

Impact of precession on aligned-spin searches for neutron-star–black-hole binaries

Tito Dal Canton,^{*} Andrew P. Lundgren,[†] and Alex B. Nielsen[‡]

Max-Planck-Institut für Gravitationsphysik, Callinstrasse 38, D-30167 Hannover, Germany

(Received 15 February 2015; published 31 March 2015)

The inclusion of aligned-spin effects in gravitational-wave search pipelines for neutron-star–black-hole binary coalescence has been shown to increase the astrophysical reach with respect to search methods where spins are neglected completely, under astrophysically reasonable assumptions about black-hole spins. However, theoretical considerations and population synthesis models suggest that many of these binaries may have a significant misalignment between the black-hole spin and the orbital angular momentum, which could lead to precession of the orbital plane during the inspiral and a consequent loss in detection efficiency if precession is ignored. This work explores the effect of spin misalignment on a search pipeline that completely neglects spin effects and on a recently developed pipeline that only includes aligned-spin effects. Using synthetic but realistic data, which could reasonably represent the first scientific runs of advanced-LIGO detectors, the relative sensitivities of both pipelines are shown for different assumptions about black-hole spin magnitude and alignment with the orbital angular momentum. Despite the inclusion of aligned-spin effects, the loss in signal-to-noise ratio due to precession can be as large as 40%, but this has a limited impact on the overall detection rate: even if precession is a predominant feature of neutron-star–black-hole binaries, an aligned-spin search pipeline can still detect at least half of the signals compared to an idealized generic precessing search pipeline.

DOI: [10.1103/PhysRevD.91.062010](https://doi.org/10.1103/PhysRevD.91.062010)

PACS numbers: 04.80.Nn

I. INTRODUCTION

Astrophysical black holes (BHs) can be simply described by their mass and spin angular momentum. A large number of BH properties can be derived from knowing just these two values. Searching for faint gravitational wave (GW) signals from BHs in coalescing binary systems using ground-based interferometers such as advanced LIGO [1], advanced Virgo [2] and KAGRA [3] requires a bank of potential inspiral signals (templates) over a range of possible parameter values. Such template banks have traditionally been built ranging over possible BH masses, but in most cases assuming the absence of spin (see e.g. Refs. [4–8]). This was recently extended to enable searches for gravitational radiation from coalescing neutron-star–black-hole (NSBH) binaries which include aligned-spin effects [9–11]. Under reasonable distributions of binary configurations and the assumption that the BH spin and orbital angular momentum are aligned, including spin effects has been shown to improve the sensitivity of the search method with respect to methods neglecting the effect of spin altogether [11].

However, NSBH binaries are generally expected to have a BH spin misaligned with respect to the orbital angular momentum. The formation of NSBH binaries coalescing in

a time useful for detection seems to require a dissipative common-envelope phase [12,13], which tends to align the angular momenta and spin up the BH, followed by the supernova explosion of the smaller object, which imparts a significant kick to the resulting neutron star (NS). Binaries surviving the second supernova explosion turn out to have a NS kick that effectively tilts the orbital plane with respect to the BH spin [14].

Nevertheless, tilting the orbit by large angles turns out to be very hard; simulations suggest that most NSBH binaries will generally have a misalignment angle smaller than about 60° , with a large fraction having an angle smaller than 45° [15–17]. Unfortunately, x-ray studies of BH spin values are not able to measure directly the alignment of the spin; instead, they rely on the assumption that the accretion disk should align with the spin, and observational constraints on spin magnitudes can thus be influenced by misalignment [18].

A misalignment between the BH spin and the orbital angular momentum breaks a symmetry of the system and leads to a precessing orbital plane. The time-varying orientation of the orbital angular momentum then causes a characteristic phase and amplitude modulation of the chirping GW signal observed on Earth [19,20] which is absent in a nonprecessing waveform.

Including the effect of precession in a GW search pipeline has been attempted before, but it did not yield a significant gain in sensitivity [21,22]. The best strategy for

^{*}tito.dalcanton@aei.mpg.de

[†]andrew.lundgren@ligo.org

[‡]alex.nielsen@ligo.org

including such effects is still unknown, as are the resources required by a fully precessing search pipeline. Thus, even if the spin-orbit misalignment may not be very large, it is important to assess its impact on a search pipeline that completely neglects precession. In light of the results of Ref. [22], it is possible that an aligned-spin search pipeline performs no better than a nonspinning one. It is therefore important to determine how precession affects the sensitivity of including aligned-spin effects over a cheaper method that completely neglects spin. This problem has been recently investigated using bank simulations and assuming the final design sensitivity of advanced LIGO interferometers [10,23]. Here the investigation is extended; in particular it is applied to realistic data rather than idealized noise, and an “early-advanced-LIGO” sensitivity curve, which more realistically represents the first (2015) scientific runs of advanced LIGO, is assumed. The curve is available in Ref. [24] and plotted in Fig. 1 of Ref. [11] (black solid line).

Using bank simulations, we first study the loss in signal-to-noise ratio (SNR) imparted by precession when using template banks which (i) neglect spin altogether and (ii) only include aligned-spin effects. For both cases we explore the dependency of the loss on the parameters which mainly affect precession, i.e. the magnitude and tilt of the BH spin. We also use the bank simulations to estimate the loss of detections produced by considering aligned-spin effects but neglecting precession. Then, using the approach described in Ref. [11], we search for a population of simulated precessing NSBH binaries in synthetic data, and we present the sensitivity of the search pipeline for different possible distributions of BH spins. The synthetic data we analyze consist of 60 days of real data (of which about 25 are useful for the analysis) from the sixth scientific run of the initial LIGO Hanford and Livingston interferometers. However, the power spectral density (PSD) of the data is modified to resemble the aforementioned early advanced-LIGO sensitivity curve, in order to reasonably represent data expected from the first scientific runs of advanced LIGO.

Although the merger and ringdown parts of NSBH signals happen within the bandwidth of the detectors for some regions of the NSBH parameter space, particularly antialigned systems, we neglect them in this work and focus instead on purely precessional effects, which we expect to be somewhat complementary. We reserve a comprehensive study of merger and ringdown effects to a future article.

The paper is organized as follows. In Sec. II we present bank simulations showing the effectualness of a nonspinning template bank for precessing NSBH signals and of a spinning but nonprecessing bank. In Sec. III we present the results of running our search pipeline, using the above banks, on a stretch of realistic data containing simulated precessing NSBH binaries. Section IV summarizes conclusions and future work.

II. EFFECT OF PRECESSION ON SNR LOSS

The first step in studying the effect of precession is calculating the fitting factor between the population of precessing signals we want to observe and the template banks we intend to use, which is equivalent to asking how good, on average, our bank is at recovering the signals we target. This can be done via bank simulations. For a full description of the basic ideas behind template banks and bank simulations, as well as more details of the banks we use in this section, see Ref. [11].

The source population we are interested in contains NSBH binaries with BH mass m_{BH} between $3M_{\odot}$ and $15M_{\odot}$ and NS mass m_{NS} between $1M_{\odot}$ and $3M_{\odot}$. The dimensionless BH spin has magnitude uniformly distributed within the bounds imposed by the Kerr solution, $0 \leq \chi_{\text{BH}} \leq 1$. The tilt angle ϑ of the BH spin with respect to the orbital angular momentum is defined by a uniform distribution of $\kappa \equiv \cos \vartheta$ in the range ± 1 , such that $\hat{\chi}_{\text{BH}}$ is uniformly distributed on the sphere. Although the resulting spin distribution is a reasonable choice if no information is available about BH spins, a number of studies and observations exist suggesting that χ_{BH} may be large [12,25]. At the same time, models based on population synthesis suggest that ϑ is likely peaked at 0, with only a small fraction of NSBH binaries having $\vartheta > 60^{\circ}$ [15,16]. In other words, there is a small region of the $(\chi_{\text{BH}}, \vartheta)$ plane which may be much more common astrophysically than the rest. The NS spin is known to have a small effect on a search for NSBH coalescence [11]. Nevertheless we take a uniform distribution of its magnitude in the range $0 \leq \chi_{\text{NS}} \leq 0.05$, which should include the fastest spinning NSs observed in compact binaries (see e.g. Ref. [26]). The distribution of $\hat{\chi}_{\text{NS}}$ is also uniform on the unit sphere.

