



Shouts and Murmurs: Combining Individual Gravitational-wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers

Tom Callister¹ , Maya Fishbach² , Daniel E. Holz³ , and Will M. Farr^{1,4} 

¹ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA; tcallister@flatironinstitute.org

² Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

³ Enrico Fermi Institute, Department of Physics, Department of Astronomy and Astrophysics, and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁴ Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794, USA

Received 2020 March 26; revised 2020 May 24; accepted 2020 May 28; published 2020 June 17

Abstract

One of the goals of gravitational-wave astronomy is to quantify the evolution of the compact binary merger rate with redshift. The redshift distribution of black hole mergers would offer considerable information about their evolutionary history, including their progenitor formation rate, the dependence of black hole formation on stellar metallicity, and the time delay distribution between formation and merger. Efforts to measure the binary redshift distribution are currently limited, however, by the detection range of existing instruments, which can individually resolve compact binary merger events only out to $z \lesssim 1$. We present a novel strategy with which to measure the redshift distribution of binary black hole mergers well beyond the detection range of current instruments. By synthesizing direct detections of individually resolved mergers with *indirect* searches for the stochastic gravitational-wave background due to unresolved distant sources, we can glean information about the peak redshift, z_p , at which the binary black hole merger rate attains its maximum, even when this redshift is beyond the detection horizon. Using data from Advanced LIGO and Virgo's first and second observing runs, we employ this strategy to place joint constraints on z_p and the slope α with which the binary merger rate increases at low redshifts, ruling out merger rates that grow faster than $\alpha \gtrsim 7$ and peak beyond $z_p \gtrsim 1.5$. Looking ahead, we project that approximately one year of observation with design-sensitivity Advanced LIGO will further break remaining degeneracies, enabling a direct measurement of the peak redshift of the binary black hole merger history.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Black holes (162); Compact binary stars (283); LIGO (920); Gravitational wave sources (677)

1. Introduction

The Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015) gravitational-wave experiments are rapidly transitioning between low- and high-statistics regimes. With the LIGO–Virgo detections of eleven compact binary mergers during the past O1 and O2 observing runs (Abbott et al. 2019b) and tens more anticipated in the present O3 run (Abbott et al. 2019c), we can now begin to understand the ensemble properties of compact binaries, including the distributions of their component masses and spins (Farr et al. 2017; Fishbach & Holz 2017, 2020; Talbot & Thrane 2017, 2018; Abbott et al. 2019a; Wysocki et al. 2019). Beyond the distributions of these intrinsic binary parameters, we might also seek to understand the *redshift distribution* of binary black hole mergers—how the merger rate evolves as we look back to earlier times in the universe's history. If measured, the redshift distribution of compact binary mergers would offer substantial insight into the birth and evolution of compact binaries, encoding such properties as the time delay distribution between black hole formation and merger (see, e.g., Adhikari et al. 2020), the dependence of black hole production on stellar metallicity (e.g., Belczynski et al. 2016), and perhaps even the relative contributions from competing binary formation channels, including field binaries, hierarchical triples, dynamical capture, or primordial black holes (Dominik et al. 2015; Mandic et al. 2016; Mandel & Farmer 2018; Rodriguez & Loeb 2018).

Study of the binary black hole redshift distribution, however, is made difficult by the limited range of existing gravitational-wave

detectors. Figure 1, for example, shows a typical model for the source-frame rate of binary black hole mergers as a function of redshift. To obtain this figure, we assume progenitor formation following the star formation rate of Madau & Dickinson (2014) weighted by the fraction of stellar formation occurring at metallicities $Z \leq 0.3Z_\odot$ (Langer & Norman 2006). We further adopt a $p(t_d) \propto t_d^{-1}$ probability distribution for the time delay t_d between binary formation and merger, with $50 \text{ Myr} \leq t_d \leq 13.5 \text{ Gyr}$. Within this simple model, the binary black hole merger rate peaks at $z \sim 2$, while more sophisticated models generally predict merger rates peaking between redshifts $z \sim 2$ to 4, depending on the specific formation channel presumed (Dominik et al. 2013; Mapelli et al. 2017; Rodriguez & Loeb 2018; Baibhav et al. 2019; Santoliquido et al. 2020).

In contrast, design-sensitivity Advanced LIGO is expected to successfully detect optimally oriented $30 + 30 M_\odot$ binary black holes only out to redshifts $z \lesssim 1.2$ (Chen et al. 2017; Abbott et al. 2019c). Current efforts to study the redshift distribution of compact binary mergers therefore attempt only to measure the leading-order, low-redshift evolution of the binary merger rate (Fishbach et al. 2018; Abbott et al. 2019a); observation of the peak and subsequent turnover of the black hole redshift distribution is a challenge left to future third-generation detectors (Vitale et al. 2019).

In this Letter we demonstrate that present-day gravitational-wave observatories *can* provide meaningful measurements of the high-redshift evolution of the compact binary merger rate. We achieve these measurements by synthesizing the direct detections of compact binaries in the local universe with an

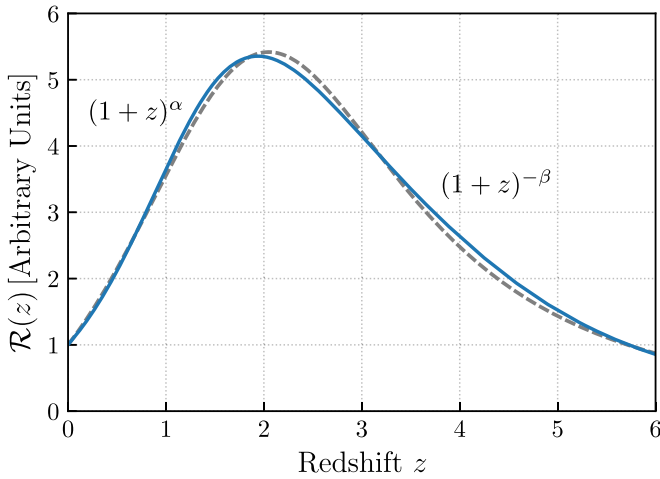


Figure 1. Example prediction of the source-frame rate density of binary black hole mergers (solid blue), assuming progenitor formation that follows the rate of stellar formation at metallicities $Z \leq 0.3 Z_{\odot}$ (Langer & Norman 2006; Madau & Dickinson 2014), and time delays t_d between binary formation and merger distributed as $p(t_d) \propto t_d^{-1}$, with $0.05 \text{ Gyr} \leq t_d \leq 13.5 \text{ Gyr}$. In this work, we will adopt a phenomenological model for the binary black hole merger rate (Equation (2)) that allows for the same qualitative behavior as the prediction shown here, rising as $\mathcal{R}(z) \propto (1+z)^{\alpha}$ at $z \lesssim z_p$ and falling as $\mathcal{R}(z) \propto (1+z)^{-\beta}$ at redshifts $z \gtrsim z_p$. The specific prediction plotted here, for instance, is well fit by Equation (2) using $\alpha = 1.9$, $\beta = 3.4$, and $z_p = 2.4$ (the dashed gray curve).

additional piece of information: the astrophysical stochastic gravitational-wave background (Romano & Cornish 2017; Christensen 2019). Arising from the superposition of all distant individually undetectable compact binaries, the stochastic gravitational-wave background manifests as excess correlated noise shared among a network of gravitational-wave detectors (Allen & Romano 1999). The strength of the present-day gravitational-wave background is determined by the cumulative merger history of binary black holes, integrated across all redshifts (Phinney 2001). The observation of (or even upper limits on) the gravitational-wave background can therefore be leveraged to place powerful constraints on the redshift distribution of binary mergers, complementary to those constraints gleaned from the direct detection of binaries in the local universe (Section 2).

We apply our approach to existing data, finding that the synthesis of binary black hole detections (Abbott et al. 2019b) and gravitational-wave background constraints (Abbott et al. 2016a, 2017a, 2017b, 2018, 2019d, 2019e) from Advanced LIGO and Advanced Virgo’s first two observing runs already yields nontrivial constraints on the peak of the binary black hole redshift distribution (Section 3). With additional data gathered from future observing runs, our method may enable a *measurement* of this peak redshift within the next five years (Section 4).

2. High-redshift Constraints from the Gravitational-wave Background

In their O1 and O2 observing runs, Advanced LIGO and Virgo confidently detected 10 binary black hole mergers, the most distant of which (GW170729) may have occurred at $z \approx 0.5$ (Abbott et al. 2019b; Chatziioannou et al. 2019). Together, these 10 events have recently allowed for the first exploration of the binary black hole merger rate’s evolution

with redshift. Adopting a model

$$\mathcal{R}(z) = \mathcal{R}_0 (1+z)^{\alpha}, \quad (1)$$

for the source-frame merger rate per comoving volume (Fishbach et al. 2018), Abbott et al. (2019a) find $\alpha = 6.5_{-9.3}^{+9.1}$ at 90% credibility. Thus, in the local universe, the binary black hole merger rate (probably) increases with redshift.

