise to theorized prompt radio emission. Chapter 8 first reviewed the myriad models predicting the emission of a prompt radio precursor (or immediate "post-cursor") to stellar-mass compact binary mergers. In Ch. 9, I then asked if prompt radio emission might *already* have been detected in the form of fast radio bursts (FRBs). While the properties of FRBs are remarkably similar to those predicted of prompt radio emission, we found that the rate of FRBs is vastly greater than the rate of compact binary mergers, largely ruling out an association between the two. Then, in Ch. 10, I presented a direct search for prompt radio emission, following up the binary black hole GW170104. While I (perhaps unsurprisingly) detected no statistically significant radio emission from GW170104, this study served to demonstrate the feasibility of prompt radio follow-up, illustrating that we can place sensitive limits on radio transients arriving in coincidence with or even *before* a gravitational-wave signal.

Although we've made significant headway in exploring both of these questions, we are certainly far from finished. With respect to the stochastic gravitational-wave background, I believe it will be crucial to continue the work begun in Ch. 7. As we saw there, there exist known terrestrial effects – Schumann resonances, correlated combs, and perhaps others – that can mimic a gravitational-wave signal. While we have more or less known for decades how to *detect* a stochastic gravitational-wave signal, until now virtually no work has been done to explore how to *differentiate* a gravitational-wave background from one of these terrestrial effects. While I proposed one technique in Ch. 7, there may certainly be others that are more effective. Whichever technique(s) are ultimately used, once a detection is indeed made of the gravitational-wave background we will need thoughtful and well-developed methods for convincing ourselves (and the broader scientific community) that the detection is legitimate.

Detection (and validation) of the gravitational-wave background, of course, is only a means to and end. That end is to learn about compact binary mergers at cosmological distances. Practically speaking, the LIGO/Virgo community has a great deal of infrastructure and methodology devoted to the prediction and detection of the stochastic background. Relatively little attention, in contrast, has been paid to *inference* on a stochastic signal – using the detection of a gravitational-wave background to test concrete properties of distant compact binaries. In the near term, there is a vast amount of work to be done exploring this latter question, particularly since the first detection of the stochastic background may soon be drawing near. This work will be hard. In Ch. 4, we showed that the current generation of gravitational-wave detectors can generally only hope to measure the background's *amplitude*, not its shape. Nevertheless, the amplitude of the gravitational-wave background contains, indirectly at least, a great deal of information that is complementary to studies of individually-resolvable signals in the local Universe. One of the most fruitful avenues of research, I believe, will be the development of ways to synthesize these two sources of information – measurements of the stochastic background and observations of individual, nearby events – to understand how the properties of compact binary mergers evolve as we look back in time.

The search for prompt radio emission also has a long list of clear next steps. Using the OVRO-LWA, we now have radio data coincident with eight gravitational-wave candidates from O3, including one high-significance binary neutron star candidate. This data is currently undergoing analysis. In parallel, we should explore ways to further optimize our analysis. As discussed in Ch. 10, the binary black hole GW170104 was localized to 1,600 sq. degrees on the sky; at the time we considered this to represent a "worst-case scenario." We were, invariably, proven wrong. The binary neutron star candidate S190425z has a localization region spanning 10,000 sq. degrees, of which approximately 60% is covered by the OVRO-LWA's field of view. Poorly localized events like S190425z currently represent a significant computational challenge, requiring both a vast amount of storage and long computing times. We would therefore benefit immensely from any methods allowing us to search faster over larger localization regions.

Clearly, the most exciting outcome of the OVRO-LWA follow-up efforts would be the unambiguous detection of a prompt radio counterpart. Pragmatically, a perhaps more likely outcome is a set of non-detections associated with a population of binary black hole and neutron star candidates. Given a sufficiently large set of non-detections, though, we will be able to begin placing meaningful constraints on the properties of prompt radio emission – its beaming angle, energetics, etc. It will be valuable to anticipate this endeavor, learning how to account for non-trivial selection effects and identifying specific emission models that we may be able to rule out.

Finally, as we saw in Ch. 10, the sensitivity of current OVRO-LWA follow-up is limited by our 13 s integration times. If the duration of prompt radio emission is much less than a second, then we would gain considerable sensitivity by moving to shorter integrations. The ultimate goal is to do away with integrations altogether – to instead buffer and search the raw voltages output by the radio antennas. In addition to specialized hardware and a vast amount of computing memory, this daunting task will require us to develop and deploy methods with which to coherently dedisperse and search these raw voltages in real time.

Bibliography

[1] J. Weber, *Detection and Generation of Gravitational Waves*, Phys. Rev. **117**, 306 (1960).

[2] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, *Observation of a Rapidly Pulsating Radio Source*, Nature **217**, 709 (1968).

[3] J. Weber, *Gravitational Radiation from the Pulsars*, Phys. Rev. Lett. **21**, 395 (1968).

[4] J. J. Giganti, J. V. Larson, J. P. Richard, and J. Weber, NASA Report No. SP-330, 1973 (unpublished).

[5] J. Weber, *Evidence for Discovery of Gravitational Radiation*, Physical Review Letters **22**, 1320 (1969).

[6] J. Weber, *Anisotropy and Polarization in the Gravitational-Radiation Experiments*, Physical Review Letters **25**, 180 (1970).

[7] R. Weiss, Quarterly Progress Report 105: Electromagnetically coupled broadband gravitational antenna, 1972.

[8] Interview with Ronald W. P. Drever (*Caltech Oral Histories*), 1997, http://resolver.caltech.edu/CaltechOH:OH_Drever_R.

[9] Interview with Rainer Weiss (*Caltech Oral Histories*), 2002, http://resolver.caltech.edu/CaltechOH:OH_Weiss_R.

[10] Interview with Barry C. Barish (*Caltech Oral Histories*), 2001, http://resolver.caltech.edu/CaltechOH:OH_Barish_B.

[11] Interview with Stanley E. Whitcomb on LIGO (*Caltech Oral Histories*), 1997, http://resolver.caltech.edu/CaltechOH:OH_Whitcomb_S.

[12] Interview with Stanley E. Whitcomb (*Caltech Oral Histories*), 2017, http://resolver.caltech.edu/CaltechOH:OH_Whitcomb_S_Ligo_2017.

[13] R. Vogt, E., R. Drever, W. P., F. J. Raab, K. S. Thorne, and R. Weiss, *A Laser Interferometer Gravitational-Wave Observatory (LIGO)*, Proposal to the National Science Foundation (1989), https://dcc.ligo.org/LIGO-M890001/public/.

[14] A. Abramovici *et al.*, *LIGO: The Laser Interferometer Gravitational-Wave Observatory*, Science **256**, 325 (1992).

[15] LIGO Scientific Collaboration, B. Abbott *et al.*, *Detector description and performance for the first coincidence observations between LIGO and GEO*, Nucl. Instrum. Meth. **A517**, 154 (2004).

[16] LIGO Scientific Collaboration, B. Abbott *et al.*, *Analysis of LIGO data for gravitational waves from binary neutron stars*, Phys. Rev. D **69**, 122001 (2004).

[17] LIGO Scientific Collaboration, B. Abbott *et al.*, *Analysis of first LIGO science data for stochastic gravitational waves*, Phys. Rev. D **69**, 122004 (2004).

[18] LIGO Scientific Collaboration, B. Abbott *et al.*, *First upper limits from LIGO on gravitational wave bursts*, Phys. Rev. D **69**, 102001 (2004).

[19] LIGO Scientific Collaboration, B. Abbott *et al.*, *Setting upper limits on the strength of periodic gravitational waves from PSR* J1939 + 2134 *using the first science data from the GEO 600 and LIGO detectors*, Phys. Rev. D **69**, 082004 (2004).

[20] LIGO Scientific, B. P. Abbott *et al.*, *LIGO: The Laser interferometer gravitational-wave observatory*, Rept. Prog. Phys. **72**, 076901 (2009).

[21] VIRGO, T. Accadia *et al.*, *Virgo: a laser interferometer to detect gravitational waves*, JINST **7**, P03012 (2012).

[22] LIGO Scientific Collaboration, G. M. Harry *et al.*, *Advanced LIGO: the next generation of gravitational wave detectors*, Class. Quant. Grav. **27**, 084006 (2010).

[23] LIGO Scientific Collaboration, J. Aasi *et al.*, *Advanced LIGO*, Class. Quant. Grav. **32**, 074001 (2015).

[24] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW150914: The Advanced LIGO Detectors in the Era of First Discoveries*, Phys. Rev. Lett. **116**, 131103 (2016).

[25] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Observation of Gravitational Waves from a Binary Black Hole Merger*, Phys. Rev. Lett. **116**, 061102 (2016).

[26] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW150914: Implications for the Stochastic Gravitational-Wave Background from Binary Black Holes*, Phys. Rev. Lett. **116**, 131102 (2016).

[27] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914*, Astrophys. J. **833**, L1 (2016).

[28] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Properties of the Binary Black Hole Merger GW150914*, Phys. Rev. Lett. **116**, 241102 (2016).

[29] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Astrophysical implications of the Binary Black Hole Merger GW150914*, Astrophys. J. **818**, L22 (2016).

[30] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Directly comparing GW150914 with numerical solutions of Einstein's equations for binary black hole coalescence*, Phys. Rev. D **94**, 064035 (2016).

[31] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW150914: First results from the search for binary black hole coalescence with Advanced LIGO*, Phys. Rev. D **93**, 1 (2016).

