

Discovering the Underlying Distributions of Black Hole Populations

Caltech LIGO SURF 2019, Final Report

Phoebe McClincy (phoebemclincy@gmail.com)
Department of Astronomy and Astrophysics, Pennsylvania State University

Mentors: Alan Weinstein, Jonah Kanner, Liting Xiao
LIGO Laboratory, California Institute of Technology
(Dated: 14 October 2019)

The LIGO and Virgo detectors have been observing the cosmos in search of gravitational waves (GW) since 2000. All three detectors were upgraded to Advanced versions, which for LIGO began observing in 2015 and for Virgo in 2017. In Advanced LIGO's first (12 September 2015 to 19 January 2016) and second (30 November 2016 to 25 August 2017) observing runs (O1 and O2, respectively), the detectors found 10 GW signals from binary black hole (BBH) mergers, and 1 from a binary neutron star (BNS) merger, all with high significance, or low probability of being due to instrumental noise fluctuations. Already in the first several months of O3, which began in April 2019, dozens of candidates have been seen with such high significance. The two aforementioned categories, along with neutron star/black hole mergers (NSBH), are collectively known as compact binary coalescence (CBC). In the coming years, as the detectors' sensitivities are improved, we expect to accumulate tens, hundreds, or thousands of CBC events. From such large samples, we expect to be able to infer the underlying population of CBC systems as a function of their masses, component black hole spins, and redshift. This, in turn, will allow us to better understand the astrophysical processes governing the formation, evolution, and final fate of such systems, as tracers of the most massive stars. In this project, we aim to develop tools and techniques to accomplish this through detailed simulation and Bayesian inference.

I. INTRODUCTION

At present, LIGO has catalogued over 30 CBC events, most of which have been detected solely in the third observing run. As our detectors improve, the volume of spacetime in which LIGO is able to observe increases. This quantity is known as the sensitive spacetime volume

$$\langle VT \rangle = \frac{4}{3}\pi D_{avg}^3 T, \quad (1)$$

where D_{avg} represents LIGO's sensitive distance and T represents the observation time of the LIGO detectors. This represents the sensitive spacetime volume in Euclidean space, and this generalized equation no longer holds in the presence of cosmological effects. It is important to note that D_{avg} is a strong function of mass; systems with a larger overall mass produce louder GW, and result in a larger D_{avg} . Because $\langle VT \rangle$ is proportional to the sensitivity of our detectors, we are able to observe larger regions of spacetime as we produce higher-sensitivity detectors.

As larger regions of spacetime are observed, LIGO is expected to recover events in larger and larger numbers. A larger sample size of BBHs offers the unique opportunity to reveal the underlying naturally-occurring relationship between merger rate and BBH masses, spins, and redshift. By studying populations of BBHs, we may better understand the relationship between the progenitor star initial mass function and the mass function

that governs BBH. We also stand to uncover information about models that describe formation channels for BBHs, and which formation channel is most prevalent in nature. We have, and will continue to carry out detailed simulation and Bayesian inference to do so.

II. MOTIVATIONS

A. Measuring binaries with gravitational waves

Gravitational radiation comes from rapidly-changing gravitational fields, and is observed in the form of a wave which propagates at the speed of light. When a GW propagates, it distorts spacetime; we thus observe a change in the length of the arms of our Michelson laser interferometer. This quantity is called strain:

$$h = \frac{\Delta L}{L}, \quad (2)$$

where L is the original length of the detector arm, and ΔL is the change in that length as induced by a GW of strain h .

General relativity allows us to compute the strain of a GW if we know the binary's parameters. Strain depends primarily on the intrinsic parameters, mass and spin, which directly affect the shape of the observed waveform. The other parameters, which are all extrinsic (observer-dependent), govern only the strength of the signal. We can thus analyze the evolution of the waveform's shape

by using Bayesian inference techniques to determine the chirp mass \mathcal{M} , as well as the symmetric mass ratio η . The component masses are implicit in η [1]:

$$\eta = \frac{m_1 m_2}{M_{tot}^2} = \left(\frac{\mathcal{M}}{M_{tot}} \right)^{\frac{5}{3}}. \quad (3)$$

The ultimate goal of this project is to learn about the universe's more massive stars, which lead to the formation of BBHs. The component masses of a BBH are related to the masses of its progenitor stars. More massive stars, such as those that form BBHs, have lower metallicity and were formed early in the universe's evolution (before heavier elements existed). Younger, higher-metallicity stars have masses too small to form BBHs; as such, it is important to determine the underlying distribution of the high-mass progenitor stars so that we may be informed about their evolution into BBHs.