If the binary black holes observed with LIGO and Virgo are born from stellar progenitors, then the black hole merger rate cannot continue to increase out to arbitrarily high redshifts. Instead, it must reach a maximum at some peak redshift, z_p , and then decay to zero as star formation ceases in the very early universe. Generically, we can describe this complete merger history with a phenomenological model of the form (Madau & Dickinson 2014; Madau & Fragos 2017)

$$\mathcal{R}(z) = \mathcal{C}(\alpha, \beta, z_p) \frac{\mathcal{R}_0 (1+z)^{\alpha}}{1 + \left(\frac{1+z}{1+z_p} \right)^{\alpha+\beta}}, \quad (2)$$

allowing a source-frame merger rate that initially evolves as $\mathcal{R}(z) \propto (1+z)^{\alpha}$, reaches a maximum near z_p , and subsequently falls as $\mathcal{R}(z) \propto (1+z)^{-\beta}$. The example binary black hole merger rate shown in Figure 1, for example, is well fit by this phenomenological model using $\alpha = 1.9$, $\beta = 3.4$, and $z_p = 2.4$, shown via a dashed gray curve. The normalization constant $\mathcal{C}(\alpha, \beta, z_p) = 1 + (1+z_p)^{-\alpha-\beta}$ ensures that $\mathcal{R}(0) = \mathcal{R}_0$.

At present the direct detection of binary black holes with Advanced LIGO and Virgo can offer no meaningful constraints on z_p or β . In O2, the range within which Advanced LIGO could detect a typical $30 + 30 M_{\odot}$ binary black hole (averaging over sky location and binary orientation; see Chen et al. 2017) was $z \lesssim 0.5$; in the future O5 observing run this range may be pushed to $z \lesssim 1.2$ (Abbott et al. 2019c). Meanwhile, if the black hole merger rate roughly follows the star formation rate, it should peak at $z_p \gtrsim 2$, well beyond our ability to probe with direct detections.

We have another piece of information at our disposal, however. Although individually undetectable, the superposition of *all* distant binary black holes gives rise to a stochastic gravitational-wave background, detectable in the form of excess cross-power between widely separated detectors (Romano & Cornish 2017; Christensen 2019). The stochastic gravitational-wave background is conventionally described by a dimensionless energy–density spectrum (Allen & Romano 1999)

$$\Omega(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}, \quad (3)$$

where $\frac{d\rho_{\text{GW}}}{d \ln f}$ is the present-day energy density in gravitational-waves per logarithmic frequency interval and $\rho_c = \frac{3H_0^2 c^2}{8\pi G}$ is the universe’s critical energy density. Here, c is the speed of light, G is Newton’s constant, and H_0 is Hubble’s constant; we adopt $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

The energy density arising from the population of binary black hole mergers is given by (Phinney 2001)

$$\Omega(f) = \frac{f}{\rho_c} \int_0^{z_{\text{max}}} dz \frac{\mathcal{R}(z) \left\langle \frac{dE_c}{df} \Big|_{f(1+z)} \right\rangle}{(1+z)H(z)}. \quad (4)$$

Here, $\langle dE_s/df_s \rangle$ is the source-frame energy spectrum radiated by a single binary (Ajith et al. 2008), averaged over the binary black hole population. If the intrinsic parameters of individual binary black holes (e.g., their masses and spins) are denoted by ϕ and have distribution $p(\phi)$, then

$$\left\langle \frac{dE_s}{df_s} \right\rangle = \int d\phi p(\phi) \frac{dE_s}{df_s}(\phi). \quad (5)$$

Note that in Equation (4) we evaluate $\langle dE_s/df_s \rangle$ at the source-frame frequency $f(1+z)$. Meanwhile, $H(z) = H_0 \sqrt{\Omega_M(1+z)^3 + \Omega_\Lambda}$ is the Hubble parameter at redshift z (neglecting radiation density). We take the energy densities of matter and dark energy to be $\Omega_M = 0.3$ and $\Omega_\Lambda = 0.7$, respectively. Finally, the integral in Equation (4) is taken up to a cutoff redshift z_{\max} ; we fix $z_{\max} = 10$, beyond which we expect virtually no star formation and hence no black hole mergers (assuming stellar progenitors). Alternatively, allowing z_{\max} itself to vary as another free parameter may help to provide constraints on binary black holes of non-stellar origin, like the mergers of primordial black holes (Mandic et al. 2016; Wang et al. 2018; Koushiappas & Loeb 2017).

The energy density, $\Omega(f)$, measured by stochastic searches is, in essence, a weighted integral over the binary black hole merger history $\mathcal{R}(z)$, sensitive to the *total* number of past mergers. Thus, if the local rate \mathcal{R}_0 is independently fixed by direct detections, then knowledge of $\Omega(f)$ provides strong bounds on the possible values of α , β , and z_p . This is true even given a *nondetection* of the gravitational-wave background. To illustrate this, we can consider how the signal-to-noise ratio (S/N) of the gravitational-wave background varies with α and z_p .

Given a model $\Omega_M(f)$ for the true energy-density spectrum, the S/N of the gravitational-wave background is (Allen & Romano 1999)

$$\text{S/N} = \frac{\langle \hat{C} | \gamma \Omega_M \rangle}{\sqrt{\langle \gamma \Omega_M | \gamma \Omega_M \rangle}}. \quad (6)$$

Here,

$$\hat{C}(f) = \frac{1}{T} \frac{20\pi^2}{3H_0^2} f^3 \tilde{s}_1(f) \tilde{s}_2^*(f) \quad (7)$$

is the cross-correlation statistic between the strains $\tilde{s}_1(f)$ and $\tilde{s}_2(f)$ measured by two gravitational-wave detectors (Callister et al. 2017; Romano & Cornish 2017), and we have defined an inner product

$$\langle A | B \rangle = 2T \left(\frac{3H_0^2}{10\pi^2} \right)^2 \int_0^\infty df \frac{\tilde{A}(f) \tilde{B}^*(f)}{f^6 P_1(f) P_2(f)}, \quad (8)$$

where $P_{i(f)}$ is the one-sided noise power spectral density of detector i and T is the total observation time. In the presence of a gravitational-wave background, the expectation value of $\hat{C}(f)$ is

$$\langle \hat{C}(f) \rangle = \gamma(f) \Omega(f) \quad (9)$$

and its variance is $\langle \hat{C}(f) \hat{C}(f') \rangle = \delta(f - f') \sigma^2(f)$, with

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

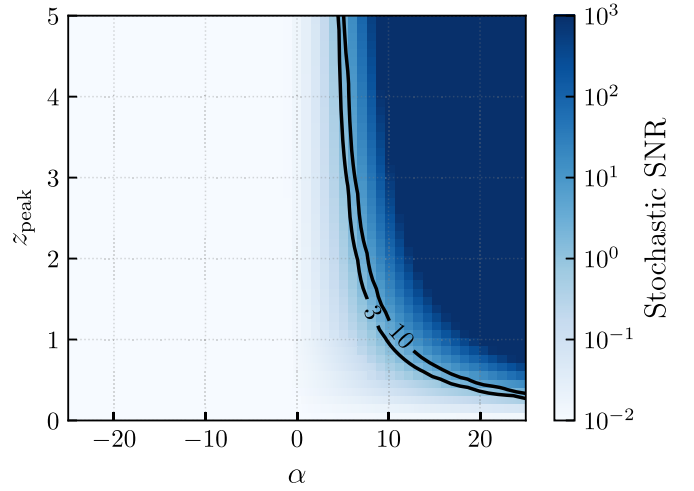


Figure 2. Optimal signal-to-noise ratio with which the binary black hole stochastic background should be visible in Advanced LIGO’s O1 and O2 observing runs, as a function of the leading slope, α , and peak redshift, z_p , of the merger rate $\mathcal{R}(z)$; see Equation (2). For purposes of illustration, we have fixed $\mathcal{R}_0 = 30 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $\beta = 3$, and assumed equal mass binaries with chirp masses $\mathcal{M}_c = 30 M_\odot$. The two black curves trace contours of constant signal-to-noise ratios, at $(\text{S/N})_{\text{opt}} = 3$ and 10. Given our choices of \mathcal{R}_0 , β , and binary mass distribution, the nondetection of a stochastic gravitational-wave background in O1 and O2 (Abbott et al. 2017b, 2019e; Renzini & Contaldi 2019) excludes values of α and z_p at which $\text{S/N}_{\text{opt}} \gtrsim 3$, ruling out a large fraction of the α - z_p parameter space.

