



# CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO

Carl L. Rodriguez, Meagan Morscher, Bharath Pattabiraman, Sourav Chatterjee, Carl-Johan Haster, and Frederic A. Rasio

Phys. Rev. Lett. **115**, 051101 — Published 30 July 2015

DOI: [10.1103/PhysRevLett.115.051101](https://doi.org/10.1103/PhysRevLett.115.051101)

# Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO

Carl L. Rodriguez,<sup>1</sup> Meagan Morscher,<sup>1</sup> Bharath Pattabiraman,<sup>1,2</sup>  
Sourav Chatterjee,<sup>1</sup> Carl-Johan Haster,<sup>1,3</sup> and Frederic A. Rasio<sup>1</sup>

<sup>1</sup>*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Dept. of Physics and Astronomy, Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208, USA*

<sup>2</sup>*Dept. of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL, USA*

<sup>3</sup>*School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT, United Kingdom*

(Dated: July 6, 2015)

The predicted rate of binary black hole mergers from galactic fields can vary over several orders of magnitude and is extremely sensitive to the assumptions of stellar evolution. But in dense stellar environments such as globular clusters, binary black holes form by well-understood gravitational interactions. In this letter, we study the formation of black hole binaries in an extensive collection of realistic globular cluster models. By comparing these models to observed Milky Way and extragalactic globular clusters, we find that the mergers of dynamically-formed binaries could be detected at a rate of  $\sim 100$  per year, potentially dominating the binary black hole merger rate. We also find that a majority of cluster-formed binaries are more massive than their field-formed counterparts, suggesting that Advanced LIGO could identify certain binaries as originating from dense stellar environments.

## INTRODUCTION

By the end of this decade, the Advanced LIGO and Virgo detectors are expected to observe gravitational waves (GWs), ushering in a new post-electromagnetic era of astrophysics [1, 2]. The most anticipated sources of observable GWs will be the signals generated by mergers of binaries with compact object components, such as binary neutron stars (NSs) or binary black holes (BHs). While coalescence rates of NS-NS or BH-NS systems can be constrained from observations, it is not currently possible to produce observationally-motivated rate predictions for BH-BH mergers [3]. Typical detection rates of binary BH (BBH) mergers in galaxies can span several orders of magnitude from  $0.4 \text{ yr}^{-1}$  to  $1000 \text{ yr}^{-1}$  with a fiducial value of  $\sim 20 \text{ yr}^{-1}$  [4]; however, these estimates typically ignore the large numbers of BBHs that are formed through dynamical interactions in dense star clusters [5, 6].

The dynamical formation of BBHs is a probabilistic process, requiring a very high stellar density. These conditions are believed to exist within the cores of globular clusters (GCs), very old systems of  $\sim 10^5 - 10^6$  stars with radii of a few parsecs. Approximately 10 Myr after the formation of a GC, the most massive stars explode as supernovae, forming a population of single and binary BHs with individual masses from  $\sim 5M_\odot$  to  $\sim 25M_\odot$  [7]. The BHs, being more massive than the average star in the cluster, sink to the center of the GC via dynamical friction, until the majority of the BHs reside in the cluster core [8]. After this “mass segregation” is complete, the core becomes sufficiently dense that three-body encounters can frequently occur [9], producing BBHs at high rates. In effect, GCs are dynamical factories for BBHs: producing large numbers of binaries within their cores and ejecting them via energetic dynamical encounters.

In this letter, we use an extensive and diverse collection of GC models to study the population of BBHs that Advanced LIGO can detect from GCs. We explore how the observed parameters of a present-day GC correlate with the distribution of BBH inspirals it has produced over its lifetime. We then compare our models to the observed population of Milky Way GCs (MWGCs) and use recent measurements of the GC luminosity function to determine a mean number of BBH inspirals per GC. Finally, we combine these estimates with an updated estimate of the spatial density of GCs in the local universe [10] into a double integral over comoving volume and inspiral masses to compute the expected Advanced LIGO detection rate. We assume cosmological parameters of  $\Omega_M = 0.309$ ,  $\Omega_\Lambda = 0.691$ , and  $h = 0.677$ , consistent with the latest combined Planck results [11].

