



LIGO SURF Interim Report 2: Constraining the Precession of Binary Black Hole Systems Using New Parameters

CHARLES F. A. GIBSON ^{1,2} AND JAVIER ROULET ^{1,3}

¹*LIGO, California Institute of Technology, Pasadena, CA 91125, USA*

²*Department of Physics, Allegheny College, Meadville, Pennsylvania 16335, USA*

³*TAPIR, Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125, USA*

ABSTRACT

The precession of Binary Black Holes (BBHs) can be informative of the formation channel of the system; weakly/non-precessing systems are likely to have formed through binary stellar evolution, while strongly precessing systems may have formed dynamically. Despite the growing number of LIGO sources, evidence of precession is strongly debated in the literature. The parameter χ_p is currently used to evaluate the precession of observed BBH systems. However, χ_p is difficult to constrain to a narrow range of values for most events and can yield vanishing prior probability density at the aligned-spin configuration. We present an alternative spin precession parameter, the cosine of the angle between the total spin and the orbital angular momentum $\cos\theta_{LS}$, that provides better localization of a precession value and allows a non-zero probability of aligned spins. We begin by testing $\cos\theta_{LS}$, χ_p , and other parameters against synthetic data with known values to determine the best statistical measurement of precession. We then use $\cos\theta_{LS}$ to evaluate the precession in events from O3, the third observing run of LIGO and Virgo.

1. MOTIVATION

The properties of binary black hole (BBH) mergers observed from LIGO-Virgo-KAGRA detections can be informative of the formation channel of the system (e.g., Mandel & Farmer 2022). Two primary theories of the origin of BBHs exist. The first is that the systems were formed through stellar evolutionary channels. Namely, as a binary system between two intermediately massive stars evolved, both stars remained in orbit, with the resulting black holes (BHs) surviving the supernovae at the end of the stars' lives. Eventually, due to the emission of gravitational waves (GWs), the two coalesced into a single BH through a BH-BH merger.

Alternatively, the BBH system may have been formed dynamically. Through the gravitational interactions of stars and black holes in dense stellar environments such as globular clusters and galactic nuclei, scattering events can place two, previously unrelated BHs into orbit around each other. This would most likely be from a three-body interaction in which an intruding BH kicks a less massive companion from a the binary the other BH is in, yielding a BBH system.

One way to potentially differentiate between these two formation channels is through analyzing the precession of the orbit. The BHs in BBH systems that formed from binary stellar evolution likely have spins \vec{S} that

are aligned with the orbital angular momentum \vec{L} . This stems from the preferential alignment of stellar rotation axes with the \vec{L} of the binary, initialized by the angular momentum in stellar nurseries. Additional complications such as kicks from the supernovae of companion stars in the binary may misalign spins. However, the details of these processes are still not well modeled, so approximations to the effects must be taken into account. The simplest of approximations neglect these kicks, claiming that \vec{S} and \vec{L} remain aligned through the entire binary evolution process through the BBH merger.

Conversely, dynamically formed BBH systems are much more likely to have isotropic spin distributions. Because there is no initial relationship between \vec{S} and \vec{L} , the alignment of \vec{S} and \vec{L} is just as likely as the misalignment of the two vectors. This assumption leads to the prediction that the orbits of dynamically formed BBH systems are more likely to precess than the orbits of binary stellar evolution remnants.

By understanding the precession of a BBH system, information regarding the formation channel of the binary can be gleaned. In particular, analyzing the precession found in LIGO-Virgo data from O1, O2, and O3 can help inform predictions of the origins of known BBH merger candidates. With just under 100 candidates of BBH systems as of O3 (Abbott et al. 2023; Mehta et al.

2023; Nitz et al. 2023), statistical conclusions can begin to be made about the nature of BBH precession and, therefore, the origin of the BBH systems.

These conclusions may be especially useful in understanding the nature of binary evolution, dense stellar environments, and dynamical interactions.

Currently, there exists a parameter χ_p that has been used to describe the precession of the orbit. However, claims of individual precession candidates are controversial (Hannam et al. 2022; Payne et al. 2022). That is because this parameter is not necessarily very informative of the individual precession of a BBH system. The issues with χ_p are described in detail in Section 2. This summer, we focus on defining a new parameter that can better constrain orbital precession of BBH systems. The progress made so far is detailed in Section 3, and the challenges encountered are outlined in Section 4

2. PROBLEM

The effective precession parameter currently used to describe the precession of a BBH system, χ_p , is defined as

$$\chi_p = \max \left(\chi_1 \sin \theta_{S_1 L}, \frac{q(4q+3)}{3q+4} \chi_2 \sin \theta_{S_2 L} \right), \quad (1)$$

where χ_i is the dimensionless spin parameter of the BH i , q is the mass ratio m_2/m_1 (where $m_1 > m_2$), and $\theta_{S_i L}$ is the angle between the spin \vec{S} of BH i and the orbital angular momentum \vec{L} (Schmidt et al. 2015). When $\chi_p = 0$, the system is not precessing, and when $\chi_p = 1$, the system is strongly precessing.

This parameter has two main issues that make it difficult to analyze precession: Both can be seen by the posterior distributions in Figure 1. Most of the posterior distributions for χ_p are very broad. A broad posterior distribution is not very informative on the true value associated with the data, as it makes it difficult to constrain the value to a reasonable range. The second is displayed by the prior distribution in Figure 1: the prior distribution of χ_p sharply approaches 0 as χ_p approaches 0. In mathematical terms, $\pi(\chi_p = 0) = 0$.

The first issue makes χ_p a poor parameter statistically. The second issue fails to address a fundamental goal of the χ_p parameter: to put to test the hypothesis that the spins are aligned. However, by initially assuming the spins are misaligned (as the probability of alignment is 0 in the prior in χ_p), the parameter fails to reject aligned spins. This is because the posterior distribution is defined as the prior distribution times the likelihood, so if the prior is 0 at a value, then the posterior will always be 0 at that value.