The factor $\gamma(f)$ in Equations (6) and (9), known as the overlap reduction function, quantifies the geometrical sensitivity of a given detector pair to an isotropic gravitational-wave background (Christensen 1992; Flanagan 1993). The optimal S/N occurs when we choose a model $\Omega_M(f) = \Omega(f)$ matching the true energy density; the expected S/N in this case is (Allen & Romano 1999)

$$\langle \text{S/N} \rangle_{\text{opt}} = \sqrt{\langle \gamma \Omega | \gamma \Omega \rangle}. \quad (11)$$

In Figure 2 we plot the optimal S/N with which the gravitational-wave background *would* have appeared in O1 and O2 as a function of possible values for α and z_p . In this example we fix $\mathcal{R}_0 = 30 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $\beta = 3$, and assume a population of equal mass binaries with chirp mass $\mathcal{M}_c = 30 M_\odot$. If $\alpha \lesssim 5$, virtually no stochastic signal is expected, consistent with the nondetection of the gravitational-wave background in O1 (Abbott et al. 2017b) and O2 (Abbott et al. 2019e; Renzini & Contaldi 2019). However, the expected S/N rises sharply towards the upper right corner of Figure 2. In particular, if $\alpha \gtrsim 5$ and $z_p \gtrsim 1$, we should have seen an extraordinarily loud stochastic gravitational-wave signal. The fact that no such background was detected means that we can already reject this portion of parameter space, ruling out binary black hole backgrounds rising faster than $\alpha \sim 5$ and peaking beyond $z_p \sim 1$. We note, though, that these exact limits depend strongly on the assumed local merger rate \mathcal{R}_0 and black hole mass distribution (and to a lesser extent on β), and so the results in Figure 2 should be taken as an example only. In Section 3 below, we will instead seek to *simultaneously* measure these different properties, leveraging both the observational limits on the stochastic gravitational-wave background and the current catalog of direct binary black hole detections.

So far, our argument has implicitly assumed that the distribution of binary black hole parameters is independent of redshift, such that the average energy radiated by a given binary (Equation (5)) does not vary with z . This is not necessarily the case. It is possible, for instance, that black holes born at high redshifts are preferentially more massive, due to the increased stellar masses predicted to occur at low metallicities (Belczynski et al. 2010; Spera et al. 2015; Abbott et al. 2016b), although more recent work suggests that the mass distribution of *merging* binaries may be approximately constant (Mapelli et al. 2019). By neglecting the possibility increased masses at higher redshifts, the constraints we obtain on α and z_p are *conservative*. Given a fixed observational limit on $\Omega(f)$, any presumed increase in the average radiated energy ($\langle dE_s/df_s \rangle$) must be balanced by a decrease in the merger rate $\mathcal{R}(z)$ at high redshifts, yielding stricter limits than those shown in Figure 2. Nevertheless, one could incorporate effects like metallicity-dependent masses in this analysis by amending Equation (5) to additionally include integration over distributions of formation redshifts or progenitor metallicities (Abbott et al. 2016a, 2017b).

3. Peak Redshift Constraints from O1 and O2

The best constraints on $\mathcal{R}(z)$ will come from neither the direct detection of binary black holes nor the gravitational-wave background searches considered separately, but instead from a joint analysis that self-consistently synthesizes both sources of information. In this Letter we perform the first such joint analysis, synthesizing stochastic data and direct black hole observations to hierarchically measure the redshift distribution of binary black hole mergers. We take as inputs the integrated cross-correlation spectrum $\hat{C}(f)$ measured between the LIGO Hanford and Livingston detectors (Abbott et al. 2017b, 2019e) during O1 and O2, as well as parameter estimation results for each of the 10 binary black hole mergers comprising the LIGO and Virgo GWTC-1 catalog (Abbott et al. 2019b).

In order to robustly constrain $\mathcal{R}(z)$, it will also be important to simultaneously fit for the *mass* distribution of binary black holes. In Figure 2, for instance, the exact exclusion region depends on our presumed black hole mass distribution: heavier or lighter black holes would increase or decrease the expected energy density $\Omega(f)$, leading us to draw different conclusions about $\mathcal{R}(z)$ in the case of a stochastic nondetection. Strong degeneracies also exist between the inferred mass and redshift distributions of directly detected black hole mergers (Fishbach et al. 2018). A dearth of detections at large redshifts, for example, simply implies a low merger rate for high-mass binaries, since low-mass binaries go undetected at large distances. This can be explained either by a low overall rate at high redshifts, or by a mass distribution that prefers low-mass binaries.

Consider a population of binary black hole mergers, with a local merger rate per unit comoving volume \mathcal{R}_0 and whose mass and redshift distributions are characterized by parameters Λ . The likelihood of obtaining data $\{d_i\}_{i=1}^{N_{\text{obs}}}$ from N_{obs} direct detections, as well as a stochastic cross-correlation spectrum $\hat{C}(f)$, is

$$p(\hat{C}, \{d_i\} | \Lambda, \mathcal{R}_0) = p_{\text{BBH}}(\{d_i\} | \Lambda, \mathcal{R}_0) p_{\text{stoch}}(\hat{C} | \Lambda, \mathcal{R}_0), \quad (12)$$

which has been factored into a direct-detection and a stochastic term.

The likelihood, $p_{\text{BBH}}(\{d_i\} | \Lambda, \mathcal{R}_0)$, of our direct binary black hole detections is given by (Loredo 2004; Taylor & Gerosa 2018; Mandel et al. 2019)

$$p_{\text{BBH}}(\{d_i\} | \Lambda, \mathcal{R}_0) \propto [N(\Lambda, \mathcal{R}_0) \xi(\Lambda)]^{N_{\text{obs}}} e^{-N(\Lambda, \mathcal{R}_0) \xi(\Lambda)} \times \prod_{i=1}^{N_{\text{obs}}} \frac{\int p(d_i | \phi) p(\phi | \Lambda) d\phi}{\xi(\Lambda)}. \quad (13)$$

Here, $p(d_i | \phi)$ is the likelihood for event i given its component masses m_1 and m_2 and redshift z , together abbreviated as $\phi = \{m_1, m_2, z\}$. Meanwhile, $p(\phi | \Lambda)$ is the ensemble distribution of these source parameters. The quantity $N(\Lambda, \mathcal{R}_0)$ is the total number of binary black hole mergers (both observed and unobserved) expected to occur during our observation time; see Equation (18). Observational selection effects are captured by the factor $\xi(\Lambda)$, the fraction of all binary black holes that we expect to successfully detect. If $P_{\text{det}}(\phi)$ is the probability of successfully detecting an event with parameters ϕ , then

$$\xi(\Lambda) = \int P_{\text{det}}(\phi) p(\phi | \Lambda) d\phi. \quad (14)$$

In our analysis we precompute $P_{\text{det}}(\phi)$ over a grid of masses and redshifts, using the semi-analytic prescription of Finn & Chernoff (1993), and requiring detections to have a matched filter S/N of $\rho > 8$ in a single detector.

In practice, we do not have direct access to the likelihoods, $p(d_i | \phi)$, needed to compute Equation (13). Instead, we have discrete samples $\{\phi_i\}$ drawn from each event's *posterior* distribution $p(\phi | d_i)$, obtained via parameter estimation with Monte Carlo integration or nested sampling (Veitch et al. 2015). Parameter estimation itself is performed while assuming some default prior, $p_{\text{pe}}(\phi)$, that is generally *not* equal to the population prior $p(\phi | \Lambda)$ appearing in Equation (13). To evaluate Equation (13), we must therefore replace the integral with an average over discrete samples, weighting each sample with $p_{\text{pe}}^{-1}(\phi)$ to undo the influence of the prior used in parameter estimation:

$$p_{\text{BBH}}(\{d_i\} | \Lambda, \mathcal{R}_0) \propto [N(\Lambda, \mathcal{R}_0) \xi(\Lambda)]^{N_{\text{obs}}} e^{-N(\Lambda, \mathcal{R}_0) \xi(\Lambda)} \times \prod_{i=1}^{N_{\text{obs}}} \frac{1}{\xi(\Lambda)} \left\langle \frac{p(\phi_i | \Lambda)}{p_{\text{pe}}(\phi_i)} \right\rangle_{\text{samples}}. \quad (15)$$

The stochastic cross-correlation spectrum $\hat{C}(f)$, meanwhile, is generally obtained through the weighted combination of a large number of measurements performed over short $\mathcal{O}(100 \text{ s})$ time segments (Allen & Romano 1999; Romano & Cornish 2017), and so the likelihood $p_{\text{stoch}}(\hat{C} | \Lambda, \mathcal{R}_0)$ is well approximated as a Gaussian (Mandic et al. 2012; Callister et al. 2017):

$$p_{\text{stoch}}(\hat{C} | \Lambda, \mathcal{R}_0) \propto \exp \left[-\frac{1}{2} (\hat{C} - \gamma \Omega_M(\Lambda, \mathcal{R}_0) | \hat{C} - \gamma \Omega_M(\Lambda, \mathcal{R}_0)) \right], \quad (16)$$

where $\Omega_M(\Lambda, \mathcal{R}_0; f)$ is our model energy–density spectrum and we have used the inner product defined in Equation (8).