The nonspinning bank (NSB) we use here covers BH masses between $3M_{\odot}$ and $15M_{\odot}$ and NS masses between $1M_{\odot}$ and the smallest of either the equal-mass boundary or the line corresponding to a total mass of $18M_{\odot}$. The aligned-spin bank (ASB) covers the same BH masses, NS masses between $1M_{\odot}$ and $3M_{\odot}$ only, dimensionless BH spin projection along the orbital angular momentum $\vec{\chi}_{\text{BH}} \cdot \hat{L} \in [-1, 1]$ and NS spin projection $\vec{\chi}_{\text{NS}} \cdot \hat{L} \in [-0.4, 0.4]$. The NSB and ASB contain $\sim 2.8 \times 10^4$ and $\sim 1.5 \times 10^5$ templates respectively; including aligned spin leads to a $\sim 5\times$ larger computing cost. Template waveforms are computed in the frequency domain using the TaylorF2 approximant. Their inspiral phasing contains orbital terms up to 3.5 post-Newtonian (pN) order and spin-orbit terms up to 2.5 pN. Waveforms start at 30 Hz and terminate at the frequency of the innermost stable circular orbit (ISCO).

To separate the effect of spin from mass-related issues of the banks, we first restrict our attention to binaries with fixed masses $m_{\text{BH}} = 7.8M_{\odot}$ and $m_{\text{NS}} = 1.35M_{\odot}$, which represent typical mass values for BH and NSs in binaries [27,28]. This allows us to study the fitting factor of the NSB

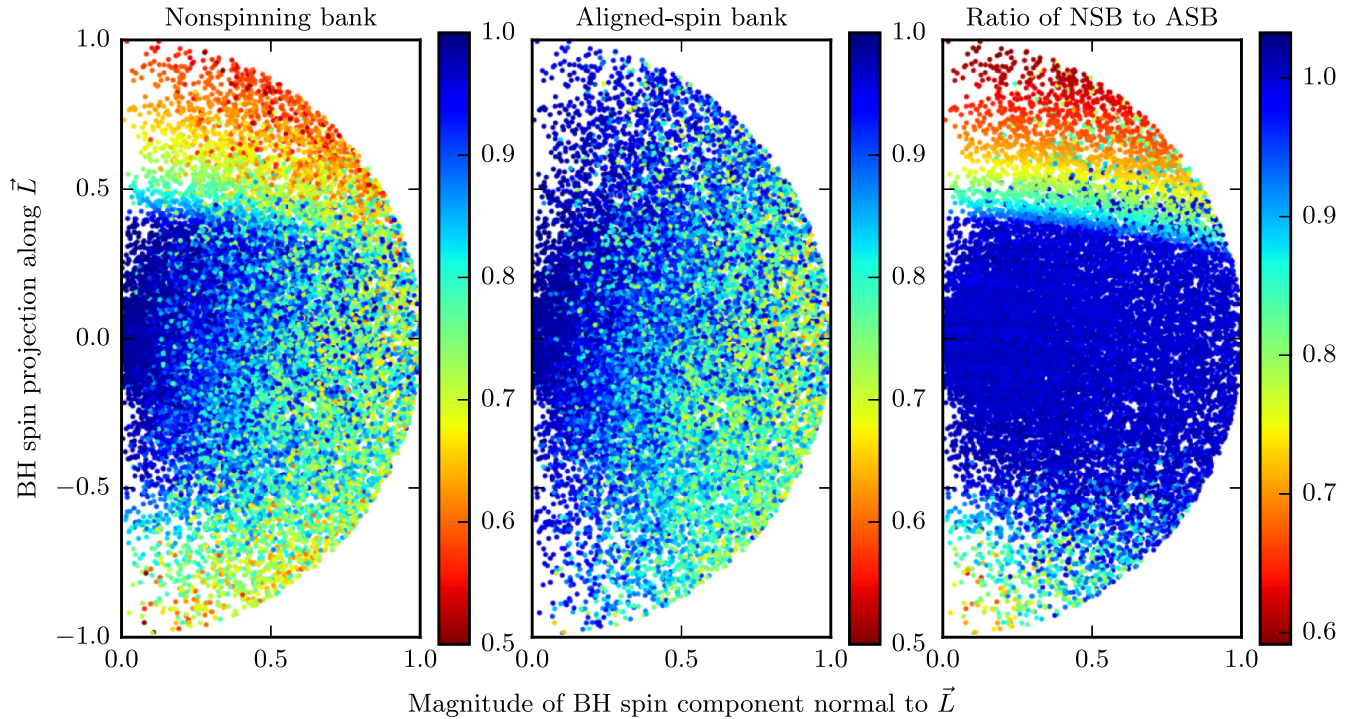


FIG. 1 (color online). Fitting factor of our nonspinning bank (left), aligned-spin bank (middle) and their ratio (right) for a population of 20,000 precessing NSBH binaries with fixed masses $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$. The horizontal axis is the amount of BH spin orthogonal to the orbital angular momentum, while the vertical axis is its projection along the orbital angular momentum; thus, the origin corresponds to nonspinning signals, and the vertical axis corresponds to the aligned-spin case. The spin magnitude and orientation on the sphere are both uniformly distributed.

and ASB as a function of the BH spin parameters only and with high statistics. The result of this first set of simulations is shown in Fig. 1, which displays the fitting factor as a function of the amount of BH spin orthogonal to the orbital angular momentum and parallel to it,

$$\chi_{\text{BH}}^{\perp} \equiv \|\vec{\chi}_{\text{BH}} - (\vec{\chi}_{\text{BH}} \cdot \hat{L})\hat{L}\| \quad (1)$$

$$\chi_{\text{BH}}^{\parallel} \equiv \vec{\chi}_{\text{BH}} \cdot \hat{L}. \quad (2)$$

The behavior of the two banks allows us to roughly divide this space in three approximately defined regions.

The first region is defined by low values of the black hole (BH) spin magnitude, approximately $\chi_{\text{BH}} \lesssim 0.4$. Not surprisingly, both banks perform well here. The spin-orbit terms in the signal waveforms are small enough that the consequent dephasing can be accommodated by a small bias in the symmetric mass ratio η [29] allowing the NSB to recover signals well. Moreover, even if the spin is tilted, its magnitude is too small for precession to induce significant modulation.

The second region is defined by a small perpendicular component of the BH spin but a large parallel component (positive or negative), i.e. $\chi_{\text{BH}}^{\perp} \lesssim 0.5$ and either $\chi_{\text{BH}}^{\parallel} \gtrsim 0.4$ or $\chi_{\text{BH}}^{\parallel} \lesssim -0.5$. Here the NSB has a sudden and severe loss of effectualness, while the ASB performs significantly better.

