Constraining Astrophysical Parameter Bias due to Glitches in Third-Generation Gravitational-Wave Detectors

Interim Report #2

Samantha Walker
Barnard College, Columbia University, New York, NY 10027, USA

Mentors: Sophie Bini and Jake Summers
LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

(Dated: August 1, 2026)

Transient detector noise or "glitches" that coincide in time and frequency with gravitational-wave signals can bias the inferred astrophysical parameters of compact binary mergers, particularly black hole spin measurements. We investigate how this bias evolves as detector sensitivity increases by simulating binary black hole signals contaminated by glitches under projected sensitivities for LIGO A♯ and Cosmic Explorer. We consider two limiting scenarios in which the glitch amplitude or glitch signal-to-noise ratio (SNR) is held fixed relative to present-day detectors and perform Bayesian parameter estimation on each simulated data set. We find that glitches with fixed SNR produce progressively smaller parameter biases as the gravitational-wave signal SNR increases, whereas glitches with fixed amplitude continue to induce biases of comparable magnitude despite substantially narrower posterior distributions. These results suggest that the impact of glitches in future gravitational-wave detectors depends critically on how glitch populations scale with detector sensitivity and provide insight into when computationally expensive glitch-mitigation methods remain necessary.

I. INTRODUCTION

Gravitational waves (GWs) are traveling-wave perturbations of spacetime generated by an accelerating mass-energy quadrupole moment. Since 2015, the Laser Interferometer Gravitational-Wave Observatory (LIGO) has detected over 200 gravitational-wave events from binary black hole (BBH) mergers and binary neutron star (BNS) mergers to date, with many more observations still to come from the most recent observing run [1].

Using GW waveform models and Bayesian inference, the properties of the source can be inferred from the GW signal, a process known as parameter estimation (PE). For BBH mergers, these parameters include the masses of the black holes (BHs), their spins, the sky position, and the distance to the source, among other parameters. Accurate parameter estimation is crucial to characterizing GW sources, which enables scientific advances in gravitational-wave physics, tests of general relativity, and multi-messenger astronomy [2]. Measuring the BH spins is of particular interest, as they provide insight into the origin and evolution of the binary system, which may help address open questions regarding binary stellar evolution and the existence of BBH subpopulations [3–5].

The broad characteristics of the BH spins are encoded in two parameters, $\chi_{\rm eff}$ and $\chi_{\rm p}$, which describe the BH spin aligned with the orbital angular momentum and the BH spin parallel to the orbital plane, respectively. Previous observations (GW191109_010717 and GW200129_065458) have exhibited spins not aligned with the angular momentum of the orbital system, or spins in the plane of the binary [6, 7]. This suggests that the systems could have been formed through dynamical interactions rather than through isolated binary evolution, where the progenitor stars are born and evolve together as a binary system [5].

However, these conclusions may be biased by transient noise, known as "glitches." Typically, the process of parameter estimation assumes as a baseline a transient signal and stationary Gaussian noise [8]. If a glitch is suspected, one method is to model it as a separate component of the data and subtract the best fitting estimate before proceeding with analysis [9]. A more robust approach is joint compact binary coalescence (CBC)-glitch inference, in which the properties of the CBC signal and the glitch are inferred simultaneously. This allows the analysis to marginalize over multiple possible glitch realizations rather than relying on a single best-fit subtraction [10]. Despite these mitigation strategies, Udall *et al.* [11] showed that glitches can still significantly bias spin parameter estimates, even for glitches with signal-to-noise ratio (SNR) below the threshold of detectability.

This paper focuses on how the issue of glitch mitigation and parameter bias will evolve with third-generation detectors. Upcoming detectors such as Cosmic Explorer (CE) will be able to detect an unprecedented number of gravitational-wave signals [12]. CE is expected to observe 10^5 BBH mergers (one every 5 minutes) and 3×10^5 BNS mergers (one every 100 seconds) per year of observation [13].

In the transition toward these next-generation observatories, planned sensitivity upgrades to current LIGO detectors will provide an intermediate step toward improved detector performance. LIGO A♯ is the final planned upgrade to the current LIGO detectors and is expected to maximize their capabilities [14].

Currently, parameter estimation for a single source can take days, or longer if a glitch overlaps with the sig-

nal (which occurs in about 10% of cases). Moreover, the process of joint BBH-glitch inference is not fully automated and must be manually assessed at each step. With the increased detection rates of third-generation detectors, it will no longer be practical to treat every glitch-contaminated event with the same level of manual intervention. It is therefore crucial to understand how glitches scale with detector sensitivity in order to identify when full joint BBH-glitch inference will be necessary, and to develop scalable glitch mitigation pipelines to ensure accurate parameter estimates across a high volume of events.

We aim to simulate increased sensitivity of GW detectors in order to extend past studies of glitch-induced parameter bias such as [11]. In doing so, we hope to clarify whether certain glitch scenarios are expected to continue to pose a significant problem in producing accurate parameter estimates with third-generation detectors.

