The Astrophysical Journal, 999:149 (15pp), 2026 March 1 https://doi.org/10.3847/1538-4357/ae40b3
© 2026. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

Identifying Compact Chirping SMBHBs in LSST Using Bayesian Analysis

Chengcheng Xin[1,2], Maximiliano Isi[3,4], Will M. Farr[3,5], and Zoltán Haiman[4,6,7]
[1] Black Hole Initiative, Harvard University, 20 Garden Street, Cambridge, MA 02138, USA; chengcheng.xin@cfa.harvard.edu
[2] Center for Astrophysics, Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA
[3] Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA
[4] Columbia University, Department of Astronomy, 550 West 120th Street, New York, NY 10027, USA
[5] Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA
[6] Columbia University, Department of Physics, 550 West 120th Street, New York, NY 10027, USA
[7] Institute of Science and Technology Austria (ISTA), Am Campus 1, Klosterneuburg 3400, Austria
Received 2025 July 9; revised 2026 January 27; accepted 2026 January 29; published 2026 February 27

Abstract

The Legacy Survey of Space and Time (LSST) is expected to observe up to ∼100 million quasars in the next decade. In this work, we show that it is possible to use such data to measure the characteristic frequency evolution of a "chirp" induced by gravitational waves, which can serve as robust evidence for the presence of a compact supermassive black hole binary. Following the LSST specifications, we generate mock lightcurves consisting of (i) a post-Newtonian chirp produced by orbital motion through, e.g., relativistic Doppler boosting, (ii) a damped random walk representing intrinsic quasar variability, and (iii) Gaussian photometric errors, while assuming nonuniform observations with extended gaps over a period of 10 yr. Through a fully Bayesian analysis, we show that we can simultaneously measure the chirp and noise parameters with little degeneracy between the two. For chirp signals with an amplitude of $A = 0.5$ mag and a range of times to merger ($t_m = 15$–10^4 yr), we can typically measure a nonzero amplitude and positive frequency derivative with over 5σ credibility. For binaries with $t_m = 50$ yr, we achieve 3σ (5σ) confidence that the signal is chirping for $A \gtrsim 0.1$ ($A > 0.2$). The median runtime of our analysis is 5.6 minutes, with a minimum as low as 35 s, making it scalable to a large number of lightcurves. This implies that LSST could, on its own, establish the presence of a compact supermassive black hole binary, and thus discover gravitational wave sources detectable by LISA and by Pulsar Timing Arrays.

Unified Astronomy Thesaurus concepts: Quasars (1319); Supermassive black holes (1663); Time domain astronomy (2109); Gravitational waves (678)

1. Introduction

The Legacy Survey of Space and Time (LSST) by the Vera C. Rubin Observatory is expected to observe tens to over 100 million quasars during its 10 yr lifespan, which will begin at the end of 2025 (C. Xin & Z. Haiman 2021, hereafter XH21). Assuming that most of these quasars are associated with coalescing supermassive black hole binaries (SMBHBs) triggered by galaxy mergers (M. C. Begelman et al. 1980; G. Kauffmann & M. Haehnelt 2000; D. M. Alexander & R. C. Hickox 2012; A. D. Goulding et al. 2018), we can expect that a small fraction are compact inspiraling SMBHBs driven by gravitational wave (GW) emission. Electromagnetic (EM) observations of these binaries are ideal for multimessenger studies when combined with their GW emission, which is expected in the nanohertz to millihertz band. This makes them prime targets for the upcoming Laser Interferometer Space Antenna (LISA; M. Colpi et al. 2024) and pulsar timing arrays (PTAs; G. Agazie et al. 2023a). XH21 estimates that up to ∼100 of the GW-driven binaries in LSST can be ultracompact (1–2 days periods), meaning that they can merge and be detected as mHz GWs within the LISA band 5–15 yr later. The lightcurves of compact binaries in LSST are expected to exhibit distinct periodicity and period evolution, namely an EM "chirp" (Z. Haiman 2017). Detecting such a chirp signal can serve as smoking-gun evidence for a GW-driven SMBHB. This is a novel task that is expected to be challenging unless the GWs of the binary are also observed (C. Xin & Z. Haiman 2024, hereafter XH24); XH24 shows that LISA data can make it possible to identify the signal in archival LSST data.

Generally, as long as they are luminous, SMBHBs can be identified in time-domain surveys as quasars with periodic variability (Z. Haiman et al. 2009). In the past decade, SMBHB candidates have been identified in time-domain surveys, including a total of ∼400 candidates identified in the Catalina Real-Time Transient survey (CRTS; M. J. Graham et al. 2015a), the Zwicky Transient Facility (CRTS+ZTF; V. Foustoul et al. 2025), the Palomar Transient Factory (PTF; M. Charisi et al. 2016), the Sloan Digital Sky Surveys and Dark Energy Survey (SDSS+DES; Y.-C. Chen et al. 2020), and the Gaia space observatory (P. Huijse et al. 2025). The identification of these candidates suffers from false-positive contamination due to quasar "red noise" (S. Vaughan et al. 2016) and observational limitations, including the cadence, photometric uncertainty, baseline, and depth of the surveys (C. A. Witt et al. 2022). LSST will observe a large quasar catalog, including a never-before-seen population of faint active galactic nuclei (AGNs; down to $i = 26$ magnitude) with high cadence and a long baseline (LSST Science Collaboration et al. 2009). The better "data quality" of LSST should allow us to (i) identify rare, short-lived binaries with shorter periods ($P \sim$ months to days) that exhibit a chirp, and (ii) distinguish SMBHBs from stochastic quasar variability.

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

The Lomb–Scargle (LS) periodogram (N. R. Lomb 1976; J. D. Scargle 1982) is commonly used in studies to detect periodicity in quasar lightcurves. However, the LS periodogram is suboptimal for identifying chirping signals and inferring multiple binary and noise parameters beyond the orbital period. Periodograms also struggle to quantify the significance of detected peaks and to identify significant rates of change in frequency.[8] XH24 have shown that the LS periodogram appears noisy, and any significant peak vanishes if a chirp component is present. Some past works have employed Bayesian methods to identify periodic SMBHBs in existing time-domain data (e.g., D. J. D'Orazio et al. 2015; V. Foustoul et al. 2025; P. Huijse et al. 2025) and in mock LSST data (C. A. Witt et al. 2022). However, they do not address our question: whether chirp and noise properties, in addition to periodicity, are measurable—and for which binaries are they measurable—within the real LSST quasar population as predicted by the quasar luminosity function.

Our goal is to identify chirping SMBHBs in LSST quasar lightcurves and to assess the detectability of various properties of these binaries. We develop a fully Bayesian framework to analyze LSST data and to extract the properties of the EM chirp, if present. We model the lightcurves as composed of (i) quasi-sinusoidal variability induced by relativistic Doppler boost (DB), with GW-driven frequency evolution evaluated to leading post-Newtonian order, (ii) a "damped random walk" (DRW) representing stochastic variability intrinsic to the quasar, and (iii) Gaussian photometric errors; we also allow for both nonuniform observing cadences and extended gaps in observations. An efficient likelihood parameterization and the use of a gradient-assisted, hardware-accelerated sampler allow us to efficiently analyze a large number of lightcurves, and show that LSST will be able to measure the chirp in sufficiently compact SMBHB sources.

We find that chirp and noise parameters are measurable for a wide range of compact SMBHBs with $M = 10^5$–$10^9\,M_\odot$, $z = 0$–6 and $t_m = 15$–10^4 yr, assuming the chirp amplitude (A) is $\gtrsim$0.5 mag, with periods 2–300 days. The statistical confidence depends on A and other system parameters. We perform our analysis for fiducial binaries ($t_m = 50$ yr), and we find that we can achieve at least 3σ confidence in measuring binary parameters when $A > 0.1$. This paper is organized as follows. In Section 2, we discuss our selection of simulated LSST source parameters (Section 2.1), as well as our method to simulate mock lightcurves (Section 2.2) and to analyze them (Section 2.3). We report our results in Section 3. Finally, in Section 4, we summarize our findings and discuss future plans, including expanding this study beyond the DB and DRW assumptions.

2. Methods

2.1. Sampling LSST Binary Quasars

To define a potential set of LSST sources, we follow XH21 to estimate the number of quasars as a function of their brightness and redshift. Using the extrapolated quasar luminosity function (G. Kulkarni et al. 2019), we estimate that LSST is expected to observe up to $N_{\rm total} \sim 100$ million quasars during its 10 yr survey, assuming that the detection limit reaches magnitude $m_i = 26$ in the i-band (LSST Science Collaboration et al. 2009). The binary mass ($M_{\rm bin}$) can be estimated using m_i; assuming a typical quasar spectral energy distribution (Z. Haiman et al. 2009, see also XH24, Equation (10)) and a fiducial Eddington ratio $f_{\rm Edd} = 0.3$, then

$$m_i = 26 + 2.5\log\left[\left(\frac{f_{\rm Edd}}{0.3}\right)^{-1}\left(\frac{M_{\rm bin}}{3\times10^7 M_\odot}\right)^{-1} \times\left(\frac{d_L(z)}{d_L(z=2)}\right)^2\right], \quad (1)$$

where $d_L(z)$ is the luminosity distance as a function of redshift z in the Λ − CDM cosmology we adopted with $H_0 = 70$ km s^{-1} Mpc^{-1}, $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$. In Figure 1, we show the distribution of the number of quasars N_Q as a function of $M_{\rm bin}$ and z. The quasar density peaks around $z \approx 2$ and $M_{\rm bin} \approx 10^5 M_\odot$, consistent with the measured quasar luminosity function. The left edge of the contour represents the LSST magnitude limit of $m_i = 26$, per the relationship in Equation (1). We use Figure 1 to draw masses and redshifts of putative LSST sources. Using the same quasar luminosity function, and assuming that all quasars are associated with SMBHBs and reasonable choices of Eddington ratios, Á. Kis-Tóth et al. (2025) have found that PTA sources are dominated by compact SMBHBs with $M_{\rm bin} \sim 10^9\,M_\odot$ and $z \sim$ 2–3, at higher redshifts than those predicted by other binary population models (e.g., G. Agazie et al. 2023b).