## COMPUTING THE RATE

We use a collection of 48 GC models generated by our Cluster Monte Carlo (CMC) code, an orbit-averaged Hénon-type Monte Carlo code for collisional stellar dynamics [12]. The models span a range of initial star numbers ( $2 \times 10^5$  to  $1.6 \times 10^6$ ), initial virial radii (0.5 pc to 4 pc), and consider low stellar metallicities ( $Z = 0.0005, 0.0001$ ) and high stellar metallicities ( $Z = 0.005$ ). In addition, the code implements dynamical binary formation via three-body encounters, strong three and four-body binary interactions, and realistic single and binary stellar evolution. See [13] for a complete description of our code and the models used.

Previous studies have explored the contribution of BBHs from GCs to the Advanced LIGO detection rate [14–19]; however, the majority of these studies have relied on either approximate analytic arguments or simpli-

fied  $N$ -body models with  $N \lesssim 10^5$  particles and have assumed a single black hole mass of  $10M_\odot$ . The one exception is [17], which used a Monte Carlo approach to model GCs of a realistic size ( $N = 5 \times 10^5$ ). However, their study only considered GCs of a single mass, and did not extrapolate that result to the observed distribution of GCs in the local universe. Ours is the first study to compare models with all the relevant physics over a range of masses to observed GCs. This comparison enhances our BBH merger rate by more than an order of magnitude over previous results.

We express the rate of detectable mergers per year,  $R_d$ , as the following double integral over binary chirp mass ( $\mathcal{M}_c \equiv (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ ) and redshift:

$$R_d = \iint \mathcal{R}(\mathcal{M}_c, z) f_d(\mathcal{M}_c, z) \frac{dV_c}{dz} \frac{dt_s}{dt_0} d\mathcal{M}_c dz. \quad (1)$$

This equation is similar to that found in [20, 21]. The components of Eqn. 1 are as follows:

- $\mathcal{R}(\mathcal{M}_c, z)$  is the rate of merging BBHs from GCs with chirp mass  $\mathcal{M}_c$  at redshift  $z$ .
- $f_d(\mathcal{M}_c, z)$  is the fraction of sources with chirp mass  $\mathcal{M}_c$  at redshift  $z$  that are detectable by a single Advanced LIGO detector.
- $dV_c/dz$  is the comoving volume at a given redshift [22].
- $dt_s/dt_0 = 1/(1+z)$  is the time dilation between a clock measuring the merger rate at the source versus a clock on Earth.

This letter focuses on estimating the rate,  $\mathcal{R}(\mathcal{M}_c, z)$ , using our collection of GC models. We assume the rate can be expressed as the product of the mean number of inspirals per GC, the distribution of those sources in  $\mathcal{M}_c - z$  space, and the density of GCs in the local universe, i.e.  $\mathcal{R}(\mathcal{M}_c, z) = \langle N \rangle \times P(\mathcal{M}_c, z) \times \rho_{GC}$ . The spatial density of GCs in the local universe is taken to be  $\rho_{GC} = 0.77 \text{ Mpc}^{-3}$ , based on recent measurements of extragalactic GC systems [23] and modern near-infrared Schechter functions [24]. Note that this estimate, computed in the Supplemental Materials [10], is substantially lower than the previous estimate of  $\rho_{GC} = 8.4 h^3 \text{ Mpc}^{-3}$  from [14] that has been used in previous studies. We now estimate the values of  $\langle N \rangle$  and  $P(\mathcal{M}_c, z)$ .

## MEAN NUMBER OF MERGERS PER CLUSTER

To determine the mean number of BBH inspirals produced by a GC, we use the collection of models to explore how the present-day observable parameters of GCs relate to the number of BBHs it has produced over its lifetime. To quantify the realism of a particular model,




FIG. 1. The number of BBH inspirals per model at 12 Gyr for each of our 48 GC models as a function of final cluster mass. We show the weighted linear regression (with  $1\sigma$  uncertainties on the slope) for the low and high metallicity models.

we compare the total masses and concentrations of our models to GCs observed in the Milky Way. The concentrations are measured by considering the ratio of a cluster’s core radius to its half-light radius,  $R_c/R_h$ . This mass-concentration space is similar to the “fundamental plane” of GCs described in [25], with  $R_c/R_h$  in place of the King concentration [26].