At present, there exist well-supported models of the star formation rate density as a function of m_1 , m_2 , and redshift (z), such as the Salpeter Initial Mass Function (Fig. 1) [2] and the Madau-Dickinson model (Fig. 2) [3]. These models are based on observations from electromagnetic radiation [4, 5]. For BHs, we postulate a relation between BH event rate density \mathcal{R} , z , and mass. We describe \mathcal{R} , or the number of events per unit comoving volume per unit time, as

$$\mathcal{R}(m_1, m_2, z) = \frac{dN}{dV_c dt_s}. \quad (4)$$

We aim to determine the dependence of \mathcal{R} on z , m_1 , and m_2 . Such a measurement will allow us to infer the underlying mass distribution and track the evolution of BBHs in cosmic spacetime.

Although there are six parameters used to describe spin, only two combinations of them dominate in the phase and amplitude evolution of a gravitational waveform: χ_{eff} and χ_p . χ_{eff} is the mass-weighted combination of the component spins along the binary's orbital angular momentum vector [6]. χ_p refers to the components of spin perpendicular to the orbital angular momentum which cause the binary's orbit to precess. These quantities carry information about the mechanisms by which a binary was formed, and are described [7] by

$$\chi_{\text{eff}} = \frac{a_1 m_1 \cos\theta_1 + a_2 m_2 \cos\theta_2}{m_1 + m_2}, \quad (5)$$

and

$$\chi_p = \max\left(a_1 \sin\theta_1, \left(\frac{4m_1 + 3m_2}{4m_2 + 3m_1}\right)\left(\frac{m_1}{m_2}\right)a_2 \sin\theta_2\right), \quad (6)$$

where θ_1 and θ_2 are the angles between the angular momentum vectors of each component and the binary's total angular momentum vector.




FIG. 1. Various initial mass functions following different distributions. We are interested in the Salpeter IMF, shown in dark blue. These curves describe the initial mass distribution for a stellar population. It appears linearly on a log-log plot, as its true nature is a power law with index -2.35. BHs that are detected by LIGO tend to have masses greater than $10 M_{\odot}$, thus our work begins where this plot ends. (Johannes Buchner)




FIG. 2. Madau-Dickinson star formation rate density (rate/unit volume) as a function of redshift. The data points come from many other bodies of work. [3] For BHs, this distribution will be shifted left (to lower redshifts) because it takes time for stars to evolve and collapse into BHs, and even more time for BHs in binaries to merge and be detected in GWs. Additionally, we are limited in the maximum redshift at which we can observe; generally speaking, the cutoff is below $z = 1$. This will improve with next-generation detectors.

Because spin is a higher-order effect in gravitational

waveform evolution, it is more difficult to measure, whereas chirp mass is of first order and its range of values can be more accurately constrained. As such, we will primarily focus on mass for now.

B. Channels of formation

If we are able to determine the underlying distribution of masses, we are then able to gain valuable insights into possible methods by which a binary in question was formed. There are many proposed formation channels, but two (shown in Fig. 3) that are of particular interest to us [8]. The first main channel is common evolution. In one common evolution sub-channel marked by a common envelope phase, the BBH evolves via traditional collapse of both components from a progenitor stellar binary, within a single cloud of gas. In another sub-channel, the BBH forms via chemically homogeneous evolution, in which the orbit of the binary components does not expand traditionally during main sequence helium production, but rather remains compact [9].

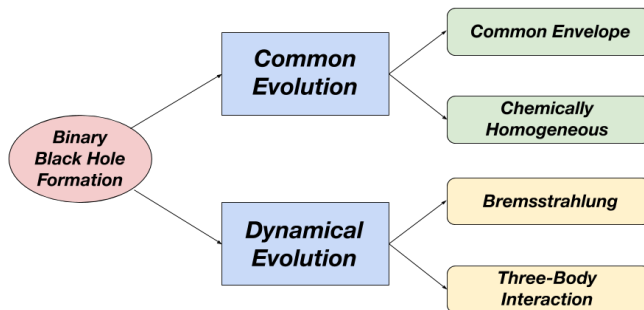


FIG. 3. Formation channels of interest for binary black holes. The first such channel is common evolution, in which binaries are either formed from common envelope evolution or chemically-homogeneous evolution. The second such channel is dynamical formation, in which black holes formed in separate environments interact, in either two-body or three-body interactions.