[32] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Observing gravitational-wave transient GW150914 with minimal assumptions*, Phys. Rev. D **93**, 122004 (2016).

[33] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Tests of General Relativity with GW150914*, Phys. Rev. Lett. **116**, 221101 (2016).

[34] Virgo Collaboration, F. Acernese *et al.*, *Advanced Virgo: a second-generation interferometric gravitational wave detector*, Class. Quant. Grav. **32**, 024001 (2015).

[35] Y. Aso *et al.*, *Interferometer design of the KAGRA gravitational wave detector*, Phys. Rev. D **88**, 043007 (2013).

[36] T. Akutsu *et al.*, *Construction of KAGRA: an underground gravitational-wave observatory*, Progress of Theoretical and Experimental Physics **2018** (2018).

[37] KAGRA Collaboration, T. Akutsu *et al.*, *First cryogenic test operation of underground km-scale gravitational-wave observatory KAGRA*, Preprint (2019), arXiv:1901.03569.

[38] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA*, Living Reviews in Relativity **21**, 1 (2018).

[39] B. Iyer *et al.*, *Proposal of the Consortium for Indian Initiative in Gravitational-wave Observations (IndIGO)*, LIGO-India Technical Report No. LIGO-M1100296 (2011).

[40] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence*, Phys. Rev. Lett. **116**, 241103 (2016).

[41] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Binary Black Hole Mergers in the First Advanced LIGO Observing Run*, Phys. Rev. X **6**, 041015 (2016).

[42] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2*, Phys. Rev. Lett. **118**, 221101 (2017).

[43] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170608: Observation of a 19 Solar-mass Binary Black Hole Coalescence*, Astrophys. J. **851**, L35 (2017).

[44] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence*, Phys. Rev. Lett. **119**, 141101 (2017).

[45] The LIGO Scientific Collaboration and the Virgo Collaboration, B. P. Abbott *et al.*, *GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs*, Preprint (2018), arXiv:1811.12907.

[46] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, Phys. Rev. Lett. **119**, 161101 (2017).

[47] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Properties of the Binary Neutron Star Merger GW170817*, Phys. Rev. X **9** (2019).

[48] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences*, Phys. Rev. Lett. **120**, 091101 (2018).

[49] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170817: Measurements of Neutron Star Radii and Equation of State*, Phys. Rev. Lett. **121**, 161101 (2018).

[50] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *On the Progenitor of Binary Neutron Star Merger GW170817*, Astrophys. J. **850**, L40 (2017).

[51] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Constraining the p-Mode–g-Mode Tidal Instability with GW170817*, Phys. Rev. Lett. **122**, 061104 (2019).

[52] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Tests of General Relativity with GW170817*, Preprint (2018), arXiv:1811.00364.

[53] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A*, Astrophys. J. **848**, L13 (2017).

[54] LIGO Scientific Collaboration and Virgo Collaboration and others, B. P. Abbott *et al.*, *Multi-messenger Observations of a Binary Neutron Star Merger*, Astrophys. J. **848**, L12 (2017).

[55] D. Sobral *et al.*, *A large Hα survey at z = 2.23, 1.47, 0.84 and 0.40: the 11 Gyr evolution of star-forming galaxies from HiZELS*, Mon. Not. R. Astron. Soc. **428**, 1128 (2013).

[56] A. G. Riess *et al.*, *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant*, Astrophys. J. **116**, 1009 (1998).

[57] C. Will, *Theory and Experiment in Gravitational Physics*, 2nd edition ed. (Cambridge University Press, 2018).

[58] M. Isi and A. J. Weinstein, *Probing gravitational wave polarizations with signals from compact binary coalescences*, Preprint (2017), arXiv:1710.03794.

[59] A. Goldstein *et al.*, *An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi -GBM Detection of GRB 170817A*, Astrophys. J. **848**, L14 (2017).

[60] D. A. Coulter *et al.*, *Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source*, Science **358**, 1556 (2017).

[61] M. Soares-Santos *et al.*, *The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. I. Discovery of the Optical Counterpart Using the Dark Energy Camera*, Astrophys. J. **848**, L16 (2017).

[62] G. Hallinan *et al.*, *A radio counterpart to a neutron star merger*, Science **358**, 1579 (2017).

[63] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Estimating the Contribution of Dynamical Ejecta in the Kilonova Associated with GW170817*, Astrophys. J. **850**, L39 (2017).

[64] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *A gravitational-wave standard siren measurement of the Hubble constant*, Nature **551**, 85 (2017).

[65] E. E. Flanagan and S. A. Hughes, *The basics of gravitational wave theory*, New Journal of Physics **7** (2005).

[66] S. Carroll, *Spacetime and Geometry: An Introduction to General Relativity* (Addison Wesley, 2004).

[67] M. Maggiore, *Gravitational Waves: Volume 1: Theory and Experiments* (Oxford University Press, 2008).

[68] T. Moore, *A General Relativity Workbook* (University Science Books, 2012).

[69] P. R. Saulson, *If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves?*, Am. J. Phys. **65**, 501 (1997).

[70] P. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors*, 2nd ed. (World Scientific, 2017).

[71] B. F. Schutz, *Gravitational waves on the back of an envelope*, Am. J. Phys. **52**, 412 (1984).

[72] S. Bonazzola and E. Gourgoulhon, *Gravitational waves from pulsars: emission by the magnetic field induced distortion*, Preprint (1996), arXiv:astro-ph/9602107.

[73] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run*, Preprint (2019), arXiv:1902.08442.

[74] The LIGO Scientific Collaboration and the Virgo Collaboration, B. P. Abbott *et al.*, *Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data*, Preprint (2019), arXiv:1902.08507.

[75] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data*, Preprint (2019), arXiv:1903.01901.

[76] S. W. Hawking and W. Israel, *General Relativity: An Einstein Centenary Survey* (Cambridge University Press, 1979).

[77] V. Morozova, D. Radice, A. Burrows, and D. Vartanyan, *The Gravitational Wave Signal from Core-collapse Supernovae*, Astrophys. J. **861**, 10 (2018).

[78] J. Powell and B. Müller, *Gravitational Wave Emission from 3D Explosion Models of Core-Collapse Supernovae with Low and Normal Explosion Energies*, Mon. Not. R. Astron. Soc. **487**, 1250 (2019).

[79] D. Radice, V. Morozova, A. Burrows, D. Vartanyan, and H. Nagakura, *Characterizing the Gravitational Wave Signal from Core-collapse Supernovae*, Astrophys. J. **876**, L9 (2019).

[80] D. V. Martynov *et al.*, *The Sensitivity of the Advanced LIGO Detectors at the Beginning of Gravitational Wave Astronomy*, Phys. Rev. D **93**, 112004 (2016).

[81] P. B. Covas *et al.*, *Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO*, Phys. Rev. D **97** (2018).

[82] P. Madau and M. Dickinson, *Cosmic Star-Formation History*, Annu. Rev. Astron. Astrophys. **52**, 415 (2014).

[83] B. Allen and J. D. Romano, *Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities*, Phys. Rev. D **59**, 102001 (1999).

[84] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *A search for the isotropic stochastic background using data from Advanced LIGO's second observing run*, Preprint (2019), arXiv:1903.02886.

[85] S. Crowder, R. Namba, V. Mandic, S. Mukohyama, and M. Peloso, *Measurement of parity violation in the early universe using gravitational-wave detectors*, Phys. Lett. B **726**, 66 (2013).

[86] R. Jackiw and S.-Y. Pi, *Chern-Simons modification of general relativity*, Phys. Rev. D **68**, 104012 (2003).

[87] S. Alexander and J. Martin, *Birefringent gravitational waves and the consistency check of inflation*, Phys. Rev. D **71**, 063526 (2005).

[88] S. Alexander, L. S. Finn, and N. Yunes, *Gravitational-wave probe of effective quantum gravity*, Phys. Rev. D **78**, 066005 (2008).

[89] S. Alexander and N. Yunes, *Chern-Simons modified general relativity*, Phys. Rep. **480**, 1 (2009).

[90] K. Yagi and H. Yang, *Probing gravitational parity violation with gravitational waves from stellar-mass black hole binaries*, Phys. Rev. D **97**, 104018 (2018).

[91] Planck Collaboration, R. Adam *et al.*, *Planck 2015 results. I. Overview of products and scientific results*, Astronomy & Astrophysics **594**, A1 (2016).

[92] N. Mazumder, S. Mitra, and S. Dhurandhar, *Astrophysical motivation for directed searches for a stochastic gravitational wave background*, Phys. Rev. D **89**, 084076 (2014).

[93] G. Cusin, C. Pitrou, and J.-P. Uzan, *Anisotropy of the astrophysical gravitational wave background: Analytic expression of the angular power spectrum and correlation with cosmological observations*, Phys. Rev. D **96** (2017).

[94] G. Cusin, C. Pitrou, and J.-P. Uzan, *The signal of the gravitational wave background and the angular correlation of its energy density*, Phys. Rev. D **97** (2018).

[95] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, *First Predictions of the Angular Power Spectrum of the Astrophysical Gravitational Wave Background*, Phys. Rev. Lett. **120** (2018).

[96] A. C. Jenkins, M. Sakellariadou, T. Regimbau, and E. Slezak, *Anisotropies in the astrophysical gravitational-wave background: Predictions for the detection of compact binaries by LIGO and Virgo*, Phys. Rev. D **98** (2018).

[97] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, *Properties of the stochastic astrophysical gravitational wave background: astrophysical sources dependencies*, Preprint (2019), arXiv:1904.07797.