II. BACKGROUND

A. Expressing Detector Sensitivity

The sensitivity of a GW detector is reflected in its power spectral density (PSD) or amplitude spectral density (ASD) function, where the ASD is the square root of the PSD. The PSD and ASD are functions of frequency and are measured in units of Hz^{-1} and $\mathrm{Hz}^{-1/2}$ respectively, with lower values indicating greater sensitivity. Figure 1 illustrates the ASD for current LIGO detectors and the projected ASD for LIGO A♯ detectors and CE. In order to simulate increased sensitivity of the detector, we vary the PSD throughout our analysis.

The PSD also defines the relationship between "colored" and "whitened" strain data. Colored strain refers to the detector strain before whitening, in which the frequency content is shaped by the detector's frequency-dependent noise. Whitening divides out this frequency dependence through the relation

$$\tilde{h}_{\mathrm{white}}(f) = \frac{\tilde{h}(f)}{\sqrt{S_n(f)}}, \tag{1}$$

where $\tilde{h}(f)$ is the Fourier transform of the colored strain and $S_n(f)$ is the PSD. This suppresses frequency components in regions of high detector noise and emphasizes those in regions of low detector noise.

B. Quantifying Parameter Bias

The GW strain amplitude is fixed by the physics of the source, but with increased CE detector sensitivity, we expect the noise to decrease, thus increasing the signal SNR. However, it is unclear how the glitches will be impacted. In this project, we will consider two limiting cases: 1) the glitch amplitude remains constant and 2) the glitch SNR remains constant, relative to glitches observed in current detectors.

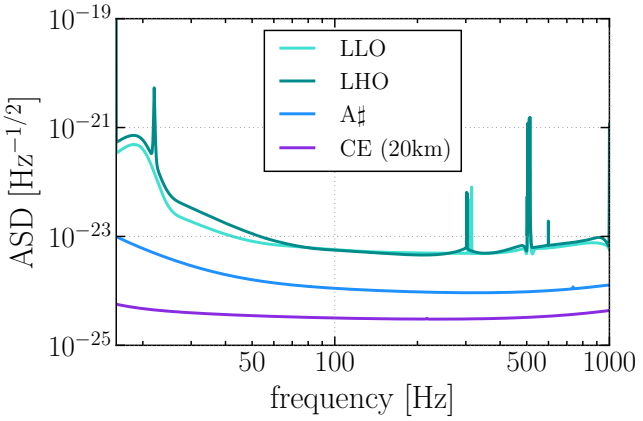


FIG. 1. Amplitude spectral density (ASD) as a function of frequency for the LIGO Livingston Observatory (LLO) and LIGO Hanford Observatory (LHO) during the third observing run (O3), shown in green and teal respectively, along with the projected sensitivities of LIGO A♯ (blue) and the 20 km Cosmic Explorer (CE) detector (purple).

To quantify the parameter bias in these two cases, we consider two different measures, the first being the raw or absolute bias, which is defined as

$$\delta\theta^{\alpha} \equiv \hat{\theta}^{\alpha} - \theta_0^{\alpha}, \tag{2}$$

where $\hat{\theta}^{\alpha}$ are the maximum-likelihood parameters and θ_0^{α} are the injection parameters. The second is the standardized bias, which is defined as $\delta\theta^{\alpha}/\sigma_{\alpha}$ where σ_{α} is the standard deviation of the parameter θ^{α}.

Based on an approximate Fisher matrix scaling argument outlined in Udall *et al.* [11], we expect the two glitch scaling cases to affect the bias differently. Our simulations aim to test whether these expectations hold. Consider an increase in detector sensitivity which scales the BBH signal SNR from $\rho \to \kappa\rho$. In the case of fixed glitch SNR, the raw bias is expected to scale approximately as

$$\delta\theta^{\alpha} \propto 1/\kappa. \tag{3}$$

As the standard deviation follows the same scaling, the standardized bias then remains constant with increasing SNR.

Conversely, for the fixed glitch amplitude case, the raw bias is predicted to remain approximately constant, leaving the standardized bias to scale accordingly with the standard deviation:

$$\delta\theta^{\alpha}/\sigma_{\alpha} \propto \kappa. \tag{4}$$

Thus, in the fixed-SNR case, absolute bias is expected to decrease, while the statistical deviation of the maximum likelihood parameters from truth will remain the same. In the fixed-amplitude case, the absolute bias is

expected to remain the same, while the statistical deviation from truth will increase. Note that these scalings assume the validity of the linear-signal approximation, which generally applies in the high-SNR regime [11, 15–17]. Our simulations use a variety of signal SNR values to test whether these expectations hold.

C. Modeling the Glitch

Glitches will be modeled as a single Morlet-Gabor wavelet, given by the expression

$$\Psi(t) = Ae^{-(t-t_c)^2/\tau^2}\cos(2\pi f_c(t-t_c)+\phi_0), \quad (5)$$

where $\tau = Q/(2\pi f_c)$ and A is the amplitude, Q is the quality factor, t_c and f_c are the central time and frequency, and ϕ_0 is the phase offset [18].