The second main channel is dynamical formation, which entails the interaction of components formed independently of each other. In one such case, a binary system interacts with a third body in a dense stellar cluster, which results in the ejection of the binary’s less massive component and the capture of the third, more massive body into the binary. In another dynamical formation case, one single body captures a second body via the gravitational Bremsstrahlung radiation caused by the acceleration of the second body through the first body’s gravitational field. We aim to be able to distinguish between these formation channels, primarily from mass, then eventually from spin.

C. Populations, mass distributions, and mass gaps

Within the mass distribution of CBC events (shown in Fig. 4), there exist three proposed regions of scarcity (mass gaps): one below $1 M_{\odot}$, one between $\sim 2-5 M_{\odot}$, and another between $\sim 50-150 M_{\odot}$. The first is proposed to exist due to the small likelihood that traditional stellar collapse mechanisms would produce any BHs in this region; if there were BHs in this region, they could be of primordial nature [10]. The second is thought to be caused by a disparity in NS and BH masses; NS masses gather between $1-2 M_{\odot}$, whereas BH masses tend to begin around $5 M_{\odot}$ [11]. The third hypothesized mass gap is proposed to exist due to pulsational pair-instability supernovae [12], in which progenitor binary stars with component masses between $100-150 M_{\odot}$ eject a significant amount of their mass upon going supernova. This theoretically causes the subsequent BBH component masses to settle around $\sim 40-65 M_{\odot}$ [13]. We can thus infer truths about this particular mass gap by measuring the underlying mass distribution of BBH.

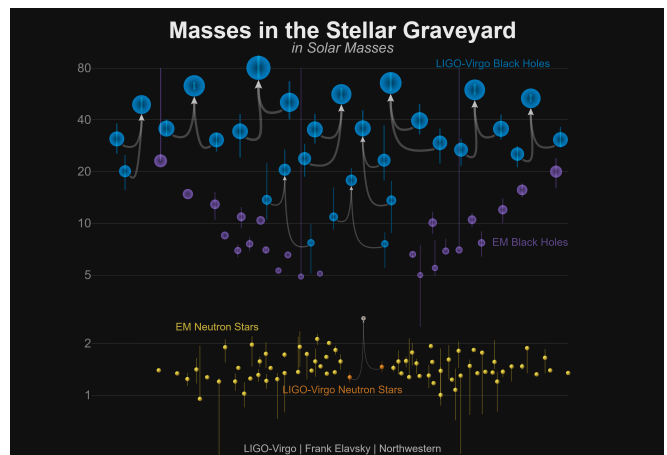


FIG. 4. Distribution of NS and BH masses. Yellow represents NS detected via electromagnetic radiation, and orange represents NS detected via Advanced LIGO. Purple represents BHs detected via electromagnetic radiation, and blue represents BHs detected via Advanced LIGO. There are regions of scarcity below $1 M_{\odot}$ and between $\sim 2-5 M_{\odot}$, as well as a possible mass gap between $\sim 50-150 M_{\odot}$. The x-axis does not have significance in this figure.

III. METHODS

A. Bayesian inference

We have ultimately used the method of Bayesian inference for this project. The first step in this method is to construct a posterior distribution, given by:

$$p(\theta|d) = \frac{\mathcal{L}(d|\theta)\pi(\theta)}{\mathcal{Z}}. \quad (7)$$

The posterior distribution is the probability density of θ , which describes the parameters of the model, given the strain data from the detectors, d . $\mathcal{L}(d|\theta)$ represents the likelihood function of the strain data given the parameters of the model, $\pi(\theta)$ represents the prior distribution of the model parameters, and \mathcal{Z} is the normalization factor, also called the evidence:

$$\mathcal{Z} \equiv \int d\theta \mathcal{L}(d|\theta) \pi(\theta) \quad (8)$$

We can thus define evidences for both the signal and the noise:

$$\mathcal{Z}_{signal} \equiv \int d\theta \mathcal{L}(d|\theta) \pi(\theta), \quad (9)$$

$$\mathcal{Z}_{noise} \equiv \mathcal{L}(d|n). \quad (10)$$

We can then define the Bayes factor, or the ratio of evidence, for signal and noise to be:

$$BF_N^S = \frac{\mathcal{Z}_{signal}}{\mathcal{Z}_{noise}}. \quad (11)$$

We can use this technique to determine preferred models, as well. For example, given two models 1 and 2, we can produce a Bayes factor:

$$BF_2^1 = \frac{\mathcal{Z}_1}{\mathcal{Z}_2}. \quad (12)$$

We can thus compare different models and determine which model produces the best fit for our actual data [14].