[98] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, *Comment on the article "Anisotropies in the astrophysical gravitational-wave background: The impact of black hole distributions" by A.C. Jenkins et al. [arXiv:1810.13435]*, Preprint (2018), arXiv:1811.03582.

[99] A. C. Jenkins *et al.*, *Response to Cusin et al's comment on arXiv:1810.13435*, Preprint (2019), arXiv:1901.01078.

[100] D. Meacher, E. Thrane, and T. Regimbau, *Statistical Properties of Astrophysical Gravitational-Wave Backgrounds*, Phys. Rev. D **89**, 084063 (2014).

[101] A. C. Jenkins and M. Sakellariadou, *Shot noise in the astrophysical gravitational-wave background*, Preprint (2019), arXiv:1902.07719.

[102] S. Drasco and E. E. Flanagan, *Detection methods for non-Gaussian gravitational wave stochastic backgrounds*, Phys. Rev. D **67**, 082003 (2003).

[103] E. Thrane, *Measuring the non-Gaussian stochastic gravitational-wave background: a method for realistic interferometer data*, Phys. Rev. D **87**, 043009 (2013).

[104] L. Martellini and T. Regimbau, *Semiparametric approach to the detection of non-Gaussian gravitational wave stochastic backgrounds*, Phys. Rev. D **89**, 124009 (2014).

[105] L. Martellini and T. Regimbau, *Efficiency of the cross-correlation statistic for gravitational wave stochastic background signals with non-Gaussian noise and heterogeneous detector sensitivities*, Phys. Rev. D **92**, 104025 (2015).

[106] N. J. Cornish and J. D. Romano, *When is a gravitational-wave signal stochastic?*, Phys. Rev. D **92**, 042001 (2015).

[107] R. Smith and E. Thrane, *Optimal Search for an Astrophysical Gravitational-Wave Background*, Phys. Rev. X **8**, 021019 (2018).

[108] F. Hernandez Vivanco, R. Smith, E. Thrane, and P. D. Lasky, *Accelerated detection of the binary neutron star gravitational-wave background*, Preprint (2019), arXiv:1903.05778.

[109] J. C. N. de Araujo and O. D. Miranda, *Star formation rate density and the stochastic background of gravitational waves*, Phys. Rev. D **71**, 127503 (2005).

[110] E. S. Phinney, *A practical theorem on gravitational wave backgrounds*, Preprint (2001), arXiv:0108028.

[111] The LIGO Scientific Collaboration and the Virgo Collaboration, B. P. Abbott *et al.*, *Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo*, Preprint (2018), arXiv:1811.12940.

[112] T. Regimbau and V. Mandic, *Astrophysical Sources of a Stochastic Gravitational-Wave Background*, Class. Quant. Grav. **25**, 184018 (2008).

[113] P. A. Rosado, *Gravitational wave background from binary systems*, Phys. Rev. D **84**, 084004 (2011).

[114] X.-J. Zhu, E. Howell, T. Regimbau, D. Blair, and Z.-H. Zhu, *Stochastic Gravitational Wave Background from Coalescing Binary Black Holes*, Astrophys. J. **739**, 86 (2011).

[115] C.-J. Wu, V. Mandic, and T. Regimbau, *Accessibility of the Stochastic Gravitational Wave Background from Magnetars to the Interferometric Gravitational Wave Detectors*, Phys. Rev. D **87**, 042002 (2013).

[116] X.-J. Zhu, E. J. Howell, D. G. Blair, and Z.-H. Zhu, *On the gravitational wave background from compact binary coalescences in the band of ground-based interferometers*, Mon. Not. R. Astron. Soc. **431**, 882 (2013).

[117] E. Vangioni *et al.*, *The impact of star formation and gamma-ray burst rates at high redshift on cosmic chemical evolution and reionization*, Mon. Not. R. Astron. Soc. **447**, 2575 (2015).

[118] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Upper Limits on the Rates of Binary Neutron Star and Neutron Star–Black Hole Mergers From Advanced Ligo'S First Observing Run*, Astrophys. J. **832**, L21 (2016).

[119] V. Ferrari, S. Matarrese, and R. Schneider, *Stochastic background of gravitational waves generated by a cosmological population of young, rapidly rotating neutron stars*, Mon. Not. R. Astron. Soc. **303**, 258 (1999).

[120] T. Regimbau and J. A. de Freitas Pacheco, *Cosmic background of gravitational waves from rotating neutron stars*, Astronomy & Astrophysics **376**, 381 (2001).

[121] T. Regimbau and J. A. de Freitas Pacheco, *Gravitational wave background from magnetars*, Astronomy & Astrophysics **447**, 1 (2006).

[122] E. Howell, T. Regimbau, A. Corsi, D. Coward, and R. Burman, *Gravitational wave background from sub-luminous GRBs: prospects for second- and third-generation detectors: GW background from SL-GRBs*, Mon. Not. R. Astron. Soc. **410**, 2123 (2011).

[123] S. Marassi, R. Ciolfi, R. Schneider, L. Stella, and V. Ferrari, *Stochastic background of gravitational waves emitted by magnetars: Stochastic background of gravitational waves emitted by magnetars*, Mon. Not. R. Astron. Soc. **411**, 2549 (2011).

[124] X.-J. Zhu, X.-L. Fan, and Z.-H. Zhu, *Stochastic gravitational-wave background from neutron star r-mode instability revisited*, Astrophys. J. **729**, 59 (2011).

[125] P. A. Rosado, *Gravitational wave background from rotating neutron stars*, Phys. Rev. D **86**, 104007 (2012).

[126] L. Warszawski and A. Melatos, *Gravitational-wave bursts and stochastic background from superfluid vortex avalanches during pulsar glitches: Gravitational-waves bursts from vortex avalanches*, Mon. Not. R. Astron. Soc. **423**, 2058 (2012).

[127] C.-J. Wu, V. Mandic, and T. Regimbau, *Accessibility of the stochastic gravitational wave background from magnetars to the interferometric gravitational wave detectors*, Phys. Rev. D **87**, 042002 (2013).

[128] D. Talukder, E. Thrane, S. Bose, and T. Regimbau, *Measuring neutron-star ellipticity with measurements of the stochastic gravitational-wave background*, Phys. Rev. D **89**, 123008 (2014).

[129] Q. Cheng, Y.-W. Yu, and X.-P. Zheng, *Stochastic gravitational wave background from magnetic deformation of newly born magnetars*, Mon. Not. R. Astron. Soc. **454**, 2299 (2015).

[130] M. Surace, K. D. Kokkotas, and P. Pnigouras, *The stochastic background of gravitational waves due to the f-mode instability in neutron stars*, Astronomy & Astrophysics **586**, A86 (2016).

[131] Q. Cheng, S.-N. Zhang, and X.-P. Zheng, *Stochastic gravitational wave background from newly born massive magnetars: The role of a dense matter equation of state*, Phys. Rev. D **95**, 083003 (2017).

[132] D. Blair and L. Ju, *A cosmological background of gravitational waves produced by supernovae in the early Universe*, Mon. Not. R. Astron. Soc. **283**, 648 (1996).

[133] V. Ferrari, S. Matarrese, and R. Schneider, *Gravitational wave background from a cosmological population of core-collapse supernovae*, Mon. Not. R. Astron. Soc. **303**, 247 (1999).

[134] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Stochastic Background of Gravitational Waves*, Phys. Rev. D **61**, 124015 (2000).

[135] R. Schneider, A. Ferrara, B. Ciardi, V. Ferrari, and S. Matarrese, *Gravitational wave signals from the collapse of the first stars*, Mon. Not. R. Astron. Soc. **317**, 385 (2000).

[136] D. M. Coward, R. R. Burman, and D. G. Blair, *The stochastic background of gravitational waves from neutron star formation at cosmological distances*, Mon. Not. R. Astron. Soc. **324**, 1015 (2001).

[137] D. M. Coward, R. R. Burman, and D. G. Blair, *Simulating a stochastic background of gravitational waves from neutron star formation at cosmological distances*, Mon. Not. R. Astron. Soc. **329**, 411 (2002).

[138] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Gravitational wave background from Population III black hole formation*, Mon. Not. R. Astron. Soc. **330**, 651 (2002).

[139] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Background of gravitational waves from pre-galactic black hole formation*, Class. Quant. Grav. **19**, 1335 (2002).

[140] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Gravitational wave background from black holes of the Springel & Hernquist star formation epoch*, Mon. Not. R. Astron. Soc. **348**, 1373 (2004).

[141] J. C. N. de Araujo, O. D. Miranda, and O. D. Aguiar, *Gravitational wave background from stellar black holes in a scenario of structure formation*, Class. Quant. Grav. **21**, S545 (2004).

[142] E. Howell, D. Coward, R. Burman, D. Blair, and J. Gilmore, *The gravitational wave background from neutron star birth throughout the cosmos*, Mon. Not. R. Astron. Soc. **351**, 1237 (2004).

[143] A. Buonanno, G. Sigl, G. G. Raffelt, H.-T. Janka, and E. Müller, *Stochastic gravitational-wave background from cosmological supernovae*, Phys. Rev. D **72**, 084001 (2005).

[144] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni, *Gravitational waves from the first stars*, Phys. Rev. D **73**, 104024 (2006).

[145] S. Marassi, R. Schneider, and V. Ferrari, *Gravitational wave backgrounds and the cosmic transition from Population III to Population II stars*, Mon. Not. R. Astron. Soc. **398**, 293 (2009).

