The Astrophysical Journal, 825:116 (10pp), 2016 July 10 doi:10.3847/0004-637X/825/2/116

PARAMETER ESTIMATION ON GRAVITATIONAL WAVES FROM NEUTRON-STAR BINARIES WITH SPINNING COMPONENTS

Ben Farr[1], Christopher P. L. Berry[2], Will M. Farr[2], Carl-Johan Haster[2], Hannah Middleton[2], Kipp Cannon[3], Philip B. Graff[4,5], Chad Hanna[6], Ilya Mandel[2], Chris Pankow[7], Larry R. Price[8], Trevor Sidery[2], Leo P. Singer[9], Alex L. Urban[7], Alberto Vecchio[2], John Veitch[2], and Salvatore Vitale[10]
[1] Enrico Fermi Institute and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA; farr@uchicago.edu
[2] School of Physics & Astronomy, University of Birmingham, Birmingham, B15 2TT, UK
[3] Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, Ontario, M5S 3H8, Canada
[4] Department of Physics, University of Maryland–College Park, College Park, MD 20742, USA
[5] Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[6] The Pennsylvania State University, University Park, PA 16802, USA
[7] Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA
[8] LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA
[9] Astrophysics Science Division, NASA Goddard Space Flight Center, Code 661, Greenbelt, MD 20771, USA
[10] Massachusetts Institute of Technology, 185 Albany Street, Cambridge, MA 02139, USA
Received 2015 August 21; revised 2016 March 9; accepted 2016 March 10; published 2016 July 11

ABSTRACT

Inspiraling binary neutron stars (BNSs) are expected to be one of the most significant sources of gravitational-wave signals for the new generation of advanced ground-based detectors. We investigate how well we could hope to measure properties of these binaries using the Advanced LIGO detectors, which began operation in September 2015. We study an astrophysically motivated population of sources (binary components with masses 1.2 $M_{\odot}$–1.6 $M_{\odot}$ and spins of less than 0.05) using the full LIGO analysis pipeline. While this simulated population covers the observed range of potential BNS sources, we do not exclude the possibility of sources with parameters outside these ranges; given the existing uncertainty in distributions of mass and spin, it is critical that analyses account for the full range of possible mass and spin configurations. We find that conservative prior assumptions on neutron-star mass and spin lead to average fractional uncertainties in component masses of ∼16%, with little constraint on spins (the median 90% upper limit on the spin of the more massive component is ∼0.7). Stronger prior constraints on neutron-star spins can further constrain mass estimates but only marginally. However, we find that the sky position and luminosity distance for these sources are not influenced by the inclusion of spin; therefore, if LIGO detects a low-spin population of BNS sources, less computationally expensive results calculated neglecting spin will be sufficient for guiding electromagnetic follow-up.

Key words: gravitational waves – methods: data analysis – stars: neutron – surveys

1. INTRODUCTION

As we enter the advanced-detector era of ground-based gravitational-wave (GW) astronomy, it is critical that we understand the abilities and limitations of the analyses we are prepared to conduct. Of the many predicted sources of GWs, binary neutron-star (BNS) coalescences are paramount; their progenitors have been directly observed Lorimer (2008), and the advanced detectors will be sensitive to their GW emission up to ∼400 Mpc away (Abbott et al. 2016c).

When analyzing a GW signal from a circularized compact binary merger, strong degeneracies exist between parameters describing the binary (e.g., distance and inclination). To properly estimate any particular parameter(s) of interest, the marginal distribution is estimated by integrating the joint posterior probability density function (PDF) over all other parameters. In this work, we sample the posterior PDF using software implemented in the LALInference library (Veitch et al. 2015). Specifically, we use results from LALInfernce_nest (Veitch & Vecchio 2010), a nest sampling algorithm (Skilling 2006), and LALInference_MCMC (Christensen et al. 2004; Röver et al. 2006; van der Sluys et al. 2008), a Markov-chain Monte Carlo algorithm (Gregory 2005, chapter 12).

Previous studies of BNS signals have largely assessed parameter constraints assuming negligible neutron-star (NS) spin, restricting models to nine parameters. This simplification has largely been due to computational constraints, but the slow spin of NSs in short-period BNS systems observed to date (e.g., Mandel & O'Shaughnessy 2010) has also been used as justification. However, proper characterization of compact binary sources *must* account for the possibility of non-negligible spin; otherwise, parameter estimates will be biased (Buonanno et al. 2009; Berry et al. 2015). This bias can potentially lead to incorrect conclusions about source properties and even misidentification of source classes.

Numerous studies have looked at the BNS parameter estimation abilities of ground-based GW detectors such as the Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO; Aasi et al. 2015) and Advanced Virgo (AdV; Acernese et al. 2015) detectors. Nissanke et al. (2010, 2011) assessed localization abilities on a simulated non-spinning BNS population. Veitch et al. (2012) looked at several potential advanced-detector networks and quantified the parameter-estimation abilities of each network for a signal from a fiducial BNS with non-spinning NSs. Aasi et al. (2013) demonstrated the ability to characterize signals from non-spinning BNS sources with waveform models for spinning sources using Bayesian stochastic samplers in the LALInference library (Veitch et al. 2015). Hannam et al. (2013) used approximate methods to quantify the degeneracy between spin and mass estimates, assuming the compact objects' spins are aligned with

the orbital angular momentum of the binary (cf. Haster et al. 2015). Rodriguez et al. (2014) simulated a collection of loud signals from non-spinning BNS sources in several mass bins and quantified parameter estimation capabilities in the advanced-detector era using non-spinning models. Chatziioannou et al. (2015) introduced precession from spin–orbit coupling and found that the additional richness encoded in the waveform could reduce the mass–spin degeneracy, helping BNSs to be distinguished from NS–black hole (BH) binaries. Littenberg et al. (2015) conducted a similar analysis of a large catalog of sources and found that it is difficult to infer the presence of a mass gap between NSs and BHs (Özel et al. 2010; Farr et al. 2011; Kreidberg et al. 2012), although, this may still be possible using a population of a few tens of detections (Mandel et al. 2015). Finally, Singer et al. (2014) and the follow-on Berry et al. (2015) represent an (almost) complete end-to-end simulation of BNS detection and characterization during the first 1–2 years of the advanced-detector era. These studies simulated GWs from an astrophysically motivated BNS population, then detected and characterized sources using the search and follow-up tools that are used for LIGO–Virgo data analysis (Aasi et al. 2014; Abbott et al. 2016b). The final stage of analysis missing from these studies is the computationally expensive characterization of sources while accounting for the compact objects' spins and their degeneracies with other parameters. The present work is the final step of BNS characterization for the Singer et al. (2014) simulations using waveforms that account for the effects of NS spin.

We begin with a brief introduction to the source catalog used for this study and Singer et al. (2014) in Section 2. Then, in Section 3 we describe the results of parameter estimation from a full analysis that includes spin. In Section 3.1 we look at mass estimates in more detail and spin-magnitude estimates in Section 3.2. In Section 4 we consider the estimation of extrinsic parameters: sky position (Section 4.1) and distance (Section 4.2), which we do not expect to be significantly affected by the inclusion of spin in the analysis templates. We summarize our findings in Section 5. A comparison of computational costs for spinning and non-spinning parameter estimation is given in the Appendix.

2. SOURCE SIMULATION AND SELECTION

We have restricted our study to the first year of the advanced-detector era, using the same 250 simulations that Singer et al. (2014) analyzed with non-spinning parameter estimation. For these, Gaussian noise was generated using the "early" 2015 aLIGO noise curve found in Barsotti & Fritschel (2012). Approximately 50,000 BNS sources were simulated, using the SpinTaylorT4 waveform model (Buonanno et al. 2003, 2009), a post-Newtonian inspiral model that includes the effects of precession, to generate the GW signals. Component masses were uniformly distributed between 1.2 $M_\odot$ and 1.6 $M_\odot$, which reflects the range of observed BNS masses (Özel et al. 2012). Component spins were isotropically oriented, with magnitudes $\chi_{1,2} = c|\boldsymbol{S}_{1,2}|/Gm_{1,2}^2$ drawn uniformly between 0 and 0.05; here, $|\boldsymbol{S}_{1,2}|$ are the NSs' spin angular momenta and $m_{1,2}$ are their masses (the indices 1 and 2 correspond to the more and less massive components of the binary, respectively). The range of simulated spin magnitudes was chosen to be consistent with the observed population of short-period BNS systems, currently bounded by PSR J0737−3039 A (Burgay et al. 2003; Brown et al. 2012) from above. Finally, sources were distributed uniformly in volume (i.e., uniform in distance cubed) to a maximum distance at which the loudest signal would produce a network signal-to-noise ratio (S/N) of $\rho_{\rm net} = 5$ (Singer et al. 2014), where $\rho_{\rm net} = \sum_i \rho_i^2$ is the individual detector S/Ns ρ_i combined in quadrature.