This summer, we aim to propose an alternative parameter that addresses these two issues. Namely, we




Figure 1. The χ_p distributions of several observations from Abbott et al. (2023). Note that most of the posteriors (upper curves) are very broad, only marginally differing from the prior distribution (lower curves). For the more localized posteriors, the localization only occurs at low values of χ_p where the peak in the prior occurs, and these events have high levels of uncertainty of astrophysical origin.

want a parameter that has a narrow, well-constrained distribution and contains the true precession value, and we want a parameter that does not reject aligned spins in the prior.




Figure 2. The geometry of a BBH system. The spins of each black hole are denoted by \vec{S}_i (with the total spin $\vec{S}_{\text{tot}} = \vec{S}_1 + \vec{S}_2$), the orbital angular momentum is expressed as \vec{L} , and the total angular momentum ($\vec{L} + \vec{S}_1 + \vec{S}_2$) is \vec{J} . θ_{LS} is the angle between \vec{L} and \vec{S}_{tot} . β is the angle between \vec{J} and \vec{L} .

3. ACCOMPLISHMENTS SO FAR

Using the geometry of the BBH merger outlined in Figure 2, two alternative parameters were initially selected based on the geometry of the system. First is θ_{LS} , the angle between the \vec{L} and total spin $\vec{S}_{\text{tot}} = \vec{S}_1 + \vec{S}_2$. This angle provides a direct geometric understanding of the relationship between \vec{S}_{tot} and \vec{L} , fundamentally relating to the orbital precession. The second is β , the angle between \vec{L} and the total angular momentum $\vec{J} = \vec{S}_{\text{tot}} + \vec{L}$. β is especially promising because, as a precession indicator, it impacts the magnitude of the amplitude modulations in the waveform (Fairhurst et al. 2020). Particularly, the parameter $b = \tan(\beta/2)$ is directly used to compute the waveform. However, unlike β , θ_{LS} , and χ_p , b has infinite bounds, making it more difficult to constrain a “maximum” precession. Regardless, θ_{LS} and β share the same issue with χ_p in the sense that (under an isotropic spin prior) their probability densities both tend towards zero when \vec{S} and \vec{L} are aligned. To combat this issue, we consider the cosine of the angles, $\cos \theta_{\text{LS}}$ and $\cos \beta$. This coordinate shift to cosine is chosen because it yields a non-zero probability of aligned spins in the prior.

In order to measure how informative the three parameters (χ_p , $\cos \theta_{\text{LS}}$, and $\cos \beta$) are, we needed to test them on known values. As the exact values of the main parameters from LIGO-Virgo-KAGRA sources are not known, we instead used synthetic data with posteriors formed from known injections. These injections were generated assuming an isotropic spin distribution. In other words,

all true angles between \vec{L} and \vec{S} are equally likely in the synthetic data. We used roughly 3000 posteriors with known injection values in this data set that was used.¹

Using the posterior distributions obtained from the injected samples, we constructed several functions to convert the raw (posterior) data into a posterior distribution of the parameters $\cos \beta$, $\cos \theta_{\text{LS}}$, and χ_p . Initial examinations of the effectiveness of each parameter for a randomly selected case are shown in Figure 3.

However, in order to evaluate how informative each parameter is most effectively, we ran a statistical significance test, starting from the Neyman-Pearson Lemma. This is the strongest test for comparing two hypotheses, \mathcal{H}_0 and \mathcal{H}_1 , against each other (in this case, having the \vec{S}_{tot} isotropically misaligned (\mathcal{H}_0) or aligned (\mathcal{H}_1) with \vec{L}). This test is defined the likelihood ratio of two hypotheses, expressed as

$$\Lambda = \frac{p(d | \mathcal{H}_1)}{p(d | \mathcal{H}_0)} \quad (2)$$

(Neyman & Pearson 1933). If the probability of \mathcal{H}_1 is greater than the probability of \mathcal{H}_0 , the ratio is greater than 1. A threshold to eliminate the null hypothesis \mathcal{H}_0 can be set.

Each hypothesis can be a set of parameters that yield some result. For the 1-dimensional case,

$$\mathcal{H}_0: \theta \sim \pi(\theta) \quad (3)$$

$$\mathcal{H}_1: \theta = \theta_*. \quad (4)$$

In other words, \mathcal{H}_0 is the initial estimate of the probability distribution isotropic spin alignment, and \mathcal{H}_1 is the value corresponding to \vec{S}_{tot} and \vec{L} alignment. The aligned-spin model is said to be nested in the isotropic spin model because it has fewer parameter and it can be recovered as a particular case of the isotropic spin model, in which its internal parameters (spin tilts) have been set to a specific value. This allows us to rewrite the likelihood ratio (2) as follows.

Using the relationship between the likelihood \mathcal{L} , the posterior \mathcal{P} , the prior π , and the evidence \mathcal{Z} ,

$$\mathcal{P} = \frac{\mathcal{L}\pi}{\mathcal{Z}} \quad (5)$$

and the definitions of \mathcal{H}_0 and \mathcal{H}_1 , we can express the likelihood $p(d | \mathcal{H}_1)$ as

$$p(d | \mathcal{H}_1) = p(d | \theta_*) \quad (6)$$

$$p(d | \theta_*) = \mathcal{L} \quad (7)$$

$$\mathcal{L} = \frac{p(\theta_* | d, \mathcal{H}_0) p(d | \mathcal{H}_0)}{\pi(\theta_* | \mathcal{H}_0)}. \quad (8)$$

¹ The injections and posterior distributions can be found at <https://zenodo.org/records/10910135>.