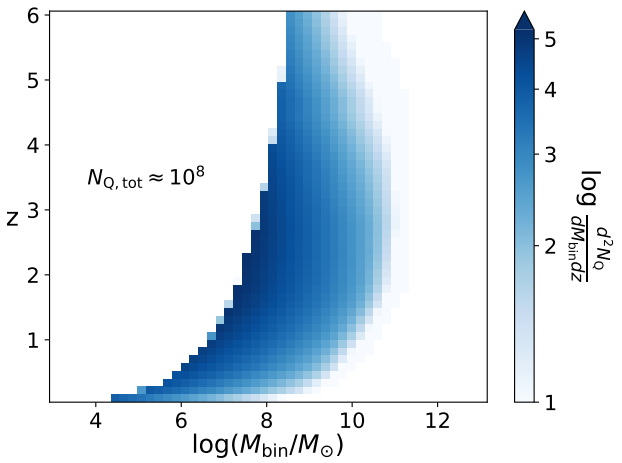


Figure 1. Expected number density of binary quasars in LSST (color) as a function of total binary mass (abscissa) and redshift (ordinate), using the extrapolated quasar luminosity function from G. Kulkarni et al. (2019) and assuming that the LSST coadded i-band magnitude threshold is $m_i = 26$ (left edge of the contour). The total number of expected LSST quasars is ∼120 million.

To simulate an LSST observation, we further need to specify the orbital separation or, equivalently (for a fixed mass ratio $q \equiv m_2/m_1 \leqslant 1$; see below), the remaining time to merger. By a simple counting argument, the abundance of binaries with (detector-frame) time to merger t_m must be proportional to t_m/t_Q, where we take $t_Q = 10^7$ yr to be the typical observable lifetime of a luminous quasar (P. Martini 2004). For instance, given the results in Figure 1, the total number of binaries with $t_m = 50$ yr is approximately $t_m/t_Q \times f_{\rm bin} \times N_{\rm Q,tot} = (50\,{\rm yr}/10^7\,{\rm yr}) \times 10^8 \approx 500$, assuming that the fraction of quasars associated with SMBHBs is

[8] J. Robnik et al. (2024) used a "null-signal template" as an alternative approach to more efficiently test periodic signals.

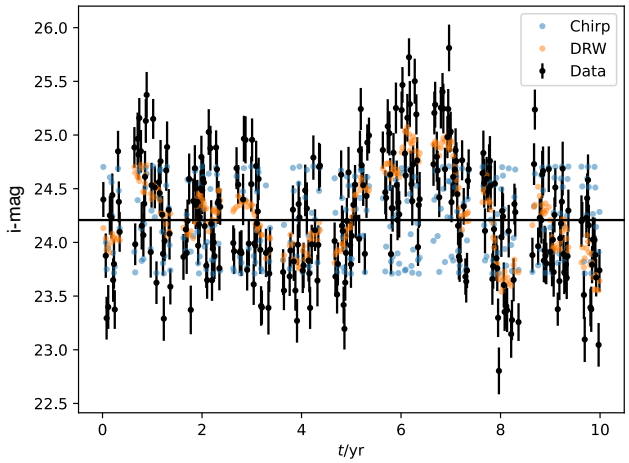


Figure 2. Example mock LSST lightcurve for an SMBHB with $M_{\rm bin} = 10^{7.8} M_\odot$ and $z = 2.9$. The lightcurve (black points) has a mean magnitude $m_i = 24.2$ per Equation (1) (line), modulated by the binary through relativistic DB with a chirping sinusoid, Equation (3) with $A = 0.5$ mag (blue points); additionally, the quasar varies stochastically following a DRW, Equation (5) with $\sigma = 0.5$ mag and $\tau = 136$ days (orange points); finally, the photometric errors have amplitudes of ∼0.2 mag (error bars). The cadence is 6 days, with seasonal gaps that skip 1/3 of observations each year.

$f_{\rm bin} \approx 1$. By the same token, we expect up to ∼150, 10^3, 10^4, and 10^5 binaries in LSST with times to merger of $t_m = 15$, 10^2, 10^3, and 10^4 yr, respectively. Note that the effects of gas drag on the inspiral are neglected, but may be important in some cases (see Section 4.1 for discussion).

In order to study different types of signals, we construct five sets of fiducial sources, each corresponding to a different choice of t_m in [15, 50, 10^2, 10^3, 10^4] yr. For each value of t_m, we draw 100 sources from the ($M_{\rm bin}$, z) distribution in Figure 1. The resulting sets of binaries represent a realistic expectation of LSST sources conditioned on each t_m. For simplicity, we assign all binaries a mass ratio of $q = m_2/m_1 = 0.1$, a commonly expected value for compact SMBHBs (K. Li et al. 2020, 2022).

2.2. Modeling Lightcurves

We next simulate $T = 10$ yr of LSST data for each of the selected binaries. Each lightcurve is centered around a mean luminosity m_i given by Equation (1), with modulations composed of (i) a DB chirp signal $y_{\rm chirp}$, (ii) a DRW signal $y_{\rm drw}$, and (iii) Gaussian photometric errors y_n, such that the overall data $y(t)$ are given by

$$y(t) = m_i + y_{\rm chirp}(t) + y_{\rm drw}(t) + y_n(t). \quad (2)$$

Note that the first two terms on the right-hand side are deterministic, and the last two are stochastic. We describe each component below, and present an example mock lightcurve in Figure 2, for which we can measure periodicity and chirp.

A distinct periodicity in quasar lightcurves containing compact SMBHBs is expected from hydrodynamical modulations on orbital timescales, which can persist all the way through merger. The relativistic DB modulation is associated with the Doppler brightening of the emission from the minidisks around each SMBH (D. J. D'Orazio et al. 2015; M. Charisi et al. 2022). The DB effect introduces a modulation in flux with an amplitude of order v/c, where v is the binary's orbital velocity. In this work, we adopt the DB model for binary variability in the LSST lightcurves. For simplicity, we assume that the SMBHB orbit is circular; this leads to DB modulations with a sinusoidal profile (D. J. D'Orazio et al. 2015). However, hydrodynamical effects (B. D. Farris et al. 2015; Y. Tang et al. 2018; J. Zrake et al. 2021; J. R. Westernacher-Schneider et al. 2022; F. Cocchiararo et al. 2024) and lensing (D. J. D'Orazio & R. Di Stefano 2018; A. Ingram et al. 2021; J. Davelaar & Z. Haiman 2022a, 2022b; L. M. Krauth et al. 2024) can also cause periodic variability in SMBHBs, notably with different (nonsinusoidal) shapes and amplitudes. See Section 4.1 for more details.

Assuming pure GW-driven inspiral during the last ∼10^4 yr of merger, the modulation phasing is then given by a simple chirp, which we evaluate to leading post-Newtonian order (e.g., C. Cutler & E. E. Flanagan 1994). Concretely, the observed (orbital) modulation frequency f is related to the source-frame GW frequency $f_{\rm gw}$ by $f = f_{\rm gw}/2(1+z)$; for a binary with chirp mass $\mathcal{M} = [q/(1+q)^2]^{3/5} M_{\rm bin}$, $f_{\rm gw}$ itself evolves following $\dot{f}_{\rm gw} \propto \mathcal{M}^{5/3} f_{\rm gw}^{11/3}$. We derive a reference value of the orbital frequency f_0 from the chosen time to merger through $t_m \propto [(1+z)\mathcal{M}]^{-5/3} f_0^{-8/3}$; this then determines the entire evolution $f(t)$. Finally, we assume that the modulation amplitude (A) remains constant over the observation period, so that the form of the chirp is

$$y_{\rm chirp}(t) = A \cos\left[\Phi(t; z, \mathcal{M}, t_m) - \phi\right], \quad (3)$$

where the *chirp phase* Φ is given by

$$\Phi(t; \mathcal{M}, t_m) = -\left(\frac{t_m - t}{5 t_{\mathcal{M}}}\right)^{5/8}, \quad (4)$$

where $t_{\mathcal{M}} \equiv (1+z) G\mathcal{M}/c^3$ and ϕ is the phase at $t = t_m$. Note that we define t_m, f, and $\dot{f}$ in the observer frame, but $\mathcal{M}$ and $f_{\rm gw}$ in the source frame.

In real signals, we expect the amplitude to increase as the binary approaches merger (D. J. D'Orazio et al. 2015; Z. Haiman 2017). Although we could build this effect into our model, we neglect it here for simplicity; an increasing amplitude should facilitate detection.

The blue points in Figure 2 show an example of this DB-induced signal for one of our binaries in the $t_m = 50$ yr set. The initial observed frequency is $f = 3.7 \times 10^{-7}$ Hz ($P = 1/f \approx 30$ days), which increases to $f \approx 4.2 \times 10^{-7}$ Hz ($P = 27.6$ d) over 10 yr; this corresponds to a "chirping rate" $\gamma \equiv \dot{f}\, T/f_0 \approx 7.5\%$. For our other systems, P ranges from ∼1.7 to 300 days, with most of them falling between ∼3 weeks and ∼3 months. The γ rates are 25%, 7.5% 3.75%, 0.375% and 0.0375% over 10 yr for $t_m =$ 15, 50, 100, 10^3, and 10^4 yr. These binaries span the majority of the frequency domain of GW-driven SMBHBs that can be identified in LSST.

In addition to the signal induced by orbital modulations, we expect the quasar to exhibit intrinsic stochastic variability. A commonly used model to describe this is a DRW (B. C. Kelly et al. 2009; C. L. MacLeod et al. 2010; Z. Ivezic & C. L. MacLeod 2013). Other quasar variability models have been investigated in recent years, including a "damped harmonic oscillator" (W. Yu et al. 2022). We follow XH24 to implement the DRW model. The characteristic amplitude SF_∞ and timescale τ of the DRW are empirically related to the BH mass, redshift, absolute i-band magnitude, and rest-frame wavelength as in Equation (7) in C. L. MacLeod et al. (2010) or Equation (11) in XH24. We model $y_{\rm drw}(t)$ as drawn from a

Gaussian process with an exponential kernel given by

$$K(t_1, t_2) = \sigma^2 \exp\left(-\frac{|t_1 - t_2|}{\tau}\right), \quad (5)$$

where $\sigma^2 = \mathrm{SF}_\infty^2/2$ (orange points in Figure 2).