Two trends emerge in our models. First, the total number of BBH inspirals over 12 Gyr is nearly linearly proportional to the final cluster mass. Second, the number of inspirals is higher for more compact clusters (those with smaller  $R_c/R_h$ ). Since the model coverage of the  $R_c/R_h$  space is poorer than the coverage of the mass, and since there are no detailed observations of extragalactic GC concentrations, we elect to focus on the linear relationship between a GC’s mass and the number of inspirals it has produced. We perform a weighted linear regression for both low-metallicity and high-metallicity GCs (Fig. 1). The weights are created by generating a kernel density estimate (KDE) [27] of the MWGCs in the fundamental plane, then measuring the probability of each model as reported by the KDE. In other words, GC models that are more likely to represent draws from the distribution of MWGCs are more heavily weighted.

We compute the mean number of inspirals per GC by multiplying the linear relationships from Fig. 1 by the mass distribution of GCs. Recent work [28] has suggested that the distribution of GC luminosities is universal and well-described by a log-normal distribution:

$$\frac{dN}{d \log L} = N_0 \exp\left(-\frac{(\log L - \log L_0)^2}{2\sigma_L^2}\right) \quad (2)$$

with  $\log L_0 = 5.24$  and  $\sigma_L = 0.52$ . Assuming a mass-to-light ratio of 2 in solar units [28, 29], we convert this luminosity function to a mass function. We then compute

the average number of inspirals per GC by integrating the linear relationship over the normalized GC mass function. The results for both metallicities and different high-mass cutoffs are shown in Table I.

	Mass Cutoff		
Metallicity	$4 \times 10^6 M_\odot$	$2 \times 10^7 M_\odot$	$2 \times 10^8 M_\odot$
Low	430	967	1512
High	830	1954	3103

TABLE I. The mean number of inspirals per GC over 12 Gyr. The result depends on our choice of maximum GC mass. We consider cutoffs of  $4 \times 10^6 M_\odot$  (the approximate mass of the most massive MWGC,  $\omega$  Cen),  $2 \times 10^7 M_\odot$  (the approximate cutoff used in [28]), and  $2 \times 10^8 M_\odot$  (the mass of the ultra-compact dwarf M60-UCD1 [30]).

### DISTRIBUTION OF INSPIRALS

The numbers quoted in Table I provide us with the mean number of BBH inspirals from a GC over 12 Gyr. We could use this average rate to compute a detection rate for Advanced LIGO. However, it is qualitatively obvious that the mass distribution of BBH sources is not constant in time (Fig. 2).

Therefore, we must use the distribution of BBH inspiral events over time from GCs to compute the rate. We select inspirals randomly from each of our models, drawing more inspirals from models with higher weights according to the following scheme:

$$W(M, R_c/R_h) = \frac{KDE_{MWGC}(M, R_c/R_h)}{KDE_{Models}(M, R_c/R_h)} \quad (3)$$

where the weight,  $W(M, R_c/R_h)$ , of a model with mass  $M$  and compactness  $R_c/R_h$  at 12 Gyr is defined by the ratio of the MWGC KDE at  $M, R_c/R_h$  divided by the KDE of the models themselves, evaluated at  $M, R_c/R_h$ . The reason for these weights is as follows: we wish to draw more samples from models that are more likely to represent MWGCs, but because our collection of models is drawn from a different distribution (the initial conditions from [5]), we cannot simply draw inspirals at random from each model according to how well it represents real GCs. To do so would bias our samples with the distribution that results from our initial conditions. By dividing the probability of a model representing a MWGC by the probability density of our collection of models, our scheme naturally corrects for this. Models unlikely to represent MWGCs have small numerators and low weights. Models with no neighboring models that are likely to represent MWGCs have large numerators and small denominators, yielding high weights. Conversely, models *with* neighbors that are likely to represent

MWGCs will have large numerators and large denominators, yielding smaller weights; however, as we will select some number of inspirals from each of these neighboring models, the cumulative effect is the same.