This loss happens because the spin-orbit terms in the waveform phase acquire their most extreme values, so neglecting them causes the largest possible dephasing. The resulting bias in η when ignoring spin is too large to fit into the NSB: unphysical templates with $\eta > 1/4$ would be required to recover positively aligned signals, and templates with $m_{\text{NS}} < 1$ would be needed to match negatively aligned ones. In other words, signals in this region “fall off” the NSB [11]. However, the signal modulation due to precession is still small in this region, so the ASB is able to recover almost all the signal power.

The third region is the highly precessing case, roughly identified by $\chi_{\text{BH}}^{\perp} \gtrsim 0.5$. The main features here are the very similar performance of both banks and the significant spread of the fitting factor with respect to the other two regions, with some sources being recovered well and others poorly. The former effect happens because spin-orbit terms are small, so the NSB is still able to compensate them by using templates with a biased symmetric mass ratio. At the same time, the modulation induced by precession is strong, and both banks are equally unable to recover the power going into the modulation sidebands. This suggests that the effects of the modulation and of $\chi_{\text{BH}}^{\parallel}$ effectively decouple, as predicted for example in Refs. [30,31].

The second effect visible in this region, i.e. the scatter of the fitting factor, is due to the different possible orientations




FIG. 2 (color online). Fitting factor of the aligned-spin bank for strongly precessing NSBH binaries as a function of the orientation of the total angular momentum with respect to the detector. The south and north poles of the projection correspond to face-on and face-off orientations, while the equatorial line is the edge-on case. Sources have fixed masses $m_{\text{BH}} = 7.8M_{\odot}$ and $m_{\text{NS}} = 1.35M_{\odot}$ and spin parameters $\chi_{\text{BH}} > 0.8$ and $|\kappa| < 0.2$. Each source is shown as a black dot.

of the total angular momentum \vec{J} with respect to the detector, as demonstrated in Fig. 2. In fact, the orientation changes the fraction of signal power in the modulation sidebands, which nonprecessing templates cannot recover. In particular, face-on and face-off binaries almost look like nonprecessing systems when viewed from the detector, so their waveforms have a smaller modulation, and thus larger fitting factor, than edge-on systems. Note that edge-on systems can also have good fitting factors at four particular orientations of \vec{J} ; however, those orientations produce quiet signals at the detector. Because of the almost-linear polarization of their radiation, edge-on systems are in fact generally quieter than face-on or face-off ones at the same distance. Thus, although recovering them is challenging even with the ASB, they are also the least detectable even in the ideal case.

One can therefore ask whether the poor performance of the ASB in the high-precession case is really affecting the overall sensitive volume to a population of binaries with uniform spin and angle distribution as we assume here. The bank simulation enables a rough estimate of the sensitive volume of an aligned-spin search pipeline relative to a hypothetical ideal generic precessing pipeline:

$$\mathcal{V} \equiv \frac{V_{\text{ASB}}}{V_{\text{prec}}} \approx \frac{\sum_i (m_i \rho_i)^3}{m_{\text{prec}}^3 \sum_i \rho_i^3}. \quad (3)$$

Here the sums are over the simulated signals, m_i is the fitting factor of signal i obtained from the bank simulation, ρ_i is the optimal SNR of signal i at a fixed reference distance, and $m_{\text{prec}} = 98.5\%$ is a guess for the average fitting factor of the hypothetical precessing bank. This

estimate does not include the likely increased false-alarm background of the precessing pipeline, which would increase \mathcal{V} . Evaluating this increase requires constructing a generic precessing template bank, which is an open problem and is outside the scope of this paper. The estimation also neglects the effect of signal-based vetoes, which would instead reduce \mathcal{V} . The effect of vetoes, coincidence and realistic data can be included by running a full aligned-spin search pipeline, which will be described in Sec. III. With this caveat in mind, the fixed-mass bank simulation gives, for no restriction on spin,

$$\mathcal{V}_{\text{FM}}^{\text{all}} \approx 81\%. \quad (4)$$

If we restrict our attention to highly precessing binaries, say with $\chi_{\text{BH}} > 0.7$ and $\frac{1}{4}\pi < \vartheta < \frac{3}{4}\pi$, we obtain instead

$$\mathcal{V}_{\text{FM}}^{\text{HP}} \approx 61\%. \quad (5)$$

We now repeat the bank simulation for the full distribution of masses described earlier in this section. The spin parameters are distributed as before. Figure 3 shows the results, and it can be seen that the match variation is qualitatively consistent with the fixed-mass case. A noticeable difference, however, is the much worse mismatch of the NSB for some of the strongly spinning and antialigned systems (lower region of first and last plots). These happen to be heavy systems, with BH mass larger than about $11M_{\odot}$. Their poor match is not due to precession but again to the bias in the symmetric mass ratio produced when attempting to recover spinning signals with zero-spin templates. In this case the spin is antialigned, so the bias is negative; signals are well matched by templates associated with a lighter NS and a heavier BH. For signals with small enough mass, the bias is within the mass space covered by the NSB, and a good fitting factor is obtained. Signals with negatively spinning BHs heavier than $\sim 11M_{\odot}$, however, fall off the high-BH-mass edge of the NSB; templates with BH masses higher than $15M_{\odot}$ would be needed to recover these signals. This effect is less dramatic in the fixed-mass bank simulation because the masses are far from the boundary of the bank.

As done for the fixed-mass bank simulation, we can again estimate the loss of detections of the aligned-spin pipeline relative to a hypothetical generic precessing pipeline. Using Eq. (3) with the result of the varying-mass bank simulation yields

$$\mathcal{V}_{\text{VM}}^{\text{all}} \approx 86\%. \quad (6)$$

High-mass systems have a larger weight in this estimate due to the higher intrinsic loudness of their signals, and they also exhibit the largest precession effects due to the high mass ratio. \mathcal{V}_{VM} may therefore overestimate precession effects, but nevertheless the resulting loss is still quite small. The relative volume for varying-mass, highly precessing systems ($\chi_{\text{BH}} > 0.7$ and $\frac{1}{4}\pi < \vartheta < \frac{3}{4}\pi$) is