Of this simulated population, detectable sources were selected using the GSTLAL_INSPIRAL matched-filter detection pipeline (Cannon et al. 2012) with a single-detector S/N threshold $\rho > 4$ and false alarm rate (FAR) threshold of FAR $< 10^{-2}$ yr^{-1}. The FAR for real detector noise is largely governed by non-stationary noise transients in the data that can mimic GWs from compact binary mergers, which, as Berry et al. (2015) demonstrate, make negligible difference to parameter estimation for the (low-FAR, BNS) signals considered here. Because our simulated noise is purely stationary and Gaussian with no such artifacts, FAR estimates are overly optimistic. To compensate, an additional threshold on the network S/N of $\rho_{\rm net} > 12$ was applied. This S/N threshold is consistent with the above FAR threshold when applied to data similar to previous science runs (Berry et al. 2015; Abbott et al. 2016c). A random subsample of 250 detections was selected for parameter estimation with LALINFERENCE.[11] The mass and spin distributions of this subset are statistically consistent with those that the sources were drawn from (Berry et al. 2015). See Singer et al. (2014) for more details regarding the simulated data and GSTLAL_INSPIRAL analyses.

3. SPINNING ANALYSIS

Singer et al. (2014) detail the detection, low-latency localization, and medium-latency (i.e., non-spinning) follow-up of the simulated signals in 2015. In this work, we perform the expensive task of full parameter estimation that accounts for non-zero compact-object spin. Whereas Singer et al. (2014) used the (non-spinning) TaylorF2 waveform model, we make use of the SpinTaylorT4 waveform model (Buonanno et al. 2003, 2009), parameterized by the 15 parameters that uniquely define a circularized compact binary inspiral.[12]

We assume the objects to be point masses with no tidal interactions. The estimation of tidal parameters using post-Newtonian approximations is rife with systematic uncertainties that are comparable in magnitude to statistical uncertainties (Wade et al. 2014; Yagi & Yunes 2014). Though marginalizing over uncertainties in tidal parameters can affect estimates of other parameters, the fact that tidal interactions only impact the evolution of the binary at late times (only having a measurable impact at frequencies above ∼450 Hz; Hinderer et al. 2010) limits both their measurability and the resulting biases in other parameter estimates caused by ignoring them (Damour et al. 2012).

The simulated population of BNS systems contains slowly spinning NSs with masses between 1.2 $M_\odot$ and 1.6 $M_\odot$ and spin magnitudes $\chi < 0.05$. This choice was motivated by the

[11] The mean (median) $\rho_{\rm net}$ of the set of 250 events is 16.7 (14.6).

[12] The 15 parameters are two masses (either component masses or the chirp mass and mass ratio); six spin parameters describing the two spins (magnitudes and orientations); two coordinates for sky position; distance; an inclination angle; a polarization angle; a reference time, and the orbital phase at this time (see Veitch et al. 2015, for more details). The masses and spins are intrinsic parameters that control the evolution of the binary, while the others are extrinsic parameters that describe its orientation and position.

characteristics of NSs found thus far in Galactic BNS systems expected to merge within a Hubble time through GW emission. However, NSs *outside* of BNS systems have been observed with spins as high as $\chi = 0.4$ (Hessels et al. 2006; Brown et al. 2012) and, depending on the NS equation of state (EOS), could theoretically have spins as high as $\chi \lesssim 0.7$ (Lo & Lin 2011) without breaking up. For these reasons, the prior assumptions used for Bayesian inference of source parameters are broader than the spin range of the simulated source population.

To simulate a real analysis scenario where the class of compact binary and the NS EOS are not known, we use uniform priors in component masses between 0.6 $M_\odot$ and 5.0 $M_\odot$ to avoid any prior constraints on mass posteriors, and our standard BH spin prior: uniform in spin magnitudes $\chi_{1,2} \sim U(0, 1)$ and isotropic in spin orientation. Prior distributions for the location and orientation of the binary match that of the simulated population, i.e., isotropically oriented and uniform in volume (out to a maximum distance of 218.9 Mpc, safely outside the detection horizon, which is ~137 Mpc for a 1.6 $M_\odot$–1.6 $M_\odot$ binary).[13] Choosing any particular upper bound for spin magnitude would require either assuming hard constraints on NS spin-up, which are based upon observations with hard-to-quantify selection effects, or making assumptions regarding the unknown EOS of NSs. For these reasons, we choose not to rule out compact objects with high spin a priori by using an upper limit of $\chi < 1$, encompassing all allowed NS and BH spins. In Section 3.3 we look at more constraining spin priors, and, particularly, how such choices can affect mass estimates.

We describe parameter-estimation accuracy using several different quantities, depending upon the parameter of interest:

1. The simplest is the fractional uncertainty $\sigma_x/\langle x\rangle$, where σ_x and $\langle x\rangle$ are the standard deviation and mean of the distributions for parameter x, respectively. This is particularly useful for showing how uncertainty scales with S/N: in the limit of high S/N, the standard deviation can be approximated from the (inverse) Fisher matrix and scales inversely with the S/N (Vallisneri 2008).
2. The credible interval CI_p^x in the range that contains the central p of the integrated posterior, with $(1-p)/2$ falling both above and below the limits (Aasi et al. 2013). Specifying the credible interval for several values of p gives information about the shape of the posterior.
3. As an alternative to credible intervals, we use credible upper or lower bounds. These are the one-sided equivalents of credible intervals and are useful for distributions that are peaked toward one end of the parameter range or for parameters we are interested in putting a limit upon (the spin magnitude satisfies both of these criteria).
4. For sky localization, we use credible regions (the two-dimensional generalization of the credible interval), which are the smallest sky areas that encompass a given total posterior probability. The credible region for a total posterior probability p is defined as

$$\mathrm{CR}_p = \arg\max_A \int_A d\mathbf{\Omega} P_\Omega(\mathbf{\Omega}), \quad (1)$$

where $P_\Omega(\mathbf{\Omega})$ is the posterior PDF over sky position $\mathbf{\Omega}$, and A is the sky area integrated over (Sidery et al. 2014).

We also consider the searched area A_*, the area of the smallest credible region that includes the true location.

To check that differences between our spinning and non-spinning analyses were a consequence of the inclusion of spin and not because of a difference between waveform approximants, we also ran SpinTaylorT4 analyses with spins fixed to $\chi_{1,2} = 0$. There were no significant differences in parameter estimation between the non-spinning TaylorF2 and zero-spin SpinTaylorT4 results for any of the quantities we examined.[14] Therefore, we only use the TaylorF2 results to illustrate the effects of neglecting spin.

3.1. Mass Estimates

To maximize sampling efficiency, model parameterizations are chosen to minimize degeneracies between parameters. To leading order, the post-Newtonian expansion of the waveform's phase evolution depends on the *chirp mass*, $\mathcal{M}_\mathrm{c} = (m_1 m_2)^{3/5}(m_1 + m_2)^{-1/5}$, making it a *very* well constrained parameterization of binary mass. The second mass parameter used is the mass ratio $q = m_2/m_1$, where $0 < q \leqslant 1$. Detectors are much less sensitive to the mass ratio, and strong degeneracies with spin make constraints on q even worse (Cutler & Flanagan 1994). It is primarily the uncertainty in q that governs the uncertainty in component masses m_1 and m_2.

Figure 1 shows the superimposed one-dimensional marginal posterior PDFs and cumulative density functions (CDFs) for the chirp mass (centered on each mean) and mass ratio for all 250 events. As a representation of a typical event's posterior distribution, we show the average PDFs and CDFs, where the average is taken over all 250 posterior PDFs and CDFs at each point. Chirp-mass distributions are usually well approximated by normal distributions about the mean, while mass ratio estimates have broad support across most of the prior range. The simulated population had a narrower range between 0.75 and 1.