Finally, we assume that the LSST observations are subject to some uncorrelated Gaussian measurement errors $y_n(t) \sim \mathcal{N}[0, \sigma_n^2(t)]$, where the variance of the Gaussian σ_n^2 is allowed to vary for each timestamp t. For each data point, we set σ_n randomly by drawing $\log \sigma_n \sim \mathcal{N}(\log(0.2), 0.01)$, consistent with LSST specifications (LSST Science Collaboration et al. 2009). We assume that LSST's i-band cadence allows one visit every 6 days; to account for irregularity in the observations, we mimic a nonuniform cadence by jittering observation times by $\delta t \sim \mathcal{N}(0, (3\,\mathrm{hr})^2)$. We also simulate expected "seasonal gaps" due to the position of the Sun, which results in $\sim 1/3$ of the data being blocked each year.

2.3. Bayesian Inference

We apply Bayesian inference to simultaneously measure the parameters of both the chirp and DRW from simulated data, as in Figure 2. Our likelihood is defined by the data-generating model of Equation (2), but with some key reparameterizations to facilitate sampling. We parameterize the post-Newtonian chirp in terms of the orbital frequency f_0 and frequency derivative $\dot{f}_0$ defined at the median LSST observation time, which are equivalent to the binary's $M_{\rm bin}$ and t_m. We set priors on f_0 and $\dot{f}_0$ based on the expected uncertainty in the observed binary mass, which reflect the range of Eddington ratios for luminous quasars: $\Delta \log(\mathcal{M}) \approx 1$ (J. A. Kollmeier et al. 2006; P. F. Hopkins & L. Hernquist 2009; C. S. Kochanek et al. 2012), which implies

$$\frac{\Delta f_0}{f_0} = -\frac{5}{11}\frac{\Delta \mathcal{M}}{\mathcal{M}} = 0.39, \quad (6)$$

$$\frac{\Delta \dot{f}_0}{\dot{f}_0} = \frac{5}{3}\frac{\Delta \mathcal{M}}{\mathcal{M}} = 1.43. \quad (7)$$

It follows that Δf is bound between $\min(f) = 0.61 f_t$ and $\max(f) = 1.39 f_t$, and $\Delta \dot{f}$ is bound between $\min(\dot{f}) = \max(0, -0.43\dot{f}_t)$ and $\max(\dot{f}) = 2.43\dot{f}_t$, where f_t and $\dot{f}_t$ are the true values of f_0 and $\dot{f}_0$ respectively. In general, binaries with higher orbital frequency will have wider priors, making the computation more costly. Besides f_0 and $\dot{f}_0$, we infer the chirp phase ϕ and amplitude A; the model reduces to a pure DRW for $A \to 0$. We parameterize the DRW itself by σ and τ, which are also measured from the data.

The chirp likelihood is challenging to sample due to degeneracies that lead to multiple false maxima; without dedicated treatment, this can make the sampling prohibitively inefficient. We tackle this through a custom strategy, which we dub the "frequency comb" method. This involves the creation of auxiliary f_0 and $\dot{f}_0$ grids that are used, together with analytic marginalization over A and ϕ, in a hybrid Gibbs and Metropolis–Hastings sampling scheme. We describe this, as well as our (generally broad) priors, in Appendix A.

The result of the inference is a set of posterior samples on $\{f_0, \dot{f}_0, A, \phi, \sigma, \tau, m_i\}$. This can answer whether (i) binary modulations are detectable ($A > 0$) and, if so, (ii) whether the chirp is also detectable ($\dot{f}_0 > 0$); we can also evaluate whether the binary parameters are recovered without bias and whether there are significant degeneracies among the binary parameters or between the noise and signal.

3. Results

We begin by applying our Bayesian analysis to the example lightcurve in Figure 2. Figure 3 shows the resulting posterior density, compared to the true injected values. It is apparent that the chirp is confidently detected, with both $P(A > 0)$ and $P(\dot{f}_0 > 0) > 99.98\%$, the greatest certainty achievable given our finite number of samples. In terms of the z-score, i.e., the ratio of the posterior mean to its standard deviation, we determine $A > 0$ at 29σ confidence and $\dot{f}_0 > 0$ at 64σ confidence—an unambiguous detection of the chirp. The signal is well recovered, and the posterior shows no correlations with the DRW sector.

Assuming a linear chirp ($\dot{f} = \mathrm{const}$), the contribution of the chirp term to the accumulated phase over an observation time T is $\Delta\Phi = \dot{f} T^2/2$. Basic scaling arguments suggest that the overall phase accuracy we can obtain in our measurement is $\sigma_\Phi \simeq \pi/\rho$, where ρ is the signal-to-noise ratio; in this example, $\rho = \sigma_A/A \simeq 30$. Ignoring correlations between $\dot{f}$ and other variables (f, A, etc.), appropriate for very loud signals, this suggests a back-of-envelope estimate of the uncertainty on $\dot{f}$ of $\sigma_{\dot{f}} \simeq 2\pi/\rho/T^2 \simeq 2.3 \times 10^{-8}\,\mathrm{d}^{-2}$ in this example. For an average $\dot{f} \simeq 1.2 \times 10^{-6}\,\mathrm{d}^{-2}$, this corresponds to an estimated z-score of $z_{\dot{f}} \simeq 52$, in close agreement with the actual measured value.

3.1. Effect of Merger Time

We next repeat the analysis on the 100 lightcurves simulated for each choice of t_m, still with true amplitude $A_t = 0.5$ mag. The most compact binaries considered here ($t_m = 15$ yr) have periods ranging from days to weeks, making them potential future sources for LISA; the widest binaries ($t_m = 10^4$ yr) have periods of years.

Figure 4 shows that in all cases we detect the binary signal with $>5\sigma$ (amplitude z-score $\gtrsim 5$), and typically 10–30σ confidence. The widest binary in this set whose signal we can detect (at 8σ) has $P_0 \approx 5$ yr ($f_0 \approx 6.2$ nHz, star in bottom panel of Figure 4). Its chirp rate is $\gamma = 3.75\%$, corresponding to an extremely slow $\dot{f} = 2.4 \times 10^{-3}$ nHz yr^{-1}; we can detect this PTA-like source with only ~ 2 visible cycles. As t_m increases, binaries with a given mass become wider and P_0 increases; this leads to a reduction in the number of observed cycles and a corresponding degradation in detection confidence as t_m increases from 15 to 10^4 yr (especially for $t_m = 10^3$ and 10^4 yr). On the other hand, ultrashort-period binaries with $P_0 \lesssim 10$ days can become difficult to identify due to aliasing, given LSST's ~ 6 day cadence, often degrading the z-score despite the large number of visible cycles (evident from the "jitter" in z-score for such short-period sources in, e.g., the $t_m = 100$ yr panel, at $z \approx 0$–1 and $\log(M_{\rm bin}/\mathrm{M}_\odot) \sim 6 - 7$). We expect that LSST's uneven sampling strategy (e.g., a pair of points in a single night) will ameliorate this issue. However, for $P_0 \gtrsim 30$ days, finite sampling is not an issue, and the z-scores increase with smaller P_0.

Besides detecting $A > 0$, we also accurately recover the orbital frequency f_0 and confidently detect a nonnegligible frequency evolution $\dot{f}_0 > 0$ for all sources in Figure 4 (see

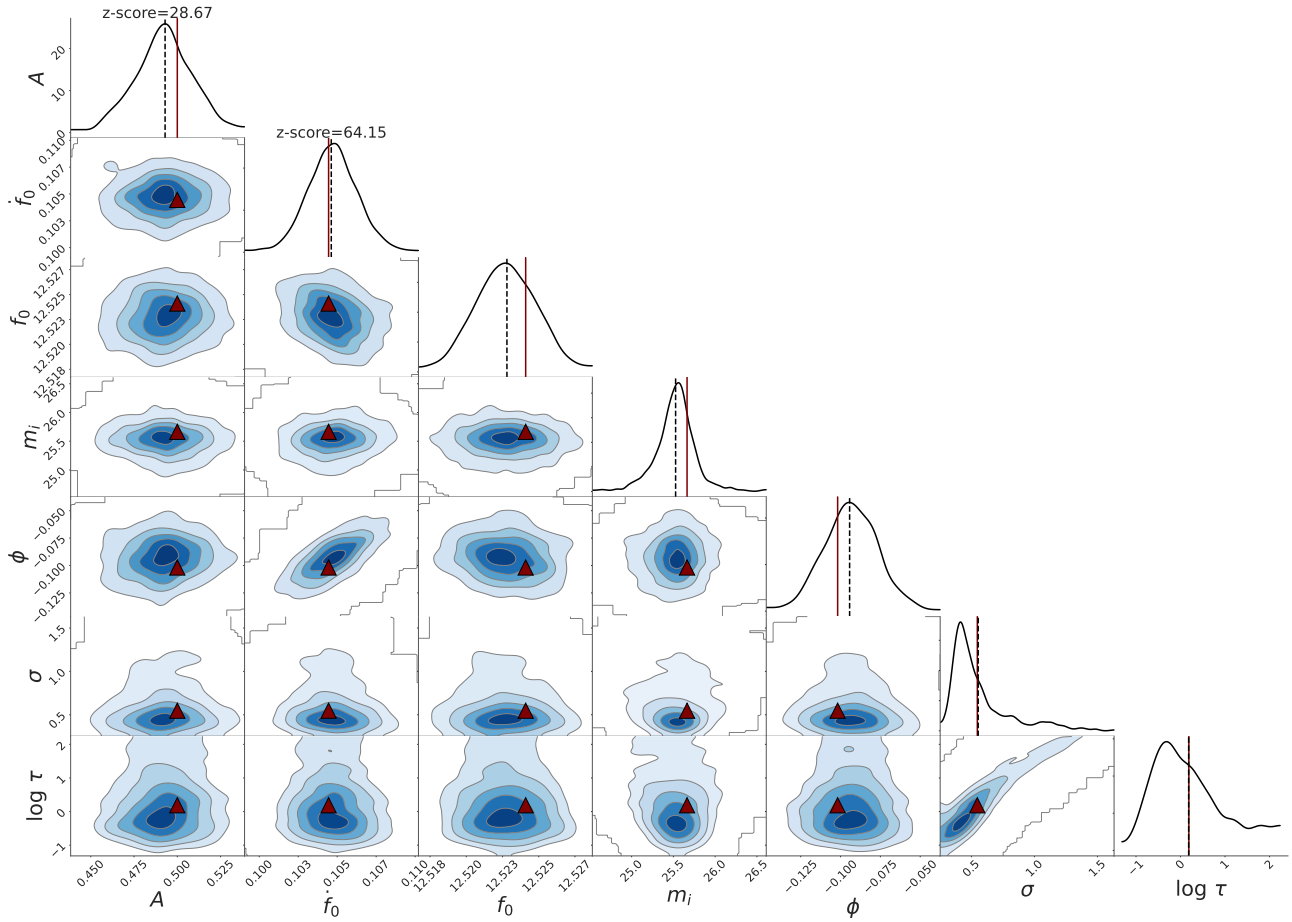


Figure 3. Posterior for all seven parameters for the binary lightcurve in Figure 2. Each contour encloses 10%, 30%, 50%, 70%, and 90% of the probability, respectively. We also show the truth (red triangles in the contours and red lines in the marginal panels) and the mean of the posterior (black lines in the marginal) for each parameter. The z-score values are shown for A and $\dot{f}_0$.