[146] X.-J. Zhu, E. Howell, and D. Blair, *Observational upper limits on the gravitational wave production of core collapse supernovae: Upper limits on the GW production of SNe*, Mon. Not. R. Astron. Soc. **409**, L132 (2010).

[147] K. Crocker *et al.*, *Model of the stochastic gravitational-wave background due to core collapse to black holes*, Phys. Rev. D **92**, 063005 (2015).

[148] F. Pacucci, A. Ferrara, and S. Marassi, *Gravitational waves from direct collapse black holes formation*, Mon. Not. R. Astron. Soc. **449**, 1076 (2015).

[149] I. Dvorkin, E. Vangioni, J. Silk, J.-P. Uzan, and K. A. Olive, *Metallicity-constrained merger rates of binary black holes and the stochastic gravitational wave background*, Mon. Not. R. Astron. Soc. **461**, 3877 (2016).

[150] K. Crocker *et al.*, *Systematic study of the stochastic gravitational-wave background due to stellar core collapse*, Phys. Rev. D **95**, 063015 (2017).

[151] L. A. Rosi and R. L. Zimmerman, *A Maximum Background Spectrum of Gravitational Radiation*, Astrophysics and Space Science **45**, 447 (1976).

[152] B. J. Carr, *Cosmological gravitational waves - Their origin and consequences*, Astronomy & Astrophysics **89**, 6 (1980).

[153] M. Zimmermann and K. S. Thorne, The Gravitational Waves That Bathe the Earth: Upper Limits Based on Theorists' Cherished Beliefs, in *Essays in General Relativity*, pp. 139–155, Elsevier, 1980.

[154] C. Caprini and D. G. Figueroa, *Cosmological backgrounds of gravitational waves*, Class. Quant. Grav. **35**, 163001 (2018).

[155] D. W. Hogg, *Distance measures in cosmology*, Preprint (1999), arXiv:astro-ph/9905116.

[156] M. Maggiore, *Gravitational Wave Experiments and Early Universe Cosmology*, Phys. Rep. **331**, 283 (2000).

[157] N. Christensen, *Stochastic gravitational wave backgrounds*, Reports on Progress in Physics **82**, 016903 (2019).

[158] A. A. Starobinsky, *Spectrum of relict gravitational radiation and the early state of the universe*, Journal of Experimental and Theoretical Physics Letters **30**, 682 (1979).

[159] R. Easther, J. T. Giblin, and E. A. Lim, *Gravitational Wave Production at the End of Inflation*, Phys. Rev. Lett. **99**, 221301 (2007).

[160] N. Barnaby, E. Pajer, and M. Peloso, *Gauge field production in axion inflation: Consequences for monodromy, non-Gaussianity in the CMB, and gravitational waves at interferometers*, Phys. Rev. D **85**, 023525 (2012).

[161] J. L. Cook and L. Sorbo, *Particle production during inflation and gravitational waves detectable by ground-based interferometers*, Phys. Rev. D **85**, 023534 (2012).

[162] E. Witten, *Cosmic separation of phases*, Phys. Rev. D **30**, 272 (1984).

[163] A. Kosowsky, M. S. Turner, and R. Watkins, *Gravitational radiation from colliding vacuum bubbles*, Phys. Rev. D **45**, 4514 (1992).

[164] A. Kosowsky and M. S. Turner, *Gravitational radiation from colliding vacuum bubbles: Envelope approximation to many-bubble collisions*, Phys. Rev. D **47**, 4372 (1993).

[165] C. Caprini, R. Durrer, and G. Servant, *Gravitational Wave Generation from Bubble Collisions in First-Order Phase Transitions: An Analytic Approach*, Phys. Rev. D **77**, 124015 (2008).

[166] M. Hindmarsh, S. J. Huber, K. Rummukainen, and D. J. Weir, *Gravitational Waves from the Sound of a First Order Phase Transition*, Phys. Rev. Lett. **112**, 041301 (2014).

[167] C. Caprini *et al.*, *Science with the Space-Based Interferometer eLISA. II: Gravitational Waves from Cosmological Phase Transitions*, J. Cosmol. Astropart. Phys. **2016**, 001 (2016).

[168] X. Siemens, V. Mandic, and J. Creighton, *Gravitational-Wave Stochastic Background from Cosmic Strings*, Phys. Rev. Lett. **98**, 111101 (2007).

[169] S. Ölmez, V. Mandic, and X. Siemens, *Gravitational-Wave Stochastic Background From Kinks and Cusps on Cosmic Strings*, Phys. Rev. D **81**, 1 (2010).

[170] T. Regimbau, S. Giampanis, X. Siemens, and V. Mandic, *Stochastic background from cosmic (super)strings: Popcorn-like and (Gaussian) continuous regimes*, Phys. Rev. D **85**, 066001 (2012).

[171] LIGO Scientific Collaboration and Virgo Collaboration, J. Aasi *et al.*, *Constraints on Cosmic Strings from the LIGO-Virgo Gravitational-Wave Detectors*, Phys. Rev. Lett. **112**, 131101 (2014).

[172] S. Henrot-Versillé *et al.*, *Improved constraint on the primordial gravitational-wave density using recent cosmological data and its impact on cosmic string models*, Class. Quant. Grav. **32**, 045003 (2015).

[173] Y. Cui, M. Lewicki, D. E. Morrissey, and J. D. Wells, *Probing the pre-BBN universe with gravitational waves from cosmic strings*, Preprint (2018), arXiv:1808.08968.

[174] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Constraints on cosmic strings using data from the first Advanced LIGO observing run*, Phys. Rev. D **97**, 102002 (2018).

[175] M. Isi and L. C. Stein, *Measuring stochastic gravitational-wave energy beyond general relativity*, Phys. Rev. D **98**, 104025 (2018).

[176] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run*, Phys. Rev. Lett. **118**, 121101 (2017).

[177] P. Ajith *et al.*, *A template bank for gravitational waveforms from coalescing binary black holes: non-spinning binaries*, Phys. Rev. D **77**, 104017 (2008).

[178] M. D. Kistler, H. Yuksel, and A. M. Hopkins, *The Cosmic Star Formation Rate from the Faintest Galaxies in the Unobservable Universe*, Preprint (2013), arXiv:1305.1630.

[179] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, *How Massive Single Stars End Their Life*, Astrophys. J. **591**, 288 (2003).

[180] K. Belczynski *et al.*, *The effect of metallicity on the detection prospects for gravitational waves*, Astrophys. J. **715**, L138 (2010).

[181] K. Belczynski *et al.*, *On the maximum mass of stellar black holes*, Astrophys. J. **714**, 1217 (2010).

[182] M. Mapelli, N. Giacobbo, F. Santoliquido, and M. C. Artale, *The properties of merging black holes and neutron stars across cosmic time*, Mon. Not. R. Astron. Soc. (2019).

[183] M. Dominik *et al.*, *Double Compact Objects I. The Significance of the Common Envelope on Merger Rates*, Astrophys. J. **759**, 52 (2012).

[184] E. Thrane and J. D. Romano, *Sensitivity curves for searches for gravitational-wave backgrounds*, Phys. Rev. D **88**, 124032 (2013).

[185] V. Mandic, E. Thrane, S. Giampanis, and T. Regimbau, *Parameter Estimation in Searches for the Stochastic Gravitational-Wave Background*, Phys. Rev. Lett. **109**, 171102 (2012).

[186] N. Christensen, *Measuring the stochastic gravitational-radiation background with laser-interferometric antennas*, Phys. Rev. D **46**, 5250 (1992).

[187] LIGO Scientific Collaboration and Virgo Collaboration, J. Aasi *et al.*, *Searching for stochastic gravitational waves using data from the two colocated LIGO Hanford detectors*, Phys. Rev. D **91**, 022003 (2015).

[188] M. Punturo *et al.*, *The Einstein Telescope: a third-generation gravitational wave observatory*, Class. Quant. Grav. **27**, 194002 (2010).

[189] B. Sathyaprakash *et al.*, *Scientific Objectives of Einstein Telescope*, Class. Quant. Grav. **30**, 079501 (2012).

[190] T. G. F. Li *et al.*, *Towards a Generic Test of the Strong Field Dynamics of General Relativity Using Compact Binary Coalescence*, Phys. Rev. D **85**, 082003 (2012).

[191] C. Brans and R. H. Dicke, *Mach's Principle and a Relativistic Theory of Gravitation*, Phys. Rev. **124**, 925 (1961).

[192] D. Gerosa, U. Sperhake, and C. D. Ott, *Numerical Simulations of Stellar Collapse in Scalar-Tensor Theories of Gravity*, Class. Quant. Grav. **33**, 135002 (2016).

[193] S. E. Gossan *et al.*, *Observing Gravitational Waves from Core-Collapse Supernovae in the Advanced Detector Era*, Phys. Rev. D **93**, 042002 (2016).

[194] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *First Targeted Search for Gravitational-Wave Bursts from Core-Collapse Supernovae in Data of First-Generation Laser Interferometer Detectors*, Phys. Rev. D **94**, 102001 (2016).

[195] S. M. Adams, C. S. Kochanek, J. F. Beacom, M. R. Vagins, and K. Z. Stanek, *Observing the Next Galactic Supernova*, Astrophys. J. **778**, 164 (2013).

[196] T. Damour and G. Esposito-Faréese, *Tensor-Multi-Scalar Theories of Gravitation*, Class. Quantum Gravity **9**, 2093 (1992).