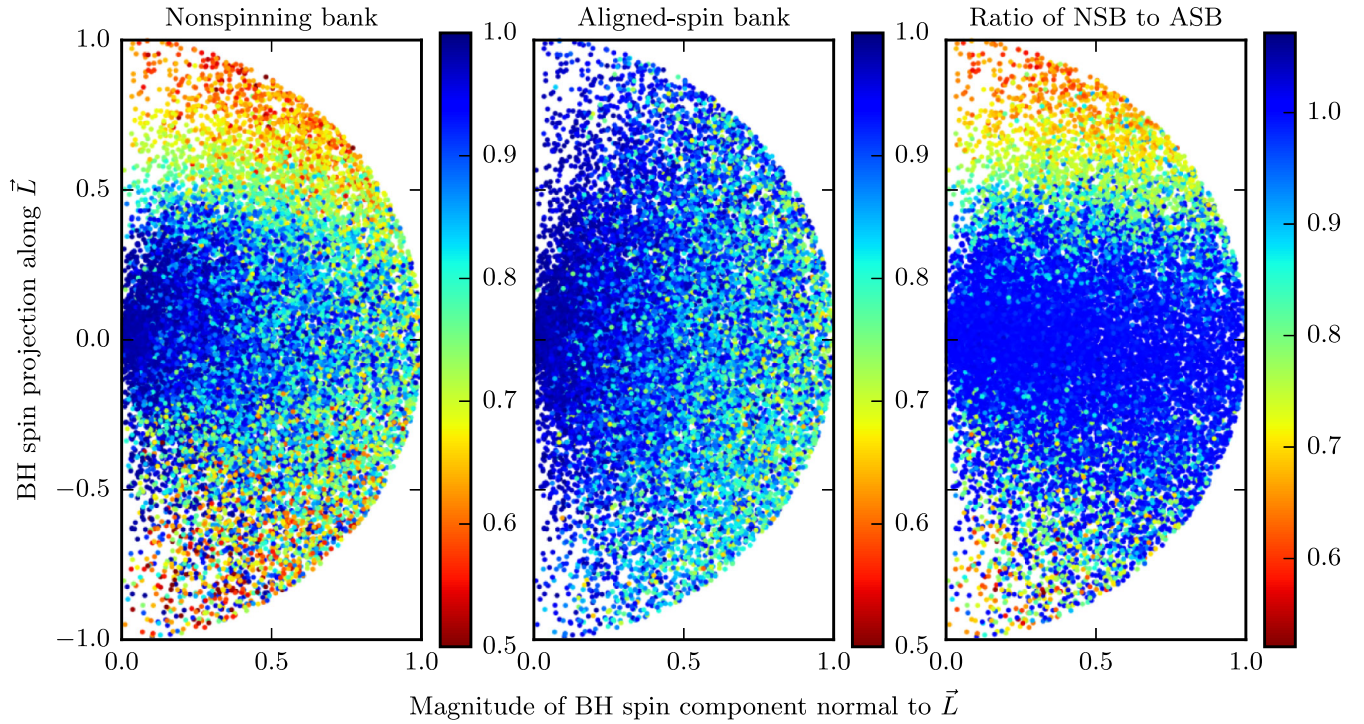


FIG. 3 (color online). Fitting factor of our nonspinning bank (left), aligned-spin bank (middle) and their ratio (right) for a population of 20,000 precessing NSBH binaries with the mass distribution described in the text. Compare with Fig. 1.

$$\mathcal{V}_{\text{VM}}^{\text{HP}} \approx 71\%. \quad (7)$$

These two first exercises show that (i) the ASB is much better exactly where population synthesis models suggest the majority of signals will be ($\chi_{\text{BH}} \rightarrow 1$ and $\vartheta \rightarrow 0$); (ii) antialigned systems ($\chi_{\text{BH}} \rightarrow 1$ and $\vartheta \rightarrow \pi$) give rise to the worst matches in the NSB, but this problem could be alleviated by adding heavier-BH templates to the NSB; (iii) both banks show a comparable inefficiency with strong precession, i.e. near maximal spin misalignments ($\vartheta \rightarrow \pi/2$), but this effect is smaller than the loss imparted by neglecting spin-orbit terms and only happens for some orientations of the total angular momentum; and (iv) even if a generic precessing template bank is not yet available, when taking into account the orientation-dependent intrinsic loudness of signals, the poor performance of the ASB with strong precession seems to reduce the detection rate by a few tens of percent at most. This result is compatible with similar existing studies [10,23]. However, the estimate only considers the loss of SNR and obviously needs to be evaluated more precisely by running a full, realistic search pipeline.

III. EFFECT OF PRECESSION ON A REALISTIC SEARCH PIPELINE

The overall efficiency of a search is determined primarily by two features. The first is the background of false alarms generated by the template bank due to the detector noise. It

has been shown in our previous study that the increase in background due to using the larger ASB is not considerable relative to the NSB when the reweighted SNR is used as a ranking statistic [11]. The second feature is the ability of the waveforms in the chosen template bank to match the targeted signals; we studied that in the previous section. Based on these results, we expect a search using the ASB to perform at least as well as the NSB and to significantly outperform it for a population of almost-aligned systems. However, it remains to be checked whether coincidence between the detectors and the inclusion of the χ^2 veto alter the result significantly in the presence of precession. As a final result, we also want to compare the search sensitivity at a fixed false-alarm rate.

In this section, we thus apply the search method described in Ref. [11], with identical parameters and data, to the fixed-mass and variable-mass precessing signal populations described in the previous section. We recall that this is a matched-filter, exact-match-coincidence pipeline based on the PyCBC toolkit, with settings similar to what was used in the last initial-LIGO searches. The data we analyze are constructed using two months of data from the sixth science run of the Hanford and Livingston LIGO detectors and modifying their PSD to resemble the early advanced-LIGO sensitivity curve we assume in this work.

Similarly to what was done in Ref. [11], we first compare, for each simulated system, the combined SNR and combined reweighted SNR observed by the pipeline




FIG. 4 (color online). Combined SNR and reweighted SNR observed by our search pipeline vs the optimal SNR one would achieve by using templates perfectly matched to each simulated system. The source population has fixed masses $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$. There is no clear difference between the left and right plots because relatively few points have large and almost aligned spins.

to the optimal combined SNR. The latter is calculated by simulating the system’s waveform and using it as its own optimal template.¹ The comparison for the fixed-mass population is shown in Figs. 4 and 5; results for the varying-mass systems are consistent with the fixed-mass case and are not shown. In Fig. 4 it can be seen that many systems have a large loss of both ranking statistics. The reweighted SNR is particularly affected, which can be explained by the χ^2 veto penalizing strongly precessing signals in both banks. As the signals shown in Fig. 4 have no constraints on spin parameters, and relatively few points have high and quasisaligned systems, the advantage of the ASB is not clearly visible in these plots. Figure 5 shows instead the observed fraction of optimal SNR across the spin parameter space and can be compared with Fig. 1. The optimal SNR is not a good predictor of the reweighted SNR when signals become too loud, because in that regime even a small residual mismatch causes the χ^2 statistic to become very large. In order to make Fig. 5 more clear, thus, the reweighted SNR plots exclude signals with observed

¹Note that, with precessing signals, waveforms observed by different detectors are not simply related by phase rotations and amplitude scalings but can be qualitatively different due to the variable relative orientations between the system and the detectors. Therefore, the optimal template in one detector is, in general, not optimal for another detector.

combined SNR larger than 50. For both statistics the results are consistent with Sec. II: the ASB performance is superior to the NSB for large spin and tilt angle close to 0° or 180° , while both banks perform similarly for small spin or tilt angle around 90° . In particular, looking at the top of the plots, it can be seen that many high-spin and small-tilt signals are not detected at all by the NSB. Thus, the χ^2 veto and exact-match coincidence preserve the features of the banks described in Sec. II.

As a final step, we turn our attention to receiver operating characteristic (ROC) curves, showing the sensitive volume of the pipeline (proportional to the detection rate) as a function of the false-alarm rate.² ROCs for fixed masses are shown in Fig. 6, where different plots compare the pipelines using the NSB (dashed lines) and ASB (solid lines) for different cuts on the BH spin magnitude χ_{BH} and tilt angle ϑ . Error intervals represent the standard deviation of 100 realizations of each curve, each constructed from a random selection of half of the background and half of the simulated signals. The first plot contains systems with no restrictions on spin parameters, and in this case the ASB gives a slightly larger sensitivity, although the difference is within the error intervals. The second plot corresponds to weakly spinning BHs with no restriction on tilt angles. Such a population would be almost “tuned” to the NSB, and in fact the ASB has a slightly lower sensitivity due to the larger background. The difference is nevertheless still comparable to the error intervals. The next three plots assume respectively large spins, large spins and small tilt angles, and small tilt angles only. In this case the ASB has a clear advantage, with a sensitivity between 50% and one order of magnitude larger. If the existing population synthesis models and spin measurements are assumed, the most realistic scenario would lie somewhere between plots 4 and 5. We can also see from Fig. 6 that the sensitivity of the ASB remains good for most choices of cuts on spin parameters, while the sensitivity of the non-spinning pipeline drops significantly when $\chi_{\text{BH}} > 0.7$. Nevertheless, if a large fraction of sources has strong and misaligned spins (plots 3 and 6) the detection rate drops noticeably also with the aligned-spin pipeline.