Appendix B). The slowest chirp we can detect with high confidence ($>5\sigma$) has $P_0 \approx 149$ days ($f_0 \approx 0.78$ nHz, diamond in middle-left panel of Figure 4). The z-score $= 6$ in $\dot{f}_0$ and 22 in A. Its chirp rate is $\gamma = 3.75\%$, corresponding to a slow chirp of $\dot{f} = 2.9$ nHz yr^{-1}.

3.2. Effect of Binary Amplitude

Thus far we have presented results assuming $A_t = 0.5$ mag. Both the DB model (D. J. D'Orazio et al. 2015) and hydrodynamical simulations (e.g., J. Zrake et al. 2021; J. R. Westernacher-Schneider et al. 2022) can induce a range of variability amplitudes $0.1 \lesssim A \lesssim 1$ mag. As A decreases relative to the DRW, our ability to recover the chirp should degrade. Here we examine different choices of true binary amplitude, in the $0.05 \text{ mag} \leqslant A_t \leqslant 0.5$ mag range, and study their detectability.

In Figure 5, we show the amplitude z-score versus A_t for binaries with $t_m = 50$ yr and varying orbital periods, 5 days $\lesssim P_0 \lesssim$ 200 days. As expected, detectability increases with the binary amplitude, reaching z-scores $\gtrsim 5$ for $A_t > 0.1$. The z-score also has a weak dependence on P_0, and is slightly higher for binaries with shorter P_0, consistent with the trends in Figure 4.

To further assess the detectability of the chirp, Figure 6 reports z-scores in $\dot{f}$, i.e., posterior mean($\dot{f}_0$)/std($\dot{f}_0$), for the same simulations as Figure 5. We find that we can measure the frequency growth with high confidence ($\geqslant 5\sigma$) for all but the weakest signals in this set. The weakest $\dot{f}_0$ we can detect corresponds to a binary with $P_0 \approx 200$ days ($f_0 = 47.9$ nHz) and $A_t \geqslant 0.3$ mag (red curve in zoom-in); this binary has $\gamma = 7.5\%$, and $\dot{f} = 0.43$ nHz yr^{-1}. We can measure $\dot{f}_0 > 0$ more confidently with decreasing P_0, achieving $\geqslant 5\sigma$ confidence for all amplitudes of interest ($A_t \geqslant 0.1$ mag), starting with $P_0 < 116$ days. This is expected from the number of visible cycles, per the $t_m = 50$ yr trend in Section 3.1.

3.3. Inferring Binary Properties

Physical properties of the SMBHBs, such as t_m and the redshifted chirp mass, $(1+z)\mathcal{M}$, can be derived from our inference of f_0 and $\dot{f}_0$ (Section 2.2). As for the signal parameters $\{A, \phi, f_0, \dot{f}_0\}$, the precision with which we can recover $(1+z)\mathcal{M}$ and t_m improves with signal-to-noise ratio, as proxied by A_t. In Figure 7, we show how the precision in both parameters improves as a function of binary amplitude A_t (left column), for the same binaries as in Figure 5. The uncertainty in the measured $(1+z)\mathcal{M}$ can be at best $\sim 10^4\, M_\odot$ (obtained for $P_0 = 5$ days binaries and $A_t > 0.15$). For some binaries ($A_t > 0.3$ and $P_0 < 200$ days), we can constrain $t_m \lesssim 5$ yr. Assuming conservatively that LISA is launched without significant delay in the mid-2030s, coinciding with the end of the 10 yr LSST observations, and that LISA operates for

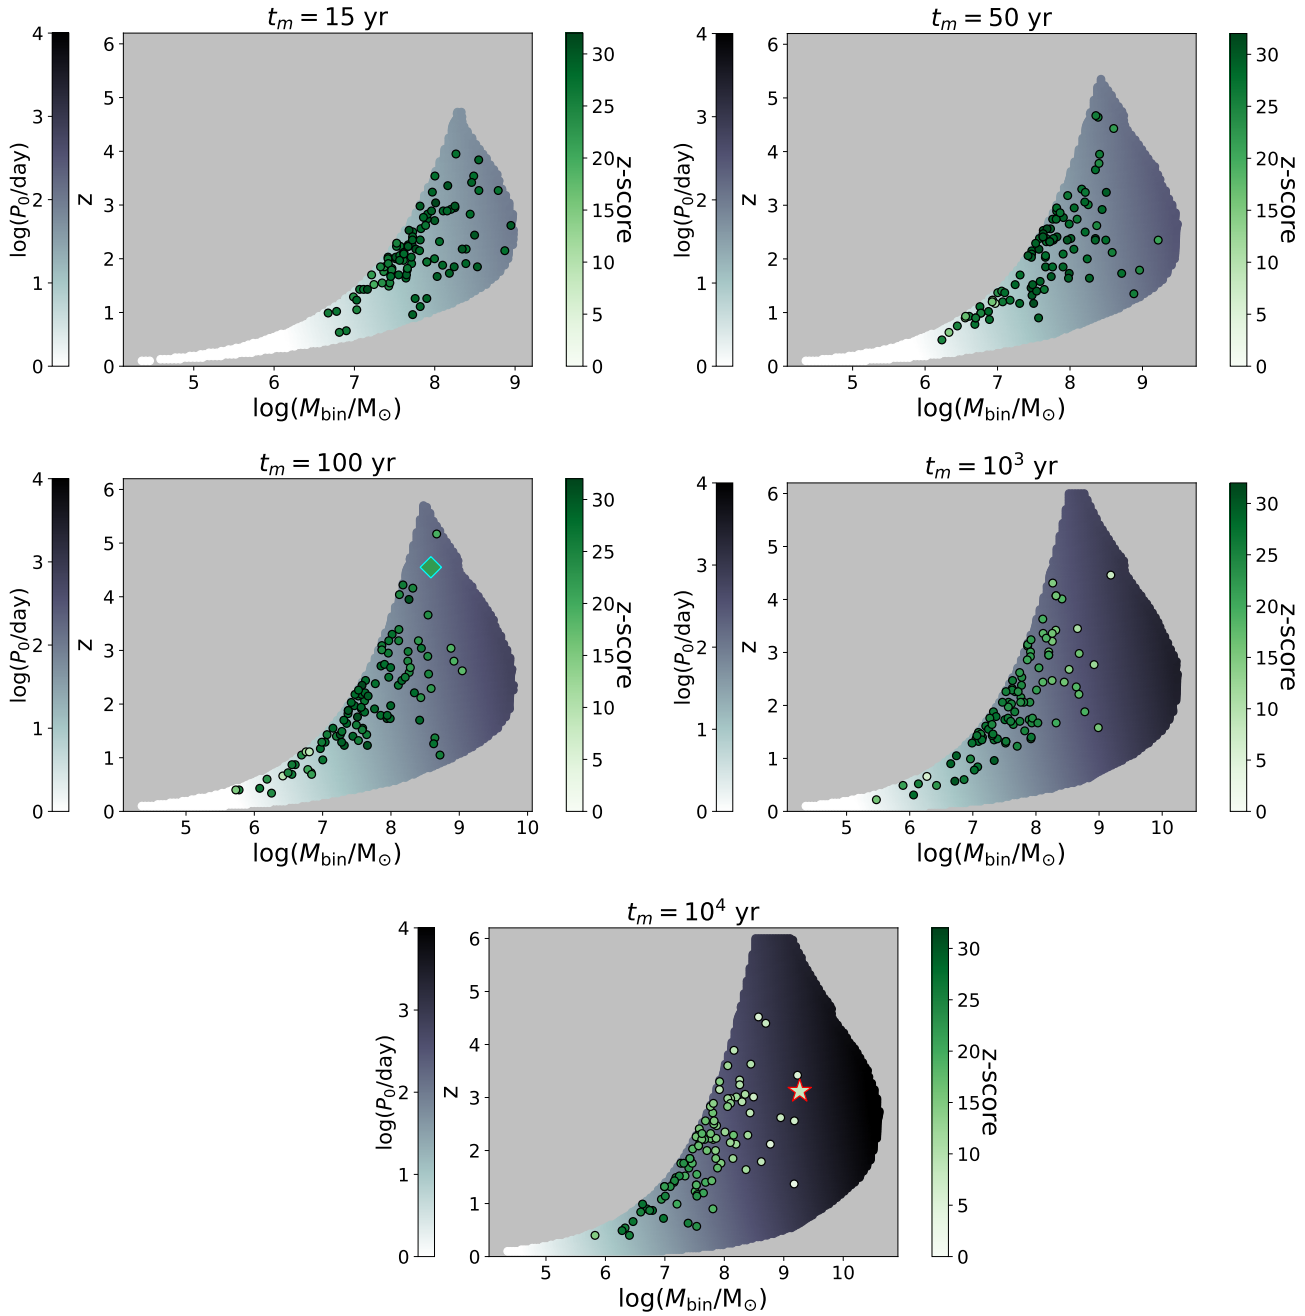


Figure 4. Summary of the inference on the 100 lightcurves simulated for $t_m = 15$, 50, 100, 10^3, and 10^4 yr (panels). We show posterior z-scores, mean(A)/std(A), for the chirp amplitude A (right color bar), as a function of binary mass (abscissa) and redshift (ordinate), as well as the corresponding observed period $P_0 = 1/f_0$ (left color bar). We fix the true amplitude to $A_t = 0.5$ mag. The z-score is a measure of how confidently the binary signal was detected, e.g., z-score = 5 implies a 5σ detection. The star highlights the widest binary in this set that we can detect with $>5\sigma$ confidence, while the diamond highlights the widest binary for which we can detect the chirp $\dot{f}_0 > 0$ with $>5\sigma$ confidence. We present a similar plot for the z-score of $\dot{f}_0$ in Appendix B.