For the fixed amplitude case, the glitch form will be modeled identically between current and third-generation detectors. For the fixed SNR case, the amplitude of the glitch must be scaled accordingly. Here, we derive the amplitude scaling relation used to fix the SNR of a glitch.

The SNR of an arbitrary GW strain time series $h(t)$ is defined as

$$\rho = \left[4\int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)}\,df\right]^{1/2}, \quad (6)$$

where $\tilde{h}(f)$ is the Fourier transform of $h(t)$, and $S_n(f)$ is the PSD of the detector. The Fourier transform of Equation (5) is given by

$$\tilde{\Psi}(f) = \frac{\sqrt{\pi}A\tau}{2}e^{-2\pi i f t_c} \times \left(e^{i\phi_0-\pi^2\tau^2(f-f_c)^2} + e^{-i\phi_0-\pi^2\tau^2(f+f_c)^2}\right). \quad (7)$$

Assuming a narrow wavelet relative to changes in the PSD, we can take $S_n(f) \approx S_n(f_c)$ over the width of the glitch. Neglecting exponentially suppressed terms, Equation (6) becomes

$$\rho^2 \approx \frac{\pi A^2\tau^2}{S_n(f_c)}\int_0^\infty e^{-2\pi^2\tau^2(f-f_c)^2}\,df. \quad (8)$$

Evaluating the Gaussian integral and simplifying yields the SNR approximation for a Morlet-Gabor wavelet:

$$\boxed{\rho \approx \frac{A\sqrt{Q}}{\sqrt{2\sqrt{2\pi}f_c S_n(f_c)}}.} \quad (9)$$

Thus, to maintain a constant glitch SNR, we rescale the glitch amplitude following

$$A' = A\sqrt{\frac{S_n'(f_c)}{S_n(f_c)}}, \quad (10)$$

Label	Detector configuration	Network SNR
O3*	LIGO O3	25.17
A100	LIGO A♯ (scaled PSD)	103.96
A♯*	LIGO A♯	138.61
A170	LIGO A♯ (scaled PSD)	173.28
CE*	CE	671.38

TABLE I. Detector configurations used in this work and the respective SNR of injection event GW200129. Configurations labeled "scaled PSD" are generated by uniformly scaling the corresponding detector PSD while preserving its spectral shape. Labels marked * indicate runs with the original detector PSD. Note that the scaled PSD runs are only currently complete in the fixed-SNR case.

where primed quantities denote the modified glitch amplitude and PSD, while all other parameters remain fixed.

III. METHODS

A. Data Generation

We consider three main detector configurations, LIGO O3, LIGO A♯, and CE, whose ASDs are pictured in Figure 1. As the future site of CE has not yet been chosen, we assume the location of two 20km CE detectors to be at the present sites of LIGO Livingston Observatory (LLO) and LIGO Hanford Observatory (LHO). Thus, all configurations are identical in location and orientation, with only their sensitivity varied.

The network SNR of the injected GW200129 signal is approximately 25 with the O3 PSD, 140 with the A♯ PSD, and 670 with the CE PSD. To study a range of signal SNRs within each detector regime, we uniformly scale the detector PSD while preserving its spectral shape. This is equivalent to rescaling the source distance, but avoids modifying the injected luminosity distance and the associated cosmological redshift effects.

Four full sets of data frames with scaled PSDs were generated in addition to the three original detector configurations: two created with a scaled A♯ PSD, and two created with a scaled CE PSD, with the baseline PSD chosen based on the SNR regime. Each of the detector configuration cases and their corresponding GW200129 network SNR are summarized in Table I.

We use Bilby [19] to simulate stationary Gaussian noise based on the PSD of the chosen detector configuration. The O3 configuration uses the PSD associated with GW200129 at the respective detector locations, while all other configurations only have projected sensitivity information, and thus assume the same PSD for both detector locations. Independent Gaussian noise frames are generated for each of the detector locations.

For each configuration, we produce 4 seconds of data with a sampling frequency of $f_{\rm samp} = 1024$ Hz, minimum

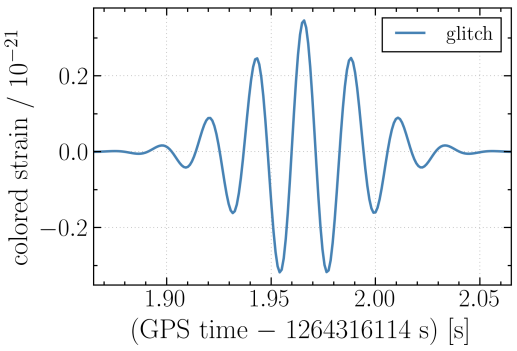


FIG. 2. Morlet-Gabor wavelet glitch, as injected into the colored strain data. The analytical form of the glitch is given by Equation (5).