To trace individual parameter uncertainties across the population we use the fractional uncertainties in chirp mass $\sigma_{\mathcal{M}_\mathrm{c}}/\langle\mathcal{M}_\mathrm{c}\rangle$ and mass ratio $\sigma_q/\langle q\rangle$. The chirp mass and mass ratio conveniently cover mass space (which is why they are used for sampling), but the total mass $M = m_1 + m_2$ is also of interest for determining the end product of the merger, so we also plot the fractional uncertainty $\sigma_M/\langle M\rangle$. The mean (median) fractional uncertainties in chirp mass, mass ratio, and total mass for the simulated population are 0.0676% (0.0642%), 28.7% (28.4%), and 6.15% (5.81%), respectively. For comparison, the mean (median) fractional uncertainties in chirp mass, mass ratio, and total mass from the non-spinning analysis are 0.0185% (0.0165%), 8.90% (8.79%), and 0.542% (0.491%), respectively. We further examine the impact of spin on mass measurements in Section 3.3.

[13] The mean (median) true distance for the set of 250 events is 52.1 Mpc (47.8 Mpc), and the maximum is 124.8 Mpc.

[14] Using, as an example, the chirp mass (the most precisely inferred parameter), we can compare the effects of switch from a non-spinning to a spinning analysis to those from switching waveform approximants by comparing the difference in the posterior means $\langle\mathcal{M}_\mathrm{c}\rangle$. The difference between means from the SpinTaylorT4 analyses with and without spin is an order of magnitude greater than the difference between the zero-spin SpinTaylorT4 and TaylorF2 analyses: defining the log ratio $\xi = \log_{10}(|\langle\mathcal{M}_\mathrm{c}\rangle^\mathrm{S} - \mathcal{M}_\mathrm{c}\rangle^0|/|\langle\mathcal{M}_\mathrm{c}\rangle^\mathrm{NS} - \mathcal{M}_\mathrm{c}\rangle^0|)$, where the superscripts S, 0, and NS indicate results of the fully spinning SpinTaylorT4, the zero-spin SpinTaylorT4, and the non-spinning TaylorF2 analyses respectively, the mean (median) value of ξ is 0.90 (1.04), and 92.4% of events have $\xi > 0$ (indicating that the shift in the mean from introducing spin is larger than the shift from switching approximants).

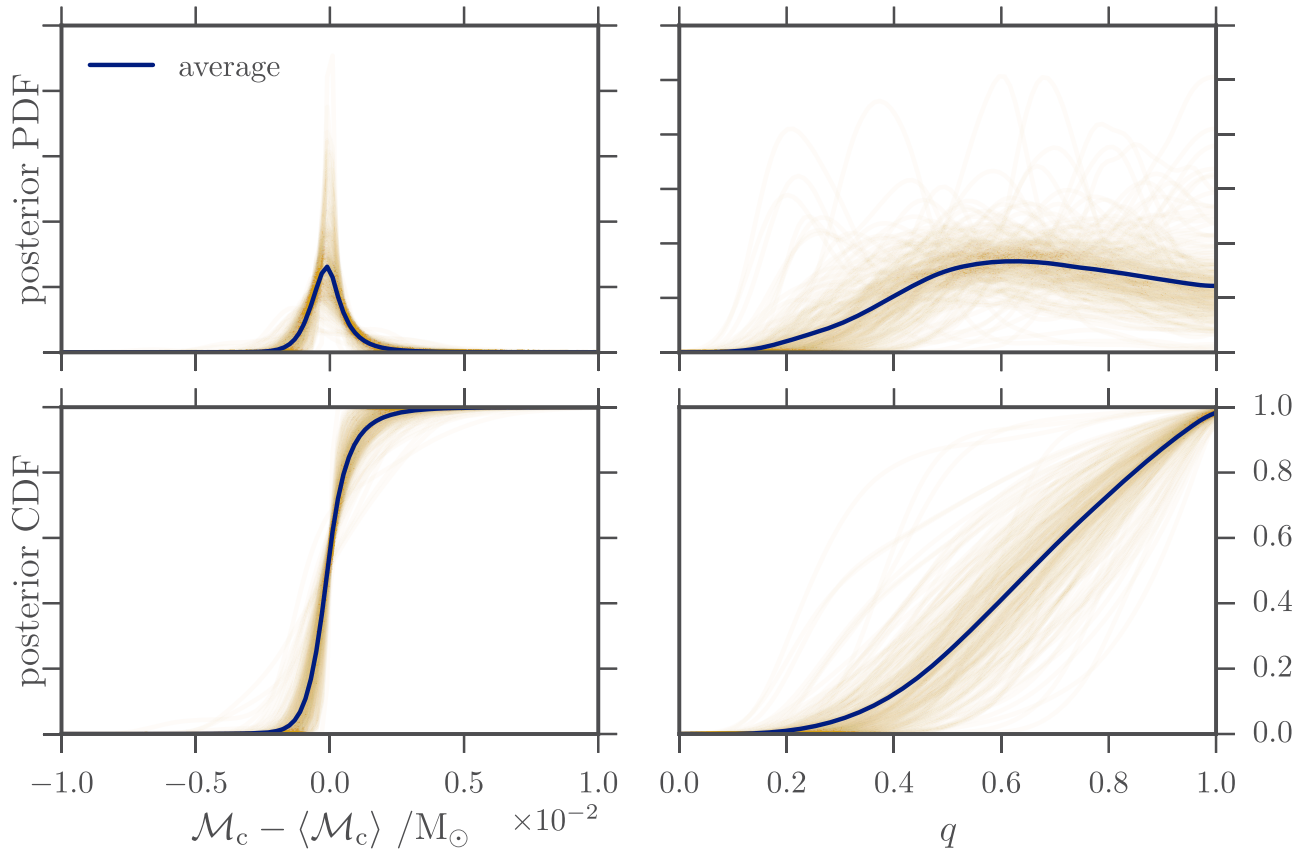

Figure 1. Superimposed posterior probability density (top) and cumulative density (bottom) functions for the chirp mass and mass ratio of all spinning analyses. The solid lines show the average distribution for the simulated population. The chirp-mass distributions have been centered on the distributions' means in order to highlight their consistent morphology.

The fractional uncertainties for the chirp mass, mass ratio, and total mass all decrease as S/N increases, as shown in Figure 2, which also shows results from the non-spinning analysis. As expected from Fisher-matrix studies (e.g., Finn & Chernoff 1993), most appear to be inversely proportional to the S/N: the exception is $\sigma_q/\langle q \rangle$ from the spinning analysis, which is better fit as $\propto \rho_{\rm net}^{-1/2}$. We do not suspect there is anything fundamental about the $\propto \rho_{\rm net}^{-1/2}$, rather, it is a useful rule-of-thumb. The behavior can still be understood from a Fisher-matrix perspective, which predicts a Gaussian probability distribution (with width $\propto \rho_{\rm net}^{-1}$). Since the mass ratio is constrained to be $0 \leqslant q \leqslant 1$, if the width of a Gaussian is large, it is indistinguishable from a uniform distribution, and the standard deviation tends to a constant $1/\sqrt{12} \simeq 0.289$. When the width of the Gaussian is small ($\lesssim$0.1), the truncation of the distribution is negligible, and the standard deviation behaves as expected, as is the case for the non-spinning results. The standard deviations obtained for the spinning runs lie in the intermediate regime, between being independent of S/N and scaling inversely with it (cf. Littenberg et al. 2015); the mean (median) standard deviation σ_q is 0.182 (0.183).[15] The mass–spin degeneracy broadens the posteriors for the chirp mass, the mass ratio, and the total mass; a consequence of the broadening for the mass ratio is that the uncertainty does not decrease as rapidly with S/N (over the range considered here).

Projecting the tightly constrained chirp mass and poorly constrained mass ratio 90% credible region from $\mathcal{M}_c$–q space into component-mass space makes it obvious how important mass-ratio uncertainties are for extracting astrophysical information. The credible regions in component-mass space are narrow bananas that lie along lines of constant chirp mass, bounded by the constraints on mass ratio (see Figure 5 for some examples posteriors).

3.2. Spin Estimates

We now look at the constraints placed on the spin of the slowly spinning simulated BNS sources. Even though the simulations occupy a small fraction of the spin-magnitude prior volume, most posterior distributions span the majority of the prior range. For non-precessing systems, where the orbital plane is stationary with respect to the line of sight, varying the spin of the compact objects has a similar effect on the phase evolution of the GW as varying the mass ratio. This results in a strong degeneracy between the two parameters. Modulation of the GWs from precession of the orbital plane can break this degeneracy (Vecchio 2004; Lang & Hughes 2006; Vitale et al. 2014; Chatziioannou et al. 2015); however, only systems with large spins that are misaligned with the orbital angular momentum significantly precess. Non-precessing systems, with either low or aligned spins, suffer the most from this degeneracy as the only information regarding the mass and spin is encoded in the phase of the GW. The simulated sources

[15] The uncertainty for the symmetric mass ratio $\eta = m_1 m_2/(m_1+m_2)^2$, which is constrained to be $0 \leqslant \eta \leqslant 1/4$, does scale approximately as $\rho_{\rm net}^{-1}$. The mean (median) standard deviation σ_η for the spinning runs is 2.00×10^{-2} (1.95×10^{-2}).