5 yr, these binaries will merge during LISA's observational window. They can thus be flagged as excellent candidates for a targeted search with LISA prior to the start of the LISA observations; using the sky location, redshift, frequency, and mass observed from the EM observation, one can assess whether an inspiraling binary can be detectable in LISA before the mission starts. Additionally, early identification of LSST compact binaries can facilitate further EM follow-up observations, see Section 4.4.

Additionally, we show the accuracy of our measurements by plotting the deviation of the posterior mean from the truth relative to the measurement precision (middle column), i.e.,

$$\mathcal{S}[x] = \frac{\text{mean}(x) - x_t}{\text{std}(x)}, \tag{8}$$

for $x = (1+z)\mathcal{M}$ or t_m. The mean displacements from the truth are typically $<2\sigma$ irrespective of A_t and P_0, and they are roughly equally spread in the positive and negative directions, suggesting our measurements are reasonably unbiased. The exceptions occur at minimum $A_t = 0.05$, where $\mathcal{S}[x]$ plunges to ~ 0 for both $(1+z)\mathcal{M}$ and t_m. This is because those binary

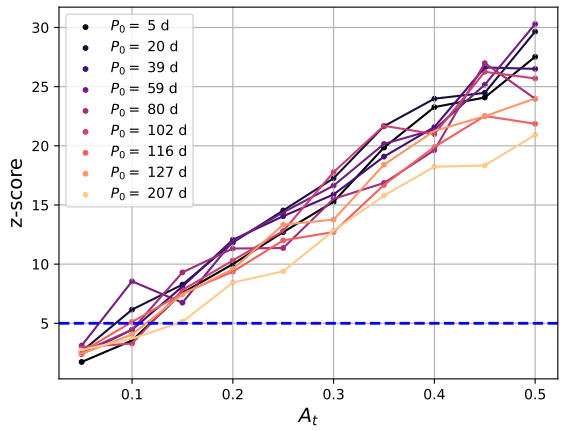


Figure 5. The detectability of binary chirp amplitude (A) as a function of the true injected amplitude A_t, for nine binaries in our sample. The z-score increases as A_t increases. It has a weak dependence on orbital periods (quoted in the legend). The chirp rate is fixed at 7.5% in all curves ($t_m = 50$ yr). The detection is generally significant; z-score > 5 (see blue horizontal line) when $A_t \gtrsim 0.1$.

signals are too weak to be detected: their posteriors are close to the prior, whose mean lands at zero, close to the truth relative to the wide variance. Finally, in the right column, we show the fractional error (defined in Equation (9)) for measuring $(1+z)\mathcal{M}$ and t_m.

$$\mathcal{D}[x] = \frac{|\mathrm{mean}(x) - x_t|}{x_t}. \quad (9)$$

For most binaries, both $(1+z)\mathcal{M}$ and t_m can be measured with $\lesssim$10% fractional errors, except the widest binaries ($P_0 = 200$ days) and the smallest amplitudes ($A < 0.2$).

3.4. LSST Multiband Data

To showcase the effects of realistic LSST cadence and photometric errors, we use the Rubin simulator tool (LSST Science Collaboration & Rubin Observatory Project 2023) to regenerate our fiducial mock lightcurve (Figure 2) using the simulated i-band cadence and photometric uncertainties (see Figure 8). During each seasonal observing window, the cadence ranges from 3 to 12 days, with a more irregular distribution compared to the uniform 6 day cadence used in our original mock lightcurves. The average photometric errors are 0.01 mag, considerably smaller than the values adopted elsewhere in this work ($\sim$0.2 mag). Note that the error bars in Figure 8 are so small that they are not visible by eye.

Using this example, we investigate how both cadence and photometric errors affect the confidence in recovering the binary parameters. We first apply the new cadence and errors to the fiducial lightcurve shown in Figure 8, which has an injected binary amplitude of 0.5 mag, an observed period of 30 days, and a time-to-merger of 50 yr. Recomputing the z-scores for measuring the amplitude A and chirp $\dot{f}$ yields values of $\sim$27 and $\sim$72, respectively, similar to the previous results of $\sim$29 and $\sim$64 and still well above the 5σ threshold.

Next, we examine the impact of the realistic cadence and smaller photometric errors on a binary with a lower injected amplitude, $A = 0.1$ mag. We test this case for a binary with $P_0 \sim 40$ days, for which the lower amplitude previously produced a z-score near the 5σ threshold (see Figure 5). The new z-scores for A and $\dot{f}$ are 5.7 and 6.8, showing no significant improvement.

Finally, we test the improved cadence and photometric errors on the two wide binaries shown in Figure 4. For the "Diamond" binary—previously detected at $>5\sigma$ in both amplitude and $\dot{f}$—the confidence for detecting A increases (with the z-score rising from $\sim$20 to $\sim$28), while the z-score for $\dot{f}$ shows only a mild increase (from 6.3 to 6.8). Similarly, the "Star" binary yields nearly unchanged statistics for measuring A, with a new z-score of 5.5.

Overall, adopting the realistic LSST cadence and photometric errors does not significantly affect the detectability of binaries.

However, LSST observes six optical bands on a rolling basis, with each band revisited every $\sim$2 weeks (although some can be visited more frequently, LSST Science Collaboration et al. 2009). In this work, we only consider the i-band, which can have a cadence as short as 6 days. Our method performs well on chirp signals for ultracompact binaries that have periods down to $P_0 \sim 1$ day, which is a few times above the Nyquist frequency. However, as we discussed in Section 3.1, some shorter-period binaries experience aliasing systematics when the orbital periods are closer to Nyquist. Combining different bands and shortening the effective cadence can potentially improve our statistics for short-period sources. This can be achieved assuming constant color correlations between each LSST band, and combining all bands can reduce the effective cadence to as low as $\sim$3 days.

4. Summary and Discussion

In this paper, we have developed a Bayesian inference framework to measure seven parameters associated with a GW-driven SMBHB chirp signal in LSST quasar lightcurves. We find that chirps from a GW-driven inspiral are typically detectable (both in amplitude and frequency derivative) with high confidence for realistic choices of the SMBHB parameters. The detection significance depends on the signal's amplitude and period, which control the signal-to-noise ratio and the chirp rate. The widest system for which we can detect the chirp signal (A) has an orbital period of $P_0 \approx 1850$ days, and chirps extremely slowly with $\dot{f}_0 \approx 2.3 \times 10^{-3}$ nHz yr^{-1}; the widest binary for which we can measure chirp ($\dot{f}_0 \neq 0$) *well* is one with $P_0 \approx 200$ days and $\dot{f} = 0.43$ nHz yr^{-1}.

Detecting a chirp represents smoking-gun evidence for a compact SMBHB. Some of the SMBHs detectable by LSST will also be potential LISA sources (day to week periods). Identifying LISA binaries as chirping quasars in LSST can place important empirical constraints on the expected binary merger rates before LISA, and enable multimessenger observations once LISA launches. We expect at least $\sim$145,000 compact SMBHBs in LSST with orbital periods of $P = 1850$ days or shorter, or with a chirp of 2.3×10^{-3} nHz yr^{-1} or faster. Our results show that we should be able to measure nonzero chirp from all of these binaries, assuming typical amplitude variations of $\sim$0.1–0.5 mag.

According to our assumption that a large fraction of luminous quasars are associated with SMBHBs, we expect that the sources making up the stochastic GW background (GWB) detected by PTAs will be in the LSST quasar catalog. Á. Kis-Tóth et al. (2025) showed that the GWB in the nHz bands is dominated by compact $\sim 10^9\,M_\odot$ binaries at $z = 2$–3 (see their Figure 5), with a nonnegligible contribution from wider binaries with masses as low as $\sim 10^6\,M_\odot$. Our results

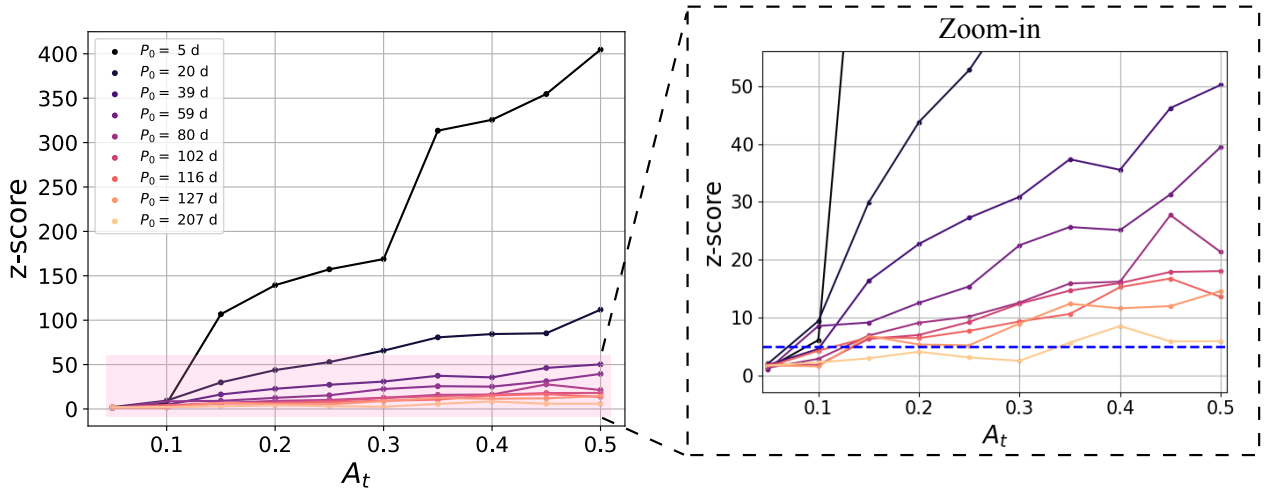


Figure 6. The significance at which the frequency chirp $\dot{f}_0$ can be measured as a function of variability amplitude A_t (abscissa) and period P_0 (color). The y-axis shows the posterior z-score for $\dot{f}_0$, i.e., mean($\dot{f}_0$)/std($\dot{f}_0$), calculated for the same set of A_t and P_0 as in Figure 5, i.e., with chirp rate $\gamma = 7.5\%$ ($t_m = 50$ yr). We achieve overall large z-scores, meaning that $\dot{f}_0$ can be inferred at high confidence, especially when $A_t \gtrsim 0.2$. The z-score is higher when P_0 is shorter, because more cycles are observed and the period changes more rapidly.