ROCs for the varying-mass population are shown in Fig. 7. The curves are qualitatively similar to the fixed-mass case, and we can thus extend our previous conclusions to a realistic distribution of masses. The curves should be compared to the final plot of Ref. [11].

Similarly to what was done in Sec. II, ROCs allow us to estimate the detection rate of an aligned-spin search pipeline relative to an ideal generic precessing pipeline, once we make the two following approximations. The first is that the precessing pipeline produces the same false-alarm background as the aligned-spin one. This is likely not the case,

²The exact definition of ROCs used here is Eq. (15) of Ref. [11].

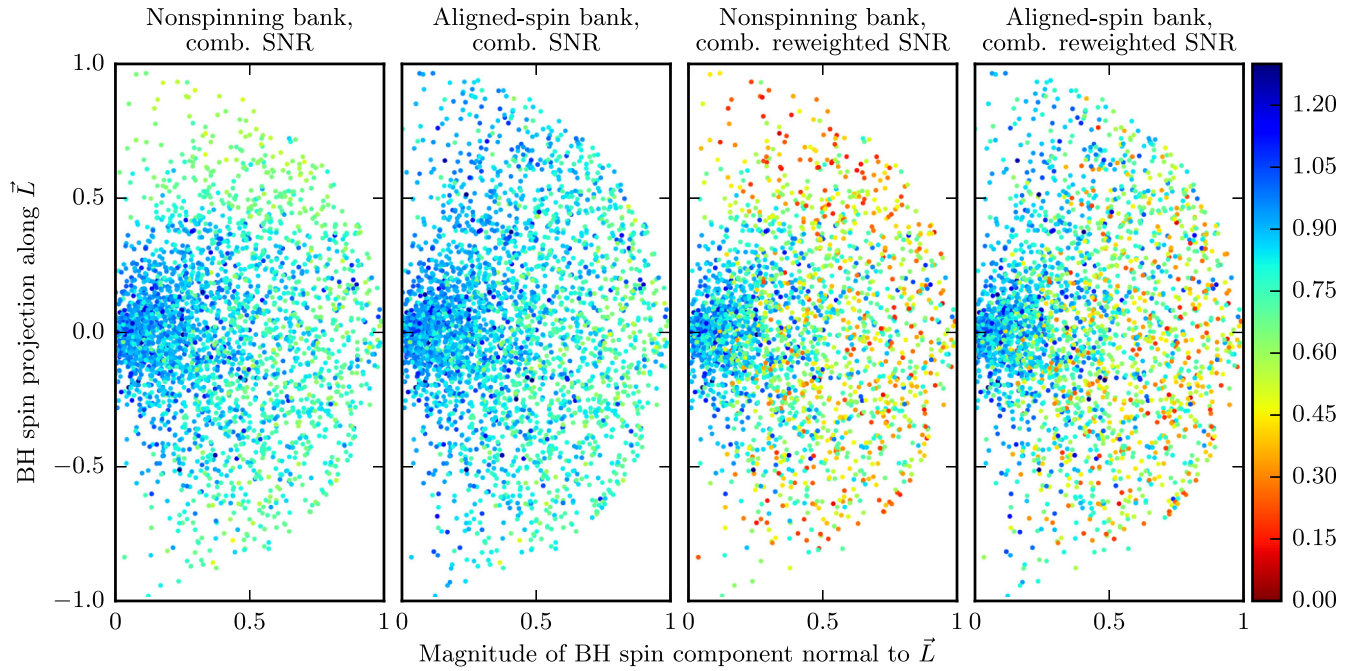


FIG. 5 (color online). Ratio between combined SNR or combined reweighted SNR observed by our search pipeline and the optimal SNR one would achieve by using templates perfectly matched to each simulated system. The source population has fixed masses $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$. In the reweighted SNR plots, only systems with an observed combined SNR smaller than 50 are shown, because for louder systems the observed reweighted SNR is much smaller than the optimal SNR.

so our result will overestimate the loss of detections due to neglecting precession; nevertheless, based on the comparison between the false alarms of the NSB and ASB in Ref. [11], we expect the error to be small. The second assumption is that the number of detections of the precessing pipeline, averaged over the orientation of \vec{J} , is independent from χ_{BH}^{\perp} ; in other words, we assume that precessing sources are not intrinsically more visible than nonprecessing ones, when all orientations of \vec{J} are considered. This assumption is reasonable because precession distributes the radiated power more evenly across the sky, but it does not make the signal significantly longer or louder [20]. If these assumptions are met, the fraction of generic sources detected by the precessing pipeline should be equal to the fraction of weakly precessing sources detected by the ASB. Since the fraction of detected sources is our definition of the ROC, the ROC of the ASB for weakly precessing sources (say, $\chi_{\text{BH}}^{\perp} < 0.4$) can then be used as a proxy for the precessing pipeline over the full spin parameter space. We can thus estimate the relative sensitivity of the aligned-spin pipeline with respect to a precessing one as

$$\mathcal{W}_{\{\text{FM,VM}\}}^{\{\text{all,HP}\}}(f) \equiv \frac{W_{\{\text{FM,VM}\}}^{\{\text{all,HP}\}}(f)}{W_{\{\text{FM,VM}\}}^{\text{LP}}(f)}, \quad (8)$$

where $W_M^S(f)$ is the ROC associated with the aligned-spin pipeline and the particular cut S of the spin parameter space

(“all”: no restriction on spin; “LP”: low precession, $\chi_{\text{BH}}^{\perp} < 0.4$; “HP”: high precession, $\chi_{\text{BH}}^{\perp} > 0.7$ and $\frac{1}{4}\pi < \vartheta < \frac{3}{4}\pi$) and mass distribution M (“FM”: fixed mass; “VM”: varying mass) and f is the false-alarm rate at which the ROC is evaluated. We obtain

$$78\% < \mathcal{W}_{\text{FM}}^{\text{all}} < 83\% \quad (9)$$

$$41\% < \mathcal{W}_{\text{FM}}^{\text{HP}} < 50\% \quad (10)$$

$$80\% < \mathcal{W}_{\text{VM}}^{\text{all}} < 87\% \quad (11)$$

$$50\% < \mathcal{W}_{\text{VM}}^{\text{HP}} < 68\%, \quad (12)$$

where the ranges include the different possible false-alarm rates. As opposed to the estimates obtained at the end of Sec. II, these estimates do include the effect of the χ^2 veto. This seems to have a small effect in the unrestricted-spin case, as $\mathcal{W}_{\text{FM}}^{\text{all}}$ and $\mathcal{W}_{\text{VM}}^{\text{all}}$ are consistent with $\mathcal{V}_{\text{FM}}^{\text{all}}$ and $\mathcal{V}_{\text{VM}}^{\text{all}}$ respectively. However, the high-precession estimates are systematically smaller than Eqs. (5) and (7), suggesting that the χ^2 veto is indeed penalizing precessing signals.