FIG. 2. A sample distribution of inspirals in redshift from the set of models. The redshift is computed by assuming that the difference between the present day and the inspiral time corresponds to the cosmological lookback time at a given redshift (e.g. [22]). The number of inspirals drawn from each model is proportional to its weight, or how similar it is to the observed distribution of MWGCs. Inspirals of BBHs that were formed primordially are indicated with stars (merged in the cluster) and diamonds (ejected before merger). Inspirals of BBHs formed dynamically are shown as squares (in-cluster) and circles (ejected). Note that there are no binaries that are formed by binary stellar evolution with chirp masses greater than  $\sim 13 M_\odot$  (dashed line). This result is consistent across all models. The blue shaded regions illustrate the regions of parameter space where 50%, 10%, and 1% of sources are detectable by Advanced LIGO.

We show a sample distribution of the chirp masses versus redshift in Fig. 2. We distinguish between two different BBH formation channels: primordial and dynamical. We define primordial BBHs as those that are formed from the supernovae of two main sequence stars in a binary, and whose components were never bound to any other star before merger; conversely, we define dynamical binaries as those that are either formed from two isolated BHs via a three-body encounter, or formed from a higher-order dynamical encounter (a binary-single or binary-binary interaction forming a new binary pair). Primordial binaries can still have their orbital parameters modified by dynamics (via a strong encounter with another BH or BBH), as long as the encounter leaves the primordial BBH intact. One immediately apparent feature is the bi-modality between primordial and dynamical BBHs. Over all of our models, the highest chirp mass that is formed by pure binary stellar evolution is  $\mathcal{M}_c \sim 13 M_\odot$ , as systems with larger progenitors

are disrupted by the supernova kick. This implies that any source from our models with a detected chirp mass greater than  $\sim 13M_{\odot}$  could *only* have formed *dynamically*.

To compare this result to BBHs formed in the field, we generated two additional CMC models, each containing  $5 \times 10^6$  binaries and different metallicities ( $Z = 0.005$  and  $Z = 0.0005$ ). These models were computed without two-body relaxation, binary formation, or strong encounters, and only considered the physics of binary stellar evolution. In this dynamics-free environment, the maximum chirp mass of any BBH was  $\mathcal{M}_c \sim 13M_{\odot}$ . Although this result depends on the metallicity and the physics of stellar evolution, it does suggest that GC dynamics forms BBHs consistently more massive than those in the field.

### DETECTION RATE

We now compute the expected rate of signals detectable by Advanced LIGO. To compute the fraction of detectable sources,  $f_d(\mathcal{M}_c, z)$ , we use gravitational waveforms that cover the inspiral, merger, and ringdown phases of a compact binary merger (known as IMRPhenomC waveforms [31]) and compute the signal-to-noise ratio (SNR) using the projected zero-detuning, high-power configuration of Advanced LIGO [32]. We then marginalize over binary orientation and sky location to determine what fraction of sources at a given chirp mass and redshift yield an  $\text{SNR} > 8$ . This approach is identical to that found in [21]. Note that we have assumed all binary components have equal masses in order to simplify the integral. This assumption is well-justified, as the dynamically-formed BBHs in our models have similar component masses. We assume all BHs to be non-spinning.

Our distribution of inspirals,  $P(\mathcal{M}_c, z)$ , is generated by creating a KDE of the inspirals in Fig. 2. Since each “draw” of inspirals will produce a slightly different distribution, we compute  $P(\mathcal{M}_c, z)$  1000 times, and then take the mean of Eqn. 1 for those 1000 draws. We find that a single Advanced LIGO detector operating at design sensitivity will detect  $\sim 100$  BBHs per year from GCs. Of those about 2/3 will originate from low-metallicity GCs and the rest from high-metallicity GCs, assuming 76% of GCs are low metallicity (consistent with the MWGC distribution). Approximately 80% of these sources will have chirp masses greater than  $13(1+z)M_{\odot}$ , meaning that the majority of BBHs detectable by Advanced LIGO from GCs could only have formed dynamically.