[197] T. Damour and G. Esposito-Farèse, *Nonperturbative Strong Field Effects in Tensor-Scalar Theories of Gravitation*, Phys. Rev. Lett. **70**, 2220 (1993).

[198] P. C. C. Freire *et al.*, *The Relativistic Pulsar-White Dwarf Binary PSR J1738+0333 - II. The Most Stringent Test of Scalar-Tensor Gravity*, Mon. Not. R. Astron. Soc. **423**, 3328 (2012).

[199] J. Antoniadis *et al.*, *A Massive Pulsar in a Compact Relativistic Binary*, Science **340**, 1233232 (2013).

[200] M. W. Horbatsch and C. P. Burgess, *Cosmic Black-Hole Hair Growth and Quasar OJ287*, J. Cosmol. Astropart. Phys. **2012**, 010 (2012).

[201] E. Berti, V. Cardoso, L. Gualtieri, M. Horbatsch, and U. Sperhake, *Numerical Simulations of Single and Binary Black Holes in Scalar-Tensor Theories: Circumventing the No-Hair Theorem*, Phys. Rev. D **87**, 124020 (2013).

[202] K. Yagi, L. C. Stein, and N. Yunes, *Challenging the Presence of Scalar Charge and Dipolar Radiation in Binary Pulsars*, Phys. Rev. D **D93**, 024010 (2016).

[203] E. Barausse, N. Yunes, and K. Chamberlain, *Theory-Agnostic Constraints on Black-Hole Dipole Radiation with Multiband Gravitational-Wave Astrophysics*, Phys. Rev. Lett. **116**, 241104 (2016).

[204] T. Damour and A. Vilenkin, *Cosmic Strings and the String Dilaton*, Phys. Rev. Lett. **78**, 2288 (1997).

[205] R. Tso, M. Isi, Y. Chen, and L. Stein, *Modeling the Dispersion and Polarization Content of Gravitational Waves for Tests of General Relativity*, Preprint (2016), arXiv:1608.01284.

[206] A. De Felice, T. Nakamura, and T. Tanaka, *Possible existence of viable models of bi-gravity with detectable graviton oscillations by gravitational wave detectors*, Prog. Theor. Exper. Phys. **2014**, 043E01 (2014).

[207] T. Narikawa *et al.*, *Detectability of Bigravity with Graviton Oscillations using Gravitational Wave Observations*, Phys. Rev. D **91**, 062007 (2015).

[208] V. A. Kostelecky, *Gravity, Lorentz Violation, and the Standard Model*, Phys. Rev. D **69**, 105009 (2004).

[209] V. A. Kostelecky and M. Mewes, *Testing Local Lorentz Invariance with Gravitational Waves*, Phys. Lett. B **757**, 510 (2016).

[210] T. Damour and K. Nordtvedt, *General Relativity as a Cosmological Attractor of Tensor-Scalar Theories*, Phys Rev. Lett. **70**, 2217 (1993).

[211] A. Joyce, B. Jain, J. Khoury, and M. Trodden, *Beyond the Cosmological Standard Model*, Physics Reports **568**, 1 (2015).

[212] C. M. Will, *The Confrontation between General Relativity and Experiment*, Living Rev. Rel. **17**, 4 (2014).

[213] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Tests of General Relativity with the Binary Black Hole Signals from the LIGO-Virgo Catalog GWTC-1*, Preprint (2019), arXiv:1903.04467.

[214] A. Nishizawa, A. Taruya, K. Hayama, S. Kawamura, and M. Sakagami, *Probing Nontensorial Polarizations of Stochastic Gravitational-Wave Backgrounds with Ground-Based Laser Interferometers*, Phys. Rev. D **79**, 082002 (2009).

[215] A. Nishizawa, A. Taruya, and S. Kawamura, *Cosmological Test of Gravity with Polarizations of Stochastic Gravitational Waves Around 0.1–1 Hz*, Phys. Rev. D **81**, 104043 (2010).

[216] A. Nishizawa and K. Hayama, *Probing for Massive Stochastic Gravitational-Wave Background with a Detector Network*, Phys. Rev. D **88**, 064005 (2013).

[217] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, *Detecting Beyond-Einstein Polarizations of Continuous Gravitational Waves*, Phys. Rev. D **91**, 082002 (2015).

[218] M. Isi, M. Pitkin, and A. J. Weinstein, *Probing dynamical gravity with the polarization of continuous gravitational waves*, Phys. Rev. D **96**, 042001 (2017).

[219] J. D. Romano and N. J. Cornish, *Detection Methods for Stochastic Gravitational-Wave Backgrounds: A Unified Treatment*, Living Rev. Rel. **20**, 2 (2017).

[220] C. R. Contaldi, J. Magueijo, and L. Smolin, *Anomalous Cosmic-Microwave-Background Polarization and Gravitational Chirality*, Phys. Rev. Lett. **101**, 141101 (2008).

[221] M. E. D. S. Alves and M. Tinto, *Pulsar timing sensitivities to gravitational waves from relativistic metric theories of gravity*, Phys. Rev. D **83**, 123529 (2011).

[222] S. J. Chamberlin and X. Siemens, *Stochastic backgrounds in alternative theories of gravity: Overlap reduction functions for pulsar timing arrays*, Phys. Rev. D **85**, 082001 (2012).

[223] L. O'Beirne, N. J. Cornish, S. J. Vigeland, and S. R. Taylor, *Constraining alternative polarization states of gravitational waves from individual black hole binaries using pulsar timing arrays*, Preprint (2019), arXiv:1904.02744.

[224] G. Cella *et al.*, *Prospects for Stochastic Background Searches Using Virgo and LSC Interferometers*, Class. Quant. Grav. **24**, S639 (2007).

[225] R. A. Isaacson, *Gravitational Radiation in the Limit of High Frequency. II. Nonlinear Terms and the Effective Stress Tensor*, Phys. Rev. **166**, 1272 (1968).

[226] T. Callister, L. Sammut, S. Qiu, I. Mandel, and E. Thrane, *Limits of Astrophysics with Gravitational-Wave Backgrounds*, Phys. Rev. X **6**, 031018 (2016).

[227] M. Vallisneri and N. Yunes, *Stealth Bias in Gravitational-Wave Parameter Estimation*, Phys. Rev. D **87**, 102002 (2013).

[228] S. Vitale and W. Del Pozzo, *How Serious Can the Stealth Bias Be in Gravitational Wave Parameter Estimation?*, Phys. Rev. D **89**, 022002 (2014).

[229] LIGO Scientific Collaboration and Virgo Collaboration, J. Aasi *et al.*, *Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009-2010 LIGO and Virgo Data*, Phys. Rev. Lett. **113**, 231101 (2014).

[230] M. Vallisneri, *Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects*, Phys. Rev. D **77**, 042001 (2008).

[231] J. Buchner *et al.*, *X-ray Spectral Modelling of the AGN Obscuring Region in the CDFS: Bayesian Model Selection and Catalogue*, Astron. Astrophys. **564**, A125 (2014).

[232] F. Feroz and M. P. Hobson, *Multimodal Nested Sampling: An Efficient and Robust Alternative to Markov Chain Monte Carlo Methods for Astronomical Data Analyses*, Mon. Not. R. Astron. Soc. **384**, 449 (2008).

[233] F. Feroz, M. P. Hobson, and M. Bridges, *MultiNest: An Efficient and Robust Bayesian Inference Tool for Cosmology and Particle Physics*, Mon. Not. R. Astron. Soc. **398**, 1601 (2009).

[234] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, Preprint (2013), arXiv:1306.2144.

[235] J. Skilling, Nested Sampling, in *AIP Conf. Proc.* Vol. 735, pp. 395–405, AIP, 2004.

[236] J. Skilling, *Nested Sampling for General Bayesian Computation*, Bayesian Analysis **1**, 833 (2006).

[237] B. Allen and A. C. Ottewill, *Detection of anisotropies in the gravitational-wave stochastic background*, Phys. Rev. D **56**, 545 (1997).

[238] E. Thrane *et al.*, *Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers*, Phys. Rev. D **80**, 122002 (2009).

[239] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Directional Limits on Persistent Gravitational Waves from Advanced LIGO's First Observing Run*, Phys. Rev. Lett. **118**, 121102 (2017).

[240] P. D. Welch, *The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms*, IEEE Trans. Audio and Electroacoustics **15**, 70 (1967).

[241] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *SUPPLEMENT: Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run*, Phys. Rev. Lett. **118**, 121101 (2017), Online Supplement.

[242] E. Thrane, N. Christensen, and R. M. S. Schofield, *Correlated magnetic noise in global networks of gravitational-wave detectors: Observations and implications*, Phys. Rev. D **87**, 123009 (2013).

[243] E. Thrane, N. Christensen, R. M. S. Schofield, and A. Effler, *Correlated noise in networks of gravitational-wave detectors: Subtraction and mitigation*, Phys. Rev. D **90**, 023013 (2014).

[244] Y. Himemoto and A. Taruya, *Impact of correlated magnetic noise on the detection of stochastic gravitational waves: Estimation based on a simple analytical model*, Phys. Rev. D **96**, 022004 (2017).

[245] C. Cahillane *et al.*, *Calibration uncertainty for Advanced LIGO's first and second observing runs*, Phys. Rev. D **96**, 102001 (2017).

[246] J. T. Whelan, E. L. Robinson, J. D. Romano, and E. H. Thrane, *Treatment of calibration uncertainty in multi-baseline cross-correlation searches for gravitational waves*, J. Phys. Conf. Ser. **484**, 012027 (2014).

