

GW190521: Tracing imprints of spin-precession on the most massive black hole binary

Simona J. Miller^{1,2}, Maximiliano Isi³, Katerina Chatziioannou^{1,2}, Vijay Varma^{4,5} and Ilya Mandel^{6,7}

¹*Department of Physics, California Institute of Technology, Pasadena, California 91125, USA*

²*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*

³*Center for Computational Astrophysics, Flatiron Institute, New York, New York 10010, USA*

⁴*Department of Mathematics, Center for Scientific Computing and Data Science Research, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

⁵*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476 Potsdam, Germany*

⁶*School of Physics and Astronomy Monash University, Clayton, Victoria 3800, Australia*

⁷*OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton, Victoria 3800, Australia*



(Received 5 October 2023; accepted 12 December 2023; published 18 January 2024)

GW190521 is a remarkable gravitational-wave signal on multiple fronts: its source is the most massive black hole binary identified to date and could have spins misaligned with its orbit, leading to spin-induced precession; an astrophysically consequential property linked to the binary’s origin. However, due to its large mass, GW190521 was only observed during its final 3–4 cycles, making precession constraints puzzling and giving rise to alternative interpretations, such as eccentricity. Motivated by these complications, we trace the observational imprints of precession on GW190521 by dissecting the data with a novel time domain technique, allowing us to explore the morphology and interplay of the few observed cycles. We find that precession inference hinges on a quiet portion of the pre-merger data that is suppressed relative to the merger ringdown. Neither premerger nor postmerger data alone are the sole driver of inference, but rather their combination; in the quasicircular scenario, precession emerges as a mechanism to accommodate the lack of a stronger premerger signal in light of the observed postmerger. In terms of source dynamics, the premerger suppression arises from a tilting of the binary with respect to the observer. Establishing such a consistent picture between the source dynamics and the observed data is crucial for characterizing the growing number of massive binary observations and bolstering the robustness of ensuing astrophysical claims.

DOI: [10.1103/PhysRevD.109.024024](https://doi.org/10.1103/PhysRevD.109.024024)

I. INTRODUCTION

At a total mass of $\sim 150M_{\odot}$, GW190521 [1,2] is the current record holder among massive black hole binaries confidently detected through gravitational waves by LIGO [3] and Virgo [4]. Such high-mass systems are essential probes of the role of hierarchical mergers [5–9] and pair-instability physics [10–15] in binary formation and evolution. Observationally, massive binaries merge toward the low edge of the detectors’ bandwidth and are only detectable for a short time. One third of the binaries in the latest gravitational-wave catalog have a median detector-frame total mass $> 100M_{\odot}$ [16], corresponding to ~ 5 signal cycles at more than half a standard deviation above the noise. The short duration makes characterizing these signals and inferring astrophysical properties such as spin challenging.

Spin is a key signature of the physics behind angular momentum transport in stellar interiors, black-hole formation, black-hole retention in dense environments, and more [17–22]. Gravitational-waves provide one of the few ways to measure spins for stellar-mass black holes directly.

Spin components parallel to the binary’s orbital angular momentum affect the signal duration [23] and are approximately conserved during the inspiral as the “effective spin” [24,25]. Spin components perpendicular to the orbital angular momentum, i.e., in the orbital plane, cause the binary to precess, leading to signal modulations as the emission pattern varies relative to the line of sight [26,27]. Although typically weak [28–30], this effect is highly sought-after: spin-induced precession and the associated in-plane spins could differentiate between dynamical and field binary formation, e.g., [17,31,32].

The elusiveness of precession is exacerbated for heavy systems. The precession timescale can be longer than the observed inspiral for large masses [33], making modulations difficult to identify. The exact imprint of precession on the ensuing merger and ringdown remains poorly understood and analytically intractable, although numerical-relativity and data-analysis studies suggest that imprints do exist [34–39]; for example, Ref. [35] suggests, based on simulations, that high-frequency data typically associated

with the merger-ringdown can constrain precession. This uncertainty makes it difficult to distinguish precession from eccentricity, another highly-valuable binary property [40–42]. Interpretation is further complicated by the high sensitivity of the measurability of precession to the system’s true parameters [34,35,43,44] and the priors [45].

In light of this, the massive system that sourced GW190521 stands out for its informative precession constraint, when analyzed assuming a quasicircular orbit. Precession can be quantified by the effective precessing spin χ_p [46–48] that is motivated by inspiral dynamics. A value of zero (one) indicates no (maximal) precession. Under the assumption of a quasicircular orbit, GW190521 has $\chi_p = 0.68^{+0.25}_{-0.37}$ at 90% credibility [1,2], the largest inferred χ_p and the one whose posterior is the most informative to date [16,49,50]¹; similar conclusions are reached under alternative parametrizations for precession [54,55].² The combined high mass and large in-plane spin make GW190521 an essential probe of hierarchical black hole mergers [7,56,57], dense stellar environments such as nuclear star clusters [19,58,59], active galactic nuclei disks [60–62], and more [2,15,63–67].

The high mass of GW190521 and its few observable cycles open the door to competing astrophysical interpretations. Romero-Shaw *et al.* [40] and Gayathri *et al.* [68] find that the data are consistent with eccentricity, though this interpretation is not supported by Iglesias *et al.* [69] and Ramos-Buades *et al.* [70]. Gamba *et al.* [71] propose a hyperbolic capture scenario. Nitz and Capano [67] suggest a highly asymmetric, but still precessing, binary interpretation. More exotic explanations include boson stars [72] and cosmic strings [73]. Any of these alternatives would have important implications if confirmed [74,75]. Additionally, random detector noise can have an outsized impact on the inference of poorly constrained effects, although Biscoveanu *et al.* [35] and Xu and Hamilton [76] show that the inference of χ_p away from zero in GW190521-like systems cannot be due to Gaussian noise alone. The fact that full-scale parameter estimation allows for competing interpretations suggests that different physical effects can result in similar observational imprints over GW190521’s few cycles. Similarly to precession and eccentricity, these imprints are often not analytically tractable.

Toward bolstering the interpretation of massive binaries, it is essential to gain intuition about the observable imprint of physical effects of interest and how their measurability is affected by mismodeling. Lacking analytical equations

¹GW200129, another massive binary [16], has a potentially comparable constraint. When averaged over waveform models, its χ_p is unconstrained [16]; when restricting to the NRSur7dq4 [51] model, GW200129 is inferred to be precessing [52]. This interpretation is however complicated by data quality issues [53].

²These and further ways to quantify precession are elaborated upon in Appendix C.

for precession in the merger phase, we introduce a novel approach that traces its imprint along the signal and identifies the role of each cycle on the χ_p constraint. We dissect the data in the time-domain and compare inference between different data subsets. We provide a cycle-by-cycle physical picture of source dynamics and explore the interplay of different data regions. Our work focuses on the data aspect that drives the inference of precession within a quasi-circular merger scenario. Extensions to further physical effects of interest such as orbital eccentricity can be tackled under a similar framework; we leave those to future work.

II. METHODS

Gravitational wave parameter estimation is typically conducted in the frequency domain for computational efficiency. Leveraging the stationarity of detector noise, noise components at different frequencies are independent which leads to a diagonal covariance matrix in likelihood calculations [77,78]. However, frequency-domain methods are nonlocal in time; thus, isolating temporal features of source dynamics and their imprint on the data requires nontrivial likelihood modifications [79,80].³ We instead adopt direct time-domain inference to isolate different signal cycles, an approach originally conceived for black hole ringdowns [84–87]. We truncate data from LIGO Livingston, LIGO Hanford, and Virgo at different times ranging from $t = -50 M_{\text{ref}}$ to $50 M_{\text{ref}}$ with respect to coalescence.⁴ We independently infer the signal properties solely from data before and after each cutoff as well as the full span of data.

We model the signal with the numerical relativity surrogate model NRSur7dq4 [51], which assumes quasicircular orbits and includes precession and higher-order modes. Within its region of validity, NRSur7dq4 displays the lowest mismatches against numerical relativity among existing models [51]. We adapt the time-domain inference code from Isi *et al.* [86] and sample the multidimensional posterior for the binary masses, spin magnitudes and tilt angles, azimuthal interspin angle, azimuthal precession cone angle, inclination, luminosity distance, and phase of coalescence. The time of coalescence, right ascension, declination, and polarization angle are fixed for computational efficiency.⁵

³Frequency truncation enables consistency checks [81,82], investigations of data-quality issues [53,83], or alternative studies of the measurability of precession in simulated data [35], but this is not equivalent to cuts in time.