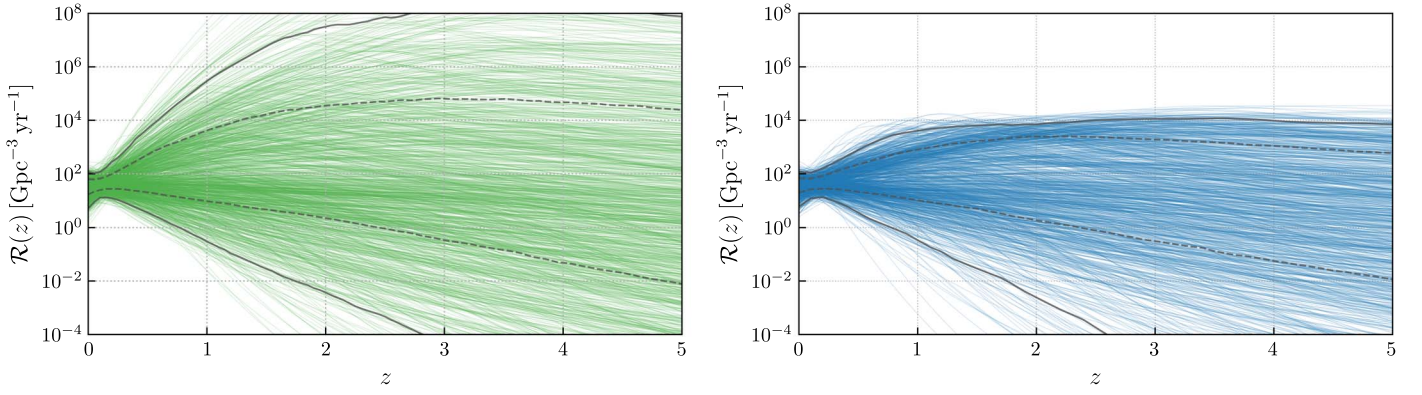


Figure 3. Posterior on the rate density $\mathcal{R}(z)$ of binary black hole mergers as a function of redshift, given the 10 binary black holes comprising GWTC-1 (left), and the joint analysis of these 10 detections with O1 and O2 searches for the stochastic gravitational-wave background (right). The rate density is parameterized as in Equation (2), and the dashed and solid gray curves show the central 68% and 95% credible bounds on $\mathcal{R}(z)$ at each redshift. The direct GWTC-1 detections alone yield a measurement of the local merger rate and marginally constrain the slope α with which the rate evolves at low redshift (see also Figure 4), but give no constraints on the high-redshift behavior of $\mathcal{R}(z)$. The nondetection of a stochastic gravitational-wave background in Advanced LIGO’s O1 and O2 observing runs, meanwhile, imposes an upper limit on the net merger rate across all redshifts. The joint analysis of direct detections and stochastic data can therefore exclude rate densities rising above $\mathcal{R} \gtrsim 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$, placing joint constraints on α and the peak redshift z_p at which $\mathcal{R}(z)$ reaches its maximum (see Figure 5).

We model the intrinsic redshift distribution of binary black hole mergers as

$$p(z|\alpha, \beta, z_p) \propto \frac{1}{1+z} \mathcal{R}(\alpha, \beta, z_p; z) \frac{dV_c}{dz}, \quad (17)$$

where $\mathcal{R}(\alpha, \beta, z_p; z)$ is given in Equation (2) and $\frac{dV_c}{dz}$ is the comoving volume per unit redshift; note that Equation (17), once normalized, is independent of the local merger rate \mathcal{R}_0 . The leading factor of $(1+z)^{-1}$ transforms between source-frame and detector-frame times. Correspondingly, the total number N of mergers expected to occur during our observation time T is

$$N(\alpha, \beta, z_p, \mathcal{R}_0) = T \int_0^{z_{\max}} dz \frac{1}{1+z} \mathcal{R}(\alpha, \beta, z_p, \mathcal{R}_0; z) \frac{dV_c}{dz}. \quad (18)$$

Following Fishbach et al. (2018) and Abbott et al. (2019a), we model the ensemble distribution of primary black hole masses as a power law

$$p(m_1|\kappa, M_{\min}, M_{\max}) \propto \begin{cases} m_1^{-\kappa} & (M_{\min} \leq m_1 \leq M_{\max}) \\ 0 & (\text{else}) \end{cases} \quad (19)$$

and assume a flat distribution

$$p(m_2|m_1, M_{\min}) = \begin{cases} \frac{1}{m_1 - M_{\min}} & (M_{\min} \leq m_2 \leq m_1) \\ 0 & (\text{else}) \end{cases} \quad (20)$$

of secondary masses.

In our analysis we fix $M_{\min} = 5 M_{\odot}$, while hierarchically inferring the parameters $\{\alpha, \beta, z_p, \mathcal{R}_0, \kappa, M_{\max}\}$ of the binary black hole redshift and mass distributions. We adopt the priors listed in Table 1, and perform inference using `emcee` (Foreman-Mackey et al. 2013). For every iteration of our sampler, we evaluate the direct-detection likelihood in Equation (15), using Equation (18) to convert the proposed event rate density \mathcal{R}_0 to a total number of mergers N . We then compute a model stochastic energy–density spectrum, integrating over the proposed mass and redshift distributions (in Equations (5) and (4), respectively) of the binary black hole

Table 1

Priors Placed on the Hyperparameters Describing the Binary Black Hole Mass and Redshift Distributions; See Equations (17)–(20)

Parameter	Prior	Minimum	Maximum
α	Uniform	−25	25
β	Uniform	0	10
z_p	Uniform	0	4
\mathcal{R}_0	Log-uniform	10^{-1}	10^3
κ	Uniform	−4	12
M_{\max}/M_{\odot}	Uniform	30	100

population, thereby evaluating the stochastic contribution (Equation (16)) to the overall likelihood.

Figure 3 shows our resulting posterior on the rate evolution of binary black hole mergers, using the direct GWTC-1 detections alone (left) and combining direct detections with existing stochastic search results (right). Each trace in these figures represents a draw from our $\{\alpha, \beta, z_p, \mathcal{R}_0\}$ posterior. The left panel of Figure 3 is directly comparable to Figure 6 of Abbott et al. (2019a). Figures 4 and 5 show the corresponding posteriors on these parameters, marginalized over κ and M_{\max} . Full parameter estimation results are listed in Table 2.

Direct detections alone allow a measurement of the local merger rate to $\mathcal{R}_0 = 30.1_{-24.9}^{+88.9} \text{ Gpc}^{-3} \text{ yr}^{-1}$ at 95% credibility (the most precise measurement actually occurs at the “waist” seen at $z \sim 0.1$). This is consistent with the results of Abbott et al. (2019a). Direct observations also allow us to roughly constrain α , with a moderate preference for $\alpha \sim 5$ shown in Figure 4. Significant uncertainties remain, however. At 95% credibility, we find $\alpha = 2.3_{-24.7}^{+13.5}$, and, since the α posterior extends all the way to our lower prior bound, we can only robustly constrain $\alpha \leq 13.7$. Direct detections offer no information about β or z_p . Correspondingly, in Figure 3 we have virtually no constraints on the merger rate beyond $z \sim 1$. At $z = 1.5$, for example, the local merger rate could plausibly lie anywhere between 10^{-4} and $10^8 \text{ Gpc}^{-3} \text{ yr}^{-1}$, a range spanning 12 orders of magnitude.