[247] T. Callister *et al.*, *Polarization-Based Tests of Gravity with the Stochastic Gravitational-Wave Background*, Phys. Rev. X **7**, 041058 (2017).

[248] J. C. Driggers *et al.*, *Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO*, Preprint (2018), arXiv:1806.00532.

[249] D. Davis *et al.*, *Improving the Sensitivity of Advanced LIGO Using Noise Subtraction*, Preprint (2018), arXiv:1809.05348.

[250] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series.* (Wiley, 1949).

[251] A. D. Viets *et al.*, *Reconstructing the calibrated strain signal in the Advanced LIGO detectors*, Class. Quant. Grav. **35**, 095015 (2018).

[252] W. O. Schumann, *Über die strahlungslosen Eigenschwingungen einer leitenden Kugel, die von einer Luftschicht und einer Ionosphärenhülle umgeben ist*, Zeitschrift Naturforschung Teil A **7**, 149 (1952).

[253] W. O. Schumann, *Über die Dämpfung der elektromagnetischen Eigenschwingungen des Systems Erde - Luft - Ionosphäre*, Zeitschrift Naturforschung Teil A **7**, 250 (1952).

[254] M. W. Coughlin *et al.*, *Subtraction of correlated noise in global networks of gravitational-wave interferometers*, Class. Quant. Grav. **33**, 224003 (2016).

[255] M. W. Coughlin *et al.*, *Measurement and subtraction of Schumann resonances at gravitational-wave interferometers*, Phys. Rev. D **97**, 102007 (2018).

[256] O. J. Sovers, J. L. Fanselow, and C. S. Jacobs, *Astrometry and geodesy with radio interferometry: experiments, models, results*, Reviews of Modern Physics **70**, 1393 (1998).

[257] H. Schuh and D. Behrend, *VLBI: A fascinating technique for geodesy and astrometry*, Journal of Geodynamics **61**, 68 (2012).

[258] N. J. Cornish and L. Sampson, *Towards robust gravitational wave detection with pulsar timing arrays*, Phys. Rev. D **93**, 104047 (2016).

[259] S. R. Taylor *et al.*, *All correlations must die: Assessing the significance of a stochastic gravitational-wave background in pulsar timing arrays*, Phys. Rev. D **95**, 042002 (2017).

[260] R. W. Hellings and G. S. Downs, *Upper limits on the isotropic gravitational radiation background from pulsar timing analysis*, Astrophys. J. Letters **265**, L39 (1983).

[261] R. Brito *et al.*, *Gravitational wave searches for ultralight bosons with LIGO and LISA*, Phys. Rev. D **96**, 064050 (2017).

[262] R. Brito *et al.*, *Stochastic and Resolvable Gravitational Waves from Ultralight Bosons*, Phys. Rev. Lett. **119**, 131101 (2017).

[263] L. Tsukada, T. Callister, A. Matas, and P. Meyers, *A first search for a stochastic gravitational-wave background from ultralight bosons*, Preprint (2018), arXiv:1812.09622.

[264] G. Cusin, R. Durrer, and P. G. Ferreira, *Polarization of a stochastic gravitational wave background through diffusion by massive structures*, Preprint (2018), arXiv:1807.10620.

[265] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, *First Predictions of the Angular Power Spectrum of the Astrophysical Gravitational Wave Background*, Phys. Rev. Lett. **120**, 231101 (2018).

[266] V. M. Lipunov and I. E. Panchenko, *Pulsars Revived by Gravitational Waves*, Astronomy & Astrophysics **312**, 937 (1996).

[267] S. Dong *et al.*, *ASASSN-15lh: A highly super-luminous supernova*, Science **351**, 257 (2016).

[268] I. Arcavi *et al.*, *Optical emission from a kilonova following a gravitational-wave-detected neutron-star merger*, Nature **551**, 64 (2017).

[269] V. M. Lipunov *et al.*, *MASTER Optical Detection of the First LIGO/Virgo Neutron Star Binary Merger GW170817*, Astrophys. J. **850**, L1 (2017).

[270] N. R. Tanvir *et al.*, *The Emergence of a Lanthanide-rich Kilonova Following the Merger of Two Neutron Stars*, Astrophys. J. **848**, L27 (2017).

[271] S. Valenti *et al.*, *The Discovery of the Electromagnetic Counterpart of GW170817: Kilonova AT 2017gfo/DLT17ck*, Astrophys. J. **848**, L24 (2017).

[272] E. Troja *et al.*, *The X-ray counterpart to the gravitational-wave event GW170817*, Nature **551**, 71 (2017).

[273] R. Chornock *et al.*, *The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. IV. Detection of Near-infrared Signatures of r - process Nucleosynthesis with Gemini-South*, Astrophys. J. **848**, L19 (2017).

[274] M. R. Drout *et al.*, *Light curves of the neutron star merger GW170817/SSS17a: Implications for r-process nucleosynthesis*, Science **358**, 1570 (2017).

[275] S. J. Smartt *et al.*, *A kilonova as the electromagnetic counterpart to a gravitational-wave source*, Nature **551**, 75 (2017).

[276] K. P. Mooley *et al.*, *A Strong Jet Signature in the Late-time Light Curve of GW170817*, Astrophys. J. **868**, L11 (2018).

[277] E. Nakar, O. Gottlieb, T. Piran, M. M. Kasliwal, and G. Hallinan, *From γ to Radio: The Electromagnetic Counterpart of GW170817*, Astrophys. J. **867**, 18 (2018).

[278] C. A. Raithel, F. Özel, and D. Psaltis, *Tidal Deformability from GW170817 as a Direct Probe of the Neutron Star Radius*, Astrophys. J. **857**, L23 (2018).

[279] M. W. Coughlin, T. Dietrich, B. Margalit, and B. D. Metzger, *Multi-messenger Bayesian parameter inference of a binary neutron-star merger*, Preprint (2018), arXiv:1812.04803.

[280] K. Hotokezaka *et al.*, *A Hubble constant measurement from superluminal motion of the jet in GW170817*, Preprint (2018), arXiv:1806.10596.

[281] J. M. Ezquiaga and M. Zumalacárregui, *Dark Energy After GW170817: Dead Ends and the Road Ahead*, Phys. Rev. Lett. **119** (2017).

[282] D. Langlois, R. Saito, D. Yamauchi, and K. Noui, *Scalar-tensor theories and modified gravity in the wake of GW170817*, Phys. Rev. D **97**, 061501 (2018).

[283] S. Jana and S. Mohanty, *Constraints on $f (R)$ theories of gravity from GW170817*, Phys. Rev. D **99** (2019).

[284] T. Totani, *Cosmological Fast Radio Bursts from Binary Neutron Star Mergers*, Publications of the Astronomical Society of Japan **65**, L12 (2013).

[285] J.-S. Wang, F.-K. Peng, K. Wu, and Z.-G. Dai, *Pre-merger Electromagnetic Counterparts of Binary Compact Stars*, Astrophys. J. **868**, 19 (2018).

[286] M. Lyutikov, *Electrodynamics of binary neutron star mergers*, Mon. Not. R. Astron. Soc. **483**, 2766 (2019).

[287] B. M. S. Hansen and M. Lyutikov, *Radio and X-ray signatures of merging neutron stars*, Mon. Not. R. Astron. Soc. **322**, 695 (2001).

[288] M. Lyutikov, *Electromagnetic power of merging and collapsing compact objects*, Phys. Rev. D **83**, 124035 (2011).

[289] M. Lyutikov, *The electromagnetic model of short GRBs, the nature of prompt tails, supernova-less long GRBs, and highly efficient episodic accretion*, Astrophys. J. **768**, 63 (2013).

[290] J.-S. Wang, Y.-P. Yang, X.-F. Wu, Z.-G. Dai, and F.-Y. Wang, *Fast Radio Bursts from the Inspiral of Double Neutron Stars*, Astrophys. J. Lett. **822**, L7 (2016).

[291] D. Lai, *DC circuit powered by orbital motion: Magnetic interactions in compact object binaries and exoplanetary systems*, Astrophys. J. **757**, L3 (2012).

[292] A. L. Piro, *Magnetic interactions in coalescing neutron star binaries*, Astrophys. J. **755**, 80 (2012).

[293] B. D. Metzger and C. Zivancev, *Pair fireball precursors of neutron star mergers*, Mon. Not. R. Astron. Soc. **461**, 4435 (2016).

[294] K. Wu, *Magnetic interaction in ultra-compact binary systems*, Research in Astronomy and Astrophysics **9**, 725 (2009).

[295] P. Goldreich and D. Lynden-Bell, *Io, a jovian unipolar inductor*, Astrophys. J. **156**, 59 (1969).

[296] J. Moortgat and J. Kuijpers, *Gravitational and magnetosonic waves in gamma-ray bursts*, Astronomy & Astrophysics **402**, 905 (2003).

[297] J. Moortgat and J. Kuijpers, *Gravitational waves in magnetized relativistic plasmas*, Phys. Rev. D **70**, 023001 (2004).

[298] J. Moortgat and J. Kuijpers, *Indirect Visibility of Gravitational Waves in Magnetohydrodynamic Plasmas*, Preprint (2005), arXiv:gr-qc/0503074.

[299] J. Moortgat and J. Kuijpers, *Scattering of magnetosonic waves in a relativistic and anisotropic magnetized plasma*, Mon. Not. R. Astron. Soc. **368**, 1110 (2006).