⁴We define t with respect to geocenter GPS time 1242442967.405764 s. Under geometric units we adopt the median detector-frame remnant mass scale $M_{\text{ref}} = 1.27 \text{ ms}$ [2]; in standard units $M_{\text{ref}} = 258.3 M_{\odot}$. The choice of remnant rather than total mass, was inspired by ringdown analyses [88].

⁵We have verified that these choices do not affect our conclusions. All parameter estimation settings, priors, and consistency checks are given in Appendix A.

We report precession constraints using the canonical effective precessing spin, χ_p [46–48],

$$\chi_p = \max \left[\chi_1 \sin \theta_1, \left(\frac{3 + 4q}{4 + 3q} \right) q \chi_2 \sin \theta_2 \right] \in [0, 1]. \quad (1)$$

Here $\chi_i \in [0, 1)$ are the dimensionless spin magnitudes and θ_i are the tilt angles between the spin and orbital angular momentum vectors. Subscripts $i \in \{1, 2\}$ denote each black hole with mass m_i and $q \equiv m_2/m_1 \leq 1$.

III. RESULTS

Figure 1 shows the inferred GW190521 properties from data before (blue) and after (orange) five representative cutoff times (vertical lines) as well as the full signal (black) for comparison.⁶ Insets in the left panels visualize the truncation in LIGO Livingston, selected as the detector in which GW190521 is the loudest. The left column shows the posterior for χ_p at a reference frequency of 11 Hz [1, 2]. For the earliest cutoff time (top row), the post- $t = -40M_{\text{ref}}$ χ_p posterior is almost identical to that of the full analysis, while the precutoff one is identical to the prior. This is due to the fact that the postcutoff analysis includes the full available signal, while none of it is contained in the precutoff data (see inset). As the cutoff moves to later times (top to bottom), the precutoff and postcutoff posteriors gradually exchange places as the data preceding each cutoff become more informative and the data following become less so.

The ~ 40 ms between $t = -40M_{\text{ref}}$ and $-10M_{\text{ref}}$ are crucial to constrain precession for GW190521. This region roughly corresponds to the final cycle before the onset of merger. The χ_p posterior obtained from data after $t = -40M_{\text{ref}}$ (first row, orange) is consistent with that from the full analysis, i.e., precession is constrained.⁷ On the other hand, data after $t = -10M_{\text{ref}}$ (third row, orange), result in a χ_p posterior that is nearly identical to the prior, i.e., uninformative. Between these times, the posterior shifts smoothly between the full measurement and the prior; e.g., the post- $t = -20M_{\text{ref}}$ analysis (second row, orange). The reduction in the signal-to-noise ratio (SNR) from excluding data is negligible between $-40M_{\text{ref}}$ and $-10M_{\text{ref}}$ suggesting this qualitative change in precession inference is not due to an SNR drop (see Fig. 8 in Appendix D).

Neither the inspiral nor the merger/ringdown data alone are fully responsible for precession constraints in GW190521. The data both pre- and post- $t = -10M_{\text{ref}}$ alone are uninformative about precession (third row, orange

and blue). Moreover, the pre- $t = 30M_{\text{ref}}$ analysis that excludes the final ringdown cycle (fifth row, blue) is consistent with the full analysis. It is therefore not solely the final premerger cycle that informs precession, but rather its combination with the subsequent 2 merger and early-ringdown cycles. This does not rule out ringdown imprints of precession that are too weak to discern at this SNR or with this waveform.

The center and right columns of Fig. 1 investigate features of the waveforms. The blue and orange waveforms are informed only by data in the unshaded regions and extended coherently into the shaded regions. As progressively less data are analyzed (center bottom to top, right top to bottom), the waveform reconstructions agree less with the full-analysis waveform, eventually becoming incoherent. The right column reveals the morphological imprint of precession on the signal during the transition from an informative χ_p posterior (first row) to the prior (third row). When the χ_p inference returns (close to) the prior, the final premerger cycle is extrapolated to be larger than when χ_p is constrained to take higher values, cf., the waveform peak at $t \sim -30M_{\text{ref}}$ and trough at $t \sim -15M_{\text{ref}}$. Again there is a progression; the post- $t = -40M_{\text{ref}}$ inferred waveforms (orange) are consistent with the full analysis (black), while the final premerger cycle subtly increases in strength toward post- $t = -10M_{\text{ref}}$ (top row to third row).

To further explore the premerger waveform suppression, we compare the full analysis (black) in which precession is constrained and the post- $t = -10M_{\text{ref}}$ analysis where the data are uninformative about precession in Fig. 2. In order to focus on waveform features that are informative compared to the noise, we plot the whitened waveform.⁸ In the top panel the grayed-out region denotes data available to the full analysis but not the post- $t = -10M_{\text{ref}}$ one. The inset focuses around $t = -15M_{\text{ref}}$, an extremum of the final premerger cycle. The reconstructions are inconsistent at the 50% credible level, with the post- $t = -10M_{\text{ref}}$ analysis resulting in a larger amplitude (in absolute value). This inconsistency only occurs at the extrema of the final premerger cycle, i.e., the peak around $t \sim -30M_{\text{ref}}$ and trough at $t \sim -15M_{\text{ref}}$, see Fig. 9 in Appendix D.

The bottom panel shows marginal posteriors for select quantities: the effective precessing spin χ_p , the absolute value of the whitened strain $|\hat{h}|$ (units of standard deviations σ of the noise), and the difference between ι —the angle between the direction of maximum signal emission and the line of sight [91]—and $\pi/2$ at $t = -15M_{\text{ref}}$. As expected from Fig. 1, χ_p and $|\hat{h}|$ are anticorrelated (albeit weakly); when the full data are analyzed (black), the final cycle is constrained to be weaker, resulting in a larger χ_p and when

⁶Results for further cutoff times are included in our accompanying Github repository [89].

⁷By “constrained,” we specifically mean “visibly different from the prior.”

⁸Whitened waveforms are obtained by dividing the Fourier-domain waveform by the noise amplitude spectral density and then inverse Fourier-transforming. See Appendix B for details.