	Fixed amplitude		Fixed SNR	
Detector	Amplitude	SNR	Amplitude	SNR
LIGO O3	3.47	7.96	3.47	7.96
LIGO A♯	3.47	42.24	0.65	7.95
CE	3.47	212.84	0.13	7.95

TABLE II. Glitch amplitude and SNR across detector configurations under two scenarios: amplitude held fixed at the LIGO O3 value, versus SNR held fixed at the LIGO O3 value. Amplitude in units of $\times 10^{-22}$. SNR values for the fixed-SNR case shift slightly due to the SNR approximation (9) used in amplitude scaling.

frequency of $f_{\min} = 20$ Hz, and maximum frequency of $f_{\max} = 448$ Hz. Then, we inject a CBC signal using the waveform model NRSur7dq4 [20] with injection parameters from the BBH merger event GW200129 [6].

For the detector located at the LLO site, we save a copy of the uncontaminated data before injecting a glitch, modeled as a single Morlet-Gabor wavelet, pictured in Figure 2. The glitch amplitude and SNR for the LIGO O3 configuration are preset to $A = 3.467 \times 10^{-22}$ and $\rho \approx 7.96$. For the other configurations, these parameters vary according to whether we are modeling fixed glitch amplitude or fixed glitch SNR. In the case of fixed glitch SNR, the amplitude is scaled with Equation (10), otherwise, the glitch form is not modified and the SNR scales accordingly as the detector noise decreases. The full set of amplitude and SNR values for each scaling case and configuration are listed in Table II. The other glitch parameters are kept constant between detector configurations and are listed in Table III. The glitch parameters were chosen to maximize the glitch bias in $\chi_{\rm p}$, as found in Udall *et al.* [11]

Parameter	Value
t_c (s)[a]	1.97
f_c (Hz)	43.56
Q	10.71
ϕ_0 (rad)	0.0

TABLE III. Parameters of the glitch which is injected into the detector located at the LLO site in all detector configurations. The glitch amplitude and SNR are not included as they vary between configurations, and are instead listed in Table II.

[a] Relative to event start time $t_0 = 1264316114$ s.

B. Parameter Estimation

We then perform parameter estimation on the generated data frames with Bilby [19] using the Dynesty sampler [21] and the default LIGO-Virgo-KAGRA priors (uniform in masses, spin magnitudes, merger time, and coalescence phase, isotropic in spin and binary orientations and sky location) [22]. NRSur7dq4 is the waveform model used for inference, matching what is used for data generation.

Probability distributions for the values of each BBH parameter, known as "posteriors," are produced from these PE runs. Due to increased detector sensitivity, we expect the width of the posteriors to decrease, scaling with $1/\rho$ where ρ is the SNR of the signal, or equivalently scaling with the square root of the detector PSD for a fixed signal, following the linear signal approximation [16].

For each of the detector configurations we perform a PE run with the uncontaminated data as a control, and then another run on the same data with the addition of the glitch. For A♯ and CE configurations, two such glitch runs are performed, one for each glitch scaling case. The PE run structure is summarized in Figure 3. Note that due to time constraints, the fixed-amplitude glitch runs for A♯ and CE configurations were conducted on data with different Gaussian noise realizations than the other runs for the same detector. However, due to the high SNR of both the GW event and the glitch in these cases, the effect is considered negligible.

IV. PE RESULTS

To validate the data generation and parameter estimation pipeline, we first examine the no-glitch runs. As the glitch parameters were chosen to maximize bias in $\chi_{\rm p}$, we will focus on that parameter. Figure 4 shows the $\chi_{\rm p}$ posterior for each of the three principal runs. Each run recovers the truth value within the 90% confidence interval. The posteriors also exhibit the expected narrowing due to increased detector sensitivity. As mentioned in

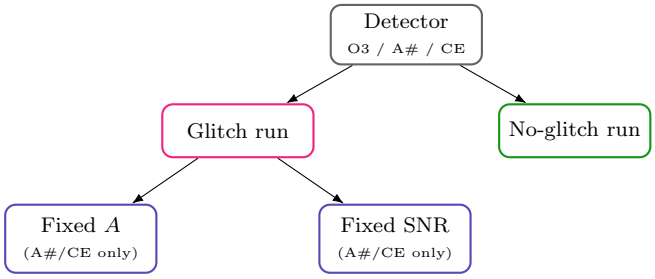


FIG. 3. Structure of the sets of PE runs for each detector configuration listed in Table I. For each, we perform a glitch-contaminated and a clean baseline run. Glitch runs are further split into fixed-amplitude and fixed-SNR injection scenarios; the resulting bias-scaling comparison is restricted to A# and CE data, with O3 serving as the fixed anchor point.