[300] H. Sotani, K. D. Kokkotas, P. Laguna, and C. F. Sopuerta, *Gravitationally driven electromagnetic perturbations of neutron stars and black holes*, Phys. Rev. D **87**, 084018 (2013).

[301] C. A. Raithel, *Constraints on the Neutron Star Equation of State from GW170817*, Preprint (2019), arXiv:1904.10002.

[302] V. Ravi and P. D. Lasky, *The birth of black holes: neutron star collapse times, gamma-ray bursts and fast radio bursts*, Mon. Not. R. Astron. Soc. **441**, 2433 (2014).

[303] H. Falcke and L. Rezzolla, *Fast radio bursts: the last sign of supramassive neutron stars*, Astronomy & Astrophysics **562**, A137 (2014).

[304] B. Zhang, *A possible connection between fast radio bursts and gamma-ray bursts*, Astrophys. J. **780**, L21 (2014).

[305] M. S. Pshirkov and K. A. Postnov, *Radio precursors to neutron star binary mergings*, Astrophysics and Space Science **330**, 13 (2010).

[306] M. Lyutikov and J. C. McKinney, *Slowly balding black holes*, Phys. Rev. D **84**, 084019 (2011).

[307] M. V. Smolsky and V. V. Usov, *Nonthermal Radiation of Cosmological Gamma-Ray Bursters*, Astrophys. J. **531**, 764 (2000).

[308] V. V. Usov and J. I. Katz, *Low Frequency Radio Pulses from Gamma-Ray Bursts?*, Preprint (2000), arXiv:astro-ph/0002278.

[309] T. Piran, *The physics of gamma-ray bursts*, Reviews of Modern Physics **76**, 1143 (2004).

[310] F. Graham-Smith, *The radio emission from pulsars*, Reports on Progress in Physics **66**, 173 (2003).

[311] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford, *A Bright Millisecond Radio Burst of Extragalactic Origin*, Science **318**, 777 (2007).

[312] E. F. Keane, M. Kramer, A. G. Lyne, B. W. Stappers, and M. A. McLaughlin, *Rotating Radio Transients: new discoveries, timing solutions and musings*, Mon. Not. R. Astron. Soc. **415**, 3065 (2011).

[313] D. Thornton *et al.*, *A Population of Fast Radio Bursts at Cosmological Distances*, Science **341**, 53 (2013).

[314] S. Burke-Spolaor and K. W. Bannister, *The galactic position dependence of fast radio bursts and the discovery of FRB011025*, Astrophys. J. **792**, 19 (2014).

[315] E. Petroff *et al.*, *A real-time fast radio burst: polarization detection and multiwavelength follow-up*, Mon. Not. R. Astron. Soc. **447**, 246 (2015).

[316] V. Ravi, R. M. Shannon, and A. Jameson, *A fast radio burst in the direction of the Carina dwarf spheroidal galaxy*, Astrophys. J. **799**, L5 (2015).

[317] D. J. Champion *et al.*, *Five new fast radio bursts from the HTRU high-latitude survey at Parkes: first evidence for two-component bursts*, Mon. Not. R. Astron. Soc. **460**, L30 (2016).

[318] E. F. Keane *et al.*, *The host galaxy of a fast radio burst*, Nature **530**, 453 (2016).

[319] L. G. Spitler *et al.*, *Fast radio burst discovered in the Arecibo pulsar ALFA survey*, Astrophys. J. **790**, 101 (2014).

[320] L. G. Spitler *et al.*, *A repeating fast radio burst*, Nature **531**, 202 (2016).

[321] C. Patel *et al.*, *PALFA Single-pulse Pipeline: New Pulsars, Rotating Radio Transients, and a Candidate Fast Radio Burst*, Astrophys. J. **869**, 181 (2018).

[322] K. Masui *et al.*, *Dense magnetized plasma associated with a fast radio burst*, Nature **528**, 523 (2015).

[323] M. Caleb *et al.*, *The first interferometric detections of fast radio bursts*, Mon. Not. R. Astron. Soc. **468**, 3746 (2017).

[324] W. Farah *et al.*, *FRB microstructure revealed by the real-time detection of FRB170827*, Mon. Not. R. Astron. Soc. **478**, 1209 (2018).

[325] W. Farah *et al.*, *Five new real-time detections of Fast Radio Bursts with UTMOST*, Preprint (2019), arXiv:1905.02293.

[326] K. W. Bannister *et al.*, *The Detection of an Extremely Bright Fast Radio Burst in a Phased Array Feed Survey*, Astrophys. J. **841**, L12 (2017).

[327] R. M. Shannon *et al.*, *The dispersion-brightness relation for fast radio bursts from a wide-field survey*, Nature **562**, 386 (2018).

[328] H. Qiu *et al.*, *A survey of the Galactic plane for dispersed radio pulses with the Australian Square Kilometre Array Pathfinder*, Mon. Not. R. Astron. Soc. **486**, 166 (2019).

[329] S. Bhandari *et al.*, *A southern sky search for repeating fast radio bursts using the Australian SKA Pathfinder*, Mon. Not. R. Astron. Soc. **486**, 70 (2019).

[330] J. P. Macquart *et al.*, *The Spectral Properties of the Bright Fast Radio Burst Population*, Astrophys. J. **872**, L19 (2019).

[331] M. Amiri *et al.*, *Observations of Fast Radio Bursts at Frequencies down to 400 Megahertz*, Nature (2019).

[332] M. Amiri *et al.*, *A Second Source of Repeating Fast Radio Bursts*, Nature (2019).

[333] E. Petroff *et al.*, *FRBCAT: The Fast Radio Burst Catalogue*, Preprint (2016), arXiv:1601.03547.

[334] E. F. Keane and E. Petroff, *Fast radio bursts: search sensitivities and completeness*, Mon. Not. R. Astron. Soc. **447**, 2852 (2015).

[335] J. M. Cordes and I. Wasserman, *Supergiant pulses from extragalactic neutron stars*, Mon. Not. R. Astron. Soc. **457**, 232 (2016).

[336] L. Connor, J. Sievers, and U.-L. Pen, *Non-Cosmological FRBs from Young Supernova Remnant Pulsars*, Mon. Not. R. Astron. Soc. **458**, L19 (2016).

[337] F. Mottez and P. Zarka, *Radio emissions from pulsar companions: a refutable explanation for galactic transients and fast radio bursts*, Astronomy & Astrophysics **569**, A86 (2014).

[338] G. E. Romero, M. V. del Valle, and F. L. Vieyro, *Mechanism for fast radio bursts*, Phys. Rev. D **93**, 023001 (2016).

[339] A. Loeb, Y. Shvartzvald, and D. Maoz, *Fast radio bursts may originate from nearby flaring stars*, Mon. Not. R. Astron. Soc. **439**, L46 (2014).

[340] D. Maoz *et al.*, *Fast radio bursts: the observational case for a Galactic origin*, Mon. Not. R. Astron. Soc. **454**, 2183 (2015).

[341] A. Barrau, C. Rovelli, and F. Vidotto, *Fast radio bursts and white hole signals*, Phys. Rev. D **90**, 127503 (2014).

[342] E. F. Keane, B. W. Stappers, M. Kramer, and A. G. Lyne, *On the origin of a highly dispersed coherent radio burst*, Mon. Not. R. Astron. Soc. **425**, L71 (2012).

[343] R. Brandenberger, B. Cyr, and A. V. Iyer, *Fast Radio Bursts from the Decay of Cosmic String Cusps*, Preprint (2017), arXiv:1707.02397.

[344] M. Lingam and A. Loeb, *Fast Radio Bursts from Extragalactic Light Sails*, Astrophys. J. Lett. **837**, L23 (2017).

[345] S. T. McWilliams and J. Levin, *Electromagnetic extraction of energy from black-hole-neutron-star binaries*, Astrophys. J. **742**, 90 (2011).

[346] C. M. F. Mingarelli, J. Levin, and T. J. W. Lazio, *Fast radio bursts and radio transients from black hole batteries*, Astrophys. J. **814**, L20 (2015).

[347] D. J. D'Orazio, J. Levin, N. W. Murray, and L. Price, *Bright transients from strongly-magnetized neutron star-black hole mergers*, Phys. Rev. D **94**, 023001 (2016).

[348] M. Lyutikov, *Schwarzschild black holes as unipolar inductors: Expected electromagnetic power of a merger*, Phys. Rev. D **83**, 064001 (2011).

[349] V. Predoi *et al.*, *Prospects for joint radio telescope and gravitational-wave searches for astrophysical transients*, Class. Quant. Grav. **27**, 084018 (2010).

[350] K. W. Bannister, T. Murphy, B. M. Gaensler, and J. E. Reynolds, *Limits on prompt, dispersed radio pulses from gamma-ray bursts*, Astrophys. J. **757** (2012).

[351] D. L. Kaplan *et al.*, *A deep search for prompt radio emission from the short GRB 150424A with the Murchison Widefield Array*, Astrophys. J. **814**, L25 (2015).

[352] C. C. Yancey *et al.*, *Multi-messenger astronomy of gravitational-wave sources with flexible wide-area radio transient surveys*, Astrophys. J. **812**, 168 (2015).

[353] D. L. Kaplan *et al.*, *Strategies for Finding Prompt Radio Counterparts to Gravitational Wave Transients with the Murchison Widefield Array*, Publications of the Astronomical Society of Australia **33**, e050 (2016).