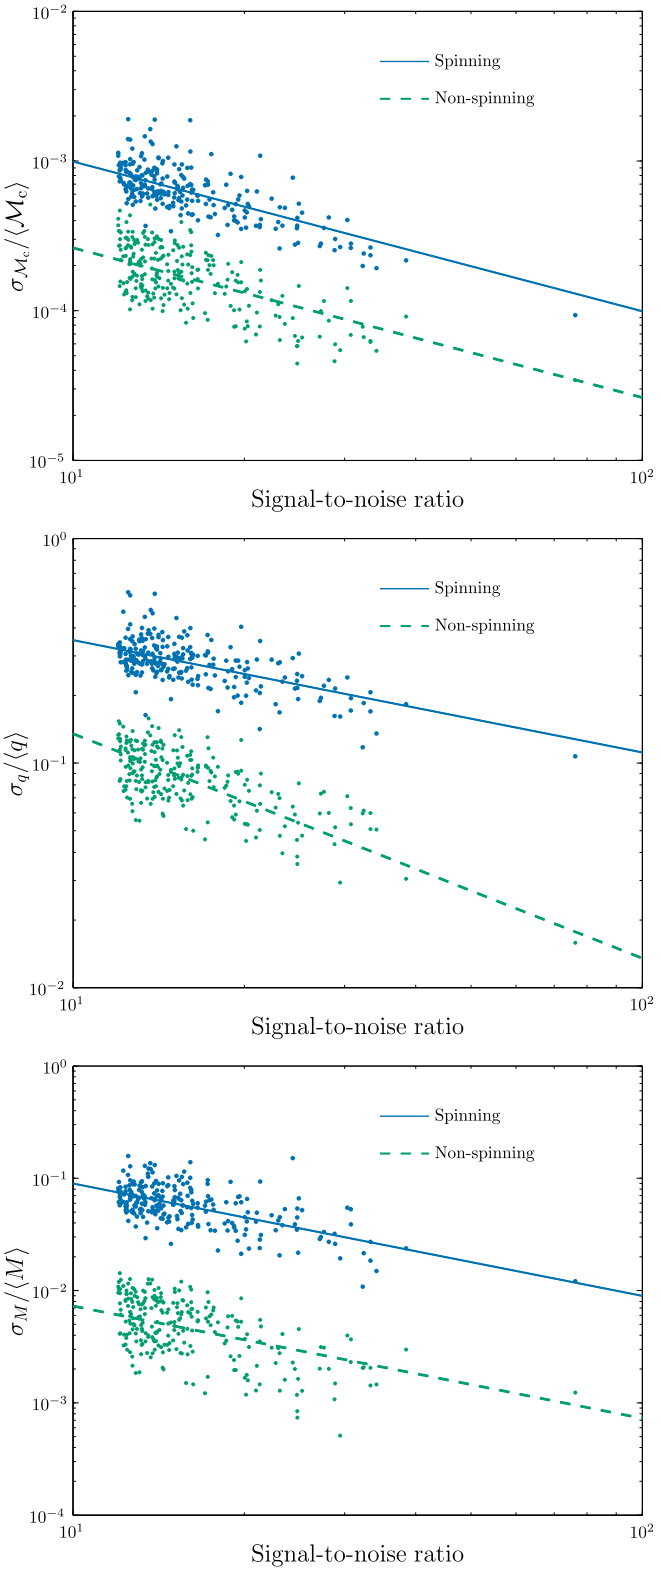

Figure 2. Fractional uncertainties in chirp mass $\mathcal{M}_c$, mass ratio q, and total mass M estimates as a function of network S/N for both the fully spinning SpinTaylorT4 analysis and the medium-latency non-spinning TaylorF2 analysis. We only show statistical uncertainties, not systematic errors (which are present when spin in not included). The lines indicate approximate power-law trends ($\propto\rho_{\rm net}^{-1/2}$ for spinning $\sigma_q/\langle q\rangle$ and $\propto\rho_{\rm net}^{-1}$ for the rest) to guide the eye.

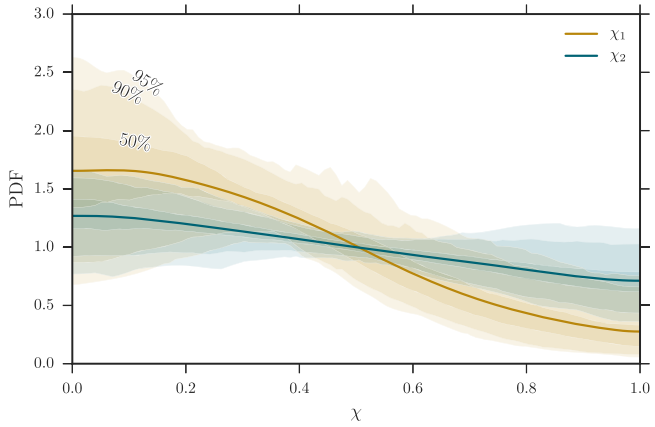

Figure 3. Distribution of one-dimensional marginalized posterior probability density functions (PDFs) of spin magnitudes of the more and less massive components (χ_1 and χ_2, respectively) for all 250 simulated sources. Shaded regions show the 90% credible boundaries for the spin distributions of the population, and the solid lines show the average of each PDF. The posteriors have consistent morphology and span the majority of the prior range. The spin of the most massive component is typically slightly more constrained toward low values, but even a maximal spin of $\chi_1 = 1$ is never ruled out with 100% certainty.

in this study fall in the latter category of low spins. Figure 3 shows the distribution of Gaussian kernel density estimates of the PDFs for the spin of the most and least massive components, χ_1 and χ_2, respectively. The labeled regions of Figure 3 bound the specified percent of PDFs as a function of spin, where the 90% region, for example, is bounded by the 5th and 95th percentiles of the PDFs at each spin value.

The spin of the more massive component has a larger effect on the GW and is, therefore, systematically better constrained, as seen in Figure 3. For both spins, however, the posterior shows slow spins to be only slightly more probable than high spins for most sources. The mean (median) 50% upper limits on χ_1 and χ_2 are 0.319 (0.302) and 0.424 (0.419), respectively; the 90% upper limits are 0.707 (0.699) and 0.855 (0.859).

3.3. Prior Constraints on Spin

Because spin is largely degenerate with mass ratio, and spin is expected to be small for BNS sources, it is interesting to ask how the mass constraints are affected by making stronger prior assumptions about the spin of NSs. First, we make the extreme assumption that NSs have negligible spin, as was done in Singer et al. (2014) and Berry et al. (2015). Figure 4 compares the distribution of (fractional) uncertainties in chirp-mass, mass-ratio, and total-mass estimates for the spinning and non-spinning analyses. The average fractional uncertainties from the non-spinning analysis are a factor of ∼3–4 smaller than the uncertainties from a spinning analysis in the case of $\mathcal{M}_c$ and q and about an order of magnitude smaller for the total mass.

Figure 5 compares cartoon 90% credible regions in component-mass space of five chosen simulated signals (cf. Hannam et al. 2013; Chatziioannou et al. 2015). As a consequence of the difficulty of estimating the narrow and nonlinearly correlated credible regions in m_1–m_2 space, we illustrate the credible regions in m_1–m_2 space as the projection of a rectangular region in $\mathcal{M}_c$–q space. To define the