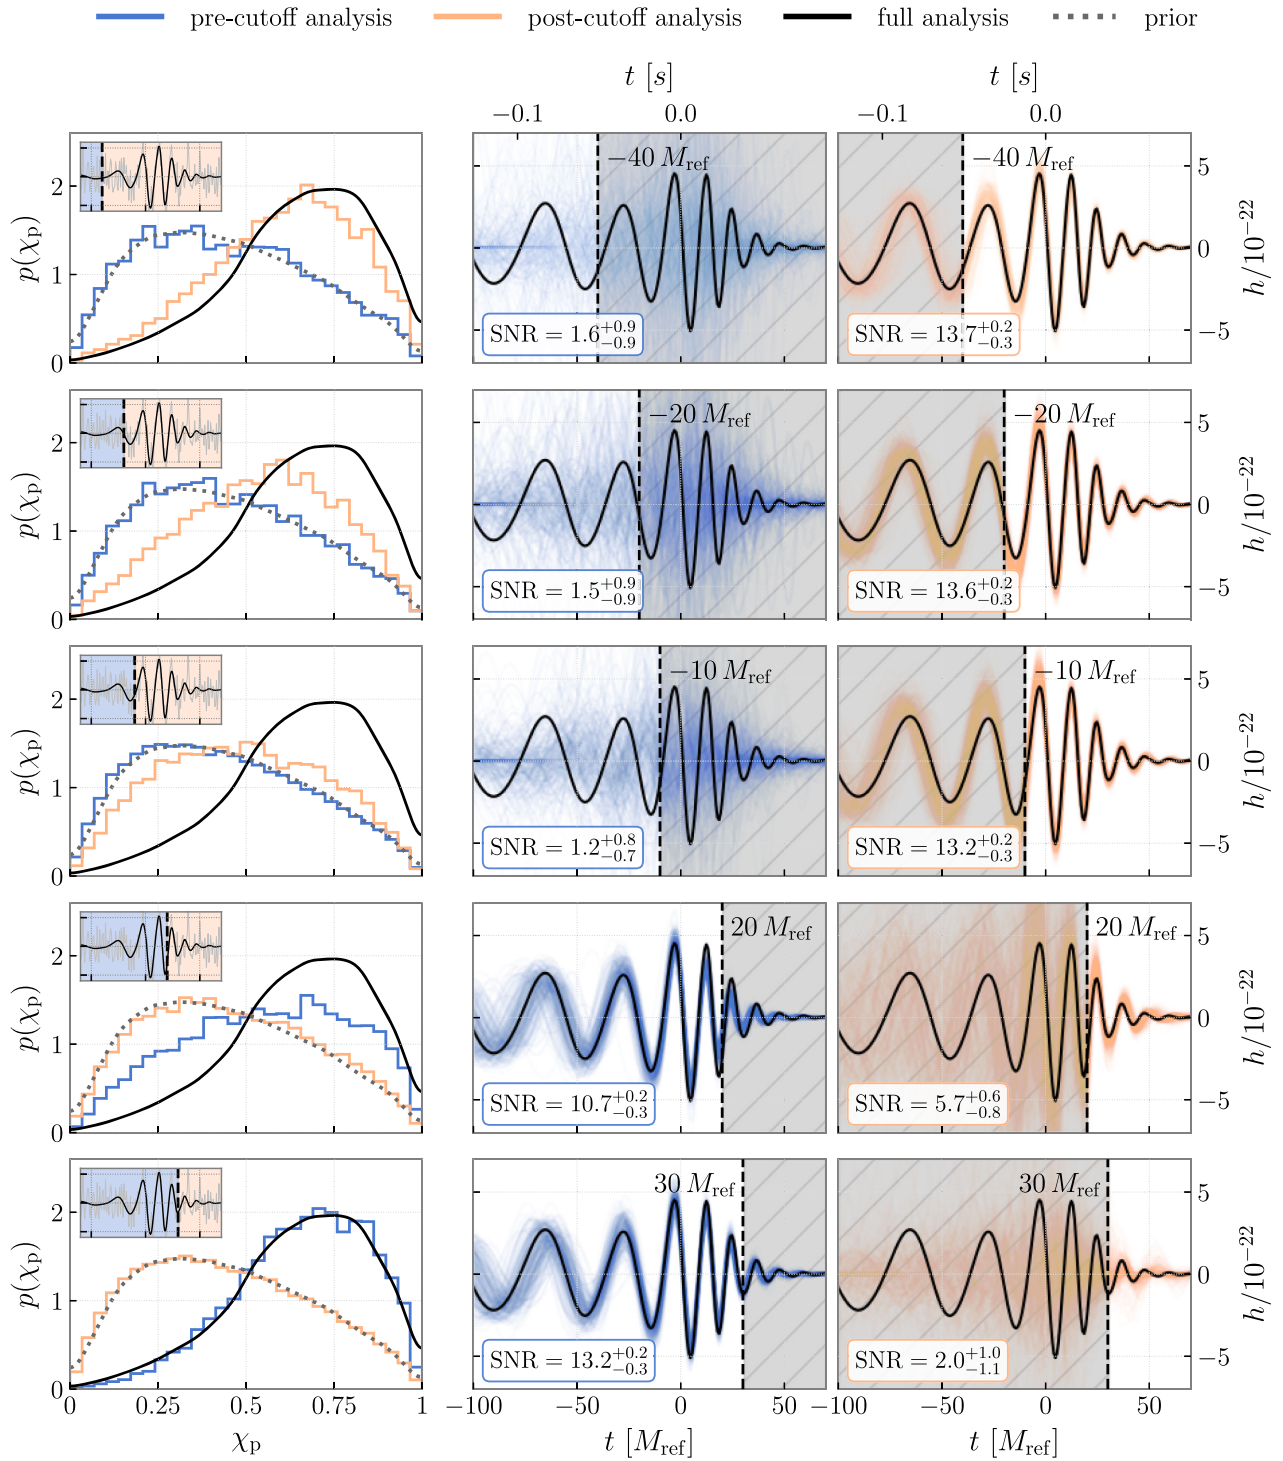


FIG. 1. Evolution of GW190521 inference for representative cutoff times $t \in \{-40, -20, -10, 20, 30\} M_{\text{ref}}$ from the time of coalescence (top to bottom; vertical black dashed lines where applicable). *Left*: Posterior for χ_p from the precutoff (blue), postcutoff (orange), and full (black solid) analysis and prior (gray dotted). The inset shows the whitened maximum-posterior waveform (with $\chi_p = 0.62$, detector-frame total mass $M = 267 M_{\odot}$, and $q = 0.89$) from the full analysis (black) along with the whitened LIGO Livingston data (gray). The blue/orange shaded regions highlight the data informing the same-color χ_p posterior. *Center and Right*: Waveform reconstruction draws for LIGO Livingston from the precutoff (center, blue) and postcutoff (right, orange) analyses and maximum-posterior waveform for the full analysis (black). Median and 90% credible intervals for the matched-filter network SNRs are given in-figure. Gray shading denotes data excluded from each analysis. See Ref. [90] for an animation of this figure including more cutoff times.

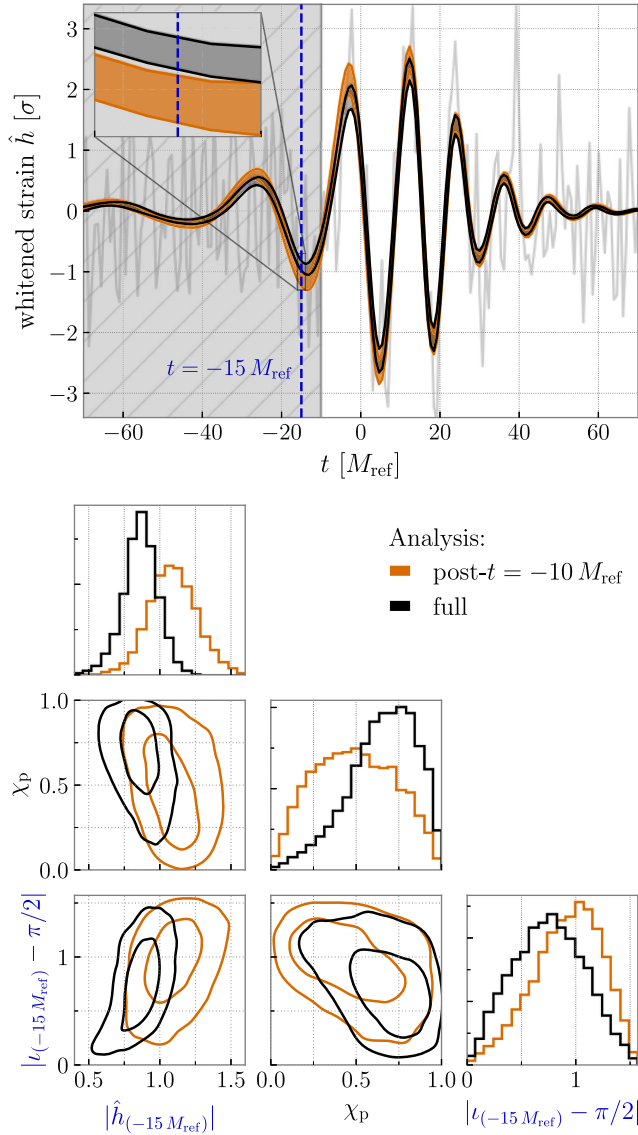


FIG. 2. Results from the full (black) and from the post- $t = -10M_{\text{ref}}$ analysis (orange) that are informative and uninformative about precession respectively. *Top*: 50% credible intervals for the whitened waveforms in LIGO Livingston in units of standard deviations of the noise. Data are plotted in gray. Gray shading denotes data excluded from the post- $t = -10M_{\text{ref}}$ analysis. The inset zooms in around $t = -15M_{\text{ref}}$ (blue dashed line), the minimum of the final premerger cycle. *Bottom*: Posteriors for χ_p , the absolute value of the whitened strain $|\hat{h}|$, and the inclination angle relative to edge-on configurations $|\iota - \pi/2|$. Quantities labeled in blue are plotted at $t = -15M_{\text{ref}}$. Contours denote 50% and 90% credible regions. The whitened strain is anticorrelated with χ_p and correlated with $|\iota - \pi/2|$. Large χ_p is paired with smaller premerger signal and more edge-on configurations.

this cycle is excluded from the analysis (orange), the extrapolated waveform is not required to have such a small value, obviating the need for a higher χ_p . In summary, precession in GW190521 is informed by a suppression of the gravitational wave in the observed waveform's last




FIG. 3. Posteriors for χ_p , M (detector frame total mass), and q (mass ratio) for the same cutoff times as Fig. 1. Posteriors from the full analysis and priors are plotted in black solid and dotted respectively. Contours denote 50% credible regions. Information about spin-precession is lost post- $t = -10M_{\text{ref}}$ (shaded light orange), while the total mass and mass ratio posteriors are informative at post- $t = 20M_{\text{ref}}$ (green). See Ref. [92] for an animation showing corner plots for mass and spin parameters for the precutoff and postcutoff analyses at more cutoff times.

cycle before merger. This is also the region in which the waveform is overall quietest; the whitened signal is less than 1σ above the noise compared to the subsequent merger cycles that are over 2σ .