Relaxing the assumption of equal background of the precessing and aligned-spin pipelines requires at least an estimate of the number of precessing templates, but it can only increase all \mathcal{W} estimates. We therefore expect that the loss of detections due to using an aligned-spin pipeline vs a precessing one is, with no restrictions on BH spin

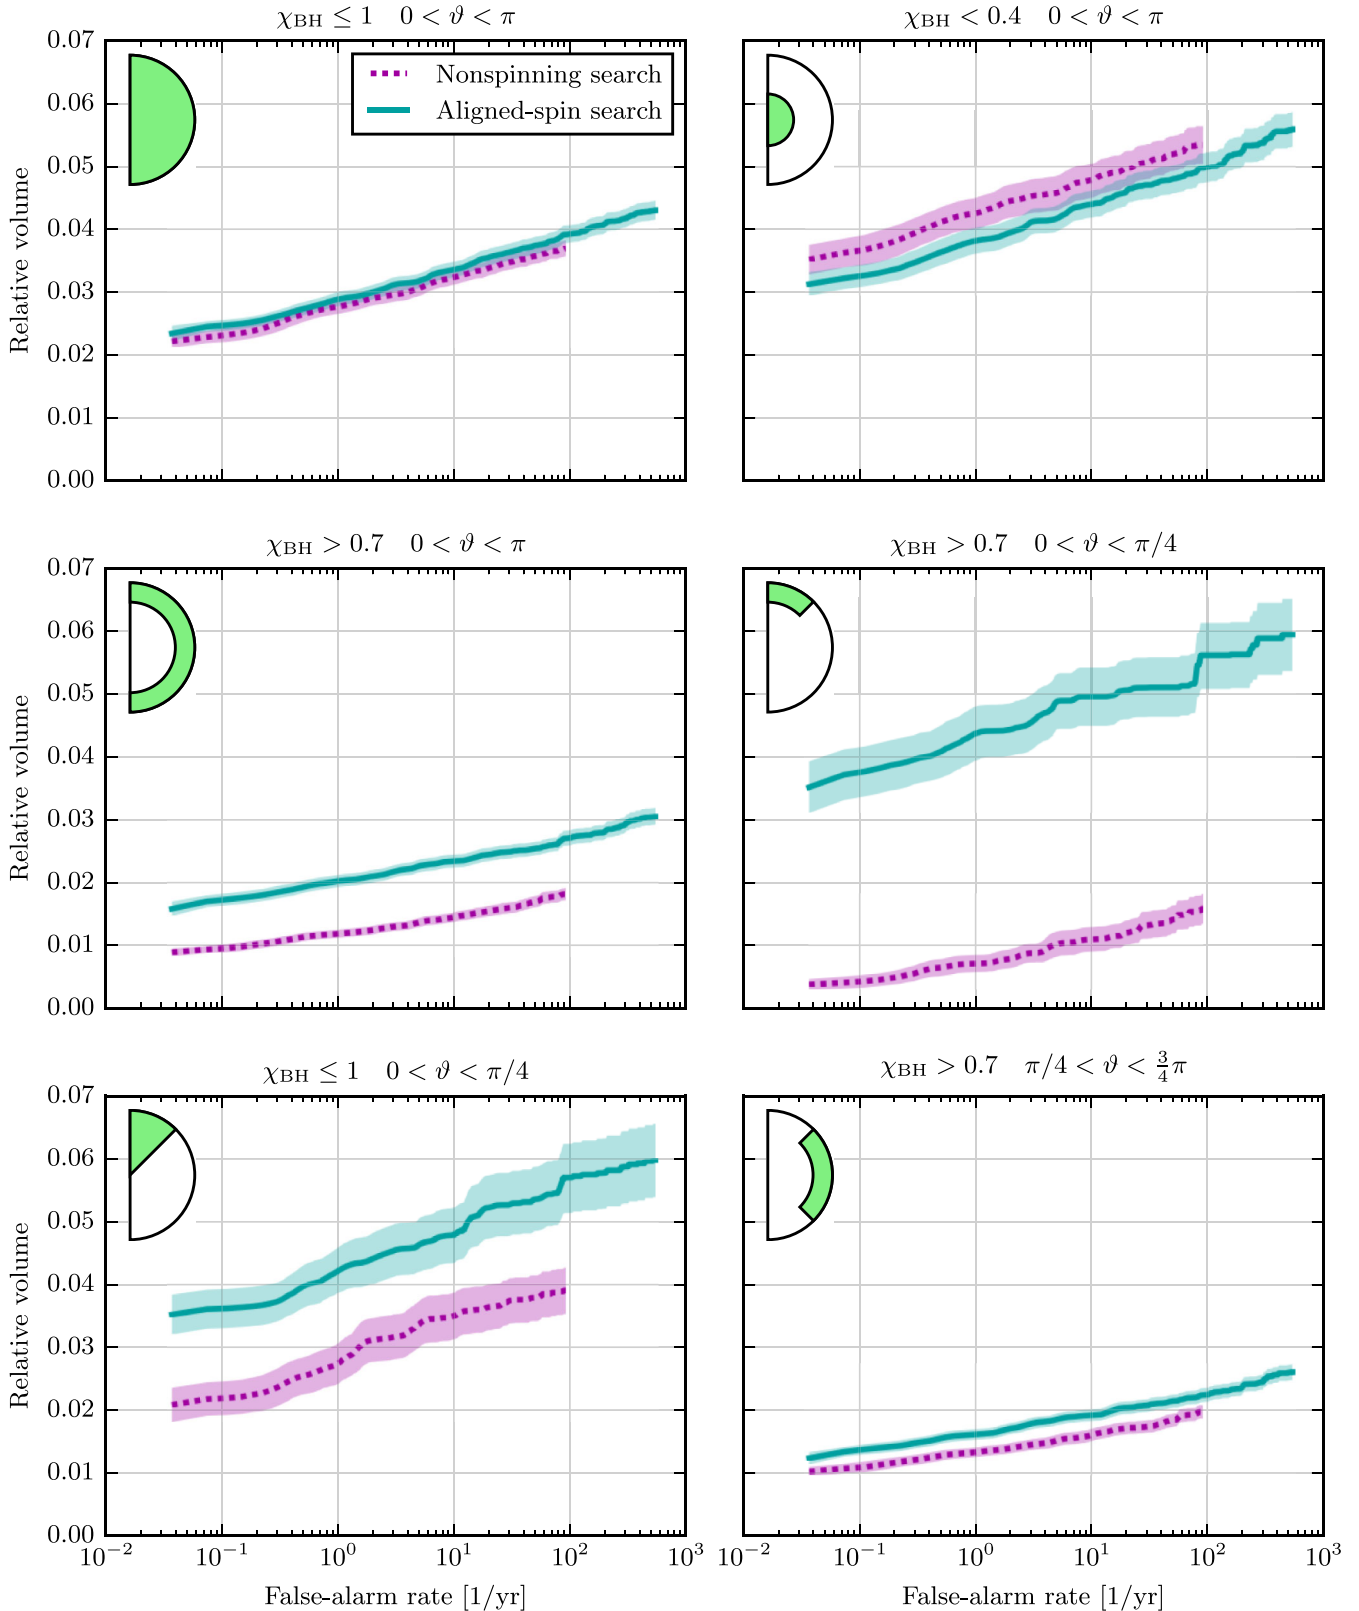


FIG. 6 (color online). ROCs associated with nonspinning and spinning search pipelines (dashed and solid lines respectively) observing populations of NSBH binaries with fixed masses $m_{\text{BH}} = 7.8M_{\odot}$, $m_{\text{NS}} = 1.35M_{\odot}$ and different constraints on the BH spin parameters ($\chi_{\text{BH}}^{\perp}, \chi_{\text{BH}}^{\parallel}$) (visualized in the insets). The lighter bands show the 68% error intervals estimated by constructing each curve 100 times from different combinations of the available data.