The majority of these BBH sources will be detected at low redshifts. For low-metallicity clusters, the distribution of detectible sources in redshift peaks at  $z \sim 0.3$ , while for high-metallicity clusters the distribution peaks at  $z \sim 0.24$ . In both cases, 90% of detectable sources are located at  $z \lesssim 0.57$ .

To obtain a rough estimate of the uncertainty on this prediction, we perform a simple error analysis that considers the optimistic and conservative rates that would be obtained by varying our assumptions and selecting the  $\pm 1\sigma$  estimates of certain quantities. For the conservative estimate, we assume that the GC mass function truncates at the mass of  $\omega$  Cen ( $4 \times 10^6 M_{\odot}$ ), and that the spatial density of GCs is  $\rho_{GC} = 0.32 \text{ Mpc}^{-3}$  (the conservative estimate from the Supplemental Materials). We also recompute the rate using the  $-1\sigma$  uncertainty from the regression in Fig. 1 and the lower standard deviation of our 1000 draws of  $\mathcal{R}(\mathcal{M}_c, z)$ . This yields a conservative estimate of  $\sim 20$  BBH inspirals per year. Conversely, if we assume the most optimistic truncation mass for GC mass function ( $2 \times 10^8 M_{\odot}$ ), the most optimistic GC spatial density ( $\rho_{GC} = 2.3 \text{ Mpc}^{-3}$ , the optimistic estimate from [10], similar to the value used in previous studies), and the  $+1\sigma$  uncertainties on the linear regression and  $\mathcal{R}(\mathcal{M}_c, z)$ , we find an optimistic rate of  $\sim 700$  BBH inspirals per year. This range is primarily influenced by the uncertainty in the GC spatial density and the truncation mass of the GC mass function.

### CONCLUSION

In this letter, we compared new GC models computed with our CMC code to the observed distributions of Milky Way and extragalactic GCs to predict the expected rate of BBH inspirals from realistic GCs. We determined a linear relationship between the present-day mass of a GC and the number of BBH inspirals produced by that cluster. By combining this with the universal GC luminosity function and a new estimate for the spatial density of GCs, we were able to predict the mean density of BBH inspirals from GCs in the local universe. Then by weighting our models according to their similarity to the observed distribution of MWGCs, we created a distribution of inspiral sources in chirp mass and redshift. Finally, by combining this with the anticipated sensitivity of Advanced LIGO, we estimated a detection rate of  $\sim 100$  BBH inspiral events per year from GC sources. With highly conservative assumptions, this rate drops to  $\sim 20$  events per year, while highly optimistic assumptions pushes the rate as high as  $\sim 700$  events per year.

We also found that no BBHs with chirp masses above  $\sim 13M_{\odot}$  were formed from a primordial binary. In other words, every inspiral with  $\mathcal{M}_c > 13M_{\odot}$  in our models was formed by dynamical processes alone. This could, in theory, provide an easy way to discriminate between binaries that were formed dynamically versus those formed by binary stellar evolution; however, this result is highly dependent on the physics of supernova kicks and the fraction of ejected supernova material which falls back onto the newly formed BH, both of which remain poorly constrained. In addition, recent work has suggested that the

mass distribution of chirp masses for BBHs produced by stellar evolution can reach as high as  $\mathcal{M}_c \sim 30M_\odot$ , depending on the physics of the common envelope [33]. As such, this result should be treated as a proof-of-principle, and not a concrete physical claim. Investigations to better understand the relationship between this formation cutoff, the distribution of supernovae kicks, and the fallback fraction, are currently underway.

Finally, the number of BHs formed is entirely dependent on the choice of the initial mass function. Although our choice of IMF is typical for studies of this type, a variation of  $1\sigma$  in the slope of the high-mass end of the IMF can produce significant differences in the number of BBHs. Investigations to quantify this effect are also underway.