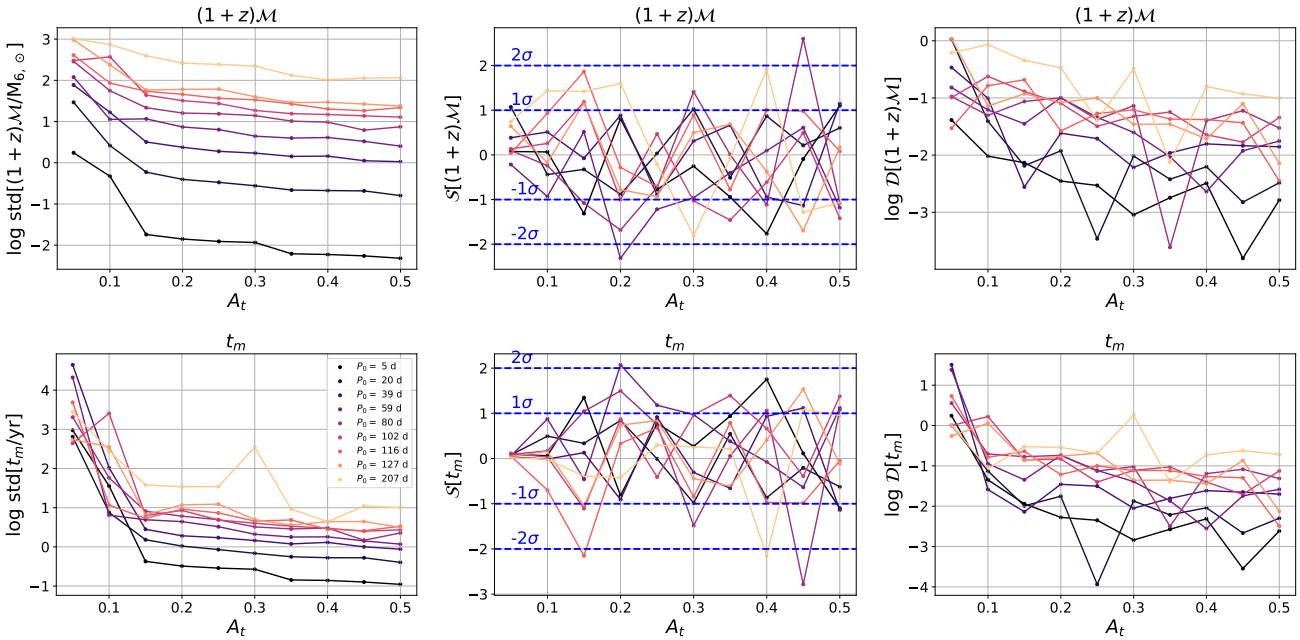


Figure 7. Precision and accuracy with which we can measure the observed chirp mass $(1+z)\mathcal{M}$ (top left panel) and time-to-merger t_m (bottom left panel), as a function of the true binary amplitude A_t (abscissa) and period P_0 (color). The left column shows precision through log-std(x) in units of $M_{6,\odot} = 10^6\,M_\odot$ for $[(1+z)\mathcal{M}]$ and yr for $[t_m]$. The middle column shows accuracy through the posterior mean distance in units of standard deviation, as in Equation (8); dashed lines mark 1σ and 2σ deviations. The right column shows the fractional error as defined in Equation (9).

imply that the bulk of the sources of the GWB will be possible to identify from EM observations in LSST alone.

Our method represents a significant improvement over standard techniques, like the Lomb–Scargle periodogram, which has difficulty with chirping signals. It achieves this at an acceptable computational cost, taking typically 10 minutes per lightcurve (Appendix C). Since our method can detect chirps across a wide range of binaries, with time-to-merger up to 10^4 yr, existing SMBHB candidates are suitable subjects for our study. In future work, we will search for periodic quasars in current time-domain surveys, including CRTS and PTF, and revisit existing candidates.

However, our analysis is subject to some key assumptions, including the binary variability shape and the noise model. In this section, we emphasize future steps to improve our analysis, particularly by incorporating other binary models, quantifying any false alarm rates, and using realistic LSST data.

4.1. Beyond the Sinusoid Profile

A notable SMBHB candidate, PG1302-102, has established periodicity that may be explained by DB modulations, both in its sinusoidal variability observed in the optical band

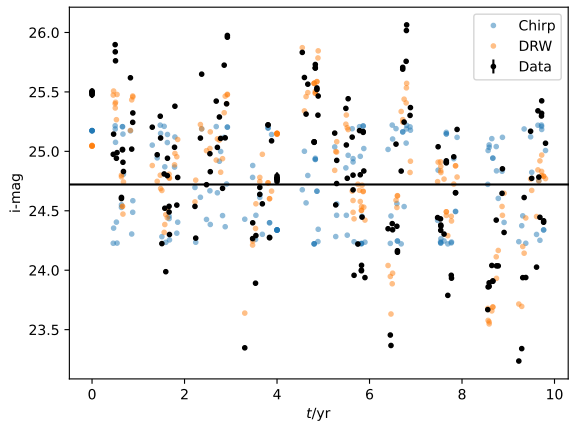


Figure 8. Mock LSST lightcurve shown in Figure 2 with realistic LSST cadence and photometric errors.

(M. J. Graham et al. 2015b) and in its features observed in multiwavelength studies (D. J. D'Orazio et al. 2015; C. Xin et al. 2020). Beyond the simple sinusoidal model motivated by DB variability for a circular orbit, the periodic variability of an SMBHB can have more complex profiles due to several effects: DB in eccentric orbits, hydrodynamical effects (B. D. Farris et al. 2015; J. Zrake et al. 2021; J. R. Westernacher-Schneider et al. 2022; F. Cocchiararo et al. 2024), self-lensing flares (D. J. D'Orazio & R. Di Stefano 2018; A. Ingram et al. 2021; J. Davelaar & Z. Haiman 2022a, 2022b; L. M. Krauth et al. 2024), or, more realistically, a combination thereof (D. J. D'Orazio et al. 2024).

Hydrodynamical simulations have shown that thermal emission from the disk can produce a saw-toothed or spikey periodicity. Self-lensing flares occur when the binary orbit is aligned with the line-of-sight within the Einstein radius, typically producing spikey-shaped periodicity as one SMBH (in the foreground) magnifies the emission from the second (background) SMBH during its passage. All these types of variability are subject to chirping during the GW-driven inspiral stage, making the chirp a robust observable.

We will extend our Bayesian framework to incorporate these different profiles, in particular, the saw-toothed and spike pulse shapes. This is similar to the work by M. C. Davis et al. (2024) and K. Park et al. (2025), who have conducted population synthesis of saw-toothed lightcurves and self-lensing flare lightcurves in LSST (however, both without chirping frequency). D. J. D'Orazio et al. (2024) have developed a public tool (`binlite`) to fit lightcurves, using LSST specifications, with a combination of DB, hydrodynamical, and lensing models. In future work, we will apply our method to such complex periodic and chirping lightcurves.

In our work, for simplicity, we assume that SMBHBs are inspiraling due to GW emission alone. It is possible that the circumbinary gas, which we are assuming produces the periodic EM emission, influences the inspiral rate. The analytic models in Z. Haiman et al. (2009), as well as recent hydrodynamical simulations (A. J. Dittmann et al. 2023; M. Garg et al. 2025; D. O'Neill et al. 2025), imply that compact binaries, even inside the LISA band, can still be coupled to their circumbinary disks, especially if the disk viscosity is high and/or the binary is retrograde with respect to the disk. In this case, the inspiral is generally expected to be accelerated compared to the pure GW-driven case, which means that it could be more observable than our results imply. On the other hand, since these gaseous effects are uncertain, additional free parameters will be required in the LSST lightcurve analyses to allow for these uncertainties. We leave investigations of these issues to a follow-up study.

4.2. False Positives from Nonchirp Variations

In this work, we follow the chirp template as predicted by the GW frequency evolution. However, other forms of chirplike frequency evolution might arise from the lightcurves that are not caused by a binary SMBH. For example, the only reported binary candidate with an observed rate of change in its orbital period is the so-called "tick-tock" binary (N. Jiang et al. 2022). However, M. Dotti et al. (2023) have ruled out "tick-tock" as a chirping SMBHB in optical follow-ups; they have found that its frequency change might be explained by AGN-disk precession around a single SMBH. For this reason, we will implement alternative templates for $\dot{f}$ that do not follow the GW frequency, and we will assess whether or not our Bayesian model can distinguish them from true GW chirps. These templates are (i) an arbitrary nonzero constant $\dot{f}_0$ deviating from $\dot{f}_{\rm gw}$, or (ii) setting the higher order ($\mathcal{O}(\geqslant 2)$) derivatives of $\dot{f}_{\rm gw}$ to be zero. Additionally, the circumbinary disk can drive an SMBHB to mildly eccentric ($e \sim 0.4$–0.5; J. Zrake et al. 2021; M. Siwek et al. 2023) or highly eccentric (C. Tiede & D. J. D. Orazio 2024) orbits, resulting in a different GW chirp compared to the circular case (P. C. Peters & J. Mathews 1963).

4.3. Stochastic AGN Variability and False Alarm Rate

We have adopted the commonly used DRW model to describe the stochastic variability of AGN in the optical band (B. C. Kelly et al. 2009; C. L. MacLeod et al. 2010) in our study. S. Vaughan et al. (2016) have shown that pure DRW noise can mimic a periodic signal, given a sufficiently large sample of pure DRW lightcurves. We emphasize that even though we produce only one DRW realization per lightcurve, our posterior statistics already demonstrate the probabilities that the truth can be found inside a certain credible region (see Figure 3); generating more DRW realizations will give very similar results.

However, we assume that in real life, the binary and noise variability conform to our model—DB + DRW. Under these assumptions, we detect a wide range of binary signals (i.e., in Figure 4) with very high confidence; z-scores are between $\sim$8–30, corresponding to extremely small false alarm probability of $\lesssim 10^{-16}$. On the other hand, both the binary model (e.g., discussion in Section 4.1) and the noise model are inherently uncertain, and these uncertainties will dominate the z-score, precluding such low false alarm probabilities. Various works have demonstrated the shortcomings of DRW in fitting the high-frequency end of quasar variability (e.g., R. F. Mushotzky et al. 2011; K. L. Smith et al. 2018) due to the sparseness of the AGN data used in the original studies. Alternative models have been proposed, including the so-called "damped harmonic oscillator" (W. Yu et al. 2022), which has found different correlations between the timescale and amplitude of AGN noise and the physical parameters (BH mass, luminosity, and wavelength).