In contrast, the inclusion of O1 and O2 stochastic search data provides a hard upper bound on the high-redshift merger rate;

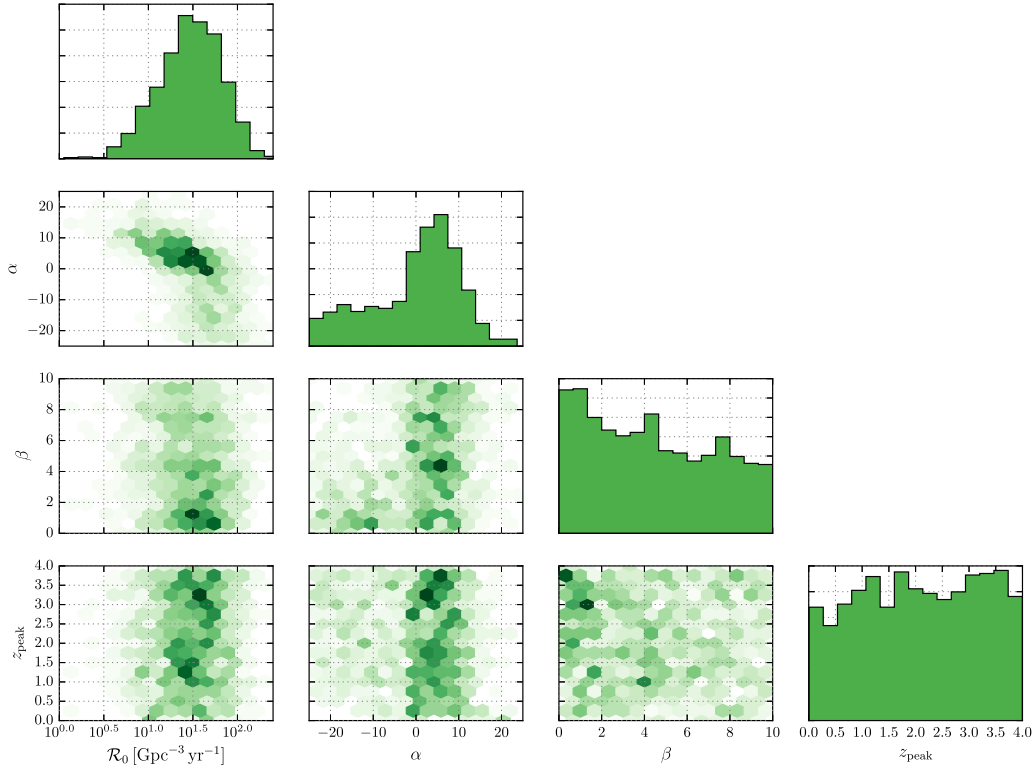


Figure 4. Posterior distribution on the local density \mathcal{R}_0 , leading slope α , trailing slope β , and peak redshift z_p of the binary black hole merger rate $\mathcal{R}(z)$ (Equation (2)), given the 10 binary black hole mergers comprising GWTC-1. We have marginalized over the parameters κ and M_{\max} governing the black hole mass distribution (Equation (19)). The GWTC-1 detections yield marginal constraints on α , but offer no information about z_p or β . This posterior is used to construct the $\mathcal{R}(z)$ samples on the left side of Figure 3. Full parameter estimation results, including bounds on κ and M_{\max} , are given in Table 2.

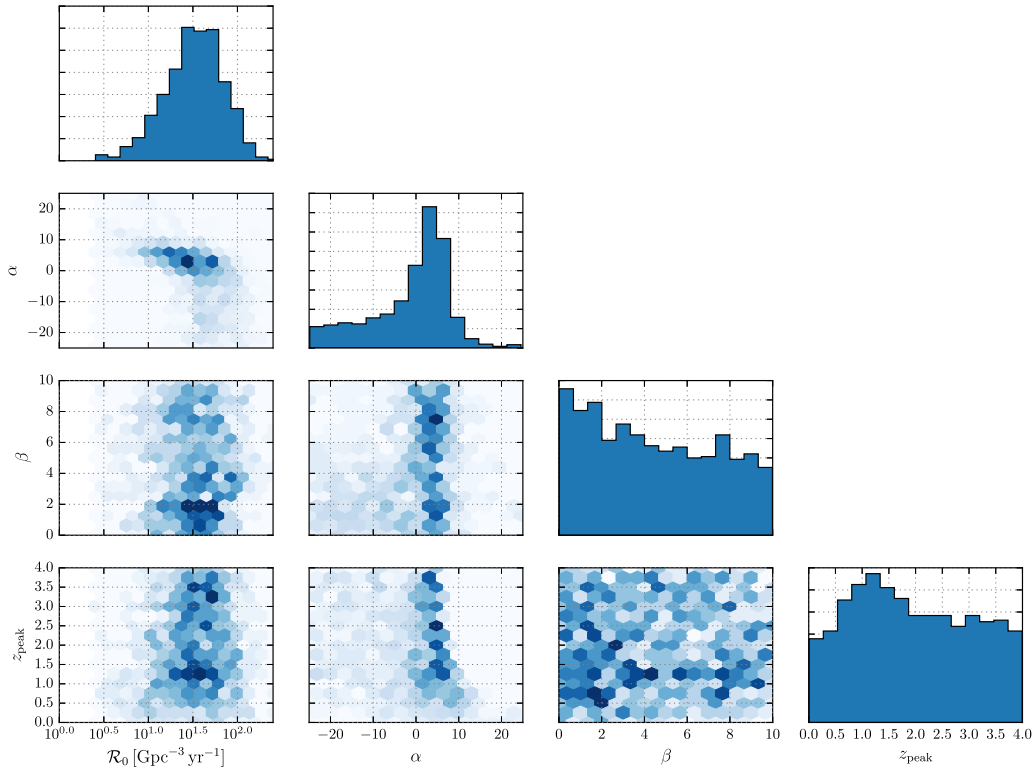


Figure 5. As in Figure 4, but incorporating a joint analysis using the GWTC-1 binary black holes as well as Advanced LIGO limits on the stochastic gravitational-wave background from O1 and O2. Although the inclusion of stochastic measurements does not affect the marginalized one-dimensional posteriors, the nondetection of a gravitational-wave background by Advanced LIGO imposes a joint constraint on α and z_p , ruling out rate densities that evolve faster than $\alpha \gtrsim 7$ and reach maxima at redshifts beyond $z_p \gtrsim 1$. Draws from this posterior are used to generate the rate evolution constraints on the right side of Figure 3.

Table 2
95% Credible Constraints on Parameters Governing the Mass and Redshift Distribution of Binary Black Hole Mergers

Run	M_{\max} (M_{\odot})	κ	α	z_p	\mathcal{R}_0 ($\text{Gpc}^{-3} \text{yr}^{-1}$)
O1–O2: Direct	$42.7^{+18.6}_{-6.4}$	$1.1^{+1.9}_{-2.3}$	≤ 13.7	...	$30.1^{+88.9}_{-24.9}$
O1–O2: Direct/Stochastic	$42.8^{+20.1}_{-6.3}$	$1.0^{+1.7}_{-2.5}$	≤ 10.1	...	$33.6^{+90.1}_{-27.5}$
Design (Mock): Direct	$44.9^{+1.2}_{-1.2}$	$1.3^{+0.3}_{-0.3}$	$3.2^{+0.8}_{-0.6}$	≥ 1.7	$29.2^{+7.2}_{-6.4}$
Design (Mock): Direct/Stochastic	$45.1^{+1.7}_{-1.0}$	$1.2^{+0.3}_{-0.3}$	$3.2^{+1.8}_{-0.9}$	$1.9^{+1.6}_{-1.1}$	$28.9^{+8.5}_{-7.5}$

Note. The second and third lines show true results given data from Advanced LIGO’s O1 and O2 observing runs, using direct observations of binary mergers alone, as well as the synthesis of direct detections with constraints on the stochastic gravitational-wave background. While both the “Direct” and “Direct/Stochastic” O1 and O2 analyses give similar one-dimensional results, the inclusion of stochastic data excludes a nontrivial portion of the joint $\alpha - z_p$ space; see Figure 5. In neither case can we measure α ; instead we place an upper limit. The fourth and fifth lines give parameter estimation results from our mock catalog corresponding to one year of Advanced LIGO observation at design sensitivity. When analyzing mock direct detections alone, we can at best place a lower limit on the peak redshift z_p , while the inclusion of simulated stochastic data allows us to directly measure z_p . None of the four cases give informative marginalized measurements of β , and so this parameter is excluded from the table.

our nondetection of the gravitational-wave background effectively excludes rate densities that rise above $\mathcal{R}(z) \gtrsim 10^4 \text{Gpc}^{-3} \text{yr}^{-1}$. This additional constraint is reflected in Figure 5. While the inclusion of O1 and O2 stochastic data does not notably alter the one-dimensional marginal posteriors, it *does* significantly alter our joint posterior on α and z_p . As argued in Section 2, the nondetection of a stochastic gravitational-wave signal provides a joint constraint on these two parameters, rejecting a large portion of the $\alpha - z_p$ parameter space. When this stochastic exclusion region is combined with the constraint on α from direct GWTC-1 detections, we can already see hints of a preferred contour in the $\alpha - z_p$ plane.