We thank Ilya Mandel for his assistance in computing the fraction of detectable sources, as well as useful discussions. We also thank Ben Farr, Claude-Andre Faucher-Giguere, William Harris, Kyle Willett, Michael Schmidt, and Tyson Littenberg for useful discussions. CR was supported by an NSF GRFP Fellowship, award DGE-0824162. This work was supported by NSF Grant AST-1312945 and NASA Grant NNX14AP92G. CJH acknowledges support from an RAS grant.

- 
- [1] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Class. Quant. Grav.* **32**, 074001 (2015).
- [2] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardini, et al., *Class. Quant. Grav.* **32**, 024001 (2015).
- [3] K. Belczynski, V. Kalogera, and T. Bulik, *Astrophys. J.* **572**, 407 (2002).
- [4] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, et al., *Class. Quant. Grav.* **27**, 173001 (2010), 1003.2480.
- [5] M. Morscher, B. Pattabiraman, C. Rodriguez, F. A. Rasio, and S. Umbreit, *Astrophys. J.* **800**, 9 (2015), 1409.0866.
- [6] A. Sadowski, K. Belczynski, T. Bulik, N. Ivanova, F. A. Rasio, and R. O’Shaughnessy, *Astrophys. J.* **676**, 1162 (2008), 0710.0878.
- [7] K. Belczynski, A. Sadowski, F. A. Rasio, and T. Bulik, *Astrophys. J.* **650**, 303 (2006).
- [8] L. Spitzer, Jr., *Astrophys. Lett.* **158**, L139 (1969).
- [9] N. Ivanova, K. Belczynski, J. M. Fregeau, and F. A. Rasio, *Mon. Not. R. Astron. Soc.* **358**, 572 (2005), 0501131.
- [10] See Supplemental Material [url], which includes Refs.[23, 24, 28, 34–40].
- [11] P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al., *ArXiv e-prints* (2015), 1502.01589.
- [12] K. J. Joshi, F. A. Rasio, S. P. Zwart, and S. Portegies Zwart, *Astrophys. J.* **540**, 969 (2000), 9909115.
- [13] See Supplemental Materials [url], which includes Refs.[3, 5, 12, 41–50].
- [14] S. P. Zwart and S. McMillan, *Astrophys. Lett.* p. 12 (1999), 9910061.
- [15] R. M. O’Leary, F. A. Rasio, J. M. Fregeau, N. Ivanova, and R. O’Shaughnessy, *Astrophys. J.* **637**, 937 (2006).
- [16] S. Banerjee, H. Baumgardt, and P. Kroupa, *Mon. Not. R. Astron. Soc.* **402**, 371 (2010).
- [17] J. M. B. Downing, M. J. Benacquista, M. Giersz, and R. Spurzem, *Mon. Not. R. Astron. Soc.* **416**, 133 (2011), 1008.5060.
- [18] Y.-B. Bae, C. Kim, and H. M. Lee, *Mon. Not. R. Astron. Soc.* **440**, 2714 (2014).
- [19] A. Tanikawa, *Mon. Not. R. Astron. Soc.* **435**, 1358 (2013).
- [20] K. Belczynski, T. Bulik, I. Mandel, B. S. Sathyaprakash, A. A. Zdziarski, and J. Mikolajewska, *Astrophys. J.* **764**, 96 (2013), 1209.2658.
- [21] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. C. Miller, and M. Walczak (2014), 1403.0677.
- [22] D. W. Hogg, *ArXiv Astrophysics e-prints* (1999), astro-ph/9905116.
- [23] W. E. Harris, G. L. H. Harris, and M. Alessi, *Astrophys. J.* **772**, 82 (2013), 1306.2247.
- [24] L. S. Kelvin, S. P. Driver, A. S. G. Robotham, A. W. Graham, S. Phillipps, N. K. Agius, M. Alpaslan, I. Baldry, S. P. Bamford, J. Bland-Hawthorn, et al., *Mon. Not. R. Astron. Soc.* **439**, 1245 (2014), 1401.1817.
- [25] D. E. McLaughlin, *Astrophys. J.* **539**, 618 (2000).
- [26] The King concentration is defined as the logarithm of the core radius over the tidal radius. We use the simpler  $R_c/R_h$ , as the tidal radius can be difficult to determine observationally.
- [27] We select the bandwidths using maximum likelihood cross-validation, since for multi-dimensional distributions, the optimal bandwidth is not necessarily equal along each dimension, particularly when using mixed units. This is implemented in the `StatsModels` Python package.
- [28] W. E. Harris, W. Morningstar, O. Y. Gnedin, H. O’Halloran, J. P. Blakeslee, B. C. Whitmore, P. Côté, D. Geisler, E. W. Peng, J. Bailin, et al., *Astrophys. J.* **797**, 128 (2014).
- [29] E. F. Bell, D. H. McIntosh, N. Katz, and M. D. Weinberg, p. 23 (2003), 0302543.
- [30] J. Strader, A. C. Seth, D. A. Forbes, G. Fabbiano, A. J. Romanowsky, J. P. Brodie, C. Conroy, N. Caldwell, V. Pota, C. Usher, et al., *Astrophys. J.* **775**, L6 (2013).
- [31] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, et al., *Phys. Rev. D* **82**, 064016 (2010).
- [32] Both the noise curve and technical reports describing it can be found under LIGO Document T0900288-v3.
- [33] K. Belczynski, M. Dominik, T. Bulik, R. O’Shaughnessy, C. Fryer, and D. E. Holz, *Astrophys. J.* **715**, L138 (2010).
- [34] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O’Neil, *ArXiv e-prints* (2014), 1403.6015.
- [35] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)* (The MIT Press, 2005), ISBN