B. Mathematical model

The actual number of events that we have observed (or our collected data) is described as N . The true number of events that occur in nature is described as

$$\hat{N}_{true} = \int \frac{dN}{dm_1 dm_2 dz dt_s} dm_1 dm_2 dz dt_s. \quad (13)$$

The component masses m_1 and m_2 , as well as source time t_s and z , can be written as a series of parameters called $\vec{\theta}$. We assume that \hat{N}_{true} depends on m_1 and m_2 , the distribution of which we will describe by hyperparameters α and β , respectively. We also assume that \mathcal{R} has a dependence on z , which is described by another hyperparameter γ . Collectively, we denote these hyperparameters $\vec{\lambda}$. Spin is also believed to be a dependent of \hat{N}_{true} ; however, we will not address spin at this point in the project. We can thus rewrite Eqn. (13) as

$$\hat{N}_{true} = \int \frac{dN(\vec{\lambda})}{d\vec{\theta}} d\vec{\theta}, \quad (14)$$

where

$$\frac{dN(\lambda)}{d\vec{\theta}} = \mathcal{R}(1+z)^\gamma f(m_1|\alpha) f(m_2|\beta) \frac{dV_c}{dz} \frac{dt_d}{dt_s} \frac{1}{T_d}, \quad (15)$$

with t_d being the time as measured at the detector (which has been dilated with respect to the time at the source due to cosmological expansion), T_d the observation time of the detector, and V_c the comoving volume, whose relation to z is determined in accordance with the Λ CDM cosmological model [15]. From this, we can construct an expression for the expected amount of events that we will observe:

$$\hat{N}_{det} = \int \frac{dN(\vec{\lambda})}{d\vec{\theta}} \mathcal{E}(\vec{\theta}) d\vec{\theta}, \quad (16)$$

where $\mathcal{E}(\vec{\theta})$ represents the efficiency of detection. Upon collecting N , we can construct a Poissonian probability distribution for N such that

$$P(N|\hat{N}_{det}, \vec{\lambda}) = \frac{\hat{N}_{det}^N e^{-\hat{N}_{det}}}{N!}. \quad (17)$$

We step through this process until we recover something that can be directly compared to observations. This enables the calculation of a single likelihood for $\vec{\lambda}$; when this cycle is repeated, the *full* likelihood in $\vec{\lambda}$ space can be recovered.

To do this, we have randomly generated a population of masses and have proposed possible values for the hyperparameters in $\vec{\lambda}$. We have developed methodology for measuring $\vec{\lambda}$ using Bayesian inference to compare how close our experimentally-recovered values are to our proposed values. In particular, we have used the dynamic nested sampler `dynesty` [16] to carry this process out in three major steps:

(1) Simulate a dataset consisting of many observed binary systems, following a distribution that is provided and known, with an arbitrary choice of $\vec{\lambda}$.

(2) Use Bayesian inference techniques to recover the most accurate posterior probability distribution for the underlying $\vec{\lambda}$.

(3) Reconstruct the distribution for \hat{N}_{true} .

This has been attempted by other groups. In [17], three models were presented for the BBH primary mass distribution, denoted Models A, B, and C (Fig. 5). Model A fixes m_{min} to be $5 M_\odot$, and allows m_{max} to vary. Model B allows both mass limits to vary. Model C allows multiple functions to describe the distribution; a second component of Gaussian nature appears due to the




FIG. 5. Differential merger rate distribution for BBHs as a function of primary mass and mass ratio (q) for proposed Models A, B, and C in [17]. At lower masses, Model C follows a power law distribution, and at higher masses, it follows a Gaussian distribution. This distribution is based on data from O1 and O2 only (10 BBH mergers). In this project, we use Model C.

pair instability in massive progenitor stars. As such, for Model C, a power law distribution fits at lower masses, and a Gaussian distribution fits at higher masses. In this case, α and β generally refer to the power law indices, and γ describes the Gaussian component. We have used Model C for the analysis done in this project.