Although we do not explore these other noise models here, we note that, if the observation baseline is long enough ($T \gg \tau$), and there are many chirp cycles in the observation,

then the precise description of the noise model should not matter much (the chirp and noise signals become uncorrelated). Previous studies analyze data from higher-mass quasars, with masses $M_{\rm bh} = 10^{8-10} M_{\odot}$, due to the observing limitations of the AGN lightcurves. With LSST, we expect low-mass AGNs, down to $M_{\rm bh} \sim 10^5 M_{\odot}$, to be the most common. Therefore, detections of low-mass AGNs with LSST will be essential for understanding AGN intrinsic variability.

We plan to investigate alternative noise models, including using realistic AGN lightcurves. However, low-mass AGN observations would be needed to empirically determine realistic noise for a significant fraction of binaries considered in our study (currently unavailable).

4.4. Follow-up Observations of LSST Binaries

SMBHBs with confirmed chirps in LSST are ideal targets for the upcoming GW missions such as LISA and the PTAs, since they can yield GWs in the mHz to nHz range. Once identified in LSST, the GWs of these binaries can be individually searched for, which can persist for $\sim$ a few years (in LISA) to thousands of years (in PTAs) before the binaries merge, during which we have simultaneous EM and GW observations for multimessenger studies. It is likely that binaries with periods $\lesssim$1 month can be identified within $\sim$1 yr of LSST observations (with $\sim$10 cycles in the lightcurves), which can facilitate GW and/or EM follow-ups shortly after their identification. Although we perform our analysis using the optical data,[9] we can apply our methods to observations in other wavelengths such as UV and X-rays. In particular, the DB predicts that the UV and X-ray lightcurves should vary in tandem with the optical lightcurve, and the amplitude boosts between bands are proportional to the ratio between the power-spectral slopes (D. J. D'Orazio et al. 2015). In general, the UV/X-ray emission from the disks is amplified as the binary approaches merger, since they come from the compact inner-disk regions near the BHs. As a result, periodic UV/X-ray signatures are expected, especially from DB and self-lensing, in addition to the optical signature. Additionally, hydrodynamical variability observed in the optical and infrared bands is also expected to appear in the UV, especially from the minidisks, all the way through merger (L. M. Krauth et al. 2023; F. Cocchiararo et al. 2024; L. M. Krauth et al. 2024). Toward the binary decoupling phase, the binary lightcurve is expected to exhibit unique "drop-off" signatures in both UV and X-ray (L. M. Krauth et al. 2023; A. J. Dittmann et al. 2023). This signature can be searched for in UV and X-ray time domain as a follow-up to LSST, or after the GW detection, e.g., by LISA (XH24), as further evidence to support the binary. Upcoming time-domain surveys such as the Ultraviolet Explorer (UVEX; S. R. Kulkarni et al. 2021), the Ultraviolet Transient Astronomy Satellite Ultrasat (ULTRASAT; Y. Shvartzvald et al. 2024), and the Athena+ X-ray Observatory (K. Nandra et al. 2013; D. Barret et al. 2020) will begin within a few years, and they can perform targeted observations of the chirping binaries that we have identified using LSST. Additionally, spectroscopic follow-up surveys of LSST AGNs allow observations of spectral features in compact SMBHBs with subparsec separations. For example, compact and accreting SMBHBs can exhibit truncations or peculiarities in their spectra due to central cavities in their circumbinary disks (H. Guo et al. 2020). The high-energy corona near the SMBHBs can also produce hard X-ray spectra that are different from those of single SMBHs (M. L. Saade et al. 2024). Early photometric identification of compact SMBHBs provides candidates for further spectroscopic examination of their features. Finally, if the SMBHB resides in a radio-loud and bright AGN, then its radio signatures can be searched for using very long baseline interferometer (VLBI) instruments, such as the next-generation Very Large Array (ngVLA; S. Burke-Spolaor et al. 2018) and Event Horizon Telescope (EHT; Event Horizon Telescope Collaboration et al. 2019a, 2019b). These instruments can directly probe the binary's orbital motion and jet evolution via targeted observations. Since these surveys require long-term monitoring, they are not suitable for surveying a large number of AGNs. Therefore, the promising binaries that we identify in LSST are valuable candidates for VLA/EHT follow-ups.

[9] Our analysis can be applied to other optical time-domain surveys, such as the Nancy Roman Telescope, CRTS, ZTF, etc.

Acknowledgments

C.X. thanks the Department of Astronomy of Columbia University, where the majority of this work was performed. Z.H. thanks the Center for Computational Astrophysics (CCA) at the Flatiron Institute for their hospitality during a sabbatical stay, where this work was initiated. The Flatiron Institute is a division of the Simons Foundation. This work was supported in part by NASA's LISA Preparatory Science Program (LPS) through grant 80NSSC24K0440 and by NASA Astrophysics Theory Program (ATP) grant 80NSSC22K0822. The authors thank the anonymous referee for helpful feedback, which improved this manuscript. We also thank Prof. Brian D. Metzger, Prof. David Schiminovich, and Prof. Kishalay De for valuable discussions.

Software: NUMPYRO (E. Bingham et al. 2019; D. Phan et al. 2019); CELERITE2 (D. Foreman-Mackey et al. 2017; D. Foreman-Mackey 2018); PYTHON (G. Van Rossum & F. L. Drake 2009); THE SOFTWARE CITATION STATION (T. Wagg & F. S. Broekgaarden 2024; T. Wagg et al. 2024).

Appendix A
Bayesian Algorithm

The probabilistic model that defines our inference is as described in Section 2.3 with the following specifications. As in Equation (2), we assume that the data are given by both deterministic and stochastic contributions. The deterministic component is comprised of the mean m_i and the DB chirp $y_{\rm chirp}(t)$, as in Equation (3). The stochastic component is given by a DRW $y_{\rm drw}(t)$ with exponential spectrum as in Equation (5), as well as uncorrelated, zero-mean Gaussian noise with variances $\sigma^2_{n,k}$ for each timestamp t_k due to photometric errors. Conditional on a choice of $\sigma^2_{n,k}$'s as well as the DRW amplitude σ and correlation length τ, the likelihood is thus a zero-mean Gaussian on the residuals,

$$r(t) = y(t) - m_i - y_{\rm chirp}(t; A, \phi, f_0, \dot{f}_0), \tag{A1}$$

with covariance matrix $C(\sigma, \tau;\sigma_n)$ given by the DRW spectrum of Equation (5), with diagonals enhanced by $\sigma^2_{n,k}$ (C. E. Rasmussen & C. K. I. Williams 2006). Equivalently, for N observations, the data-generation process is given by an

Table 1
Prior Distributions for All Explicitly Sampled Quantities; Definitions Clarify Inputs to Each Prior[a]

Parameter	Meaning	Prior Distribution	Definitions
$\log\sigma$	Log DRW amplitude	$\mathcal{N}(\log\sigma_y,\ 4)$	σ_y = std(y) lightcurve standard deviation
$\log\tau$	Log DRW timescale	$\mathcal{U}(\log\Delta t,\ \log T)$	Δt = median cadence, T = total span
$w_0 = 2\pi f_0$	Anchor for ω_0 grid	$\mathcal{N}_{[0,\infty)}(\omega_{\rm min},\ \Delta\omega^2)$	$\omega_{\rm min}$ = min. frequency, $\Delta\omega$=grid spacing
$\dot{w}_0 = 2\pi\dot{f}_0$	Anchor for $\dot{\omega}_0$ grid	$\mathcal{N}_{[0,\infty)}(\dot{\omega}_{\rm min},\ \Delta\dot{\omega}^2)$	$\dot{\omega}_{\rm min}$ = min. freq. deriv., $\Delta\dot{\omega}$=grid spacing
n_ω	Index for $\omega_0 = w_0 + n_\omega\,\Delta\omega$	$\mathcal{U}\{0,\ N_\omega - 1\}$	$N_\omega = \lfloor(\omega_{\rm max} - \omega_{\rm min})/\Delta\omega\rfloor$
$n_{\dot\omega}$	Index for $\dot{\omega}_0 = \dot{w}_0 + n_{\dot\omega}\,\Delta\dot{\omega}$	$\mathcal{U}\{0,\ N_{\dot\omega} - 1\}$	$N_{\dot\omega} = \lfloor(\dot{\omega}_{\rm max} - \dot{\omega}_{\rm min})/\Delta\dot{\omega}\rfloor$
a	Chirp cosine coefficient ($a = A\cos\phi$)	$\mathcal{N}(0,\ \sigma_y^2)$	σ_y = std(y)
b	Chirp sine coefficient ($b = A\sin\phi$)	$\mathcal{N}(0,\ \sigma_y^2)$	σ_y = std(y)
m_i	Mean flux	$\mathcal{N}(\mu_y,\ \sigma_y^2)$	μ_y = mean(y), σ_y = std(y)

Note.
[a] $\mathcal{N}(x,\ y)$ is a normal with mean x and variance y, $\mathcal{U}(x,\ y)$ is a uniform distribution between x and y, $\mathcal{N}_{[0,\infty)}(x,\ y)$ is a normal distribution truncated to the positive half-line, and $\mathcal{U}\{j,\ k\}$is a discrete uniform distribution over the integers from j to k, inclusive.

N-dimensional Gaussian,

$$y \sim \mathcal{N}[m_i + y_{\rm chirp}(A,\ \phi, f_0, \dot{f}_0),\ C(\sigma,\ \tau;\ \sigma_n)] \tag{A2}$$

where the covariance matrix C is $N \times N$ and y, $y_{\rm chirp}$, and σ_n are all of length N. We evaluate the DRW likelihood function, $p(y|m_i, A, \phi, f_0, \dot{f}_0, \sigma, \tau) = \mathcal{N}[y_{\rm chirp} + m_i, C(\sigma, \tau; \sigma_n)](y)$, in a numerically efficient way using CELERITE2 (D. Foreman-Mackey et al. 2017; D. Foreman-Mackey 2018).

Together with a choice of prior, the likelihood of Equation (A2) defines the posterior over the seven free parameters in our model, i.e., $\{f_0, \dot{f}_0, A, \phi, \sigma, \tau, m_i\}$, or any reparameterization thereof. We use NUMPYRO (D. Phan et al. 2019; E. Bingham et al. 2019) to draw from the 7D posterior through a hybrid Hamiltonian Monte Carlo (HMC) and Gibbs sampling strategy, together with analytic marginalization over linear parameters. This strategy is designed to address two major challenges:

1. The chirp phase, frequency, and frequency derivative can be near degenerate over short periods of observation.
2. It is inefficient to directly sample the chirp amplitude A and phase ϕ as written in Equation (3).