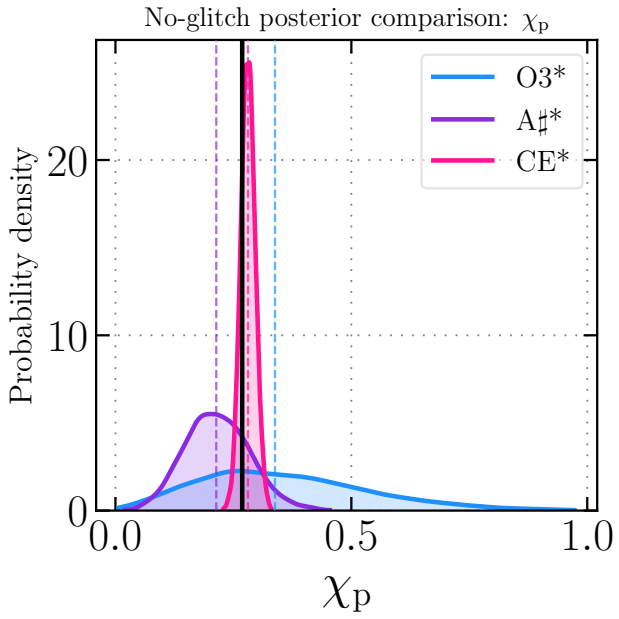


FIG. 4. $\chi_{\rm p}$ posteriors for the three principal detector configurations: O3*, and A♯* and CE*. No glitch was present in these data sets. The injected truth value is drawn in black.

Section II B, we expect the standard deviation σ_α of a parameter θ^α to follow the scaling

$$\sigma_\alpha \propto 1/\rho \tag{11}$$

where ρ is the network SNR of the signal. In Figure 5, we plot the posterior standard deviations for $\chi_{\rm p}$ for all no-glitch runs against the signal's network SNR and fit the expected relationship. The points generally agree with the expected relationship, although further data points will be necessary to confirm and quantify this agreement.

We now turn to testing the bias predictions outlined in Section II B. The only PE results available for the fixed-amplitude glitch case at the time of writing are the O3*, A♯*, and CE* configurations. The data from these runs is

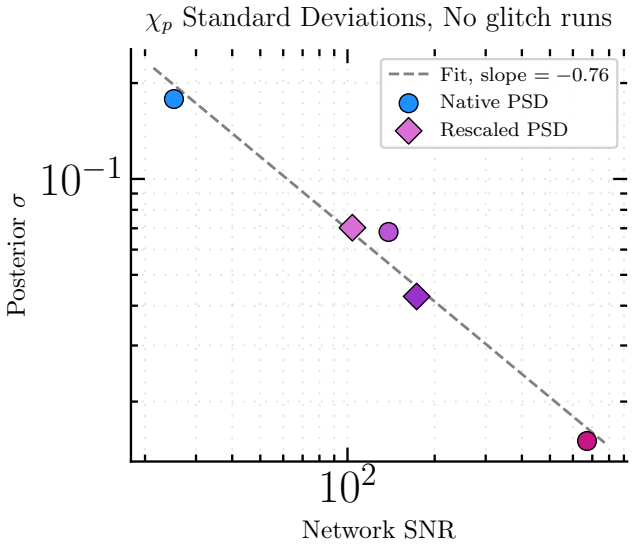


FIG. 5. Standard deviations of the $\chi_{\rm p}$ posterior against the network SNR of the signal for the detector configuration. Diamonds indicate runs with a scaled detector PSD, while circles indicate runs with the original detector PSD. Similar to Figure 4, the O3 run is shown in blue, the A♯ runs are shown in purple, and the CE runs are shown in pink. The expected relation between posterior widths and SNR given by Equation (11) is drawn as a gray dashed line.

shown in Figure 6. The expected relationships following Equation (4) are fit to the data points and the resulting line of best fit is drawn in gray dashes.

The data for the fixed-SNR glitch case includes all runs referenced in Table I. These data are similarly plotted in Figure 7.

V. CONCLUSIONS AND FUTURE WORK

These initial results allow us to visualize the general trends in the data in comparison to the expectations. Most notably, we observe a significant decrease in bias in the fixed glitch SNR case as detector sensitivity increases. This suggests that glitch mitigation may be unnecessary in the presence of low-SNR glitches for future detectors.

In the coming weeks, we will be adding error bars as well as additional data points to better assess trends. At the moment, it seems that in the fixed-SNR case, the absolute bias roughly scales with expectations, but the standardized bias generally decreases, contradicting the prediction that it would remain constant. As the standardized bias is defined as the raw bias $\delta\theta^\alpha$ divided by the posterior standard deviation σ_α, this discrepancy may either be caused by deviations from the predicted behavior in either $\delta\theta^\alpha$ or σ_α. Although we checked the behavior of σ_α for the no-glitch runs in Figure 5, some of our data suggests deviations from this behavior in the other runs. We will continue to investigate this matter