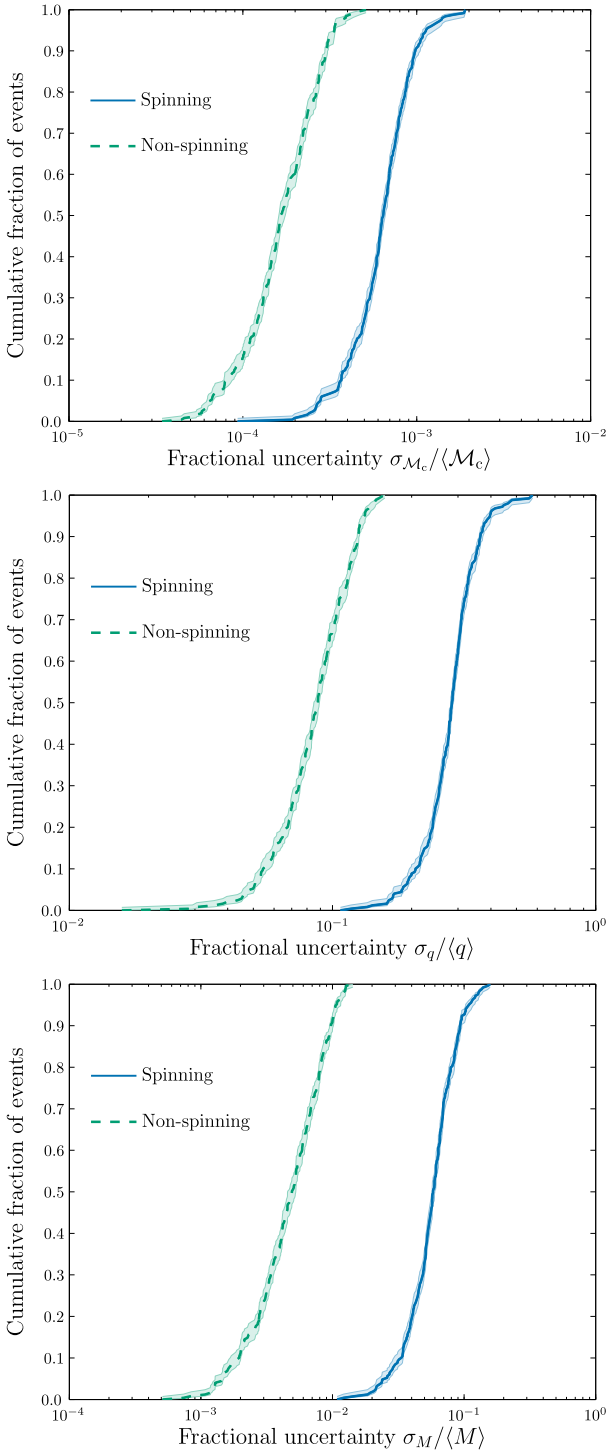

Figure 4. Fractional uncertainty in chirp mass $\mathcal{M}_c$, mass ratio q, and total mass M estimates from the non-spinning and spinning analyses. The mean fractional uncertainties from the non-spinning analysis are 0.0185%, 8.93%, and 0.542% for chirp mass, mass ratio, and total mass, respectively. These are a factor of a few smaller than those found from a spinning analysis (0.0676%, 28.7%, and 6.15% for chirp mass, mass ratio, and total mass, respectively).

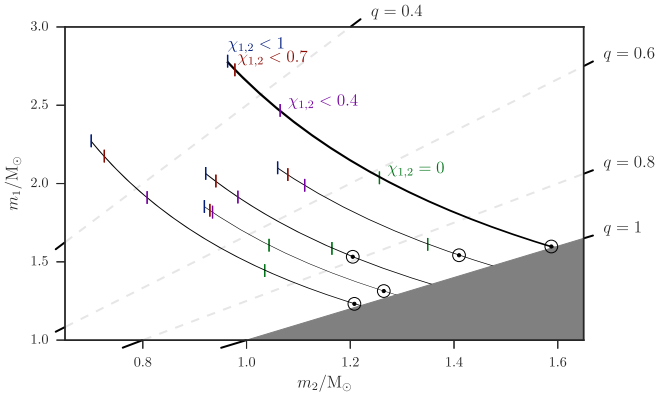

Figure 5. Approximate 90% credible regions for the component-mass estimates of five selected simulations from the spinning analysis; each region is the projection of a rectangular region of chirp-mass–mass-ratio space, bounded by the central 90% credible interval in chirp-mass and upper 90% credible interval in mass-ratio. Circles indicate the true masses of each simulation, and bars indicate the lower bounds of the upper 90% credible intervals (i.e., the 10th percentiles) on mass ratio for increasingly strict prior assumptions on the maximum spin of NSs.

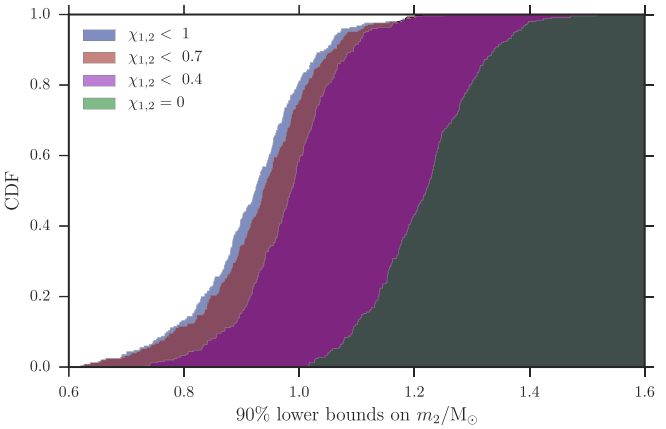

Figure 6. Cumulative distributions of the lower bounds of the upper 90% credible interval (i.e., 10th percentiles) on the estimated mass of the least massive binary components under increasingly strict prior assumptions about maximum NS spin. Restricting spins to be below break-up ($\chi \lesssim 0.7$) for non-exotic equations of state has little effect, as does restricting spin to the maximum observed NS spin ($\chi \lesssim 0.4$). Only strict prior assumptions on NS spin will significantly impact mass constraints.

rectangular region we use 90% credible intervals of the one-dimensional posterior PDFs of $\mathcal{M}_c$ and q; for $\mathcal{M}_c$ we use the central 90% credible interval (5th to 95th percentile), and for q, the upper 90% credible interval (10th to 100th percentile). These differing credible intervals were chosen to better summarize the one-dimensional posterior PDFs, which are typically normal for $\mathcal{M}_c$ and skewed toward high values for q (see Figure 1).

We can investigate the impact of stronger prior assumptions regarding the maximum spin of NSs on mass estimates by discarding posterior samples above a given spin. Figure 6 shows the cumulative distribution of lower 90% bounds on the estimates of m_2 among the 250 simulated sources for spin priors of $\chi_{1,2} \leqslant \{1, 0.7, 0.4, 0\}$. $\chi_{1,2} < 1$ and $\chi_{1,2} = 0$ correspond to the spinning and non-spinning analyses described above. $\chi < 0.7$ is consistent with the NSs remaining intact for most proposed non-exotic EOSs. $\chi < 0.4$ is consistent with the spin of observed isolated NSs to date.

From these PDFs, it is clear that fairly strong prior assumptions on NS spin are required to significantly impact mass constraints. Assuming NSs to be spinning with $\chi_{1,2} \leqslant 0.4$ a priori only constrains masses by an extra few percent compared with allowing them to have $\chi_{1,2} \leqslant 1$.

4. SOURCE LOCATION

Having discussed how GW observations can measure the intrinsic properties of their source systems, we now consider the measurement of extrinsic parameters, specifically the sky position (Section 4.1) and the distance (Section 4.2). These are central to the success of multimessenger astronomy. The sky position is required in order to direct telescopes for electromagnetic (EM) follow-up and to verify that any observed transients do coincide with the source of the gravitational waves. The distance also aids electromagnetic follow-up, as it allows cross-reference with galaxy catalogs to find the most probable source locations (Nissanke et al. 2013; Hanna et al. 2014; Fan et al. 2014; Blackburn et al. 2015). Even without an observed counterpart, the posterior for the (three-dimensional) position allows us to assign a probability that the source resides in given galaxies; combining the redshift of these galaxies (measured electromagnetically) with the GW luminosity distance gives a measure of the Hubble constant free of the usual systematics (Schutz 1986; Del Pozzo 2012). For our population of slowly spinning NSs, we do not expect the measurement of the extrinsic parameters to be affected by the inclusion of spin in the analysis.

4.1. Sky Localization

In order for EM observatories to follow-up a GW detection, they need an accurate sky location. This must be provided promptly, while there is still a visible transient. Parameter estimation while accounting for spin is computationally expensive and slow to complete (see the Appendix). There are alternative methods that can provide sky localization more quickly. The most expedient is BAYESTAR, which uses output from the detection pipeline to rapidly compute sky position (Singer & Price 2016). BAYESTAR can compute sky positions with a latency of a few seconds. Between the low-latency BAYESTAR and the high-latency full parameter estimation, there is the medium-latency option of performing non-spinning parameter estimation with computationally cheap TaylorF2 waveforms. This requires hours of wall time to complete analyses, with the exact time depending upon the degree of parallelization. Despite only using information from the detection triggers, rather than full waveforms, it has been shown that BAYESTAR produces sky areas for BNS signals fully consistent with non-spinning parameter estimation results, provided that there was a trigger from all detectors in the network (Singer et al. 2014; Berry et al. 2015; Singer & Price 2016). Having now performed a full spinning analysis, we can compare the results of high-latency parameter estimation with the more expedient methods of inferring sky position.