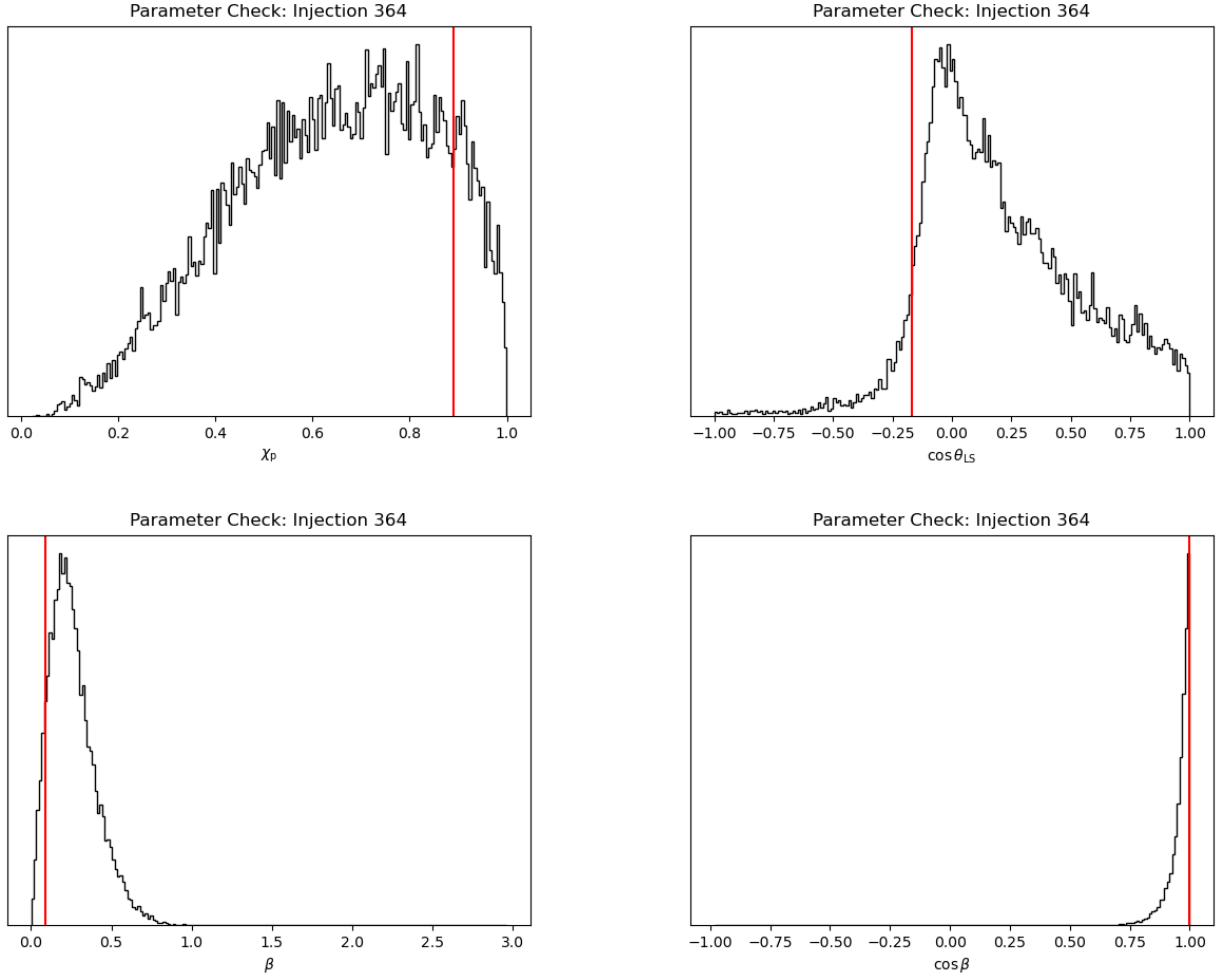


Figure 3. The posterior distributions of four parameters for a randomly selected, low-mass injection (Injection 364). Although each distribution contains the true, injected value, there are varying degrees of the breadth of each posterior distribution around the injected value. χ_p has a very broad distribution over accepted values, making it difficult to pinpoint its true value (vertical red line) without knowing it beforehand. For the same data, $\cos \theta_{LS}$ has a slightly more defined distribution around the true value. β has a narrower distribution, whereas $\cos \beta$ has the narrowest distribution of the four parameters. It is important to note that this is a general trend found across the data, but it is not necessarily the case for all posteriors. Indeed, there are some parameters in which χ_p retains a sharply defined peak in the posterior. A more quantitative value is introduced later that can statistically evaluate how good the parameter is for each injection.

198 Plugging this into Equation (2), we get

$$199 \quad \frac{p(\theta_* | d, \mathcal{H}_0)}{\pi(\theta_* | \mathcal{H}_0)}. \quad (9)$$

200 The ratio in Equation (9) is called the Savage-Dickey
 201 density ratio and can be used to estimate the likelihood
 202 ratio (Bayes factor) between two nested models. How-
 203 ever, a BBH system is not 1-dimensional, as it is defined
 204 by many parameters. Although precession is not neces-
 205 sarily based on a single parameter, our goal is to find
 206 a single parameter that can provide significant informa-
 207 tion on the precession of the system. We can express
 208 θ as a multidimensional parameter that contains a sin-
 209 gle parameter x that preserves the relevant precession

210 information and all other unrelated parameters θ' as

$$211 \quad \theta = (x, \theta') \quad (10)$$

212 We can then define a new Hypothesis $\tilde{\mathcal{H}}_1$ that remains
 213 as close to \mathcal{H}_1 as possible while only being based on one
 214 parameter. We choose $\tilde{\mathcal{H}}_1$ to differ from \mathcal{H}_0 by only one
 215 parameter as this will allow us to test simplified popu-
 216 lation models that only consider very few BBH param-
 217 eters at a time. Ideally, there should be a single value
 218 of $x = x_*$ that allows \mathcal{H}_0 to be the same as $\tilde{\mathcal{H}}_1$. That
 219 is, in an isotropic spin distribution (\mathcal{H}_0), there should
 220 be only one orientation of the vectors that gives aligned