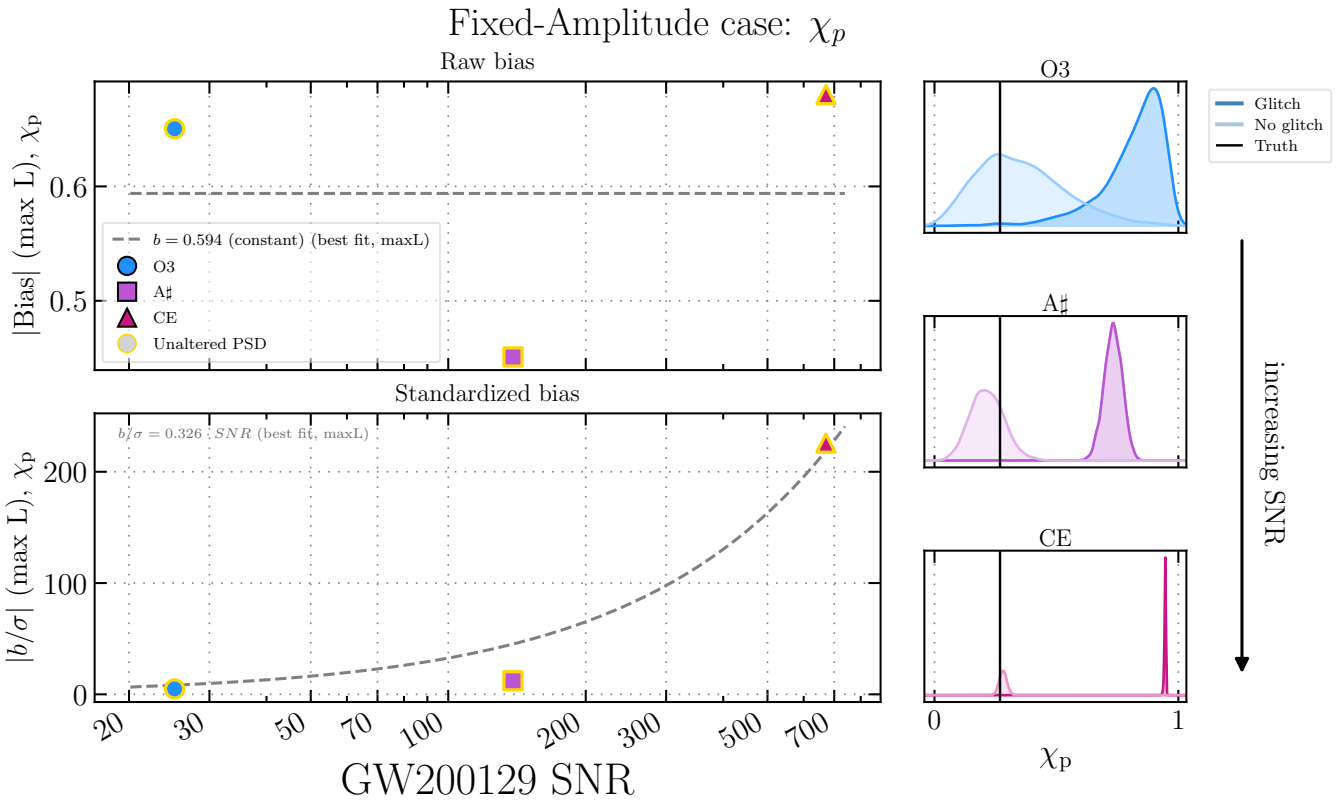


FIG. 6. Bias trends and χ_p posteriors for various detector configurations in the fixed-amplitude glitch scaling case. The top left panel shows the raw biases for the fixed-amplitude glitch runs for the three available configurations. The expected trend, fit to the data points, is drawn as a gray dashed line. Similarly, the bottom left panel shows the same points but in terms of standardized bias. The corresponding expected trendline is fit to the data points with a dashed gray line. The right panels display the posteriors for each of the runs plotted left, with the no-glitch runs in lighter shading and the truth value drawn as a black vertical line.

in the coming weeks.

Next, we plan to rerun the analysis with another injection event, GW191109, with a focus instead on bias in $\chi_{\rm eff}$. We will compare the results between the two cases and observe any differences in trends which occur. In both of these cases, we will measure changes in the direction of bias in parameter space, instead of limiting our analysis to only one parameter.

In addition to this, we are interested in exploring several variations in our analysis to see how it affects our results. The bias is defined in Equation (2) is given in terms of the maximum likelihood parameters to match the derived Fisher matrix prediction; however, in practice the parameters are typically reported at the median of the posterior. Accordingly, this may be a more practical definition of the "recovered value" if our focus is the impact on reported parameters. In the high-SNR regime, we would expect the median and maximum likelihood values to coincide, but our initial findings show that this is not always the case.

As future detectors become more sensitive in current frequency ranges, they will also lower the minimum frequency at which GWs can be detected. In the coming weeks, we will simulate this effect by lowering the minimum frequency in our data generation/parameter estimation pipeline and observing how the addition of lower frequencies in the signal affects parameter bias. An initial test performed with the A♯ PSD pictured in Figure 8 demonstrates a significant decrease in glitch bias in χ_p with a lowered frequency. More tests are needed to confirm this finding. We also plan on investigating the differences in waveform with reduced frequencies in order to attribute additional waveform features to improvements in glitch-induced bias.

Returning to the motivation of our project, we also wish to explore population models for BBHs in future detectors. If we determine a conditional threshold at which glitch mitigation is no longer necessary, these population models would be useful in determining exactly how many future events may not warrant any mitigation even in the presence of certain types of glitches.