The origin of the signal suppression can be attributed to the evolution of the emission direction. As a binary precesses, the angle between the dominant emission direction and the line-of-sight evolves, changing the amplitude of the observed signal. Systems with $\iota \sim \pi/2$ are quieter than those with $\iota \sim 0$ or $\sim \pi$. To explore these dynamics, we plot the absolute difference between the inclination angle and $\pi/2$ at the time $t = -15M_{\text{ref}}$ in Fig. 2. As expected, χ_p and $|\iota - \pi/2|$ are anticorrelated, while $|\hat{h}|$ and $|\iota - \pi/2|$ are correlated. When precession is constrained, ι is found to be closer to edge-on at the last premerger cycle than when precession is unconstrained, leading to a suppressed cycle and smaller $|\hat{h}|$.

We investigate the impact of data truncation on other source parameters in Fig. 3. Information about χ_p , M and q is not lost or gained in lockstep as a function of cutoff time. At post- $t = -20M_{\text{ref}}$ (dark orange), the χ_p posterior shifts away from the full analysis posterior; however, for M and q , this does not happen until multiple cycles later. Post- $t = -10M_{\text{ref}}$ (shaded light orange), i.e., at the end of the final premerger cycle, the χ_p posterior is close to the prior, while the detector-frame total mass M and mass ratio

q both resemble the posteriors from the full analysis. Thus, the lack of an informative χ_p posterior post- $t = -10M_{\text{ref}}$ does not simply arise from poor parameter constraints overall due to lower SNR; rather, the suppression of the final pre-merger cycle is informative specifically about precession (under the assumption that the system has a quasi-circular orbit). We also confirm that the χ_p inference is not driven by a conditional measurement based on the typically better-measured aligned spins in Appendix C.

IV. CONCLUSIONS

Precession inference for the massive, distant binary black hole signal GW190521 is subtle. It originates from contrasting a ~ 40 ms slice of data from the final pre-merger cycle between $t = -40M_{\text{ref}}$ and $-10M_{\text{ref}}$ with the loud merger cycles following it. The merger of GW190521 is 2 loud cycles that reach 2.5σ above the noise and is informative about the source’s masses. However, precession is only constrained away from the prior when the merger is observed in tandem with the final premerger cycle, which does not rise more than 1σ above the noise. The measurement is linked to a relative suppression of the aforementioned final cycle, caused by the binary tilting toward an edge-on configuration due to precession. This picture qualitatively agrees with the interpretation posited in Ref. [1] by comparing precessing and spin-aligned waveform reconstructions for the full signal, and is supported by simulations [42].

Siegel *et al.* [88] carried out a complementary study seeking a description of the GW190521 ringdown consistent with the NRSur7dq4 full analysis.⁹ Ringdown mode content encodes information about the preceding binary dynamics [38,96,97], meaning it is (in theory) possible to identify signatures of precession in the ringdown. Siegel *et al.* [88] found support for the presence of at least two modes; consistency with NRSur7dq4 suggests a configuration including the 220 and 210 fundamental modes. A large 210 mode amplitude could be expected under strong precession [88]. The fact that past GW190521 ringdown-only analyses cannot unequivocally infer precession is consistent with our finding that a postpeak analysis is not sufficient to constrain precession.

Our study highlights the delicate nature of precessional imprints on observed signals, providing a new view of spins in massive systems beyond the frequency domain [1,35,76]. For GW190521, the difference between the most informative precession inference to date and the prior boils down to a single, quiet premerger cycle that needs to be measured to better than half a standard deviation, cf., the difference between the black and orange waveforms in Fig. 2. This is quantitatively in agreement with the conclusions of Payne *et al.* [53] who

explored the impact of data quality on our ability to obtain an unbiased measurement at that level.

Our novel time-domain approach of tracing the observational imprint of interesting physical effects cycle-by-cycle can provide physical intuition about how key source properties are inferred in relation to observed data features. In anticipation of further massive observed signals, such a correspondence between source dynamics and observed data can help pinpoint the most informative data in order to assess data quality and waveform systematics, and enable us to morphologically study competing physical interpretations that are likely to keep arising.

Posterior samples from all our precutoff and postcutoff analyses—including those from additional cutoff times not included in the main text—are available on Zenodo at Ref. [98]. We release posteriors for all cutoff times in intervals of $10M_{\text{ref}}$, from $t = -50M_{\text{ref}}$ to $50M_{\text{ref}}$. Cutoff times in intervals of $2.5M_{\text{ref}}$ between $t = -30M_{\text{ref}}$ and $20M_{\text{ref}}$ were additionally explored to more finely resolve the transition between informative and uninformative χ_p posteriors.

Scripts to generate the waveform reconstructions and inclination angles are on Github at Ref. [89], as are notebooks to plot all the figures appearing in the text. The repository additionally contains animations showing results from for all pre- and post-cutoff analyses in intervals of $t = 2.5M_{\text{ref}}$:

- (i) Ref. [90]: Animation of χ_p posteriors, whitened reconstructions, and colored reconstructions for all time slices; similar to Fig. 1.
- (ii) Ref. [92]: Animation of a corner plot for $\chi_p, \chi_{\text{eff}}, M$, and q at all time slices; similar to Figs. 3 and 6.

ACKNOWLEDGMENTS

We thank Harrison Siegel for helpful discussions on ringdown analyses of GW190521, Sylvia Biscoveanu for insights about inference on high-mass gravitational-wave sources, Sophie Hourihane for assistance whitening gravitational-wave signals, Davide Gerosa for insight on alternative measures of precession, and Will Farr for insights about time-domain inference. We also extend thanks to Jacob Lange, Christopher Berry, Carlos Lousto, Juan Calderon Bustillo, and Salvatore Vitale for their helpful comments on our manuscript. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. S.M. and K.C. were supported by NSF Grant No. PHY-2110111. K.C. acknowledges support from the Sloan Foundation. The Flatiron Institute is funded by the Simons Foundation. V.V. acknowledges support from NSF Grant No. PHY-2309301, and the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement No. 896869. I.M. is a recipient of the Australian Research Council Future Fellowships (FT190100574).

⁹Capano *et al.* [80,93] provided an alternative interpretation based on waveform models of the PHENOM family [94,95].