Although the primary goal of this analysis is to measure the evolution of the binary black hole merger rate, it additionally provides a self-consistent framework for predicting the energy density $\Omega(f)$ of the binary black hole background using both the known population properties of local binary black holes and upper limits from Advanced LIGO and Advanced Virgo stochastic searches (Abbott et al. 2017b, 2019e). For every posterior sample in Figure 5 (including the mass parameters M_{\max} and κ not shown there) we compute the corresponding binary black hole energy density using Equation (4). The result, shown in Figure 6, is a prediction for the binary black hole stochastic background that is marginalized over our uncertainty in both the mass distribution and rate evolution of binary black holes, and subject to the measured upper limits from Advanced LIGO.

Within Figure 6, the dashed black curve traces the 95% credible upper limit on $\Omega(f)$ at each frequency. For comparison, the dashed blue curve shows the 2σ “power-law integrated (PI) curve” (Thrane & Romano 2013) quantifying Advanced LIGO’s integrated sensitivity to the gravitational-wave background following O1 and O2; energy–density spectra lying above this curve will generally be observed with $S/N \geq 2$. As expected, the 95% credible limit on $\Omega(f)$ lies nearly tangent to the PI curve. The solid black curve, meanwhile, marks the median predicted energy density. At 25 Hz, this median prediction gives $\Omega(25 \text{ Hz}) = 8.8 \times 10^{-10}$, comparable to the prediction made by Abbott et al. (2019e): $\Omega(25 \text{ Hz}) = 5.3 \times 10^{-10}$. The *uncertainty* on our predicted energy–density spectrum, though, is considerably larger. While the Abbott et al. (2019e) model includes uncertainty on the local rate density \mathcal{R}_0 of binary black hole mergers, it makes stringent assumptions concerning the

subsequent evolution of the merger rate with redshift, assumptions that carry considerable systematic uncertainty. In contrast, Figure 6 includes marginalization over all possible redshift distributions, making this systematic uncertainty explicit.

4. Advanced LIGO at Design Sensitivity

The continued synthesis of direct detections with stochastic search results will offer increasingly strong information regarding the leading slope, α , and peak, z_p , of the binary black hole merger history. Additional binary black holes detected in the local universe will yield ever tighter posteriors on α , while continued time integration by stochastic searches will reject a growing fraction of the joint $\alpha - z_p$ posterior space. Eventually these two effects will meet, converging to produce a true measurement of both α and z_p .

To illustrate this, here we anticipate the results that will soon be possible with design-sensitivity Advanced LIGO. We simulate a mock catalog of 500 binary black hole detections, drawn from a population whose mass distribution is characterized by $\kappa = 1.2$, $M_{\max} = 45 M_{\odot}$, and $M_{\min} = 5 M_{\odot}$. We assume a redshift distribution given by $\alpha = 3$, $\beta = 3$, $z_p = 2$, and $\mathcal{R}_0 = 30 \text{Gpc}^{-3} \text{yr}^{-1}$. With this choice of local merger rate, we would expect to detect these 500 binary black holes after $T \sim 1.2 \text{yr}$ of observation with design-sensitivity Advanced LIGO.

We construct our mock catalog following Fishbach et al. (2018). For each event, we draw an “observed” S/N

$$\rho_{\text{obs}} \sim \mathcal{N}(\rho, 1) \quad (21)$$

from a Gaussian distribution about the event’s true S/N ρ , calculated in a detector with a noise power spectral density given by the Advanced LIGO “design sensitivity” curve of Abbott et al. (2019c). We require our detected events to have $\rho_{\text{obs}} > 8$ in a single detector. For each detected event, we draw an observed maximum-likelihood chirp mass

$$\log \mathcal{M}_{\text{obs}} \sim \mathcal{N}\left(\log \mathcal{M}, \sigma_{\mathcal{M}} \frac{8}{\rho_{\text{obs}}}\right) \quad (22)$$

and symmetric mass ratio

$$\eta_{\text{obs}} \sim \mathcal{N}\left(\eta, \sigma_{\eta} \frac{8}{\rho_{\text{obs}}}\right), \quad (23)$$

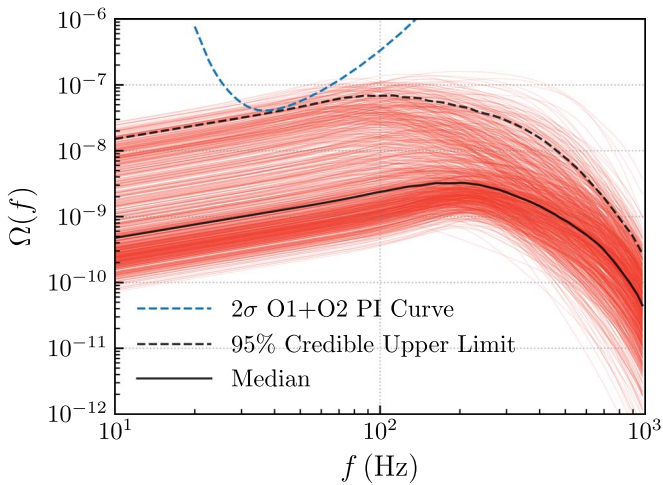


Figure 6. Posterior on the energy–density spectrum $\Omega(f)$ (see Equation (4)) of the binary black hole stochastic background, given the stochastic upper limits and direct binary black hole detections made by Advanced LIGO and Virgo during O1 and O2. Each red trace corresponds to a posterior sample drawn from Figure 5; the range of predictions shown here therefore incorporates our uncertainty in the mass and redshift distributions of binary black holes. The solid and dashed black curves mark the median and 95% credible upper limit on $\Omega(f)$, respectively. For comparison, the dashed blue curve shows Advanced LIGO’s 2σ power-law integrated curve (Thrane & Romano 2013) illustrating its sensitivity to the stochastic background following O2.

where \mathcal{M} and η are the event’s true parameters and we adopt characteristic uncertainties $\sigma_{\mathcal{M}} = 0.08$ and $\sigma_{\eta} = 0.022$. We then draw synthetic likelihood samples about $\log \mathcal{M}_{\text{obs}}$ and η_{obs} , with variances consistent with the above distributions. This prescription gives realistic uncertainties on the measured component masses and distances of binary black hole detections, matching the typical uncertainties reported in Vitale et al. (2017).

We encapsulate a binary’s inclination angle and sky location in a single Finn & Chernoff (1993) projection factor Θ , which quantifies a signal’s amplitude reduction due to suboptimal viewing angles and/or sky placement. If ρ_{opt} is a binary’s optimal S/N (i.e., face-on and directly overhead), then $\Theta\rho_{\text{opt}}$ is the event’s actual S/N. For each mock event, we draw a maximum-likelihood projection factor from

$$\Theta_{\text{obs}} \sim \mathcal{N}\left(\Theta, \sigma_{\Theta} \frac{8}{\rho_{\text{obs}}}\right), \quad (24)$$

where $\sigma_{\Theta} = 0.15$, about which we draw likelihood samples $\{\Theta\}$.

Realistic redshift samples will be strongly correlated with an event’s recovered S/N as well as its projection factor Θ . To capture these correlations, we first draw S/N samples

$$\{\rho\} \sim \mathcal{N}(\rho_{\text{obs}}, 1). \quad (25)$$

Then, noting that ρ is inversely proportional to an event’s luminosity distance D_L , we convert $\{\rho\}$ and $\{\Theta\}$ into luminosity distance samples via

$$\frac{\{D_L\}}{1 \text{ Gpc}} = \rho_{\text{opt}}(1 \text{ Gpc}) \frac{\{\Theta\}}{\{\rho\}}, \quad (26)$$

where $\rho_{\text{opt}}(1 \text{ Gpc})$ is the binary’s optimal S/N at 1 Gpc.

We additionally simulate cross-correlation measurements of the corresponding stochastic gravitational-wave background,

assuming $T = 1.2$ yr of integration with Advanced LIGO’s Hanford–Livingston baseline. Our simulated cross-correlation spectra are drawn from

$$\hat{C}(f) \sim \mathcal{N}\left(\gamma(f)\Omega(f), \sigma(f)\right), \quad (27)$$

where the gravitational-wave background’s energy density $\Omega(f)$ is calculated using Equation (4) and $\sigma(f)$ is given by Equation (10). Given the binary black hole mass and redshift distributions assumed above and a 1.2 yr integration time, the gravitational-wave background has amplitude $\Omega_0 = 2.2 \times 10^{-9}$ at $f = 25$ Hz and $\langle S/N \rangle_{\text{opt}} = 4.2$. In our particular noise realization, the binary black hole background is observed with $S/N = 3.6$, representing a marginal detection.