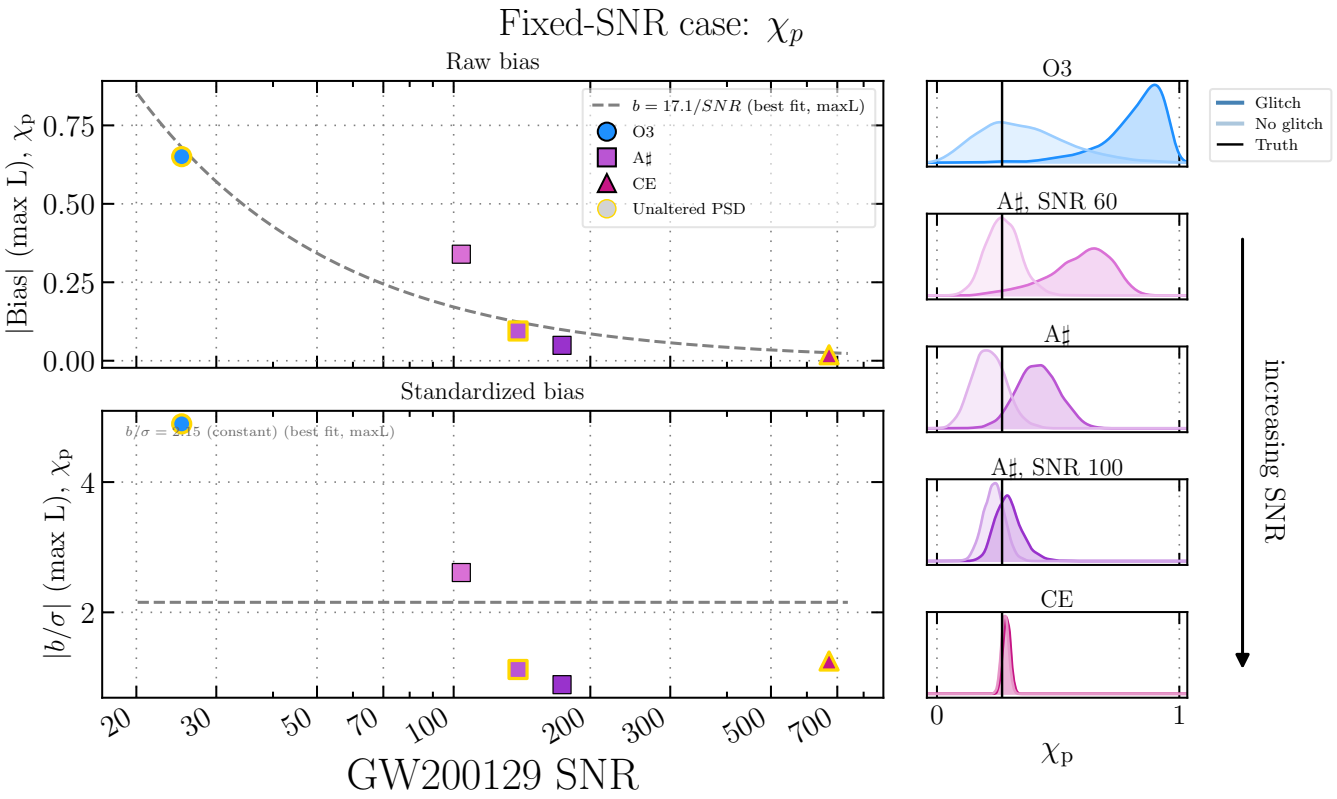


FIG. 7. Bias trends and χ_p posteriors for various detector configurations in the fixed-SNR glitch scaling case. The top left panel shows the raw biases for the fixed-SNR glitch runs for all detector configurations. The expected trend, fit to the data points, is drawn as a gray dashed line. Similarly, the bottom left panel shows the same points but in terms of standardized bias. The corresponding expected trendline is fit to the data points with a dashed gray line. For both plots, data points highlighted in yellow indicate runs conducted with the original detector PSD. The right panels display the posteriors for each of the runs plotted left, with the no-glitch runs in lighter shading and the truth value drawn as a black vertical line.

[1] A. G. Abac *et al.* (LIGO Scientific, VIRGO, KAGRA), GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run, (2025), arXiv:2508.18082 [gr-qc].

[2] N. Christensen and R. Meyer, Parameter estimation with gravitational waves, Rev. Mod. Phys. **94**, 025001 (2022).

[3] I. Mandel and A. Farmer, Merging stellar-mass binary black holes, Phys. Rept. **955**, 1 (2022).

[4] A. Ray, S. Mukherjee, M. Zevin, and V. Kalogera, On the Astrophysical Origin of Binary Black Hole Subpopulations: A Tale of Three Channels?, (2026), arXiv:2603.17987 [astro-ph.HE].

[5] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, B. Farr, and A. Vecchio, Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves, Nature **548**, 426 (2017).

[6] R. Abbott *et al.* (KAGRA, VIRGO, LIGO Scientific), GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run, Phys. Rev. X **13**, 041039 (2023).