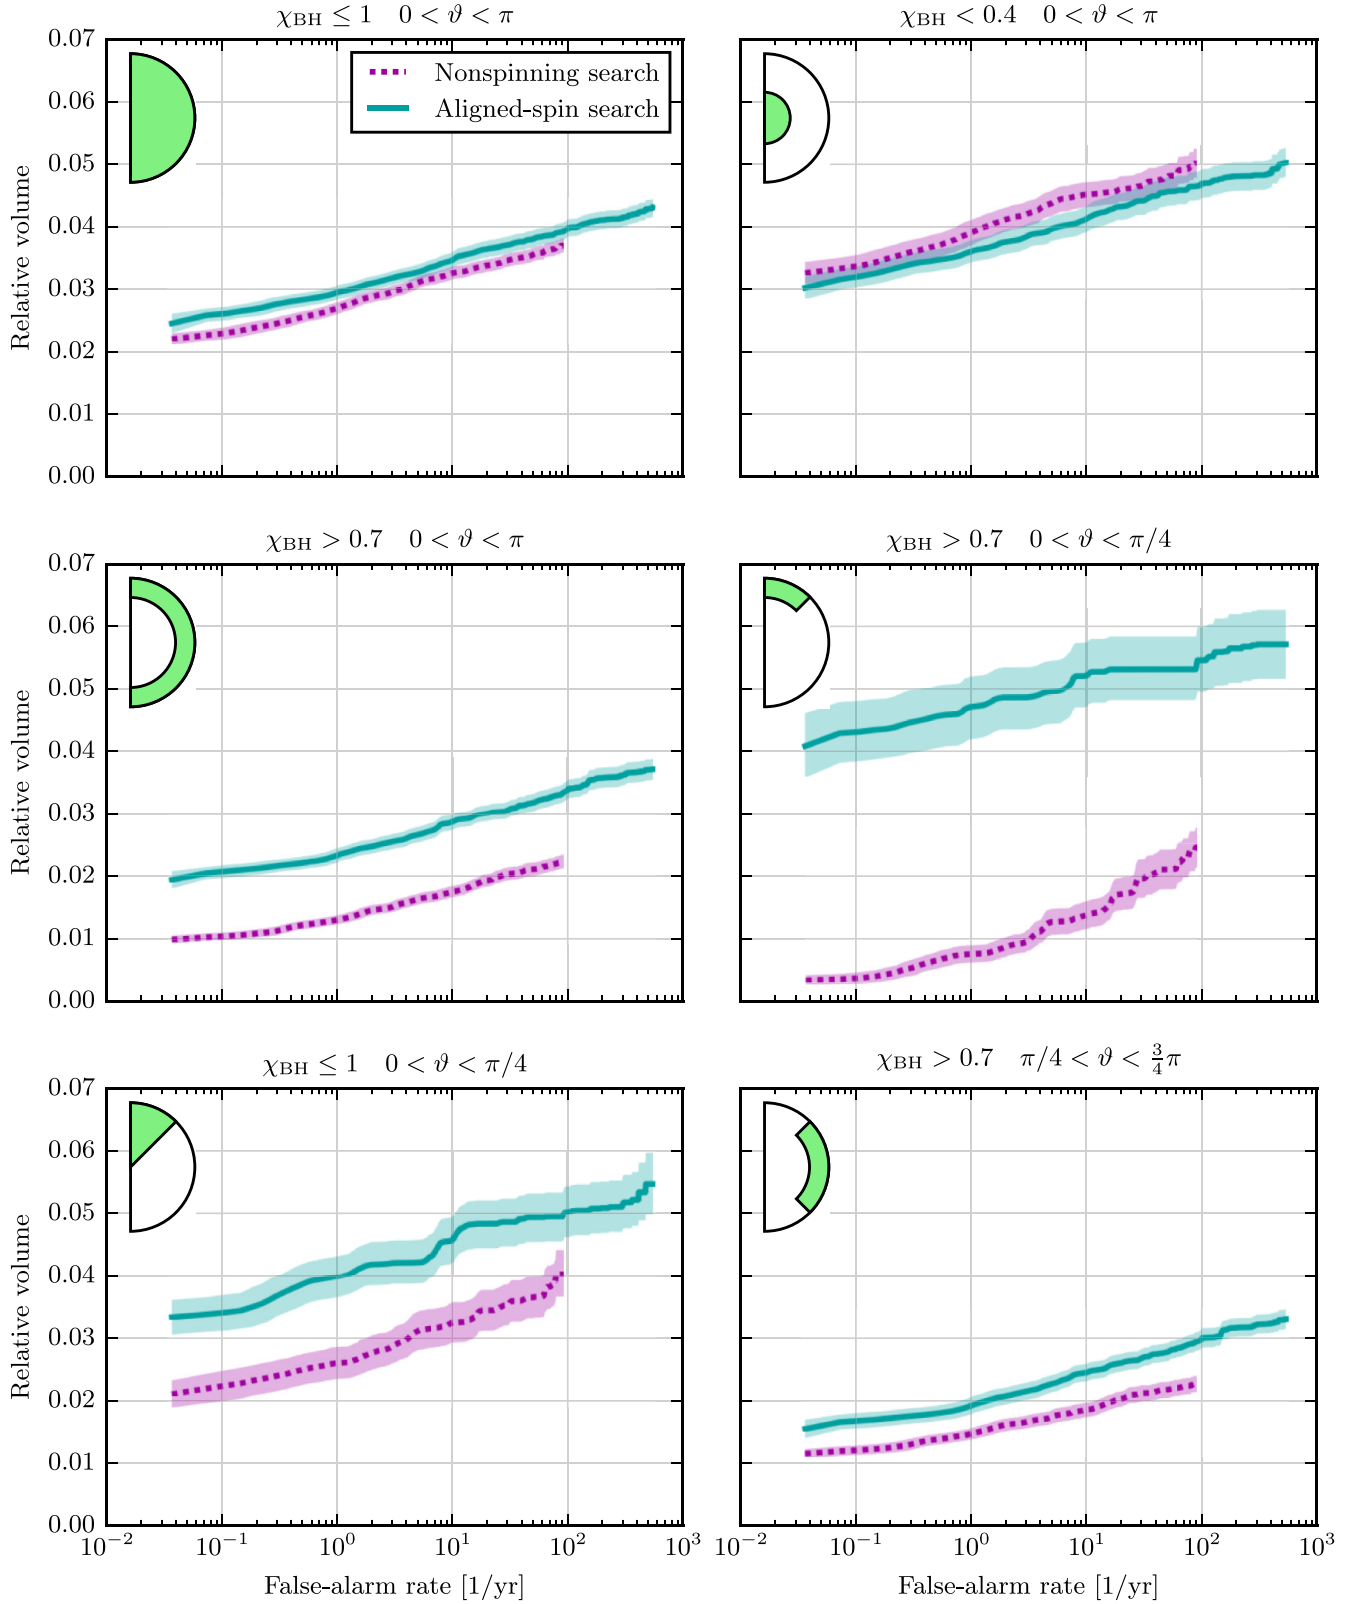


FIG. 7 (color online). ROCs associated with nonspinning and spinning search pipelines (dashed and solid lines respectively) observing populations of NSBH binaries with mass distribution as in Sec. II but different constraints on the BH spin parameters ($\chi_{\text{BH}}^{\perp}, \chi_{\text{BH}}^{\parallel}$) (visualized in the insets). The lighter bands show the 68% error intervals estimated by constructing each curve 100 times from different combinations of the available data. Compare with the final plot of Ref. [11].

parameters, no larger than $\sim 20\%$; if we make a less plausible assumption of highly precessing NSBH binaries, we expect a loss within $\sim 60\%$.

IV. CONCLUSION

In this paper we study the effect of precession on NSBH binary inspiral search pipelines using nonspinning and aligned-spin template banks.