[354] M. M. Anderson *et al.*, *A Simultaneous Search for Prompt Radio Emission Associated with the Short GRB 170112A Using the All-sky Imaging Capability of the OVRO-LWA*, Astrophys. J. **864**, 22 (2018).

[355] T. A. Callister *et al.*, *A First Search for Prompt Radio Emission from a Gravitational-Wave Event*, Preprint (2019), arXiv:1903.06786.

[356] A. Rowlinson and G. E. Anderson, *Constraining coherent low frequency radio flares from compact binary mergers*, Preprint (2019), arXiv:1905.02509.

[357] V. Kalogera *et al.*, *The Cosmic Coalescence Rates for Double Neutron Star Binaries*, Astrophys. J. **614**, L137 (2004).

[358] R. O'Shaughnessy, C. Kim, V. Kalogera, and K. Belczynski, *Constraining Population Synthesis Models via Empirical Binary Compact Object Merger and Supernova Rates*, Astrophys. J. **672**, 479 (2008).

[359] LIGO Scientific Collaboration and Virgo Collaboration, J. Abadie *et al.*, *Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors*, Class. Quant. Grab. **27**, 173001 (2010).

[360] C. Kim, B. B. P. Perera, and M. a. McLaughlin, *Implications of PSR J0737-3039B for the Galactic NS-NS binary merger rate*, Mon. Not. R. Astron. Soc. **448**, 928 (2015).

[361] M. Dominik *et al.*, *Double compact objects II. Cosmological merger rates*, Astrophys. J. **779**, 72 (2013).

[362] LIGO Scientific Collaboration and Virgo Collaboration, J. Abadie *et al.*, *Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3*, Phys. Rev. D **85**, 082002 (2012).

[363] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo*, Living Reviews in Relativity **19** (2016).

[364] B. Zhang, *Mergers of Charged Black Holes: Gravitational Wave Events, Short Gamma-Ray Bursts, and Fast Radio Bursts*, Preprint (2016), arXiv:1602.04542.

[365] B. Zhang, *On the afterglow and progenitor of FRB 150418*, Astrophys. J. **822**, L14 (2016).

[366] J.-P. Macquart and S. Johnston, *On the paucity of Fast Radio Bursts at low Galactic latitudes*, Mon. Not. R. Astron. Soc. **451**, 3278 (2015).

[367] A. Rane *et al.*, *A search for rotating radio transients and fast radio bursts in the Parkes high-latitude pulsar survey*, Mon. Not. R. Astron. Soc. **455**, 2207 (2016).

[368] C. J. Law *et al.*, *A millisecond interferometric search for fast radio bursts with the Very Large Array*, Astrophys. J. **807**, 16 (2015).

[369] K. Ioka, *The Cosmic Dispersion Measure from Gamma-Ray Burst Afterglows: Probing the Reionization History and the Burst Environment*, Astrophys. J. **598**, L79 (2003).

[370] S. Inoue, *Probing the cosmic reionization history and local environment of gamma-ray bursts through radio dispersion*, Mon. Not. R. Astron. Soc. **348**, 999 (2004).

[371] Planck Collaboration *et al.*, *Planck 2015 results. XIII. Cosmological parameters*, Preprint (2015), arXiv:1502.01589.

[372] M. Fukugita and P. J. E. Peebles, *The Cosmic Energy Inventory*, Astrophys. J. **616**, 643 (2004).

[373] Q. Chu *et al.*, *Capturing the electromagnetic counterparts of binary neutron star mergers through low-latency gravitational wave triggers*, Mon. Not. R. Astron. Soc. **459**, 121 (2016).

[374] K. S. Obenberger *et al.*, *Limits on gamma-ray burst prompt radio emission using the LWA1*, Astrophys. J. **785**, 27 (2014).

[375] T. Callister *et al.*, GCN **21680** (2017).

[376] D. Kaplan *et al.*, GCN **21927** (2017).

[377] J. Kocz *et al.*, *Digital Signal Processing Using Stream High Performance Computing: A 512-Input Broadband Correlator for Radio Astronomy*, Journal of Astronomical Instrumentation **4**, 1550003 (2015).

[378] M. W. Eastwood *et al.*, *The Radio Sky at Meter Wavelengths: m -mode Analysis Imaging with the OVRO-LWA*, Astron. J. **156**, 32 (2018).

[379] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2*, Phys. Rev. Lett. **118**, 221101 (Supplement) (2017).

[380] V. Connaughton *et al.*, *Fermi GBM observations of LIGO gravitational-wave event GW150914*, Astrophys. J. **826**, L6 (2016).

[381] N. Omodei *et al.*, GCN **21227** (2017).

[382] K. Artkop *et al.*, *Radio follow-up of a candidate gamma-ray transient in the sky localization area of GW170608*, Preprint (2019), arXiv:1903.10660.

[383] S. L. Liebling and C. Palenzuela, *Electromagnetic luminosity of the coalescence of charged black hole binaries*, Phys. Rev. D **94**, 064046 (2016).

[384] T. Liu, G. E. Romero, M.-L. Liu, and A. Li, *Fast radio bursts and their gamma-ray or radio afterglows as Kerr-Newman black hole binaries*, Astrophys. J. **826**, 82 (2016).

[385] F. Fraschetti, *Possible role of magnetic reconnection in the electromagnetic counterpart of binary black hole merger*, Journal of Cosmology and Astroparticle Physics **2018**, 054 (2018).

[386] R. Perna, D. Lazzati, and B. Giacomazzo, *Short gamma-ray bursts from the merger of two black holes*, Astrophys. J. **821**, L18 (2016).

[387] S. E. de Mink and A. King, *Electromagnetic Signals Following Stellar-mass Black Hole Mergers*, Astrophys. J. **839**, L7 (2017).

[388] A. Loeb, *Electromagnetic counterparts to black hole mergers detected by LIGO*, Astrophys. J. **819**, L21 (2016).

[389] J. M. Fedrow *et al.*, *Gravitational Waves from Binary Black Hole Mergers inside Stars*, Phys. Rev. Lett. **119** (2017).

[390] J. Greiner, J. M. Burgess, V. Savchenko, and H.-F. Yu, *On the Fermi-GBM event 0.4s after GW150914*, Astrophys. J. **827**, L38 (2016).

[391] M. Lyutikov, *Fermi GBM signal contemporaneous with GW150914 - an unlikely association*, Preprint (2016), arXiv:1602.07352.

[392] V. Savchenko *et al.*, *INTEGRAL upper limits on gamma-ray emission associated with the gravitational wave event GW150914*, Astrophys. J. **820**, L36 (2016).

[393] V. Connaughton *et al.*, *On the Interpretation of the Fermi -GBM Transient Observed in Coincidence with LIGO Gravitational-wave Event GW150914*, Astrophys. J. **853**, L9 (2018).

[394] LIGO Scientific Collaboration and Virgo Collaboration, B. P. Abbott *et al.*, *Low-latency Gravitational-wave Alerts for Multimessenger Astronomy during the Second Advanced LIGO and Virgo Observing Run*, Astrophys. J. **875**, 161 (2019).

[395] J. W. M. Baars, R. Genzel, I. I. K. Pauliny-Toth, and A. Witzel, *The Absolute Spectrum of Cas A; An Accurate Flux Density Scale and a Set of Secondary Calibrators*, Astron. Astrophys. **61**, 99 (1977).

[396] R. A. Perley and B. J. Butler, *An Accurate Flux Density Scale from 50 MHz to 50 GHz*, Astrophys. J. Supplement Series **230**, 7 (2017).

[397] M. W. Eastwood, TTCal, 2016, 10.5281/zenodo.1049160.

[398] A. R. Offringa *et al.*, *wsclean: an implementation of a fast, generic wide-field imager for radio astronomy*, Mon. Not. R. Astron. Soc. **444**, 606 (2014).

[399] G. Rybicki and A. Lightman, *Radiative Processes in Astrophysics.* (Wiley, 1991).

[400] J. Condon and S. Ransom, *Essential Radio Astronomy.* (Princeton University Press, 2016).

[401] J. M. Cordes and T. J. W. Lazio, *NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations*, Preprint (2002), arXiv:astro-ph/0207156.

[402] J. M. Cordes and T. J. W. Lazio, *NE2001. II. Using Radio Propagation Data to Construct a Model for the Galactic Distribution of Free Electrons*, Preprint (2003), arXiv:astro-ph/0301598.

[403] J. M. Cordes, R. S. Wharton, L. G. Spitler, S. Chatterjee, and I. Wasserman, *Radio Wave Propagation and the Provenance of Fast Radio Bursts*, Preprint (2016), arXiv:1605.05890.

[404] Z. Ceplecha *et al.*, *Meteor Phenomena and Bodies*, Space Science Rev. **84**, 327 (1998).

[405] J. F. Helmboldt *et al.*, *All-sky imaging of meteor trails at 55.25 MHz with the first station of the Long Wavelength Array*, Radio Science **49**, 157 (2014).

[406] B. C. Hicks *et al.*, *A Wide-Band, Active Antenna System for Long Wavelength Radio Astronomy*, Publications of the Astronomical Society of the Pacific **124**, 1090 (2012).

[407] M. Cantiello *et al.*, *A Precise Distance to the Host Galaxy of the Binary Neutron Star Merger GW170817 Using Surface Brightness Fluctuations*, Astrophys. J. **854**, L31 (2018).

[408] LIGO Scientific Collaboration and Virgo Collaboration, *LIGO/Virgo Public Alerts User Guide*, https://emfollow.docs.ligo.org/userguide/index.html (2018), Accessed: 2019-02-25.