Figure 7 illustrates the posterior we obtain on $\mathcal{R}(z)$ using our simulated direct detections (left) and direct detections plus stochastic data (right). Figure 8 shows the posterior on \mathcal{R}_0 , α , β , and z_p for this latter case; as before, we have marginalized over the parameters governing the black hole mass distribution. For reference, Figure 8 also includes the one-dimensional marginalized posteriors obtained by direct detections alone (in green). Full parameter estimation results for each case are given in Table 2.

With 500 direct detections we can very precisely measure $\alpha = 3.2_{-0.6}^{+0.8}$ at 95% credibility, yielding a tight fit to $\mathcal{R}(z)$ out to $z \sim 1$. By virtue of not directly observing a turnover of $\mathcal{R}(z)$, we can now place a lower limit $z_p \geq 1.7$. Otherwise, we are again limited by Advanced LIGO’s finite detection range. The joint analysis of our direct detections and stochastic data, meanwhile, yields a qualitatively different picture. Although the S/N of our simulated detection of the gravitational-wave background is somewhat marginal, it provides enough complementary information to rule out large z_p . While the absolute merger rate remains uncertain at large redshifts, this future data would yield a confident measurement of $z_p = 1.9_{-1.1}^{+1.6}$.

5. Conclusions

We present a powerful new constraint on the binary black hole redshift distribution, with implications for stellar evolution, and binary black hole formation and evolution. By combining detections of compact binaries in the local universe with measurements of (or upper limits on) the stochastic gravitational-wave background, we demonstrate that it is possible to explore the binary black hole redshift distribution at redshifts well beyond the present horizon of direct detections. Using existing observations from the Advanced LIGO/Virgo O1 and O2 observing runs, we have obtained novel joint constraints on the low-redshift slope α and peak z_p of the binary black hole merger rate (see Equation (2)). In particular, we can reject merger rates that grow faster than $\alpha \gtrsim 7$ and peak beyond $z_p \gtrsim 1.5$. These constraints will significantly improve with continued observation. Given an approximately year-long observation period with design-sensitivity Advanced LIGO, we have demonstrated the possibility of *directly measuring* z_p .

Although we have taken adopted a decidedly phenomenological model for the merger rate $\mathcal{R}(z)$ in this work, this is not the only possible approach. If, for instance, one were willing to assume that binary black hole *formation* is tied directly to the (potentially metallicity-dependent) star formation rate, as in Figure 1, one could instead seek to

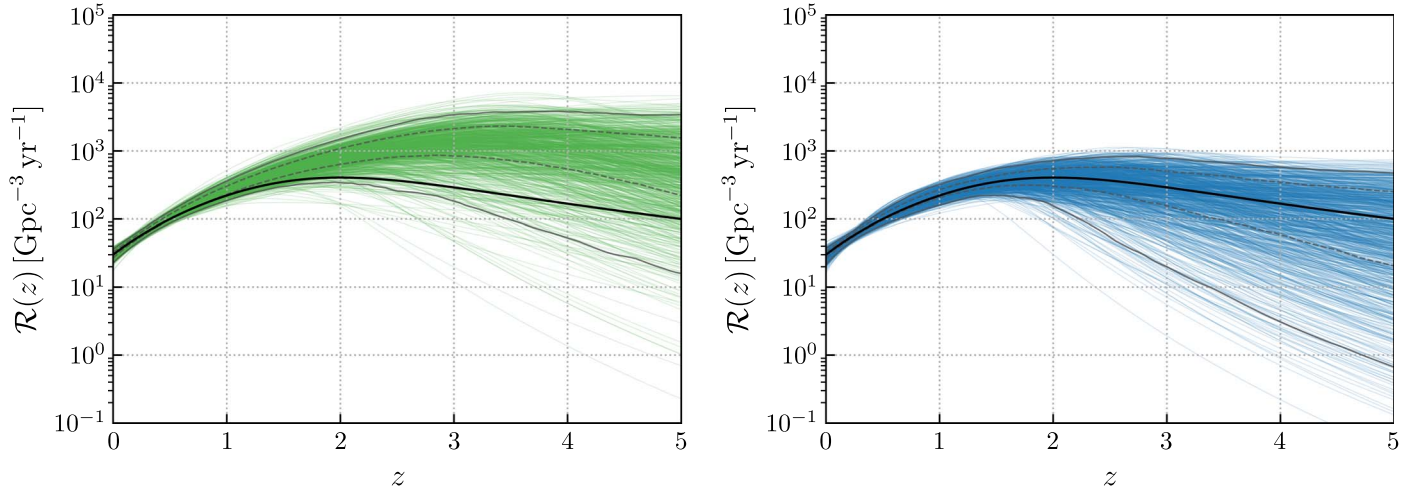


Figure 7. Expected posterior on the rate density $\mathcal{R}(z)$ of binary black hole mergers, given 1.2 yr of observation with design-sensitivity Advanced LIGO. We analyze a catalog of 500 mock detections as well as simulated measurements of the stochastic gravitational-wave background. The left subplot (green) shows results obtained from mock detections alone, while the right subplot shows results given by the synthesis of mock detections with gravitational-wave background measurements. In each case, the dashed and solid gray curves show our 68% and 95% credible symmetric bounds on the merger rate evolution, and the black trace shows the “true” injected merger rate. Although the peak of this merger rate occurs at $z_p = 2$, well beyond Advanced LIGO’s horizon, the joint analysis of direct detections with stochastic data allows us to reconstruct $\mathcal{R}(z)$, yielding the posteriors shown in Figure 8.