- 026218253X.
- [36] E. W. Peng, A. Jordan, P. Cote, M. Takamiya, M. J. West, J. P. Blakeslee, C.-W. Chen, L. Ferrarese, S. Mei, J. L. Tonry, et al., *Astrophys. J.* p. 27 (2008), 0803.0330.
  - [37] P. Schechter, *Astrophys. J.* **203**, 297 (1976).
  - [38] W. Harris, G. Harris, and M. Hudson, ArXiv e-prints (2015), 1504.03199.
  - [39] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber, *Astrophys. J.* **724**, 878 (2010).
  - [40] S. Murray, C. Power, and A. Robotham, *Astronomy and Computing* **3-4**, 23 (2013).
  - [41] K. J. Joshi, C. P. Nave, and F. A. Rasio, *Astrophys. J.* **550**, 691 (2001).
  - [42] B. Pattabiraman, S. Umbreit, W.-k. Liao, A. Choudhary, V. Kalogera, G. Memik, and F. A. Rasio, *Astrophys. J. Suppl. Ser.* **204**, 15 (2013), 1206.5878.
  - [43] J. M. Fregeau, P. Cheung, S. F. Portegies Zwart, and F. A. Rasio, *Mon. Not. R. Astron. Soc.* **352**, 1 (2004).
  - [44] J. R. Hurley, O. R. Pols, and C. A. Tout, *Mon. Not. R. Astron. Soc.* **315**, 543 (2000).
  - [45] J. R. Hurley, C. A. Tout, and O. R. Pols, *Mon. Not. R. Astron. Soc.* **329**, 897 (2002).
  - [46] S. Chatterjee, J. M. Fregeau, S. Umbreit, and F. A. Rasio, *Astrophys. J.* **719**, 915 (2010).
  - [47] P. Peters, *Phys. Rev.* **136**, B1224 (1964).
  - [48] P. Kroupa, *Mon. Not. R. Astron. Soc.* **322**, 231 (2001).
  - [49] S. Djorgovski and G. Meylan, *Astron. J.* **108**, 1292 (1994).
  - [50] W. E. Harris, *Royal Society of London Philosophical Transactions Series A* **368**, 889 (2010), 0911.0798.