In Figure 7 we show the cumulative distributions of recovered 50% credible regions, 90% credible regions, and searched areas. All three quantities show good agreement across all parameter-estimation techniques. For the slowly spinning BNSs considered here, including spin in the analysis does not change the average ability to localize sources on the sky.

We can consider sky localization in greater detail by comparing areas on an event-by-event basis and not just the

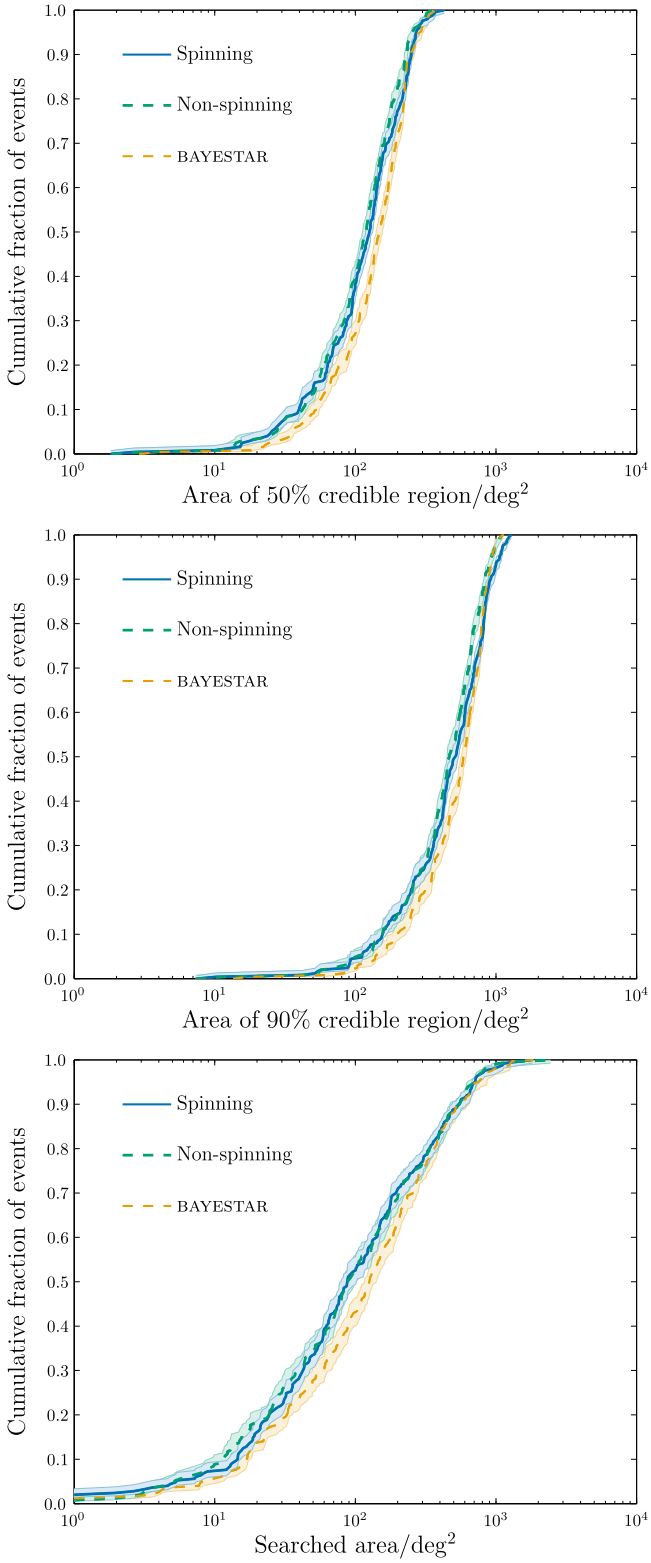

Figure 7. Cumulative fractions of events with sky localization areas smaller than the abscissa value. Top: sky area of 50% credible region $CR_{0.5}$. Middle: sky area of $CR_{0.9}$. Bottom: searched area A_*. The high-latency results including spin are indicated by the solid (blue) line. The lower latency non-spinning and BAYESTAR from Singer et al. (2014) are denoted by thicker (green) and thinner (orange) lines, respectively. The 68% confidence intervals for the cumulative distribution are denoted by the shaded areas.

Table 1
Comparison of Sky Localization Areas Produced by the Low-Latency BAYESTAR Analysis and the High-Latency Fully Spinning SpinTaylorT4 Analysis with those Produced by the Medium-Latency Non-Spinning TaylorF2 Analysis

	BAYESTAR		Spinning	
Logarithmic Ratio	Mean	Standard Deviation	Mean	Standard Deviation
$\mathcal{R}^X_{\mathrm{CR}_{0.5}}$	0.095	0.117	0.022	0.067
$\mathcal{R}^X_{\mathrm{CR}_{0.9}}$	0.075	0.094	0.028	0.063
$\mathcal{R}^X_{A_*}$	0.106	0.447	0.002	0.397

Note. Mean and standard deviation of the log ratio for the 50% credible region $\mathrm{CR}_{0.5}$, the 90% credible region $\mathrm{CR}_{0.9}$, and the searched area A_* are listed for each analysis.

cumulative distribution across the population. Doing this, we confirm that sky localization is consistent between approaches for any given event. We use the medium-latency non-spinning TaylorF2 analysis as a reference point and compare the ratio of sky areas. To summarize the variation in sky areas computed in different analyses, we use the log ratio

$$\mathcal{R}^X_A = \log_{10}\left(\frac{A^X}{A^{\mathrm{NS}}}\right), \tag{2}$$

where A^X is a credible region or the searched area as determined by method X, and A^{NS} is the same quantity from the non-spinning analysis. The log ratio $\mathcal{R}^X_A$ is zero when analysis X agrees with the non-spinning results. Considering all 250 events, the mean and standard deviation of the log ratio is given in Table 1. For the purposes of EM follow-up, there is no significant difference between analyses.[16] The computationally expensive fully spinning analysis does not improve sky localization: there is no disadvantage in using the lower-latency results for EM follow-up of slowly spinning BNSs.

4.2. Luminosity Distance

The distance is degenerate with the inclination (Cutler & Flanagan 1994; Aasi et al. 2013), and the inclination can be better constrained for precessing systems (van der Sluys et al. 2008; Vitale et al. 2014). Because we are considering a population with low spins, precession is minimal, and there should be little effect from including spin in the analysis.

The absolute size of the distance credible interval CI^D_p approximately scales with the distance; hence, we divide the credible interval by the true (injected) distance $D_\star$; this gives an approximate analogue of twice the fractional uncertainty (Berry et al. 2015). The cumulative distribution of the scaled credible intervals is plotted in Figure 8. The mean (median) values of $\mathrm{CI}^D_{0.5}/D_\star$ for the spinning and non-spinning analyses are 0.436 (0.376) and 0.426 (0.363), respectively; the values of $\mathrm{CI}^D_{0.9}/D_\star$ are 0.981 (0.845) and 0.951 (0.819), and the fractional uncertainties $\sigma_D/\langle D\rangle$ are 0.302 (0.262) and 0.245 (0.239). There is negligible difference between the spinning and non-spinning analyses, as expected.

[16] The non-spinning analysis was performed using LALInference_nest, while the spinning analysis was performed using LALInference_MCMC (Veitch et al. 2015); therefore, the consistency between analyses additionally shows the consistency of results from different sampling algorithms.

5. CONCLUSIONS

In this study, we investigated the effects of accounting for spin when estimating the parameters of BNS sources with aLIGO. We expect NSs to be only slowly spinning, and, hence, that their spins only have a small effect on the GW signature of a BNS merger. However, allowing for spins *does* have a significant effect on parameter constraints. Strong degeneracies are present in the model; not only are the spins themselves poorly constrained, but these degeneracies result in weaker constraints on other parameters, particularly masses. Excluding spin from parameter estimation results in artificially narrow and potentially inaccurate posterior distributions.

Weaker constraints are the result of accounting for broad prior assumptions on NS spins. We tested various choices for conservative prior assumptions about NS spins and found them to have little effect on mass estimates. Only strong prior assumptions, such as $\chi_{1,2} \lesssim 0.05$ (consistent with the simulated population, and NSs observed in short-period BNS binaries to date) are likely to significantly affect mass constraints. However, such strict prior assumptions are hard to justify given the small number of observed systems and possible selection effects.