IV. PROCEDURE AND RESULTS

We initially generated a distribution of masses (Fig. 9) that followed a power law (in particular, the Salpeter Initial Mass Function, $N = M^{-2.35}$). We then used `dynesty` for parameter estimation, and generated posterior distributions for both hyperparameters of the power law, the amplitude, with true value 1, and power law slope, with true value -2.35. Overall, the true values and the data values were consistent, as shown in Fig. 6.

We then began work on writing scripts for use in this project using Python 3.7 code. There is a series of four scripts that have been written to carry out the work for this project.

The first such script focuses on the generation of one “injection.” An injection is a waveform that we have created from predetermined parameters (such as mass, spin, etc.). To do this, we have used LIGO-created software called `PyCBC`. This code allows one to use already-defined programs and functions to generate waveforms, in both the time (Fig. 7) and frequency (Fig. 8) domains. As such, we used this script to first generate a waveform in the time domain.

We then applied more `PyCBC` code to translate the ideal waveform to what would be seen in the detector’s frame. Because CBC systems are usually very far away in the universe and are relativistic in nature, the waveform that is emitted from the CBC system is not what is observed




FIG. 6. Posterior probability distributions for amplitude and power law slope as recovered by `dynesty`, along with the parameters’ correlation.

on Earth. We observe the now redshifted and time-dilated event occurring at lower frequencies and higher masses, and must subsequently correct for this.

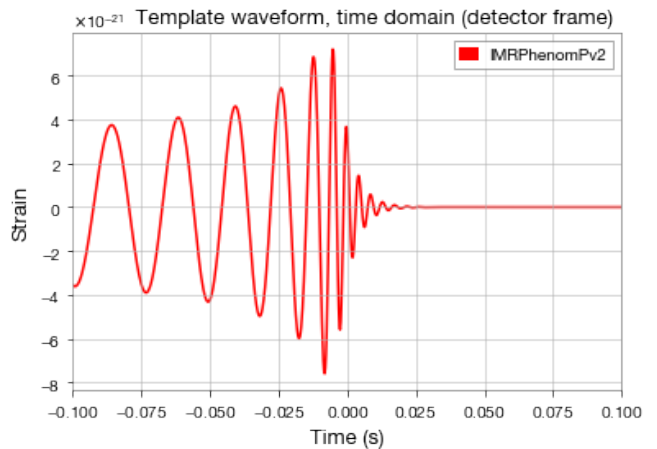


FIG. 7. A gravitational wave pictured in the time domain. Zoomed in on the merger and ringdown. Inspiral occurs until -0.01s, at which point merger causes the signal to sharply spike, and ringdown follows.

Because we are concerned with populations of BBH, we must repeat this process for many iterations. We created a second script that does just this. The second script calls the “make injection” function that was defined in the first script repeatedly. It is also here that the parameters for each waveform are randomly generated from their respective distributions. Spin is set to 0 in the x and y dimensions, but the z -dimension spin is drawn from a Gaussian distribution centered at 0.1 with a standard deviation of 0.2 (Fig. 10). The right ascension, polarization angle, coalescence phase, inclination, and time are all drawn from uniform distributions. Dec-

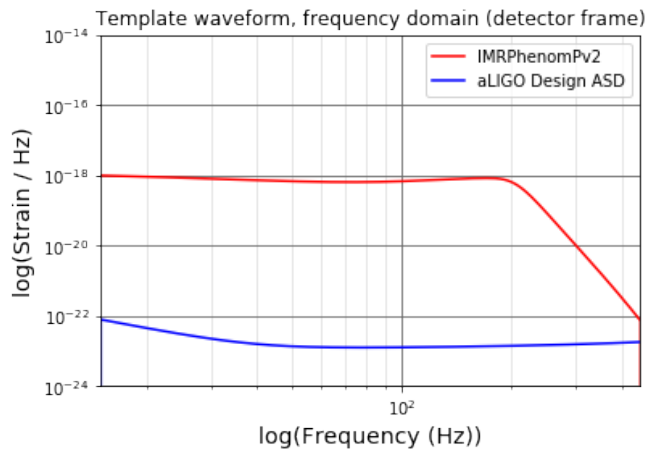


FIG. 8. A gravitational wave pictured in the frequency domain, shown in red. The blue curve represents the Advanced LIGO design sensitivity curve. The signal is well above the sensitivity curve, indicating LIGO’s ability to recover it. Merger occurs around 60 Hz, with a decrease in slope. Ring-down is shown at the “bend” in the signal curve, around 200 Hz, and continues to around 400 Hz.