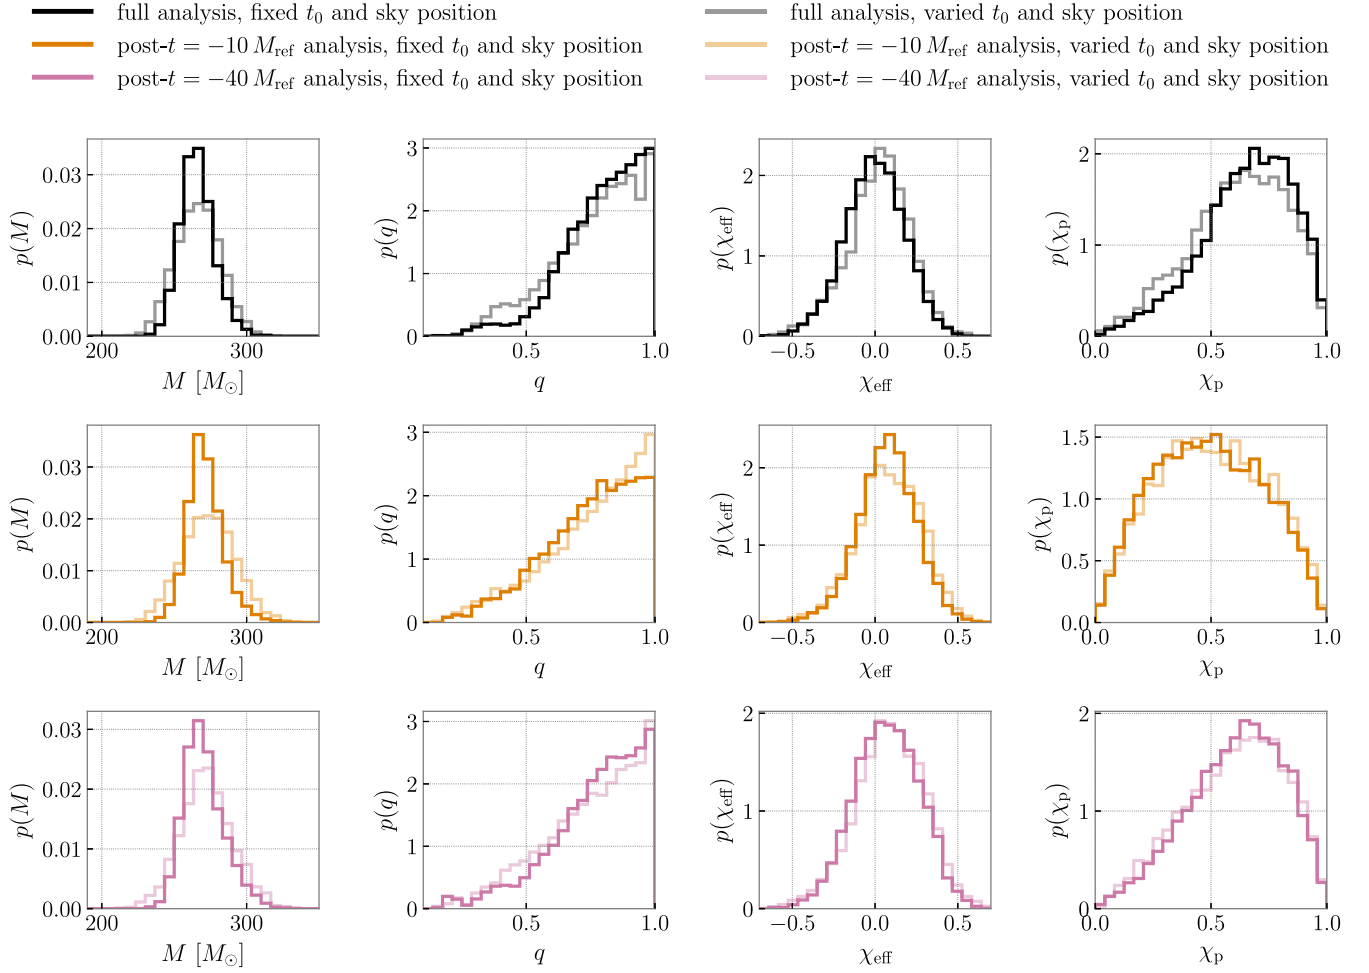


FIG. 4. Posteriors for the detector-frame total mass M , mass ratio q , effective spin χ_{eff} , and effective precessing spin χ_p for parameter estimation with (lighter color) and without (darker color) sampling over time of coalescence t_0 and sky position—referring to the right ascension, declination, and polarization angle—for the full signal (black; top row), the post- $t = -10M_{\text{ref}}$ (orange; middle row), and the post- $t = -40M_{\text{ref}}$ (pink; bottom row). All results are consistent, the only difference being a slightly wider total mass posterior.

The authors are grateful for computational resources provided by Cardiff University and supported by STFC Grant No. ST/I006285/1. Software: EMCEE [99], LALSUITE [100], NUMPY [101], SCIPY [102], H5PY [103], MATPLOTLIB [104], SEABORN [105], RINGDOWN [84,85], GWTOOLS [106].

APPENDIX A: PARAMETER ESTIMATION SETTINGS AND PRIORS

During parameter estimation, we sample the masses, spin magnitudes and tilt angles, azimuthal interspin angle, azimuthal precession cone angle, inclination angle, luminosity distance, and phase of coalescence. The time of coalescence, right ascension, declination, and polarization angle are fixed for computational efficiency; we have ensured that further sampling over these parameters does not alter our conclusions, as is shown in Fig. 4. Priors for all parameters are given in Table I. We use slightly different

TABLE I. Priors used in parameter estimation. $U(a, b)$ means uniform between a and b . An isotropic prior for an angle x means that the prior on $\cos x$ is uniform between 0 and π . Right ascension, declination, polarization, and time of coalescence are fixed to the values given in Table II.

Parameter	Symbol	Value
Detector-frame total mass	$M(1+z)$	$U(200, 500)M_{\odot}$
Mass ratio	q	$U(0.17, 1)$
Primary-spin magnitude	χ_1	$U(0, 1)$
Secondary-spin magnitude	χ_2	$U(0, 1)$
Primary-spin tilt	θ_1	Isotropic
Secondary-spin tilt	θ_2	Isotropic
Azimuthal interspin angle	ϕ_{12}	$U(0, 2\pi)$
Azimuthal cone-precession angle	ϕ_{JL}	$U(0, 2\pi)$
Inclination angle	i	Isotropic
Luminosity distance	d_L	$U(10^3, 10^4)$ Mpc
Phase of coalescence	φ	$U(-\pi, \pi)$

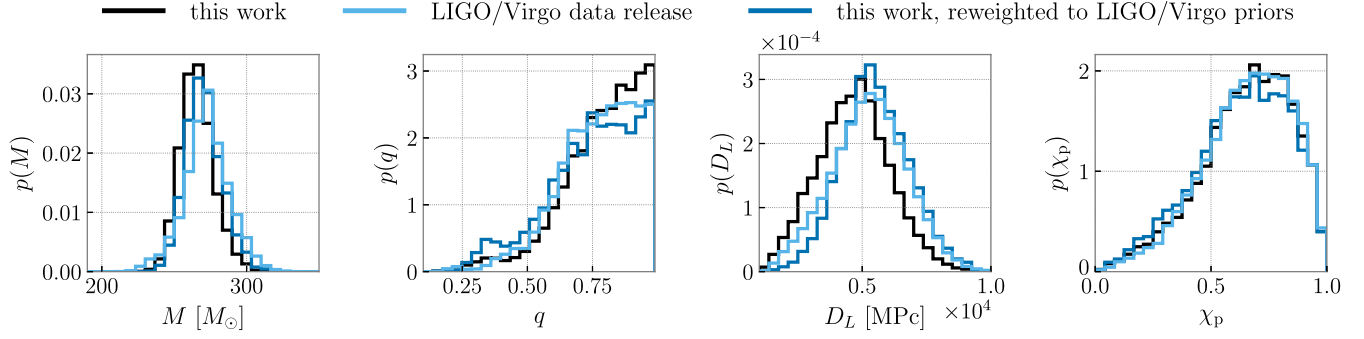


FIG. 5. Posteriors for the detector-frame total mass M , mass ratio q , luminosity distance D_L , and effective precessing spin χ_p from this work (black), Ref. [1] (LIGO/Virgo data release; light blue), and this work reweighted to the priors of Ref. [1] (dark blue). Our inference is consistent with Ref. [1] under the same priors.

TABLE II. Settings for the time-domain parameter estimation for GW190521. The precutoff analyses all begin at t_{start} , while the postcutoff ones end at t_{end} . These times are calculated with respect to the geocenter time t_0 and then shifted between detectors using the extrinsic parameters α , δ , and ψ .

Parameter	Symbol	Value
Start GPS time of “before” segments	t_{start}	1242442966.907715 s
End GPS time of “after” segments	t_{end}	1242442967.607715 s
Coalescence GPS time	t_0	1242442967.405764 s
Mass-time scaling relation	$1M_{\text{ref}}$	0.00127 s
Right ascension	α	6.075 rad
Declination	δ	-0.800 rad
Polarization angle	ψ	2.443 rad
Minimum frequency	f_{min}	11 Hz
Maximum frequency	f_{max}	1024 Hz
Sampling rate	f_{samp}	2048 Hz
Reference frequency	f_{ref}	11 Hz

mass and distance priors than Ref. [1] but, as we show in Fig. 5, obtain consistent results when analyzing the full signal.¹⁰

Parameter estimation settings are listed in Table II. All analyses make use of the three detector network of LIGO Livingston, LIGO Hanford, and Virgo, using strain and event power spectral densities publicly available on the Gravitational-wave Open Science Center [108,109].¹¹ We use a lower and reference frequency of 11 Hz, a maximum frequency of 1024 Hz, and a sampling rate of 2048 Hz.¹² Each analysis conducted before a given

¹⁰We use the NRSur7dq4 [51] posterior samples (specifically GW190521_posterior_samples.h5) released by LIGO/Virgo at Ref. [107] for this comparison.

¹¹Specifically we download the “32 sec, 16 KHz” strain data uploaded to Ref. [110].

¹²We opt for a minimum frequency of 11 Hz which is consistent with Refs. [16,49], but not the more recent Ref. [50].

cutoff time begins at the GPS time at geocenter 1242442966.9077148 s; those conducted *after* a given time all end at 1242442967.607715 s. To translate the cutoff times from geocenter to the times at the three detectors, we use a right ascension of $\alpha = 6.075$ rad, a declination of $\delta = -0.8$ rad, and polarization angle of $\psi = 2.443$ rad.

APPENDIX B: WHITENING WAVEFORMS

Whitened waveforms are obtained by dividing the Fourier-domain waveform by the noise amplitude spectral density and then inverse Fourier-transforming. In the frequency domain, the whitened waveform $\hat{h}(f)$ is obtained from the original waveform $h(f)$ by

$$\hat{h}(f) = \sqrt{\frac{f_s}{2S_n(f)}}h(f), \quad (\text{B1})$$

where $S_n(f)$ is the power spectral density and f_s is the sampling rate of the data. Though we sample the data at 2048 Hz for inference, we use 1024 Hz when plotting whitened waveforms for direct comparison to Ref. [1].