We performed parameter estimation on an astrophysically motivated population of BNS signals, assuming an aLIGO sensitivity comparable to that expected throughout its first observing run. Using a prior on spin magnitudes that is uniform from 0 to 1, spanning the range permitted for BHs and extending beyond the expected (but uncertain) upper limit for NSs, the median 90% upper limit on the spin of the more massive component is 0.70, and the limit for the less massive component is 0.86. The median fractional uncertainty for the mass ratio $\sigma_q/\langle q\rangle$ is $\sim$30%, the median fractional uncertainty for the total mass $\sigma_M/\langle M\rangle$ is $\sim$6%, and the median fractional uncertainty for the chirp mass $\sigma_{\mathcal{M}_c}/\langle\mathcal{M}_c\rangle$ is $\sim$0.06%. Despite the mass–spin degeneracy and only weak constraints on the spin magnitudes, we find that we can place precise constraints on the chirp mass for these BNS signals.

The sky-location accuracy, which is central to performing EM follow-up, is not affected by including spin in the analysis of low-spin systems; this may not be the case when spin is higher, i.e., in binaries containing a BH. For our population of BNSs, sky localization is unchanged by the inclusion (or exclusion) of spin in parameter estimation. The median $\mathrm{CR}_{0.9}$ ($\mathrm{CR}_{0.5}$) is $\sim$500 deg^2 ($\sim$130 deg^2). The luminosity distance is similarly unaffected for this population of slowly spinning NSs; the median fractional uncertainty $\sigma_D/\langle D\rangle$ is $\sim$25%. However, an analysis that includes spins requires the use of more computationally expensive waveforms (that include more physics), increasing latency by an order of magnitude. Therefore, if the population matches our current expectation of being slowly spinning, the low-latency results that could be supplied in time for EM observatories to search for a counterpart are as good as the high-latency results in this respect, and there is no benefit in waiting.

Following the submission of this article, aLIGO made its first detection (Abbott et al. 2016a). This was of a binary BH system (Abbott et al. 2016b) rather than a BNS, but much of our understanding of the abilities of the parameter-estimation analysis, such as the effects of mass–spin degeneracy, translates between sources. The era of GW astronomy has begun, and parameter estimation will play a central role in the science to come.

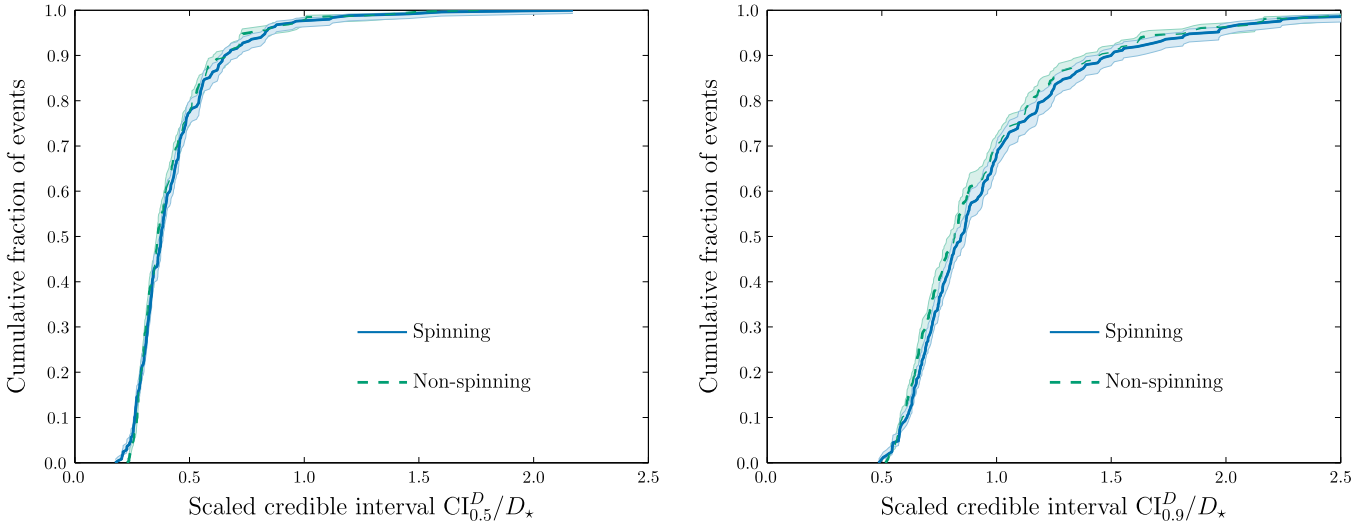

Figure 8. Cumulative fractions of events with luminosity-distance credible intervals (divided by the true distance) smaller than the abscissa value. Left: scaled 50% credible interval $\mathrm{CI}^D_{0.5}/D_\star$. Right: scaled 90% interval $\mathrm{CI}^D_{0.9}/D_\star$. Results using the spinning analysis are indicated by the solid (blue) line, and the results using the non-spinning analysis (Berry et al. 2015) are indicated by the dashed (green) line. The 68% confidence intervals for the cumulative distribution are denoted by the shaded areas.

The authors are grateful for useful suggestions from the CBC group of the LIGO Scientific–Virgo Collaboration. In particular, we are grateful to Walter Del Pozzo for useful advice, Vivien Raymond for helpful discussions, Simon Stevenson for careful reading, and Neil Cornish for beneficial comments.

B.F. was supported by the Enrico Fermi Institute at the University of Chicago as a McCormick Fellow. This work was supported in part by the Science and Technology Facilities Council. P.B.G. acknowledges NASA grant NNX12AN10G. S. V. acknowledges the support of the National Science Foundation and the LIGO Laboratory. J.V. was supported by STFC grant ST/K005014/1. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058.

This work used computing resources at CIERA funded by NSF PHY-1126812, as well as the computing facilities of the LIGO Data Grid including: the Nemo computing cluster at the Center for Gravitation and Cosmology at the University of Wisconsin–Milwaukee under NSF Grants PHY-0923409 and PHY-0600953; the Atlas computing cluster at the Albert Einstein Institute, Hannover; the LIGO computing clusters at Caltech, and the facilities of the Advanced Research Computing @ Cardiff (ARCCA) Cluster at Cardiff University.

Some results were produced using the post-processing tools of the `plotutils` library at http://github.com/farr/plotutils and `skyarea` library at https://github.com/farr/skyarea.

This paper has been assigned LIGO document reference LIGO-P1500117.

APPENDIX
COMPUTATIONAL COST

Performing a fully spinning analysis is computationally expensive. The main computational cost is generating the SpinTaylorT4 waveform, which must be done each time the likelihood is evaluated at a different point in parameter space. Progress is being made in reducing the cost of generating waveforms and evaluating the likelihood (e.g., Canizares et al. 2013; Pürrer 2014). Employing reduced order modeling can speed up the non-spinning TaylorF2 analysis by a factor of ∼30 (Canizares et al. 2015). This is still to be done for a waveform that includes the effects of two unaligned spins; however, progress has also been made in constructing frequency domain approximants using shifted uniform asymptotics, which can speed up generation of a waveform like SpinTaylorT4 by an order of magnitude (Klein et al. 2014).

In Figure 9, we present the approximate wall time taken for analyses comparable to those presented here. The low-latency BAYESTAR and the high-latency fully spinning SpinTaylorT4 results are for the 250 events considered here. The medium-latency non-spinning TaylorF2 results are from Berry et al. (2015); these are not for a different set of signals, but represent a similar population (in more realistic non-Gaussian noise) of what we hope to achieve in reality.[17] The wall times for BAYESTAR are significantly reduced compared to those in Berry et al. (2015) because of recent changes to how BAYESTAR integrates over distance (Singer & Price 2016): the mean (median) time is 4.6 s (4.5 s) and the maximum is 6.6 s. We assume that 2000 (independent) posterior samples are collected for both of the LALInference analyses. The number of samples determines how well we can characterize the posterior: ∼2000 is typically needed to calculate $\mathrm{CR}_{0.9}$ to 10% accuracy (Del Pozzo et al. 2016). In practice, we may want to collect additional samples to ensure our results are accurate, but preliminary results could also be released when the medium-latency analysis has collected 1000 samples, which would after half the time shown here with a maximum wall time of 5.87×10^4 s $\simeq$ 16 hr. We see that the fully spinning analysis is significantly (here a factor of ∼20) more expensive than the non-spinning analysis, taking a mean (median) time of 1.48×10^6 s $\simeq$ 17 days (9.19×10^5 s $\simeq$ 11 days) and a maximum of 1.48×10^7 s $\simeq$ 172 days.