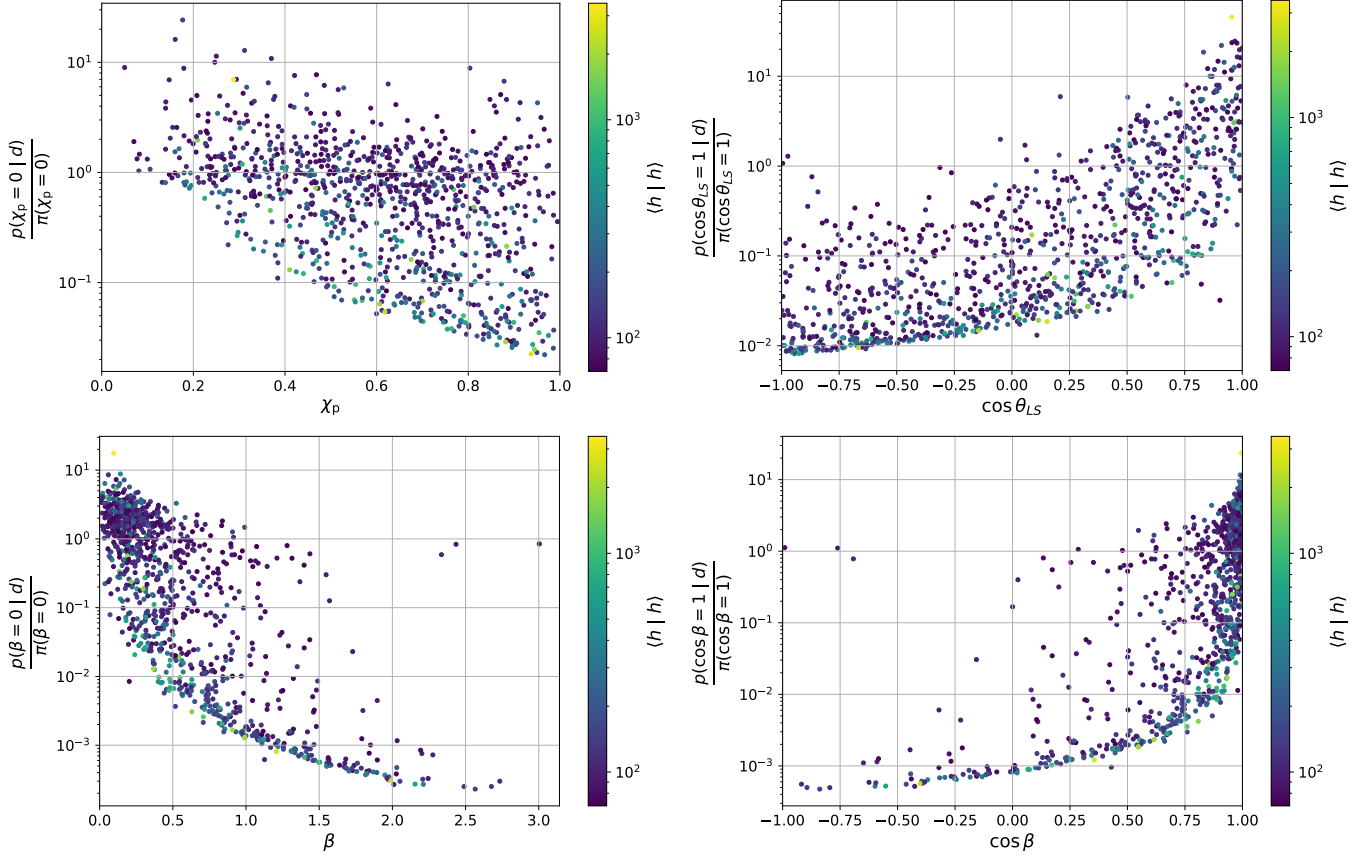


Figure 4. Savage-Dickey ratios for the three tested parameters. Each point represents an injection and associated parameter estimation. The SDR used is between the hypotheses that the spins are aligned rather than isotropic. The true, injected value is on the horizontal axis, while the Savage-Dickey ratio is on the vertical axis. The color bar symbolizes the strength of the signal h given the injected signal strength. The likelihood ratio for χ_p is strongly clustered around values ranging from ~ 1 , making it a poor test of spin alignment. Meanwhile, the likelihood ratio for $\cos \theta_{LS}$ is slightly more informative. Because it spans several orders of magnitude, strongly precessing BBH systems ($\cos \theta_{LS} \sim -1$) would be much more likely than weakly precessing systems to be ruled out as having aligned spins. Of the three parameters, β seems to be the most informative. The Savage-Dickey ratios that define β of the data span many orders of magnitude with a high accuracy of correctly interpreting the alignment of the spins. Generally, the SDRs for coordinate shifts of the same variable, such as from β to $\cos \beta$, should retain the same values. The slight discrepancies stem from deviations in the binned probability densities of the injection distribution.

spins $\tilde{\mathcal{H}}_1$ (again, based on a single parameter).² This is expressed as

$$\pi(\theta' | x_*, \tilde{\mathcal{H}}_1) = \pi(\theta' | x_*, \mathcal{H}_0), \quad (11)$$

and thus,

$$\pi(\theta | \tilde{\mathcal{H}}_1) = \delta(x - x_*) \pi(\theta' | x_*, \mathcal{H}_0). \quad (12)$$

Assuming that a single value x contains all the available information about precession, it follows that $p(d | \mathcal{H}_1) \approx p(d | \tilde{\mathcal{H}}_1)$. Using the same process as in Equations (7-

10), we can express $p(d | \tilde{\mathcal{H}}_1)$ as

$$p(d | \mathcal{H}_1) = p(d | x_*, \tilde{\mathcal{H}}_1) \quad (13)$$

$$= p(d | x_*, \mathcal{H}_0) \quad (14)$$

$$= \frac{p(x_* | d, \mathcal{H}_0) p(d | \mathcal{H}_0)}{\pi(x_* | \mathcal{H}_0)}, \quad (15)$$

and plugging Equation (15) into Equation (2), we get the revised ratio

$$\Lambda = \frac{p(d | \tilde{\mathcal{H}}_1)}{p(d | \mathcal{H}_0)} \quad (16)$$

$$= \frac{\frac{p(x_* | d, \mathcal{H}_0) p(d | \mathcal{H}_0)}{\pi(x_* | \mathcal{H}_0)}}{p(d | \mathcal{H}_0)} \quad (17)$$

$$= \frac{p(x_* | d, \mathcal{H}_0)}{\pi(x_* | \mathcal{H}_0)} \quad (18)$$

² This approximation does neglect some information. For example, if the two spins have vertical components of \vec{S} that align with \vec{L} but the horizontal components of their spins cancel, this simplification fails to identify the spin misalignment in the system.