APPENDIX C: QUANTIFYING PRECESSION

Spins aligned with the orbit are characterized at leading order by the effective spin χ_{eff} ,

$$\chi_{\text{eff}} = \frac{\chi_1 \cos \theta_1 + q\chi_2 \cos \theta_2}{1 + q} \in (-1, 1), \quad (\text{C1})$$

which is conserved to at least the second post-Newtonian order in the inspiral [24,25]. Precession—stemming from misaligned spins—is typically characterized via the “effective precessing spin” parameter [48], as defined in Eq. (1) in the main text. Posteriors for χ_{eff} versus χ_p for five representative cutoff times (compared to the full signal posterior and the prior) are shown in Fig. 6. Though typically χ_{eff} is better constrained than χ_p , this is not the

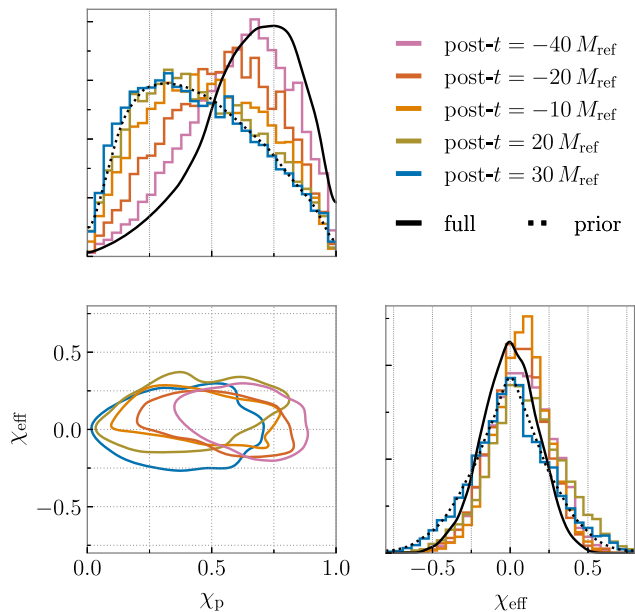


FIG. 6. Posteriors for χ_p and χ_{eff} for GW190521 informed by data after five representative cutoff times; $t = \{-40, -20, -10, 20, 30\}M_{\text{ref}}$ in pink, red, orange, green, blue respectively, compared to the full signal analysis (solid) and the prior (dotted) in black. Contours of the 50% credible regions are shown for the two-dimensional posteriors. The two spin parameters are not correlated, meaning that inferences on χ_p at the various cutoff times do not simply arise from conditional priors from χ_{eff} .

case for GW190521: χ_{eff} is never inferred to be far from its prior with only slight deviations from the prior that occur at different times from when we best constrain χ_p . Furthermore, because the χ_p and χ_{eff} are not correlated at any time slice, inferences of χ_p —and thus our main conclusions—do not arise from a conditional measurement driven by χ_{eff} [55]. In other words, we are not just identifying differences in the spins from χ_{eff} that then propagate into differences in χ_p , but rather we are directly identifying differences in the in-plane spin components' inference.

However, χ_p is a parameter motivated by inspiral precession dynamics [46–48]. The angles θ_i associated with each black hole's spin vector—used to calculate χ_p —are defined with respect to the binary's orbital angular momentum, a quantity that becomes meaningless during the merger as the two black holes cease orbiting each other. Indeed, the canonical equations describing precession dynamics [Eq. (11) of Ref. [26]] are only accurate through the second post-Newtonian order, an expansion which is only valid when the black holes are moving substantially slower than the speed of light, i.e., during the inspiral [26,27,111–113]. Once entering the merger, precession becomes no longer analytically tractable.

Alternative parametrizations of precession have been proposed, though they are all still inspired by intuition gained from the inspiral precession equations. Gerosa *et al.* [54] note that in the definition of χ_p , some—but not all—of the precession-timescale oscillations are averaged over. To rectify this and retain all variation over the precessional timescale [33], they define a “generalized χ_p ” in their Eq. (15). Defined on the range $[0, 2)$, the generalized χ_p has the advantage that it can help distinguish between binaries in which one versus both spins are precessing. While binaries with generalized $\chi_p < 1$ can have one or both spins precessing, the case in which generalized $\chi_p > 1$ only arises in the both-precessing scenario. Thomas *et al.* [114] amend the fact that χ_p does not accurately account for higher-order multipolar modes by defining the “effective precession spin vector” $\vec{\chi}_{\perp}$ —their Eq. (9). This parameter accounts for more degrees of freedom and could facilitate a better representation of precession in the strong-field regime. The magnitude of the effective precession spin vector $\|\vec{\chi}_{\perp}\|$ is a scalar measure of precession, analogous to the traditional χ_p . A third way to quantify the observability of spin precession is through the precessional SNR ρ_p , as defined in Fairhurst *et al.* [43,115]. Based on the idea that precession is only inferred when two gravitational-wave harmonics are observed, ρ_p is related to the SNR in the second most significant waveform harmonic, defined in Eq. (39) of Ref. [43].

Posteriors for generalized χ_p and $\|\vec{\chi}_{\perp}\|$ are plotted in Fig. 7 for the post- $t = -40M_{\text{ref}}$ (pink) and $-10M_{\text{ref}}$ (orange) analyses, compared to the full analysis posteriors (black). We do not plot ρ_p , as the current code to compute it [116] necessitates a frequency domain waveform, which NRSur7dq4 is not [51]; Hoy *et al.* [117] presented ρ_p constraints for GW190521 using a different set of waveform models. Both generation χ_p and $\|\vec{\chi}_{\perp}\|$ follow the same trend as the traditional χ_p , shown in the top panel of Fig. 7 for comparison. The inference of precession is closer to that from the full signal in the post- $t = -40M_{\text{ref}}$ analysis than the $t = -10M_{\text{ref}}$.

The above quantities are all defined based on the inspiral dynamics; precession representations motivated by merger dynamics are not available to the best of our knowledge. Furthermore, all of the above quantities, as well as χ_p itself, are quoted at the reference frequency of 11 Hz. Since quantifications of precession are time varying (albeit slowly [46]) we expect their measurement to change as a function of reference frequency. However, crucially we do not expect the choice of reference frequency to affect our conclusions. As long as the same reference frequency is used for all pre/postcutoff analyses, the posteriors for χ_p (and related quantities) are expected to vary with the cutoff time in a similar way to the results presented in Fig. 7.




FIG. 7. From top to bottom: posteriors for χ_p , generalized χ_p , and $\|\vec{\chi}_\perp\|$ for the full analysis (black), and the post- $t = -40M_{\text{ref}}$ (pink) and post- $t = -10M_{\text{ref}}$ (orange) analyses. All parametrizations of precession lead to qualitatively similar conclusions about spin inference and GW190521.

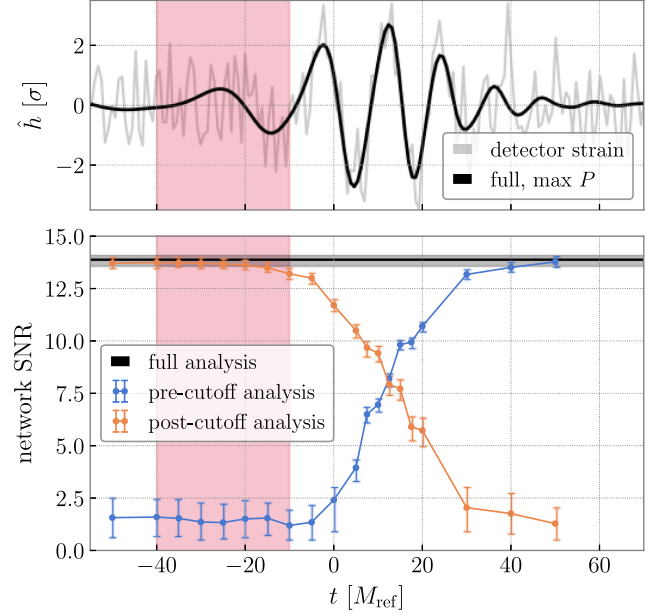


FIG. 8. (*Lower subplot*) Network matched-filter SNR for the precutoff (blue) and postcutoff (orange) analyses as a function of the cutoff time t . Points indicate the median SNR and error bars represent the 90% credible region. The horizontal black line represents the median SNR recovered by the full analysis, with the 90% credible region shaded in gray. (*Upper subplot*) Whitened strain data in LIGO Livingston (gray), and whitened waveform reconstruction for the maximum posterior (abbr. “max P ”) draw from the full analysis (black), plotted to help visualize the signal alongside the SNR evolution. In both subplots, the red shaded region—between $t = -40M_{\text{ref}}$ and $-10M_{\text{ref}}$ —is that which is identified as crucial for constraining χ_p .