The times shown in Figure 9 illustrate the hierarchy of times associated with different analyses. However, they should not be used as exact benchmarks for times expected during the first observing run of aLIGO because the version of LALInference

[17] We use the more reliably estimated figures for the LALInference runs.

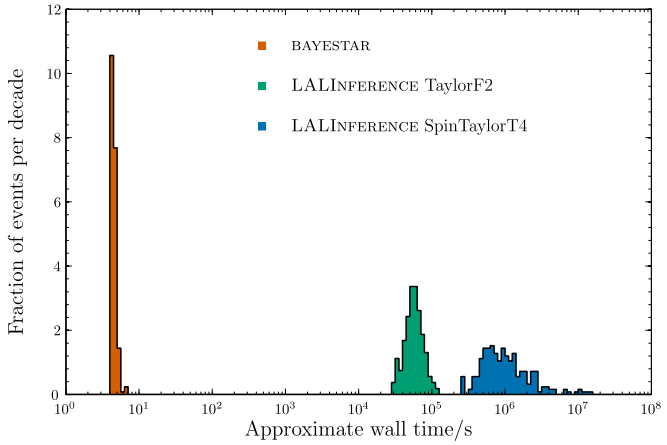

Figure 9. Distribution of wall times (per event). The left (orange–red) distribution is for low-latency sky localization with BAYESTAR; the middle (green) distribution is for medium-latency non-spinning parameter estimation (performed with LALInference_nest), and the right (blue) distribution is for low-latency fully spinning parameter estimation (performed with LALInference_MCMC). The BAYESTAR results assume parallelization across 32 cores (wall timescales inversely with the number of cores), and both LALInference results assume 2000 posterior samples (the wall time is approximately proportional to the number of samples).

used here are not the most up-to-date versions. Following the detection pipeline identifying a candidate BNS signal, we expect BAYESTAR results with latency of a few seconds, non-spinning LALInference results with a latency of a few hours, and fully spinning LALInference results only after weeks of computation.

While work is underway to improve the latency of and to optimize parameter estimation with LALInference, there is also the possibility of developing new algorithms that provide parameter estimates with lower latency (Haster et al. 2015; Pankow et al. 2015). Improving computational efficiency is important for later observing runs with the advanced-detector network: as sensitivities improve and lower frequencies can be measured, we need to calculate longer waveforms (at even greater expense).

REFERENCES

Aasi, J., Abadie, J., Abbott, B. P., et al. 2013, PhRvD, 88, 062001
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, 061102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, arXiv:1602.03840
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, LRR, 19, 1
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001
Barsotti, L., & Fritschel, P. 2012, Early aLIGO Configurations: Example Scenarios Toward Design Sensitivity, Tech. Rep. LIGO-T1200307-v4
Berry, C. P. L., Mandel, I., Middleton, H., et al. 2015, ApJ, 804, 114
Blackburn, L., Briggs, M. S., Camp, J., et al. 2015, ApJS, 217, 8
Brown, D. A., Harry, I., Lundgren, A., & Nitz, A. H. 2012, PhRvD, 86, 084017
Buonanno, A., Chen, Y.-b., & Michele, V. 2003, PhRvD, 67, 104025 (Erratum: 2006, PhRvD, 74, 029904)
Buonanno, A., Iyer, B., Ochsner, E., Pan, Y., & Sathyaprakash, B. S. 2009, PhRvD, 80, 084043
Burgay, M., D'Amico, N., Possenti, A., et al. 2003, Natur, 426, 531
Canizares, P., Field, S. E., Gair, J., et al. 2015, PhRvL, 114, 071104
Canizares, P., Field, S. E., Gair, J. R., & Tiglio, M. 2013, PhRvD, 87, 124005
Cannon, K., Cariou, R., Chapman, A., et al. 2012, ApJ, 748, 136
Chatziioannou, K., Cornish, N., Antoine, K., & Yunes, N. 2015, ApJL, 798, L17
Christensen, N., Libson, A., & Meyer, R. 2004, CQGra, 21, 317
Cutler, C., & Flanagan, E. E. 1994, PhRvD, 49, 2658
Damour, T., Nagar, A., & Villain, L. 2012, PhRvD, 85, 123007
Del Pozzo, W. 2012, PhRvD, 86, 043011
Del Pozzo, W., Berry, C. P. L., Ghosh, A., et al. 2016, Techniques for Reconstructing Sampled Probability Distributions: An Application for Gravitational-Wave Observations of Binary Neutron Stars, Tech. Rep. LIGO-P1500176
Fan, X., Messenger, C., & Heng, I. S. 2014, ApJ, 795, 43
Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, ApJ, 741, 103
Finn, L. S., & Chernoff, D. F. 1993, PhRvD, 47, 2198
Gregory, P. C. 2005, Bayesian Logical Data Analysis for the Physical Sciences (Cambridge: Cambridge Univ. Press)
Hanna, C., Mandel, I., & Vousden, W. 2014, ApJ, 784, 8
Hannam, M., Brown, D. A., Fairhurst, S., Fryer, C. L., & Harry, I. W. 2013, ApJL, 766, L14
Haster, C.-J., Mandel, I., & Farr, W. M. 2015, CQGra, 32, 235017
Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, Sci, 311, 1901
Hinderer, T., Lackey, B. D., Lang, R. N., & Read, J. S. 2010, PhRvD, 81, 123016
Klein, A., Cornish, N., & Yunes, N. 2014, PhRvD, 90, 124029
Kreidberg, L., Bailyn, C. D., Farr, W. M., & Kalogera, V. 2012, ApJ, 757, 36
Lang, R. N., & Hughes, S. A. 2006, PhRvD, 74, 122001 (Erratum: 2008, PhRvD, 77, 109901)
Littenberg, T. B., Farr, B., Coughlin, S., Kalogera, V., & Holz, D. E. 2015, ApJL, 807, L24
Lo, K.-W., & Lin, L.-M. 2011, ApJ, 728, 12
Lorimer, D. R. 2008, LRR, 11, 8
Mandel, I., Haster, C.-J., Dominik, M., & Belczynski, K. 2015, MNRAS, 450, L85
Mandel, I., & O'Shaughnessy, R. 2010, CQGra, 27, 114007
Nissanke, S., Holz, D. E., Hughes, S. A., Dalal, N., & Sievers, J. L. 2010, ApJ, 725, 496
Nissanke, S., Kasliwal, M., & Alexandra, G. 2013, ApJ, 767, 124
Nissanke, S., Sievers, J., Dalal, N., & Holz, D. 2011, ApJ, 739, 99
Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, ApJ, 725, 1918
Özel, F., Psaltis, D., Narayan, R., & Villarreal, A. S. 2012, ApJ, 757, 55
Pankow, C., Brady, P., Ochsner, E., & O'Shaughnessy, R. 2015, PhRvD, 92, 023002
Pürrer, M. 2014, CQGra, 31, 195010
Rodriguez, C. L., Farr, B., Raymond, V., et al. 2014, ApJ, 784, 119
Röver, C., Meyer, R., & Nelson, C. 2006, CQGra, 23, 4895
Schutz, B. F. 1986, Natur, 323, 310
Sidery, T., Aylott, B., Christensen, N., et al. 2014, PhRvD, 89, 084060
Singer, L. P., & Price, L. R. 2016, PhRvD, 93, 024013
Singer, L. P., Price, L. R., Farr, B., et al. 2014, ApJ, 795, 105
Skilling, J. 2006, BayAn, 1, 833
The LSC-Virgo Search Groups, the Data Analysis Software Working Group, the Detector Characterization Working Group and the Computing Committee 2014, The LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics, Tech. Rep. LIGO-T1400054-v7
Vallisneri, M. 2008, PhRvD, 77, 042001
van der Sluys, M. V., Röver, C., Stroeer, A., et al. 2008, ApJL, 688, L61
Vecchio, A. 2004, PhRvD, 70, 042001
Veitch, J., Mandel, I., Aylott, B., et al. 2012, PhRvD, 85, 104045
Veitch, J., Raymond, V., Farr, B., et al. 2015, PhRvD, 91, 042003
Veitch, J., & Vecchio, A. 2010, PhRvD, 81, 062003
Vitale, S., Lynch, R., Veitch, J., Raymond, V., & Sturani, R. 2014, PhRvL, 112, 251101
Wade, L., Creighton, J. D. E., Ochsner, E., et al. 2014, PhRvD, 89, 103012
Yagi, K., & Yunes, N. 2014, PhRvD, 89, 021303