This ratio, called the Savage-Dickey Ratio (SDR), provides a simpler way to compare the competing hypotheses with the use of a single parameter, allowing for a more quantitative way to evaluate the most informative parameter on the alignment of \vec{S}_{tot} and \vec{L} .

Using the SDR, we demonstrate in Figure 4 that χ_p is not very informative about the alignment of \vec{S} and \vec{L} , while parameters based on β (β , $\cos \beta$) are the most informative about the alignment of \vec{S} and \vec{L} out of the three parameters tested. Currently, this analysis provides a strong incentive to evaluate the strength of the orbital precession of actual LIGO-Virgo-KAGRA sources using $\cos \beta$ as the precession parameter in place of χ_p . This is because $\cos \beta$ (χ_p) yields the highest (lowest) SDR for aligned (misaligned) spins. To choose the parameter that best defines the precession, we plan to employ another quantitative test to summarize the SDR information of each parameter in a single value.

3.1. Comparisons to Other Spin Parameters

Although χ_p is currently in use as a parameter, another spin parameter, χ_{eff} , exists and is used to describe the mass-weighted average spin aligned with the orbital angular momentum. χ_{eff} is usually much better measured than χ_p , but it provides fundamentally different information from χ_p despite the two providing information on the spin. In order to ensure that the proposed alternative parameters to χ_p are also providing unique information not obtained from χ_{eff} , we also needed to compare the SDRs of the new parameters to the likelihood ratios of χ_{eff} . If there is a correlation between the two, then it means that the information present in one parameter (e.g., $\cos \beta$) can also be found from χ_{eff} , making it less unique of a parameter. Figure 5 demonstrates that $\cos \beta$ provides unique information from χ_{eff} .

4. OBSTACLES ENCOUNTERED

4.1. Indeterminate Savage-Dickey Ratios

One issue I encountered while running the analysis of the likelihood ratio involved the shape of the prior of χ_p and β . For all the isotropic injections, there was not a single Savage-Dickey Ratio for β that was greater than 1 which was unexpected behavior. As both of these values approach 0 (no precession/aligned spins), the probability of the prior also approaches 0. This property makes it difficult for the SD ratio to be evaluated at 0. Using Equation (18), we can express this as

$$\lim_{\beta, \chi_p \rightarrow 0} \frac{p(\text{Aligned Spins} | d)}{\pi(\text{Aligned Spins})} \rightarrow \frac{0}{0}. \quad (19)$$

As this expression is in indeterminate form, we can use L'Hôpital's Rule to redefine this equation as

$$\lim_{\beta, \chi_p \rightarrow 0} \frac{p'(\text{Aligned Spins} | d)}{\pi'(\text{Aligned Spins})}. \quad (20)$$

We aimed to redefine p' and π' . We began by assuming a small area ε under the curve close to 0 for both curves (where $\varepsilon_p = \varepsilon_\pi$). Each triangle then has a base q_p and q_π . This makes the height of the triangle $h = \frac{2\varepsilon}{q}$, ultimately defining $p' = \frac{2\varepsilon}{q_p^2}$ and $\pi' = \frac{2\varepsilon}{q_\pi^2}$. The ratio of the two derivatives then is defined as

$$\frac{p'}{\pi'} = \frac{q_\pi^2}{q_p^2}. \quad (21)$$

The geometry of this derivation is outlined in Figure 6

This approach of substituting Equation (21) in for Equation (18) allowed for a much more reasonable set of SDRs for both χ_p and β . We also tested it for several ε values ($\varepsilon = 0.003, 0.005, 0.007, 0.01$) and found that the shape of the resulting Savage-Dickey Ratio plots was relatively insensitive to the ε value when ε is small.

5. UPDATES SINCE INTERIM REPORT 1

5.1. Selection Effects in Synthetic Posteriors

As mentioned earlier, the synthetic data is generated based on an isotropic distribution of spin angles. However, the data is also filtered to only allow events that would have signals recognizable by LIGO. As the mass of the BHs in the system increase, selection effects begin to occur. Most notably, the frequency of the BBH merger is inversely proportional to the mass, $f_{\text{merger}} \propto 1/M$. This means that more massive BBHs merge at lower frequencies. The alignment of \vec{S} and \vec{L} also affects the frequency of the merger: highly misaligned \vec{S} and \vec{L} merge more quickly at lower frequencies. These two effects cause the final frequency of high-mass, strongly precessing mergers to occur at low frequencies, potentially being undetectable by LIGO.

5.2. Altering the Savage-Dickey Ratio of the Angle Between The Total Spin and Orbital Angular Momentum

Equation (18) is used to evaluate SDR for a hypothesis described by a single parameter. The synthetic injections used to construct the posteriors used for the analysis were based on a population with isotropic spin distributions, where all orientations of $\cos \theta_{\text{LS}}$ are equally likely. This means that $\pi(\text{Aligned Spins}) = 0.5$. However, the data is then filtered to only allow events that would be detectable by LIGO. The selection effects outlined in Section 5.1 mean that the prior of $\cos \theta_{\text{LS}}$ is not necessarily flat, as more massive, highly precessing candidates are more difficult to observe. The shape of each prior based on the mass distribution is presented in Fig-