By means of bank simulations, we first show that the two banks perform similarly over the parameter space of the dimensionless BH spin, except for a large drop in effectiveness of the nonspinning bank when the projection of the spin on the orbital angular momentum is larger than ~ 0.4 or smaller than ~ -0.5 , which corresponds to strongly spinning but weakly precessing systems. When precession is strong, both banks can lose up to $\sim 40\%$ of the SNR for particular orientations of the total angular momentum with respect to the detector. Nevertheless, the high-precession systems which are best recovered correspond to face-on and face-off orientations, which are the most likely to be observed based on their intrinsic SNR.

Using the same template banks, we then employ a realistic search pipeline to recover the same signal population in simulated noise, constructed by recoloring real initial-LIGO data to a sensitivity indicative of early advanced-LIGO detectors. The search pipeline includes the χ^2 signal-based veto and exact-match coincidence between the Hanford and Livingston LIGO detectors. We compare the resulting nonspinning and aligned-spin ROC curves for different choices of BH spin parameters.

We conclude that using an aligned-spin bank will increase the detection rate of NSBH binaries by a fraction which strongly depends on the distribution of BH spin magnitudes and tilt angles in nature. At the very minimum, assuming an extreme case of weakly spinning BHs ($0 \leq \chi_{\text{BH}} < 0.4$) and unrestricted, uniformly distributed tilt angle, a search pipeline based on the aligned-spin bank would have a larger computational cost, but it would reduce the sensitivity relative to using a nonspinning bank by $\sim 10\%$ at most. With unrestricted, uniformly distributed spin magnitude and tilt, both methods have very similar sensitivity. On the other hand, for large spin magnitude and small tilt—a distribution supported by population-synthesis models and existing BH spin measurements—the improvement in sensitivity can be as large as 1 order of magnitude. We also find a noticeable improvement assuming either small tilt and unrestricted magnitude or unrestricted tilt and large magnitude. A factor of ~ 2 in sensitivity is lost for the unlikely case of strongly precessing systems, even with the aligned-spin bank. Using these results we estimate that using a bank of generic precessing templates could increase the sensitivity by tens of percent or, under less realistic

assumptions, possibly by a factor of 2, assuming the background does not increase significantly. Quantifying this improvement more precisely requires constructing a bank of precessing waveforms (see e.g. Refs. [30,32–36]) and applying it to a full search pipeline with realistic advanced-detector data, which will be subjects of future papers.

Although we assumed a particular sensitivity curve, our conclusion is likely to remain true for the wider-band sensitivity expected for the final design advanced LIGO. Bank simulations in Ref. [11] suggest that including spin is even more important for the zero-detuned high-power design curve of advanced LIGO; similar calculations with precessing NSBH signals also produced comparable results [10,23]. The simulated precessing signals we have used are based on post-Newtonian expansions and terminate abruptly, ignoring merger and ringdown effects. For the mass range considered here, except for highly antialigned systems, these are likely to be good approximations to the true waveforms in the sensitive band of advanced LIGO detectors; thus, we do not expect the inclusion of merger and ringdown to change our results significantly. Should the search be extended to much higher masses—such that merger and ringdown provide the majority of the SNR—the inclusion of spin could be less important. Tidal interaction between the NS and BH may be important in some regions of our parameter space, for instance by altering the post-Newtonian phasing or disrupting the NS and shutting down the GW signal before merger [37]. Based on Ref. [38], we do not expect tidal deformation to significantly affect the search. Tidal disruption should also not affect our conclusions if it happens beyond the ISCO frequency used for our templates. We reserve studying the inclusion of merger and ringdown, tidal effects and comparison with numerical waveforms to a future paper.

Assuming the larger computational cost can be satisfied, we conclude with the recommendation of switching from nonspinning to aligned-spin templates for future NSBH searches in the mass range considered here, as this would provide good sensitivity across the entire BH spin parameter space, including both the aligned and precessing high-spin cases.

ACKNOWLEDGMENTS

We thank Steve Privitera, Tom Dent, and Badri Krishnan for useful discussion and comments on the paper. We also acknowledge the LIGO collaboration for providing the recolored synthetic strain data we analyzed. T.D.C. is supported by the International Max-Planck Research School on Gravitational-Wave Astronomy. This paper has LIGO Document No. P1400209.

- [1] G. M. Harry (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084006 (2010).
- [2] F. Acernese *et al.* (The Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] K. Somiya (KAGRA Collaboration), *Classical Quantum Gravity* **29**, 124007 (2012).
- [4] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **79**, 122001 (2009).
- [5] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **82**, 102001 (2010).
- [6] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **83**, 122005 (2011).
- [7] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys. Rev. D* **85**, 082002 (2012).
- [8] J. Aasi *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **87**, 022002 (2013).
- [9] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev. D* **86**, 084017 (2012).
- [10] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
- [11] T. Dal Canton, A. H. Nitz, A. P. Lundgren, A. B. Nielsen, D. A. Brown, T. Dent, I. W. Harry, B. Krishnan, A. J. Miller, K. Wette, K. Wiesner, and J. L. Willis, *Phys. Rev. D* **90**, 082004 (2014).
- [12] R. W. O’Shaughnessy, J. Kaplan, V. Kalogera, and K. Belczynski, *Astrophys. J.* **632**, 1035 (2005).
- [13] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *Astrophys. J.* **759**, 52 (2012).
- [14] V. Kalogera, *Astrophys. J.* **541**, 319 (2000).
- [15] P. Grandclément, M. Ihm, V. Kalogera, and K. Belczynski, *Phys. Rev. D* **69**, 102002 (2004).
- [16] K. Belczynski, R. E. Taam, E. Rantsiou, and M. van der Sluys, *Astrophys. J.* **682**, 474 (2008).
- [17] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski, *Astrophys. J. Lett.* **719**, L79 (2010).
- [18] P. C. Fragile, *Astrophys. J.* **706**, L246 (2009).
- [19] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [20] D. A. Brown, A. Lundgren, and R. O’Shaughnessy, *Phys. Rev. D* **86**, 064020 (2012).
- [21] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, B. S. Sathyaprakash, H. Tagoshi, and H. Takahashi, *Phys. Rev. D* **80**, 024009 (2009).
- [22] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **78**, 042002 (2008).
- [23] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, *Phys. Rev. D* **89**, 024010 (2014).
- [24] L. Barsotti and P. Fritschel, LIGO Technical Note T1200307-v4, 2014.
- [25] J. E. McClintock, R. Narayan, S. W. Davis, L. Gou, A. Kulkarni, J. A. Orosz, Robert F Penna, Ronald A Remillard, and James F Steiner, *Classical Quantum Gravity* **28**, 114009 (2011).
- [26] T. Damour, A. Nagar, and L. Villain, *Phys. Rev. D* **85**, 123007 (2012).
- [27] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, *Astrophys. J.* **725**, 1918 (2010).
- [28] F. Özel, D. Psaltis, R. Narayan, and A. Santos Villarreal, *Astrophys. J.* **757**, 55 (2012).
- [29] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, *Phys. Rev. D* **87**, 024035 (2013).
- [30] A. Lundgren and R. O’Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014).
- [31] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015).
- [32] A. Buonanno, Y.-b. Chen, and M. Vallisneri, *Phys. Rev. D* **67**, 104025 (2003).
- [33] Y. Pan, A. Buonanno, Y. Chen, and M. Vallisneri, *Phys. Rev. D* **69**, 104017 (2004).
- [34] I. Harry and S. Fairhurst, *Classical Quantum Gravity* **28**, 134008 (2011).
- [35] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **89**, 084006 (2014).
- [36] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [37] F. Foucart, *Phys. Rev. D* **86**, 124007 (2012).
- [38] F. Pannarale, L. Rezzolla, F. Ohme, and J. S. Read, *Phys. Rev. D* **84**, 104017 (2011).