lination (δ) is determined by drawing from a distribution uniform in $\sin(\delta)$. Distance is determined by drawing from a distribution uniform in comoving volume; here, we must be sure to determine what the maximum distance (rather, redshift) will be. These parameters are generated for each individual waveform. The second script stores these parameters in a file.

It is important to consider whether or not we would be able to recover these simulated waveforms in application. The efficiency referenced in Eqn.(16) quantifies our ability to recover the signal. To test this, we have written a third script that recovers the optimal signal-to-noise ratio (SNR) from each dataset. This can be done two ways; the first is to run a matched-filter-based search on each dataset using PyCBC code to recover the SNR time series, and respectively the peak SNR. The second is to calculate the optimal SNR for each dataset using

$$\rho^2_{opt} = \int \frac{\hat{h}^*(f)h(f)}{\mathcal{S}(f)} df, \quad (18)$$

where $\hat{h}^*(f)$ and $h(f)$ both represent the waveform (since we are working with the optimal SNR), and $\mathcal{S}(f)$ is the PSD (power spectral density, generated from Advanced LIGO’s colored noise curve). We have used the second method. These mathematics must be done in the frequency domain, so it is essential to first use PyCBC code to convert both the template and signal to frequency series, as the PSD is already in the frequency domain. To determine if we will recover or miss an injection, we have computed both the individual detector SNRs and the network SNR. We have accounted for noise (shown in Fig. 11) by drawing a random number from a Gaussian centered at 0 with a standard deviation of 1; we will then

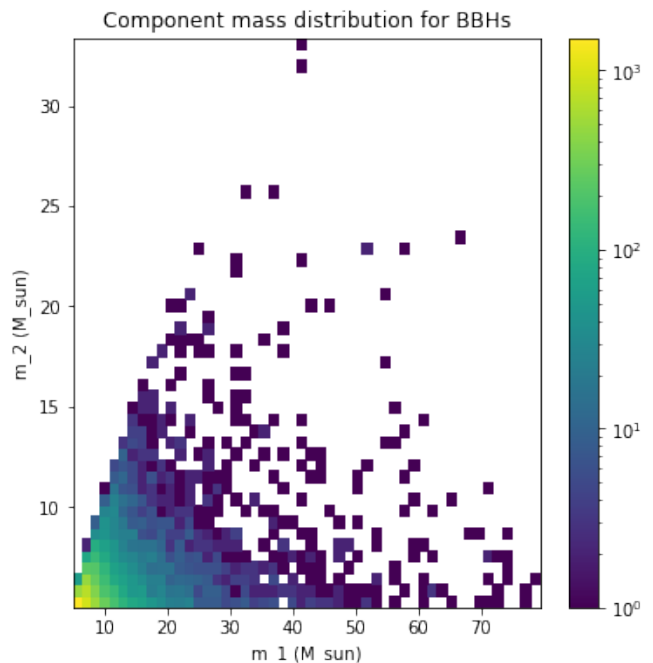


FIG. 9. 2-D histogram for primary mass vs. secondary mass. For 10,000 total events, the mass pairs were drawn from a power law distribution (following the Salpeter Initial Mass Function). Most events occur within the lower end of the mass range.

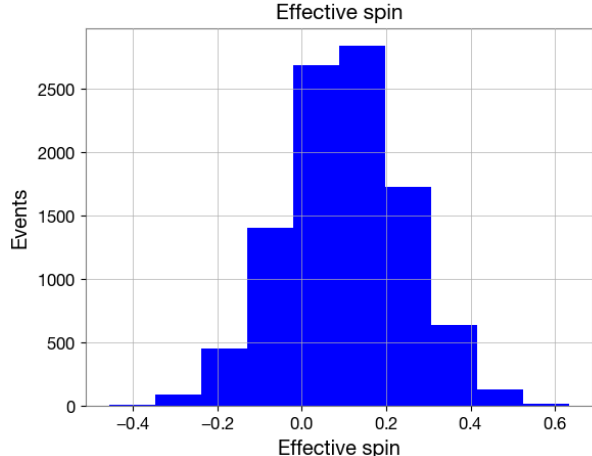


FIG. 10. Histogram of effective spin for 10,000 events, as calculated from Eqn.(4). The z-component of each BH’s spin is drawn from a Gaussian distribution centered at 0.1 with standard deviation 0.2.