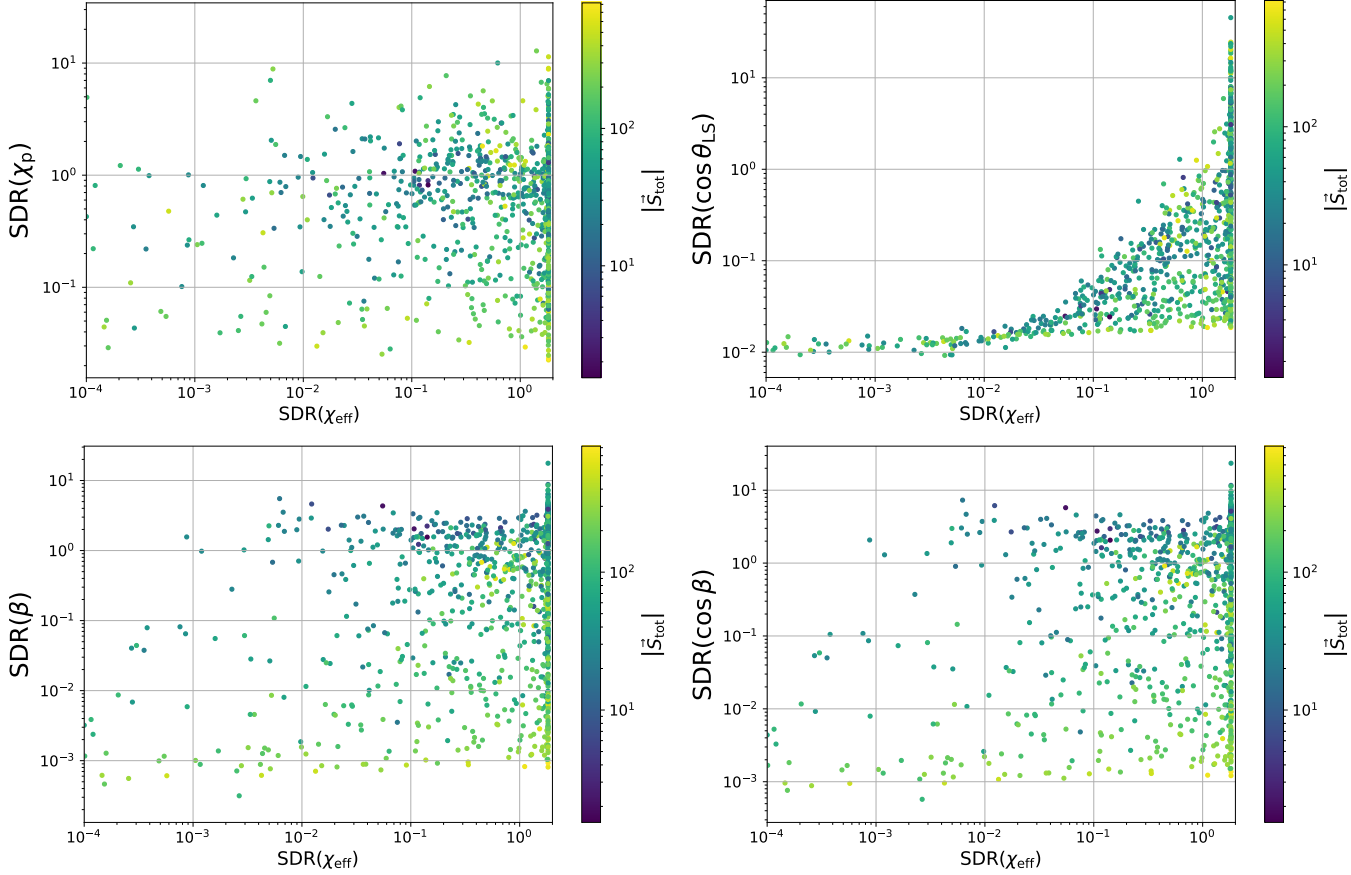


Figure 5. The Savage-Dickey ratios of the tested parameters plotted against the likelihood ratios of χ_{eff} for the low-mass injections. The data are colored based on the total spin of the system. There is little correlation between the SD ratios for $\cos \beta$ and the likelihood ratios of χ_{eff} , meaning that the information provided from $\cos \beta$ is unique from that of χ_{eff} . The same is true for β (which is expected as $\cos \beta$ is a coordinate transformation of β) and for χ_p . However, given this data, $\cos \theta_{\text{LS}}$ appears to be more informative than χ_{eff} . This is because the highest values of the SD ratio for $\cos \theta_{\text{LS}}$ strongly inform the value of χ_{eff} , while the highest likelihood ratios for χ_{eff} retain a high range of SD ratio values for $\cos \theta_{\text{LS}}$. Ultimately, it appears that parameters based on β provide unique information from χ_{eff} , giving credence to its use as an alternative precession parameter.




Figure 6. The geometry used to derive Equation (21).

332 ure 7. When initially calculating $\cos \theta_{\text{LS}}$, we used Equa-

333 tion (18) with $\pi(\text{Aligned Spins}) = 0.5$ as an analytic so-
 334 lution accounting for a flat $\cos \theta_{\text{LS}}$ prior. To get a more
 335 accurate evaluation of the SDRs for $\cos \theta_{\text{LS}}$, we altered
 336 the value of the prior of $\cos \theta_{\text{LS}}$ to be equal to the proba-
 337 bility density of $\cos \theta_{\text{LS}} = 1$, estimated from a histogram
 338 of the injected samples. When accounting for the selec-
 339 tion effects, which tend to increase $\pi(\text{Aligned Spins})$,
 340 the SDR of $\cos \theta_{\text{LS}}$ decreased by as much as a factor of
 341 2. This makes $\cos \theta_{\text{LS}}$ better at rejecting aligned spins
 342 for highly precessing candidates and about as effective
 343 at confirming aligned spins for non-precessing systems
 344 as $\cos \beta$ and χ_p .

6. CHOOSING A PARAMETER

6.1. Evaluating the Divergence of Isotropic and Aligned Spin Distributions

348 Although χ_p seems to be less informative than $\cos \beta$
 349 and $\cos \theta_{\text{LS}}$ by looking at the trend of the SDRs, it be-




Figure 7. The probability density distribution of injections for the three mass distributions. The data was initially generated to have a flat prior regardless of mass, but as the mass increases, the probability density of having aligned \vec{S} and \vec{L} increase. This is due to the filtering of events that exclusively selects events that can be detected by LIGO, as described in Section 5.1.

came difficult to compare the effectiveness of $\cos \beta$ or $\cos \theta_{LS}$ as they are both relatively effective. The synthetic posteriors for each set of mass distributions contain just over 1000 values, leaving over 3000 SDR values. One way to summarize the effectiveness of each parameter across all samples is to measure the Kullback-Leibler Divergence D_{KL} , a test that evaluates the difference between two distributions. In particular, we can use this to evaluate the difference between the likelihoods of aligned and isotropic spin distributions. By maximizing D_{KL} between the two distributions with our parameter, we could find the parameter that yields the most divergent set of distributions. The Kullback-Leibler Divergence between these two likelihoods is defined as

$$D_{KL}(p(d | \mathcal{H}_0) || p(d | \tilde{\mathcal{H}}_1)) = \int dd p(d | \mathcal{H}_0) \log_2 \frac{p(d | \mathcal{H}_0)}{p(d | \tilde{\mathcal{H}}_1)}, \quad (22)$$

which can be approximated as

$$D_{KL} \approx \frac{1}{N} \sum_{d_j \sim \mathcal{H}_0} \log_2 \frac{p(d_j | \mathcal{H}_0)}{p(d_j | \tilde{\mathcal{H}}_1)}. \quad (23)$$