APPENDIX D: SIGNAL STRENGTH OVER TIME

We track the time-bounded matched-filter signal-to-noise ratio (SNR) in the LIGO Livingston, LIGO Hanford, and Virgo detector network recovered by our analyses, as defined in Eqs. (52) and (53) of Isi and Farr [84]. Figure 8 shows the SNR evolution for the precutoff (blue) and postcutoff (orange) analyses. The change in SNR is negligible between $t = -40M_{\text{ref}}$ and $-10M_{\text{ref}}$ (red shaded region) for both the pre and postcutoff analyses, suggesting that the change in inferred χ_p between these times is not simply due to a drop in SNR.

To investigate correlations between signal strength and χ_p (see Fig. 2 in the main text), we compare the 50% credible intervals (CI) of the reconstructed waveform strain for two analyses, one in which precession is inferred (full analysis) and one in which it is not (post- $t = -10M_{\text{ref}}$ analysis). The top panel in Fig. 9 shows the 50% CIs for the




FIG. 9. Identification of regions in time (pink shaded) during which the 50% credible intervals (CI) of the reconstructed waveform strain \hat{h} are inconsistent between the full (black; labeled \hat{h}_{full}) and post- $t = -10M_{\text{ref}}$ (orange; labeled \hat{h}_{-10}) analyses. This occurs around the two extrema of the final premerger cycle ($t \sim -30M_{\text{ref}}$ and $t \sim -15M_{\text{ref}}$; indicated with vertical black dashed lines), when $\delta\hat{h}$ —the difference between the 75% CI of one analysis and the 25% CI of the other (green and blue)—is negative (enlarged in the inset).

strain \hat{h} for these two analyses in black and orange respectively. In the bottom panel, we plot the differences $\delta\hat{h}$ between the 75% CI of one analysis and the 25% CI of the other in green and blue. When either difference is *negative*, the two sets of waveform reconstructions are

inconsistent at the 50% CI. This inconsistency only occurs at the extrema of the final premerger cycle, i.e., the peak around $t \sim -30M_{\text{ref}}$ and trough at $t \sim -15M_{\text{ref}}$, as indicated by the time slices shaded in pink. These regions are enlarged in the inset.

-
- [1] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **125**, 101102 (2020).
 - [2] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **900**, L13 (2020).
 - [3] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy *et al.*, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [4] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou *et al.*, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [5] D. Gerosa and E. Berti, *Phys. Rev. D* **95**, 124046 (2017).
 - [6] D. Gerosa and M. Fishbach, *Nat. Astron.* **5**, 749 (2021).
 - [7] C. Kimball, C. Talbot, C. P. L. Berry, M. Carney, M. Zevin, E. Thrane, and V. Kalogera, *Astrophys. J.* **900**, 177 (2020).
 - [8] Z. Doctor, D. Wysocki, R. O’Shaughnessy, D. E. Holz, and B. Farr, [arXiv:1911.04424](https://arxiv.org/abs/1911.04424).
 - [9] G. Fragione, B. Kocsis, F. A. Rasio, and J. Silk, *Astrophys. J.* **927**, 231 (2022).
 - [10] S. E. Woosley, *Astrophys. J.* **836**, 244 (2017).
 - [11] S. E. Woosley and A. Heger, *Astrophys. J. Lett.* **912**, L31 (2021).
 - [12] K. Belczynski *et al.*, *Astron. Astrophys.* **594**, A97 (2016).
 - [13] L. A. C. van Son, S. E. de Mink, F. S. Broekgaarden, M. Renzo, S. Justham, E. Laplace, J. Moran-Fraile, D. D. Hendriks, and R. Farmer, *Astrophys. J.* **897**, 100 (2020).
 - [14] S. Stevenson, M. Sampson, J. Powell, A. Vigna-Gómez, C. J. Neijssel, D. Szécsi, and I. Mandel, [arXiv:1904.02821](https://arxiv.org/abs/1904.02821).
 - [15] M. Renzo, M. Cantiello, B. D. Metzger, and Y. F. Jiang, *Astrophys. J. Lett.* **904**, L13 (2020).
 - [16] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *Phys. Rev. X* **13**, 041039 (2023).
 - [17] I. Mandel and A. Farmer, *Phys. Rep.* **955**, 1 (2022).
 - [18] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, *Astrophys. J. Lett.* **832**, L2 (2016).
 - [19] C. L. Rodriguez, M. Zevin, P. Amaro-Seoane, S. Chatterjee, K. Kremer, F. A. Rasio, and C. S. Ye, *Phys. Rev. D* **100**, 043027 (2019).