add this number to the SNR value. For an injection, if any detector’s SNR is greater than 4, it is factored into the network SNR; otherwise, it is not. Then, if the injection has a network SNR above the threshold of 9, it is recovered; otherwise, the injection is missed. To differentiate between the two, recovered injections are assigned a value of 1, and missed injections are assigned a value

of 0. We can then see how the recovered and missed distributions compare.




FIG. 11. A gravitational wave in the time domain within noise. This noise is simulated by drawing a randomly-generated number from a Gaussian distribution centered at 0 with a standard deviation of 1, and adding it to the SNR.

A fourth script then reads the injections file to retrieve the mass distribution. It runs parameter estimation on the distribution using `dynesty` as the preferred sampler, and can recover the hyperparameters of the distribution (in the initial case of a power law, the amplitude and power law slope). We can then produce, for each distribution, a corner plot detailing the hyperparameters and their respective correlations with the expected values. We are also able to compare the distribution that was generated initially to the distribution that was recovered.

We have only worked with a power law distribution (Fig. 12) for the masses in order to test the scripts and be sure that the hyperparameters recovered via parameter estimation are indeed consistent with the known hyperparameters we have supplied. We have started analysis on Model C from [17], which involves a power law component at low masses, but evolves into a Gaussian distribution at higher masses. For the sake of time, we have generated the masses within the power law component, and will move on to the Gaussian component at a later time.

V. CONCLUSIONS AND FUTURE WORK

Overall, we may conclude that the known hyperparameters provided for the Salpeter Initial Mass Function testing were well-recovered by our program. Due to time constraints, we were unable to run parameter estimation on the simulated data from Model C’s power law component, but the hyperparameters should be well-recovered in this distribution, as well. Since we are able to recover

Preliminary generated mass distribution, 100,000 events

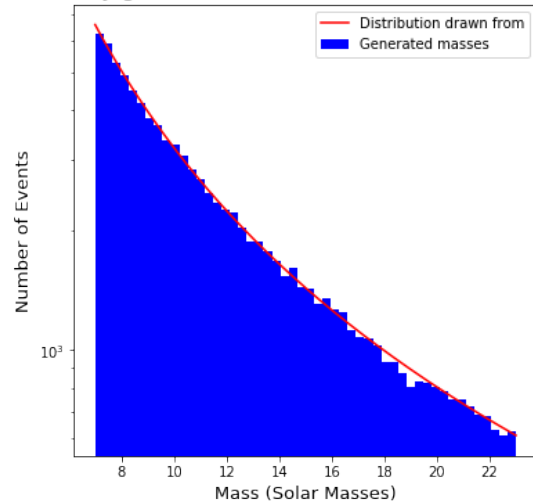


FIG. 12. Masses generated from the Model C power law component; the number of events is shown on a log-scaled axis. We selected a power law index of -2.0 (less than the Salpeter Initial Mass Function, as we know that the slope of the Model C power law component is not as steep). Masses have values between $\sim 7M_{\odot}$ and $\sim 23M_{\odot}$. There are 100,000 total events in this particular distribution.

the hyperparameters we provided well, we should also be able to recover the underlying distributions for real populations of BBHs to a good degree of accuracy.

There are still many steps yet to be taken with this project. The first is to finish the analysis of Model C’s power law component, as well as the Gaussian component. Model C offers many values for parameter estimation, such as m_{\min} , m_{\max} , and the hyperparameters that describe the underlying distribution in terms of m_1 , m_2 , and z . When the values we provide for Model C have been tested and recovered with accuracy, we then plan to move on to more complex models in hopes of determining the redshift dependence of the merger rate. In the future, we may incorporate the event rate’s possible relationship to spin into our distributions, as well.

We additionally intend to focus more heavily on incorporating the efficiency of the detectors. We did not have time to incorporate this into the mass distribution for this project, but in the future, it is important to consider whether or not our detectors would be able to recover a given event. We can simulate this by only analyzing the masses from signals that have an SNR above the threshold we set, as described above.