However, we have shown in Equations (16-18) that the inverse of this ratio of likelihoods can be expressed as the ratio of the posterior to the prior. Incorporating

Table 1. $D_{KL}(p(d | \mathcal{H}_0) || p(d | \tilde{\mathcal{H}}_1))$ (bits)

\mathcal{M}/M_\odot	χ_p	$\cos \theta_{LS}$	β	$\cos \beta$	b
(1, 5)	1.0	3.0	2.9	2.4	2.9
(5, 25)	1.7	2.9	1.8	2.1	1.8
(25, 125)	0.9	2.0	0.9	1.0	0.9

Notes: The D_{KL} values of each parameter comparing the distributions of aligned to isotropic spins. The “Mass” column corresponds to one of three mass distributions of the injections used to construct the posteriors.

this result, we can express the divergence as

$$D_{KL} = \frac{1}{N} \sum_{d_j \sim \mathcal{H}_0} \log_2 \frac{\pi(x_* | \mathcal{H}_0)}{p(x_* | d, \mathcal{H}_0)} \quad (24)$$

$$= -\frac{1}{N} \sum_{d_j \sim \mathcal{H}_0} \log_2 \text{SDR}. \quad (25)$$

The D_{KL} values for each parameter across the three distributions used in this study are reported in Table 1. χ_p consistently has the lowest D_{KL} while $\cos \theta_{LS}$ consistently has the highest. This means that χ_p is the least effective at distinguishing between the isotropic and aligned spin distributions while $\cos \theta_{LS}$ is the best. We expect $D_{KL}(\beta) = D_{KL}(\cos \beta) = D_{KL}(b)$, but this is not the case, likely for the same reason(s) that the SDR values are not consistent between the three coordinate systems of β . Regardless, none of the D_{KL} values for any coordinate of β in any mass distribution are greater than that of $\cos \theta_{LS}$. Additionally, as the distribution goes to higher masses, it becomes more difficult to distinguish between isotropic and aligned spins, especially between β and χ_p .

Given these results, it appears that $\cos \theta_{LS}$ is the strongest parameter at distinguishing misaligned spins from aligned spins, while χ_p is the weakest.

6.2. Relationship to Other Parameters

Although $\cos \theta_{LS}$ appears to most accurately diagnose the precession of a BBH system, it is important to compare it to other parameters. We have already shown that the mass of the BBH system influences which events may be detected by LIGO. However, it is also important to recognize further relationships with other parameters. Most prominently is the relationship with the total spin. Figure 8 demonstrates that $\cos \theta_{LS}$ is more effective at rejecting aligned spins for highly precessing systems with a high total spin.

7. POPULATION INFERENCE




Figure 8. The updated SDRs for $\cos \theta_{\text{LS}}$ colored by the total spin. Higher spins tend to have the lowest SDR values. This makes $\cos \theta_{\text{LS}}$ better at confirming precession for BBHs with high spins.

Now that an alternative parameter, $\cos \theta_{\text{LS}}$, has been defined, we can begin population inference. By analyzing loud signals in the LIGO strain data, called triggers, with a high likelihood of astrophysical origin, we can evaluate the precession using $\cos \theta_{\text{LS}}$.

For any N events in LIGO data, N_a of these events are of astrophysical origin, while N_b of these events come from background noise. The population has a set of hyperparameters λ that influence the detected data d . The probability that the event is of astrophysical origin is defined as

$$p_{\text{astro}} = \frac{dN_a(\lambda)}{dN_a(\lambda) + dN_b}. \quad (26)$$

By establishing a significance metric, the false-alarm rate, we can set a threshold for which triggers to analyze based on how likely they are to stem from astrophysical origin. Assuming some arbitrary set of parameters that describes a population model λ_0 , the probability of a trigger occurring is defined as

$$P(N_{\text{trig}}, d_i | \lambda) \propto e^{-N_a(\lambda)} \times \prod_{i=1}^{N_{\text{trig}}} \left[\frac{dN_a(\lambda)}{dN_a(\lambda_0)} \Big|_{d_i} p_{\text{astro},i}(\lambda_0) + (1 - p_{\text{astro},i}(\lambda_0)) \right] \quad (27)$$

This is derived in detail in Roulet et al. (2020). The likelihood for the population model can be evaluated from

posterior samples for all events while using pipeline injections to quantify the search sensitivity and the significance of each of the events above the minimum threshold. Currently, we are lacking the posterior samples for all events.

By identifying triggers with high p_{astro} through a population model, we then evaluated events that are aligned with those in the GWTC-3 catalogue. We are currently in the process of using the strain data from each detector and a flat $\cos \theta_{\text{LS}}$ prior to construct samples that model the posterior distributions of each parameter. Using these base parameters, we can then construct posterior distributions for our newly introduced spin-precession parameters, following the same analysis processes that we used to analyze the synthetic data that we used to originally constrain the precession parameters.

We have already shown that the posterior distribution of $\cos \theta_{\text{LS}}$ is better localized than that of χ_p in Figure 3. That means that any event with a localized $\cos \theta_{\text{LS}} \ll 1$ is much more likely to be precessing than a non-zero χ_p which will likely have a broad posterior distribution.

At the time of writing this report, we have analyzed a single event, GW190917, which is a marginal event that is likely a merger between a BH and a neutron star. The distributions of χ_p and $\cos \beta$ are shown in Figure 9. The distribution of χ_p is very broad, while $\cos \theta_{\text{LS}}$ is much more sharply concentrated at -1, albeit retaining a relatively broad distribution with a non-zero probability of $\cos \theta_{\text{LS}} = 1$. Although this shape of $\cos \theta_{\text{LS}}$ is promising for a precessing candidate, this event has of a relatively low p_{astro} (Abbott et al. 2023). However, we hope to analyze all events in GWTC-3, allowing for an analysis of more confident events and potentially confirm precession in other cases. We hope to find an event with a strongly localized probability of $\cos \theta_{\text{LS}} \neq 1$ with a very low/zero probability of alignment.