- [20] D. Gerosa, E. Berti, R. O’Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, *Phys. Rev. D* **98**, 084036 (2018).
- [21] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, B. Farr, and A. Vecchio, *Nature (London)* **548**, 426 (2017).
- [22] J. Fuller and L. Ma, *Astrophys. J. Lett.* **881**, L1 (2019).
- [23] M. Campanelli, C. O. Lousto, and Y. Zlochower, *Phys. Rev. D* **74**, 041501 (2006).
- [24] E. Racine, *Phys. Rev. D* **78**, 044021 (2008).
- [25] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [26] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [27] L. E. Kidder, C. M. Will, and A. G. Wiseman, *Phys. Rev. D* **47**, R4183 (1993).
- [28] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, *Phys. Rev. Lett.* **112**, 251101 (2014).
- [29] S. Vitale, R. Lynch, V. Raymond, R. Sturani, J. Veitch, and P. Graff, *Phys. Rev. D* **95**, 064053 (2017).
- [30] F. H. Shaik, J. Lange, S. E. Field, R. O’Shaughnessy, V. Varma, L. E. Kidder, H. P. Pfeiffer, and D. Wysocki, *Phys. Rev. D* **101**, 124054 (2020).
- [31] M. Zevin, S. S. Bavera, C. P. L. Berry, V. Kalogera, T. Fragos, P. Marchant, C. L. Rodriguez, F. Antonini, D. E. Holz, and C. Pankow, *Astrophys. J.* **910**, 152 (2021).
- [32] R. C. Zhang, G. Fragione, C. Kimball, and V. Kalogera, *Astrophys. J.* **954**, 23 (2023).
- [33] D. Gerosa, M. Kesden, U. Sperhake, E. Berti, and R. O’Shaughnessy, *Phys. Rev. D* **92**, 064016 (2015).
- [34] R. O’Shaughnessy, L. London, J. Healy, and D. Shoemaker, *Phys. Rev. D* **87**, 044038 (2013).
- [35] S. Biscoveanu, M. Isi, V. Varma, and S. Vitale, *Phys. Rev. D* **104**, 103018 (2021).
- [36] S. A. Hughes, A. Apte, G. Khanna, and H. Lim, *Phys. Rev. Lett.* **123**, 161101 (2019).
- [37] E. Finch and C. J. Moore, *Phys. Rev. D* **103**, 084048 (2021).
- [38] I. Kamaretsos, M. Hannam, and B. Sathyaprakash, *Phys. Rev. Lett.* **109**, 141102 (2012).
- [39] E. Hamilton, L. London, and M. Hannam, *Phys. Rev. D* **107**, 104035 (2023).
- [40] I. M. Romero-Shaw, P. D. Lasky, E. Thrane, and J. C. Bustillo, *Astrophys. J. Lett.* **903**, L5 (2020).
- [41] I. M. Romero-Shaw, D. Gerosa, and N. Loutrel, *Mon. Not. R. Astron. Soc.* **519**, 5352 (2023).
- [42] J. Calderón Bustillo, N. Sanchis-Gual, A. Torres-Forné, and J. A. Font, *Phys. Rev. Lett.* **126**, 201101 (2021).
- [43] S. Fairhurst, R. Green, C. Hoy, M. Hannam, and A. Muir, *Phys. Rev. D* **102**, 024055 (2020).
- [44] R. Green, C. Hoy, S. Fairhurst, M. Hannam, F. Pannarale, and C. Thomas, *Phys. Rev. D* **103**, 124023 (2021).
- [45] S. Olsen, J. Roulet, H. S. Chia, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, *Phys. Rev. D* **104**, 083036 (2021).
- [46] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).
- [47] P. Schmidt, M. Hannam, and S. Husa, *Phys. Rev. D* **86**, 104063 (2012).
- [48] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015).
- [49] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **11**, 021053 (2021).
- [50] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **109**, 022001 (2024).
- [51] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. Res.* **1**, 033015 (2019).
- [52] M. Hannam *et al.*, *Nature (London)* **610**, 652 (2022).
- [53] E. Payne, S. Hourihane, J. Golomb, R. Udall, R. Udall, D. Davis, and K. Chatziioannou, *Phys. Rev. D* **106**, 104017 (2022).
- [54] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, *Phys. Rev. D* **103**, 064067 (2021).
- [55] D. Gangardt, D. Gerosa, M. Kesden, V. De Renzi, and N. Steinle, *Phys. Rev. D* **106**, 024019 (2022); **107**, 109901(E) (2023).
- [56] M. C. Miller and D. P. Hamilton, *Mon. Not. R. Astron. Soc.* **330**, 232 (2002).
- [57] C. Kimball *et al.*, *Astrophys. J. Lett.* **915**, L35 (2021).
- [58] D. Gerosa and E. Berti, *Phys. Rev. D* **100**, 041301 (2019).
- [59] G. Fragione, A. Loeb, and F. A. Rasio, *Astrophys. J. Lett.* **902**, L26 (2020).
- [60] B. McKernan, K. E. S. Ford, B. Kocsis, W. Lyra, and L. M. Winter, *Mon. Not. R. Astron. Soc.* **441**, 900 (2014).
- [61] B. McKernan *et al.*, *Astrophys. J.* **866**, 66 (2018).
- [62] B. McKernan, K. E. S. Ford, R. O’Shaughnessy, and D. Wysocki, *Mon. Not. R. Astron. Soc.* **494**, 1203 (2020).
- [63] K. Belczynski, *Astrophys. J. Lett.* **905**, L15 (2020).
- [64] L. Gondán and B. Kocsis, *Mon. Not. R. Astron. Soc.* **506**, 1665 (2021).
- [65] V. De Luca, V. Desjacques, G. Franciolini, P. Pani, and A. Riotto, *Phys. Rev. Lett.* **126**, 051101 (2021).
- [66] M. Fishbach and D. E. Holz, *Astrophys. J. Lett.* **904**, L26 (2020).
- [67] A. H. Nitz and C. D. Capano, *Astrophys. J. Lett.* **907**, L9 (2021).
- [68] V. Gayathri, J. Healy, J. Lange, B. O’Brien, M. Szczepanczyk, I. Bartos, M. Campanelli, S. Klimentko, C. O. Lousto, and R. O’Shaughnessy, *Nat. Astron.* **6**, 344 (2022).
- [69] H. L. Iglesias *et al.*, arXiv:2208.01766.
- [70] A. Ramos-Buades, A. Buonanno, and J. Gair, arXiv:2309.15528.
- [71] R. Gamba, M. Breschi, G. Carullo, S. Albanesi, P. Rettegno, S. Bernuzzi, and A. Nagar, *Nat. Astron.* **7**, 11 (2023).
- [72] J. Calderón Bustillo, N. Sanchis-Gual, A. Torres-Forné, J. A. Font, A. Vajpeyi, R. Smith, C. Herdeiro, E. Radu, and S. H. W. Leong, *Phys. Rev. Lett.* **126**, 081101 (2021).
- [73] J. C. Aurrekoetxea, T. Helfer, and E. A. Lim, *Classical Quantum Gravity* **37**, 204001 (2020).
- [74] M. Zevin, I. M. Romero-Shaw, K. Kremer, E. Thrane, and P. D. Lasky, *Astrophys. J. Lett.* **921**, L43 (2021).
- [75] I. M. Romero-Shaw, P. D. Lasky, and E. Thrane, *Astrophys. J.* **940**, 171 (2022).
- [76] Y. Xu and E. Hamilton, *Phys. Rev. D* **107**, 103049 (2023).
- [77] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Classical Quantum Gravity* **37**, 055002 (2020).
- [78] M. Unser, *Signal Processing* **7**, 231 (1984).

- [79] B. Zackay, T. Venumadhav, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **104**, 063034 (2021).
- [80] C. D. Capano, M. Cabero, J. Westerweck, J. Abedi, S. Kastha, A. H. Nitz, Y.-F. Wang, A. B. Nielsen, and B. Krishnan, *Phys. Rev. Lett.* **131**, 221402 (2023).
- [81] M. Cabero, C. D. Capano, O. Fischer-Birnholtz, B. Krishnan, A. B. Nielsen, A. H. Nitz, and C. M. Biwer, *Phys. Rev. D* **97**, 124069 (2018).
- [82] A. Ghosh *et al.*, *Phys. Rev. D* **94**, 021101 (2016).
- [83] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton, *Classical Quantum Gravity* **39**, 245013 (2022).
- [84] M. Isi and W. M. Farr, arXiv:2107.05609.
- [85] M. Isi, M. Giesler, W. M. Farr, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. Lett.* **123**, 111102 (2019).
- [86] M. Isi, W. M. Farr, M. Giesler, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. Lett.* **127**, 011103 (2021).
- [87] G. Carullo, W. Del Pozzo, and J. Veitch, *Phys. Rev. D* **99**, 123029 (2019); **100**, 089903(E) (2019).
- [88] H. Siegel, M. Isi, and W. Farr, *Phys. Rev. D* **108**, 064008 (2023).
- [89] S. Miller, GW190521-timedomain-release (2023), <https://github.com/simonajmiller/gw190521-timedomain-release/tree/main>.
- [90] S. Miller, Animation of χ_p posterior and waveform reconstructions for time-domain inference of GW190521 (2023), https://github.com/simonajmiller/gw190521-timedomain-release/blob/main/gifs/gif_01.gif.
- [91] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [92] S. Miller, Animation of mass and spin posteriors for time-domain inference of GW190521 (2023), https://github.com/simonajmiller/gw190521-timedomain-release/blob/main/gifs/gif_02.gif.
- [93] C. D. Capano, J. Abedi, S. Kastha, A. H. Nitz, J. Westerweck, Y.-F. Wang, M. Cabero, A. B. Nielsen, and B. Krishnan, arXiv:2209.00640.
- [94] G. Pratten *et al.*, *Phys. Rev. D* **103**, 104056 (2021).
- [95] H. Estellés *et al.*, *Astrophys. J.* **924**, 79 (2022).
- [96] S. Borhanian, K. G. Arun, H. P. Pfeiffer, and B. S. Sathyaprakash, *Classical Quantum Gravity* **37**, 065006 (2020).
- [97] J. Abedi, C. D. Capano, S. Kastha, Y.-F. Wang, J. Westerweck, A. B. Nielsen, and B. Krishnan, *Phys. Rev. D* **108**, 104009 (2023).
- [98] S. Miller, Dataset: GW190521: Tracing imprints of spin-precession on the most massive black hole binary, Zenodo, 10.5281/zenodo.8349581, (2023).
- [99] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [100] LIGO Scientific Collaboration, LIGO Algorithm Library—LALSuite, Free software (GPL) (2019).
- [101] C. R. Harris *et al.*, *Nature (London)* **585**, 357 (2020).
- [102] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy *et al.*, *Nat. Methods* **17**, 261 (2020).
- [103] A. Collette, T. Kluyver, T. A. Caswell, J. Tocknell, J. Kieffer *et al.*, h5py: 3.2.1 (2021), 10.5281/zenodo.4584676.
- [104] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [105] M. L. Waskom, *J. Open Source Software* **6**, 3021 (2021).
- [106] C. Galley, GWTOOLS: A collection of useful gravitational wave tools (2022), <https://bitbucket.org/chadgalley/gwtools/src/master/>.
- [107] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190521 posterior samples (2021), https://dcc.ligo.org/public/0168/P2000158/004/GW190521_posterior_samples.html.
- [108] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *SoftwareX* **13**, 100658 (2021).
- [109] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), *Astrophys. J. Suppl. Ser.* **267**, 29 (2023).
- [110] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190521 (2023), https://gwosc.org/eventapi/html/O3_Discovery_Papers/GW190521/v2/.
- [111] L. Blanchet, *Living Rev. Relativity* **17**, 2 (2014).
- [112] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009); **84**, 049901(E) (2011).
- [113] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [114] L. M. Thomas, P. Schmidt, and G. Pratten, *Phys. Rev. D* **103**, 083022 (2021).
- [115] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, *Phys. Rev. D* **102**, 041302 (2020).
- [116] C. Hoy and V. Raymond, *SoftwareX* **15**, 100765 (2021).
- [117] C. Hoy, S. Fairhurst, M. Hannam, and V. Tiwari, *Astrophys. J.* **928**, 75 (2022).