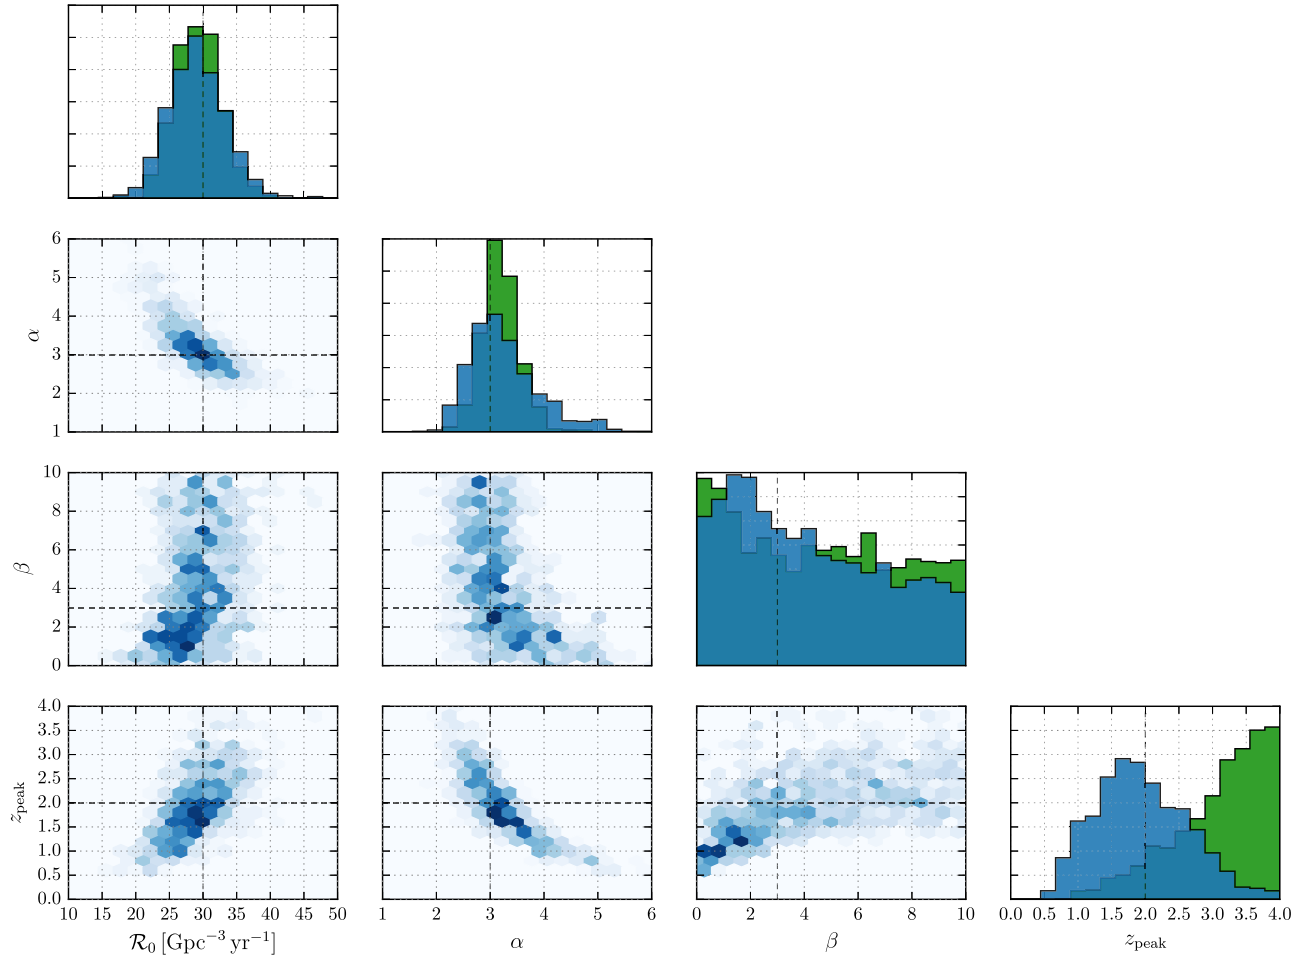


Figure 8. Expected posteriors on the local density \mathcal{R}_0 , leading slope α , trailing slope β , and peak redshift z_p of the binary black hole merger rate after 1.2 yr of Advanced LIGO observation at design sensitivity. The green marginal distributions correspond to the left-hand side of Figure 7, obtained using a mock catalog of direct BBH detections; blue distributions (both one- and two-dimensional) correspond to the right-hand side of Figure 7, given by the synthesis of the BBH catalog with simulated stochastic measurements. We have marginalized over the parameters κ and M_{\max} characterizing the black hole mass distribution. The catalog of synthetic detections provides reasonable measurements of \mathcal{R}_0 and α , but offers only a lower bound on z_p . The addition of stochastic search results imposes an upper bound on z_p ; taken together, we bound $z_p = 1.9_{-1.1}^{+1.6}$ at 95% credibility.

parameterize and measure the metallicity distribution of binary progenitors and the time delay distribution between binary formation and merger.

Looking ahead, future proposed ground-based gravitational-wave detectors like Cosmic Explorer and Voyager may be able to directly measure the rate of binary black hole mergers out to $z \gtrsim 10$ (Vitale et al. 2019). However, even a more limited ability to explore the history of binary black hole mergers with present-day instruments will allow us to ask, sooner rather than later, questions of considerable astrophysical importance: What are the progenitors of compact binary mergers, and when did they form? What is the mean time delay between binary formation and merger? How do black hole mergers across cosmic time connect to the evolution of stars and galaxies in the universe? The combination of individually resolved sources and the unresolved stochastic gravitational-wave background may soon provide answers.

We would like to thank Nelson Christensen, Andrew Matas, and others within the LIGO Scientific Collaboration and Virgo Collaboration for helpful comments and conversation. We additionally thank the anonymous referee, whose questions and feedback greatly improved the quality of this work. T.C. and W.M.F. thank the Simons Foundation for its generous support of the Flatiron Institute. T.C. was also partially supported by the Josephine de Karman Fellowship Trust. M.F. was supported by the NSF Graduate Research Fellowship Program under grant DGE-1746045. M.F. and D.E.H. were supported by NSF grant PHY-1708081. They were also supported by the Kavli Institute for Cosmological Physics at the University of Chicago through an endowment from the Kavli Foundation. DEH also gratefully acknowledges support from the Marion and Stuart Rice Award. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

ORCID iDs

Tom Callister  <https://orcid.org/0000-0001-9892-177X>
 Maya Fishbach  <https://orcid.org/0000-0002-1980-5293>
 Daniel E. Holz  <https://orcid.org/0000-0002-0175-5064>
 Will M. Farr  <https://orcid.org/0000-0003-1540-8562>

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, *CQGra*, **32**, 074001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *PhRvL*, **116**, 131102
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *ApJL*, **818**, L22

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *PhRvL*, **118**, 121102
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *PhRvL*, **118**, 121101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *PhRvL*, **120**, 091101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *ApJL*, **882**, L24
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, *PhRvX*, **9**, 031040
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019c, arXiv:1304.0670
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019d, *PhRvD*, **100**, 062001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019e, *PhRvD*, **100**, 061101
 Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *CQGra*, **32**, 024001
 Adhikari, S., Fishbach, M., Holz, D. E., Wechsler, R. H., & Fang, Z. 2020, arXiv:2001.01025
 Ajith, P., Babak, S., Chen, Y., et al. 2008, *PhRvD*, **77**, 104017
 Allen, B., & Romano, J. D. 1999, *PhRvD*, **59**, 102001
 Baibhav, V., Berti, E., Gerosa, D., et al. 2019, *PhRvD*, **100**, 064060
 Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010, *ApJ*, **714**, 1217
 Belczynski, K., Holz, D. E., Bulik, T., & O’Shaughnessy, R. 2016, *Natur*, **534**, 512
 Callister, T., Biscoveanu, A. S., Christensen, N., et al. 2017, *PhRvX*, **7**, 041058
 Chatziioannou, K., Cotesta, R., Ghonge, S., et al. 2019, *PhRvD*, **100**, 104015
 Chen, H.-Y., Holz, D. E., Miller, J., et al. 2017, arXiv:1709.08079
 Christensen, N. 1992, *PhRvD*, **46**, 5250
 Christensen, N. 2019, *RPPH*, **82**, 016903
 Dominik, M., Belczynski, K., Fryer, C., et al. 2013, *ApJ*, **779**, 72
 Dominik, M., Berti, E., O’Shaughnessy, R., et al. 2015, *ApJ*, **806**, 263
 Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, *Natur*, **548**, 426
 Finn, L. S., & Chernoff, D. F. 1993, *PhRvD*, **47**, 2198
 Fishbach, M., & Holz, D. E. 2017, *ApJL*, **851**, L25
 Fishbach, M., & Holz, D. E. 2020, *ApJL*, **891**, L27
 Fishbach, M., Holz, D. E., & Farr, W. M. 2018, *ApJL*, **863**, L41
 Flanagan, E. E. 1993, *PhRvD*, **48**, 2389
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Koushiappas, S. M., & Loeb, A. 2017, *PhRvL*, **119**, 221104
 Langer, N., & Norman, C. A. 2006, *ApJL*, **638**, L63
 Loredo, T. J. 2004, in AIP Conf. Ser. 735, Bayesian Inference and Maximum Entropy Methods in Science Engineering, ed. R. Fischer, R. Preuss, & U. von Toussaint (Melville, NY: AIP), 195
 Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
 Madau, P., & Fragos, T. 2017, *ApJ*, **840**, 39
 Mandel, I., & Farmer, A. 2018, arXiv:1806.05820
 Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, **486**, 1086
 Mandic, V., Bird, S., & Cholis, I. 2016, *PhRvL*, **117**, 201102
 Mandic, V., Thrane, E., Giampanis, S., & Regimbau, T. 2012, *PhRvL*, **109**, 171102
 Mapelli, M., Giacobbo, N., Ripamonti, E., & Spera, M. 2017, *MNRAS*, **472**, 2422
 Mapelli, M., Giacobbo, N., Santoliquido, F., & Artale, M. C. 2019, *MNRAS*, **487**, 2
 Phinney, E. S. 2001, arXiv:astro-ph/0108028
 Renzini, A. I., & Contaldi, C. R. 2019, *PhRvD*, **100**, 063527
 Rodriguez, C. L., & Loeb, A. 2018, *ApJL*, **866**, L5
 Romano, J. D., & Cornish, N. J. 2017, *LRR*, **20**, 2
 Santoliquido, F., Mapelli, M., Bouffanais, Y., et al. 2020, arXiv:2004.09533
 Spera, M., Mapelli, M., & Bressan, A. 2015, *MNRAS*, **451**, 4086
 Talbot, C., & Thrane, E. 2017, *PhRvD*, **96**, 023012
 Talbot, C., & Thrane, E. 2018, *ApJ*, **856**, 173
 Taylor, S. R., & Gerosa, D. 2018, *PhRvD*, **98**, 083017
 Thrane, E., & Romano, J. D. 2013, *PhRvD*, **88**, 124032
 Veitch, J., Raymond, V., Farr, B., et al. 2015, *PhRvD*, **91**, 042003
 Vitale, S., Farr, W. M., Ng, K. K. Y., & Rodriguez, C. L. 2019, *ApJL*, **886**, L1
 Vitale, S., Lynch, R., Raymond, V., et al. 2017, *PhRvD*, **95**, 064053
 Wang, S., Wang, Y.-F., Huang, Q.-G., & Li, T. G. F. 2018, *PhRvL*, **120**, 191102
 Wysocki, D., Lange, J., & O’Shaughnessy, R. 2019, *PhRvD*, **100**, 043012