8. FUTURE PLANS

As mentioned above, we aim to sample each event in GWTC-3 and construct corresponding posterior distributions. We then would like to re-evaluate the precession of each event using $\cos \theta_{\text{LS}}$.

This work was supported by the National Science Foundation Research Experience for Undergraduates (NSF REU) program, the LIGO Laboratory Summer Undergraduate Research Fellowship program (NSF LIGO), and the California Institute of Technology Student-Faculty Programs.

REFERENCES

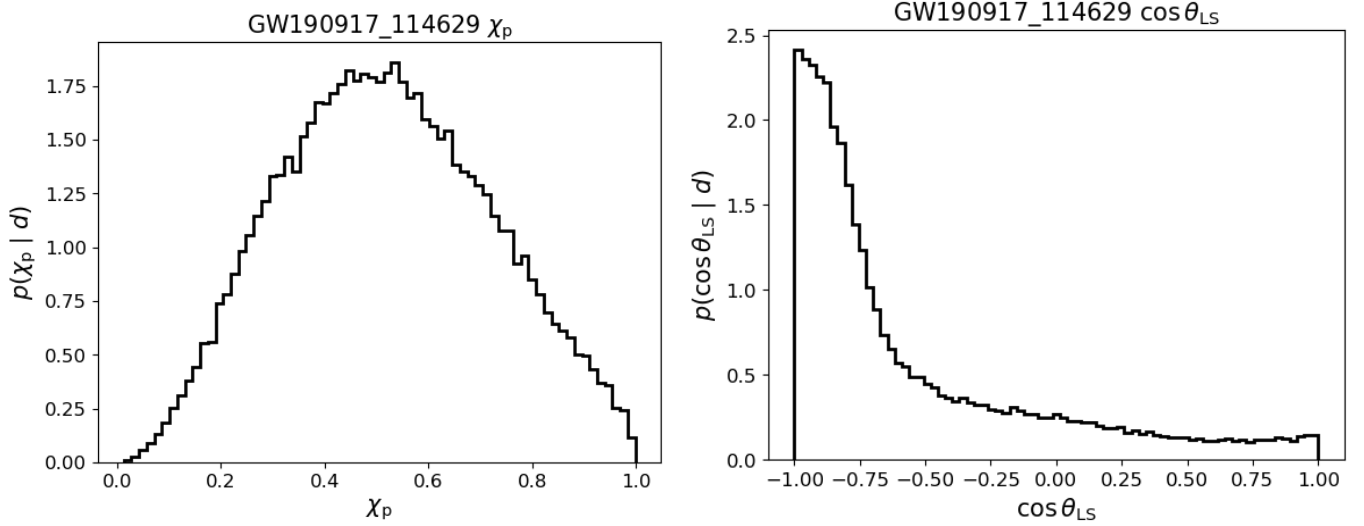


Figure 9. The probability density distribution of χ_p and $\cos \theta_{LS}$ for the marginal NSBH event GW190917. This event has high misalignment as predicted by $\cos \theta_{LS}$. Although the distribution of χ_p also peaks at a non-zero value, it retains a broad distribution, making it more difficult to evaluate the true value of precession.

472 Abbott, R., Abbott, T. D., Acernese, F., et al. 2023,
 473 Physical Review X, 13, 041039,
 474 doi: [10.1103/PhysRevX.13.041039](https://doi.org/10.1103/PhysRevX.13.041039)
 475 Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, Phys.
 476 Rev. X, 13, 011048, doi: [10.1103/PhysRevX.13.011048](https://doi.org/10.1103/PhysRevX.13.011048)
 477 Fairhurst, S., Green, R., Hoy, C., Hannam, M., & Muir, A.
 478 2020, PhRvD, 102, 024055,
 479 doi: [10.1103/PhysRevD.102.024055](https://doi.org/10.1103/PhysRevD.102.024055)
 480 Hannam, M., Hoy, C., Thompson, J. E., et al. 2022,
 481 Nature, 610, 652, doi: [10.1038/s41586-022-05212-z](https://doi.org/10.1038/s41586-022-05212-z)
 482 Mandel, I., & Farmer, A. 2022, PhR, 955, 1,
 483 doi: [10.1016/j.physrep.2022.01.003](https://doi.org/10.1016/j.physrep.2022.01.003)

484 Mehta, A. K., Olsen, S., Wadekar, D., et al. 2023, arXiv
 485 e-prints, arXiv:2311.06061,
 486 doi: [10.48550/arXiv.2311.06061](https://doi.org/10.48550/arXiv.2311.06061)
 487 Neyman, J., & Pearson, E. S. 1933, Philosophical
 488 Transactions of the Royal Society of London Series A,
 489 231, 289, doi: [10.1098/rsta.1933.0009](https://doi.org/10.1098/rsta.1933.0009)
 490 Nitz, A. H., Kumar, S., Wang, Y.-F., et al. 2023, ApJ, 946,
 491 59, doi: [10.3847/1538-4357/aca591](https://doi.org/10.3847/1538-4357/aca591)
 492 Payne, E., Hourihane, S., Golomb, J., et al. 2022, PhRvD,
 493 106, 104017, doi: [10.1103/PhysRevD.106.104017](https://doi.org/10.1103/PhysRevD.106.104017)
 494 Roulet, J., Venumadhav, T., Zackay, B., Dai, L., &
 495 Zaldarriaga, M. 2020, PhRvD, 102, 123022,
 496 doi: [10.1103/PhysRevD.102.123022](https://doi.org/10.1103/PhysRevD.102.123022)
 497 Schmidt, P., Ohme, F., & Hannam, M. 2015, PhRvD, 91,
 498 024043, doi: [10.1103/PhysRevD.91.024043](https://doi.org/10.1103/PhysRevD.91.024043)