[7] T. Islam, A. Vajpeyi, F. H. Shaik, C.-J. Haster, V. Varma, S. E. Field, J. Lange, R. O'Shaughnessy, and R. Smith, Analysis of GWTC-3 with fully precessing numerical relativity surrogate models, Phys. Rev. D **112**, 044001 (2025).

[8] E. Thrane and C. Talbot, An introduction to Bayesian inference in gravitational-wave astronomy: parameter estimation, model selection, and hierarchical models, Publ. Astron. Soc. Austral. **36**, e010 (2019), [Erratum: Publ.Astron.Soc.Austral. 37, e036 (2020)].

[9] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton, Subtracting glitches from gravitational-wave detector data during the third LIGO-Virgo observing run, Class. Quant. Grav. **39**, 245013 (2022).

[10] S. Ghonge, J. Brandt, J. M. Sullivan, M. Millhouse, K. Chatziioannou, J. A. Clark, T. Littenberg, N. Cornish, S. Hourihane, and L. Cadonati, Assessing and mitigating the impact of glitches on gravitational-wave parameter estimation: A model agnostic approach, Phys. Rev. D **110**, 122002 (2024).

[11] R. Udall, S. Bini, K. Chatziioannou, D. Davis, S. Hourihane, Y. Lecoeuche, J. McIver, and S. Miller, Inferring

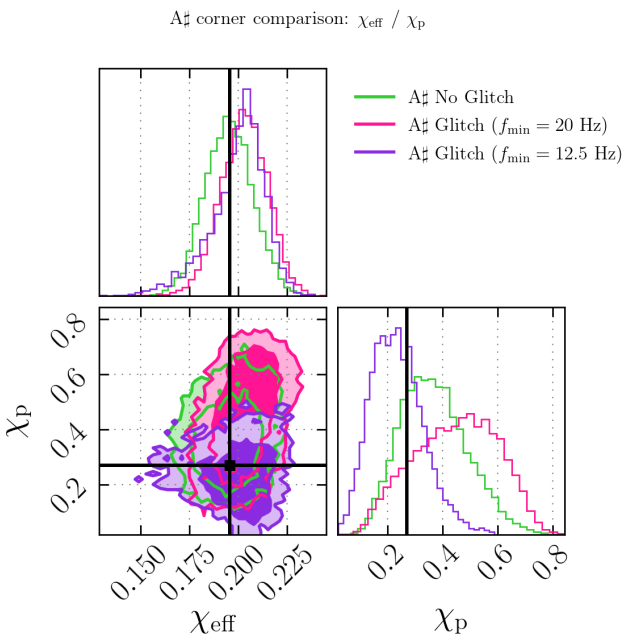


FIG. 8. Initial corner plot of the effective spin parameters, produced with posteriors from 3 A♯* runs. The green posteriors are runs conducted on data with no glitch present. The pink posteriors are the data with an injected glitch and a minimum detector frequency (in both data generation and parameter estimation) of 20 Hz. The purple posteriors are the same, but with a lower minimum frequency of 12.5 Hz.

the spins of merging black holes in the presence of data-quality issues, Phys. Rev. D **113**, 042005 (2026).

[12] D. Reitze *et al.*, Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO, Bull. Am. Astron. Soc. **51**, 035 (2019), arXiv:1907.04833 [astro-ph.IM].

[13] M. Evans *et al.*, A Horizon Study for Cosmic Explorer: Science, Observatories, and Community, (2021), arXiv:2109.09882 [astro-ph.IM].

[14] P. Fritschel, K. Kuns, J. Driggers, A. Effler, B. Lantz, D. Ottaway, S. Ballmer, K. Dooley, R. Adhikari, M. Evans, B. Farr, G. Gonzalez, P. Schmidt, and S. Raja, *Report of the LSC Post-O5 Study Group*, Tech. Rep. LIGO-T2200287 (LIGO Scientific Collaboration, 2022) v3, 2024.

[15] C. Cutler and M. Vallisneri, Lisa detections of massive black hole inspirals: Parameter extraction errors due to inaccurate template waveforms, Physical Review D **76**, 10.1103/physrevd.76.104018 (2007).

[16] C. Cutler and E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?, Physical Review D **49**, 2658–2697 (1994).

[17] M. Vallisneri, Use and abuse of the fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects, Physical Review D **77**, 10.1103/physrevd.77.042001 (2008).

[18] N. J. Cornish and T. B. Littenberg, BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, Class. Quant. Grav. **32**, 135012 (2015).

[19] G. Ashton *et al.*, BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, Astrophys. J. Suppl. **241**, 27 (2019).

[20] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations with unequal masses, Phys. Rev. Research. **1**, 033015 (2019).

[21] J. S. Speagle, dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences, Monthly Notices of the Royal Astronomical Society **493**, 3132–3158 (2020).

[22] A. G. Abac *et al.* (LIGO Scientific, VIRGO, KAGRA), GWTC-5.0: Observations from the Second Part of the Fourth LIGO-Virgo-KAGRA Observing Run and Updates to the Gravitational-Wave Transient Catalog, (2026), arXiv:2605.27225 [gr-qc].