Moving forward, we must also be cognizant of errors in the mass values. LIGO does not recover an exact value for mass, but rather a range of masses which very likely contains the true value. It is important to acknowledge the potential for error in our recovered distributions due to this.

VI. ACKNOWLEDGEMENTS

The author gratefully acknowledges financial support from the LIGO Scientific Collaboration, the California Institute of Technology, and the National Science Foundation. The author would like to extend many thanks to: Alan Weinstein for sharing his knowledge and input, his

support, and for his complete dedication to the SURF program; Liting Xiao for many hours of help debugging and answering questions; Jonah Kanner for many helpful suggestions in regard to mathematics and programming; Tom Callister for assisting with comprehension of *dynesty* and Bayesian inference; Shreya Anand and Ryan Magee for their input on the author’s presentation/report.

-
- [1] B. S. Sathyaprakash and B. F. Schutz, “Physics, astrophysics and cosmology with gravitational waves,” *Living reviews in relativity*, vol. 12, no. 1, p. 2, 2009.
 - [2] G. Chabrier, “The initial mass function: from salpeter 1955 to 2005,” in *The Initial Mass Function 50 Years Later*, pp. 41–50, Springer, 2005.
 - [3] P. Madau and M. Dickinson, “Cosmic star-formation history,” *Annual Review of Astronomy and Astrophysics*, vol. 52, pp. 415–486, 2014.
 - [4] A. E. Bauer, N. Drory, G. Hill, and G. Feulner, “Specific star formation rates to redshift 1.5,” *The Astrophysical Journal Letters*, vol. 621, no. 2, p. L89, 2005.
 - [5] F. Mannucci, G. Cresci, R. Maiolino, A. Marconi, and A. Gnerucci, “A fundamental relation between mass, star formation rate and metallicity in local and high-redshift galaxies,” *Monthly Notices of the Royal Astronomical Society*, vol. 408, no. 4, pp. 2115–2127, 2010.
 - [6] K. K. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, “Gravitational-wave astrophysics with effective-spin measurements: Asymmetries and selection biases,” *Physical Review D*, vol. 98, no. 8, p. 083007, 2018.
 - [7] C. Talbot and E. Thrane, “Determining the population properties of spinning black holes,” *Physical Review D*, vol. 96, no. 2, p. 023012, 2017.
 - [8] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, “Illuminating black hole binary formation channels with spins in advanced ligo,” *The Astrophysical Journal Letters*, vol. 832, p. L2, 2016.
 - [9] S. de Mink, M. Cantiello, N. Langer, and O. Pols, “Chemically homogeneous evolution in massive binaries,” in *AIP Conference Proceedings*, vol. 1314, pp. 291–296, AIP, 2010.
 - [10] R. Magee, A.-S. Deutsch, P. McClincy, C. Hanna, C. Horst, D. Meacher, C. Messick, S. Shandera, and M. Wade, “Methods for the detection of gravitational waves from subsolar mass ultracompact binaries,” *Physical Review D*, vol. 98, no. 10, p. 103024, 2018.
 - [11] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, “The black hole mass distribution in the galaxy,” *The Astrophysical Journal*, vol. 725, no. 2, p. 1918, 2010.
 - [12] C. Fryer, S. Woosley, and A. Heger, “Pair-instability supernovae, gravity waves, and gamma-ray transients,” *The Astrophysical Journal*, vol. 550, p. 372, 2001.
 - [13] C. Talbot and E. Thrane, “Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization,” *The Astrophysical Journal*, vol. 856, p. 173, 2018.
 - [14] E. Thrane and C. Talbot, “An introduction to bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models,” *Publications of the Astronomical Society of Australia*, vol. 36, 2019.
 - [15] D. W. Hogg, “Distance measures in cosmology,” *arXiv preprint astro-ph/9905116*, 1999.
 - [16] J. S. Speagle, “*dynesty*: A dynamic nested sampling package for estimating bayesian posteriors and evidences,” *arXiv preprint arXiv:1904.02180*, 2019.
 - [17] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, *et al.*, “Binary black hole population properties inferred from the first and second observing runs of advanced ligo and advanced virgo,” *arXiv preprint arXiv:1811.12940*, 2019.