We develop a bespoke solution for the former, and apply a preexisting linearization and factorization strategy for the latter. We detail this next, describing our priors along the way (summarized in Table 1).

A.1. Covariance Parameters

For each quasar, we assume known magnitudes of photometric errors given by the same set of $\sigma_{n,k}$'s used to generate the synthetic lightcurve. Before running the inference, we generate these by drawing $\log\sigma_{n,k} \sim \mathcal{N}(\log(0.2),\ 0.01)$, where $\mathcal{N}(x,\ y)$ is a normal distribution with mean x and variance y. Assuming N observations with timestamps t_k, this defines a $N \times N$ diagonal covariance matrix with entries $\sigma_{n,k}^2$, that we take as known and fixed during sampling.

The covariance representing the uncertainty from photometric errors is added to the diagonal of the DRW covariance matrix representing the stochasticity intrinsic to the quasar. This matrix is constructed at each sampling step contingent on a choice of DRW parameters σ and τ, per Equation (5). We draw σ and τ respectively from a lognormal and a uniform prior distributions, following

$$\log\sigma \sim \mathcal{N}(\log\sigma_y,\ 4), \tag{A3}$$

$$\log\tau \sim \mathcal{U}(\log\Delta t,\ \log T), \tag{A4}$$

where $\mathcal{U}(x,\ y)$ is a continuous uniform distribution between x and y, σ_y = std(y) is the standard deviation of the observed lightcurve, Δt is the median separation between timestamps ($\sim$6 days) and T is the length of the observation period (10 yr). We draw σ and τ at each HMC step, together with auxiliary quantities that we use to sample f_0 and $\dot{f}_0$, as we describe next.

A.2. Nonlinear Parameters and the Frequency Comb

For a given choice of σ and τ, the conditional likelihood of any $\{f_0, \dot{f}_0, A, \phi\}$ is given by Equation (A2). However, this likelihood is oscillatory as a function of f_0 and $\dot{f}_0$ (or, equivalently, $\omega_0 = 2\pi f_0$ and $\dot{\omega}_0 = 2\pi\dot{f}_0$), presenting several false maxima (sidebands) around the true peak.

The sidebands occur at semiperiodic intervals with separation close to $\Delta\omega_0 \approx 7.72525/T$ in f and $\Delta\dot{\omega}_0 \approx 15.121/T^2$. To understand this spacing, first consider the log likelihood for a periodic signal with angular frequency ω fit by a model with varying angular frequency $\omega + \Delta\omega$ in white noise observed for a time T. The log likelihood $\log\mathcal{L}$ is then proportional to

$$\log\mathcal{L}_\omega \propto -\frac{1}{T}\int_0^T {\rm d}t\ [\cos(\omega t) - \cos \times ((\omega + \Delta\omega)t)]^2 \simeq \frac{\sin(\Delta\omega T)}{\Delta\omega T} - 1, \tag{A5}$$

where we have retained only the leading order term for small $\Delta\omega$. This has a global maximum at $\Delta\omega = 0$ (naturally), and an infinite sequence of local maxima. The first such secondary local maximum occurs at $\Delta\omega \approx 7.72525/T$, which is the grid spacing we use in our analysis.

Similarly, considering a pure chirp signal with frequency derivative $\dot{\omega}$ fit by a signal with varying frequency derivative

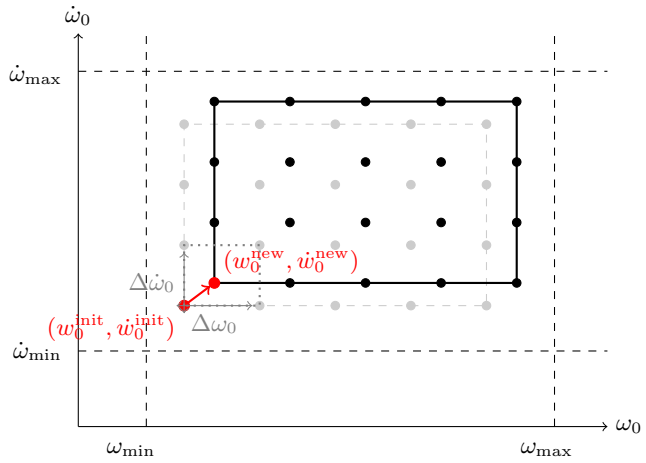


Figure 9. Illustration of the frequency comb construction in frequencies and frequency derivative space $(\omega_0, \dot{\omega}_0)$. The HMC sampler first picks a lower-left corner $(w_0^{\rm init}, \dot{w}_0^{\rm init})$ near $(\omega_{\min}, \dot{\omega}_{\min})$, following Equation (A8), which is then used to construct an $N_\omega \times N_{\dot{\omega}}$ grid with spacing $\Delta\omega_0$ and $\Delta\dot{\omega}_0$ (gray dots); this grid becomes the basis for a Gibbs step, which will sample a random location $(n_\omega, n_{\dot{\omega}})$ within it. Next, the HMC sampler picks a new lower-left corner $(w_0^{\rm new}, \dot{w}_0^{\rm new})$, which is used to construct a new grid (black dots), and the process repeats. The dashed lines indicate the minimum and maximum frequencies and frequency derivatives that will be probed during sampling.

$\dot{\omega} + \Delta\dot{\omega}$, the log likelihood is proportional to

$$\log \mathcal{L}_{\dot{\omega}} \propto -\frac{1}{T}\int_0^T \mathrm{d}t\, [\cos(\dot{\omega}t^2/2) - \cos \times ((\dot{\omega} + \Delta\dot{\omega})t^2/2)]^2 \simeq \frac{C\left(T\sqrt{\frac{\Delta\dot{\omega}}{2}}\right)}{T\sqrt{\frac{\Delta\dot{\omega}}{2}}} - 1, \quad \text{(A6)}$$

where we have again kept only the leading term in small $\Delta\dot{\omega}$, and $C(x)$ is the Fresnel cosine integral. This has a global maximum at $\Delta\dot{\omega} = 0$, and a first local maximum when $\Delta\dot{\omega} \approx 15.121/T^2$, which is the grid spacing we use in our analysis.

The local maxima in the likelihood as a function of ω and $\dot{\omega}$ can cause the HMC sampler to become locally stuck, without finding the true typical set of the posterior. To address this, we use our knowledge of the approximate separation between likelihood modes to construct a *grid* (or *comb*) of frequencies and frequency derivatives that helps the sampler hop between potential modes of the likelihood, while also smoothly sliding the grid to cover the desired frequency space. We illustrate this in Figure 9, and describe it below.

Assume that we want to look for frequencies ω_0 within some range $\omega_{\min} \leqslant \omega_0 \leqslant \omega_{\max}$ and frequency derivatives $\dot{\omega}_0$ within $\dot{\omega}_{\min} \leqslant \dot{\omega}_0 \leqslant \dot{\omega}_{\max}$. Given the above spacings $\Delta\omega_0$ and $\Delta\dot{\omega}_0$, the desired ω_0 and $\dot{\omega}_0$ ranges can be covered by a grid with a number of frequency and frequency-derivative points given by

$$N_\omega = \lfloor(\omega_{\max} - \omega_{\min})/\Delta\omega_0\rfloor, \quad \text{(A7a)}$$

$$N_{\dot{\omega}} = \lfloor(\dot{\omega}_{\max} - \dot{\omega}_{\min})/\Delta\dot{\omega}_0\rfloor \quad \text{(A7b)}$$

where $\lfloor\cdot\rfloor$ is the floor function; a location on the grid is specified by indices $(n_\omega, n_{\dot{\omega}})$ with $0 \leqslant n_\omega < N_\omega$ and $0 \leqslant n_{\dot{\omega}} < N_{\dot{\omega}}$.

To allow one of the gridpoints to snap onto the bulk of the likelihood, we free the grid to continuously slide slightly up and down in ω_0 and $\dot{\omega}_0$ space by up to an amount commensurate with the grid spacing (Figure 9). Concretely, letting $(w_0, \dot{w}_0)$ define the location of the lower-left corner of the grid in $(\omega_0, \dot{\omega}_0)$ space, i.e., the placement of the $(n_\omega = 0, n_{\dot{\omega}} = 0)$ grid point, we slide the grid by sampling

$$w_0 \sim \mathcal{N}_{[0,\infty)}(\omega_{\min}, \Delta\omega_0^2), \quad \text{(A8a)}$$

$$\dot{w}_0 \sim \mathcal{N}_{[0,\infty)}(\dot{\omega}_{\min}, \Delta\dot{\omega}_0^2), \quad \text{(A8b)}$$

where $\mathcal{N}_{[0,\infty)}(x, y)$ is a normal distribution with mean x and variance y, but truncated to positive values. That is, at each HMC step, we choose an anchor point $(w_0, \dot{w}_0)$ for the grid, which cannot fall below the minimum values of $\omega_{\min}$ and $\dot{\omega}_{\min}$ and will typically not be farther from the $(\omega_{\min}, \dot{\omega}_{\min})$ corner than the grid spacing $\Delta\omega_0$ and $\Delta\dot{\omega}_0$. Based on that anchor point and spacing, we can then construct an $N_\omega \times N_{\dot{\omega}}$ grid of frequencies and frequency derivatives.

Finally, we use a Gibbs step to select a location on the grid $(n_\omega, n_{\dot{\omega}})$ by sampling from a categorical distribution with a prior such that

$$n_\omega \sim \mathcal{U}\{0, N_\omega - 1\}, \quad \text{(A9a)}$$

$$n_{\dot{\omega}} \sim \mathcal{U}\{0, N_{\dot{\omega}} - 1\}, \quad \text{(A9b)}$$

where $\mathcal{U}\{j, k\}$ is a discrete uniform distribution over the integers j to k. With such a choice, we construct values of the frequency ω_0 and frequency derivative $\dot{\omega}_0$ at the grid point as

$$\omega_0 = w_0 + n_\omega\Delta\omega_0, \quad \text{(A10a)}$$

$$\dot{\omega}_0 = \dot{w}_0 + n_{\dot{\omega}}\Delta\dot{\omega}_0. \quad \text{(A10b)}$$

The effective prior on ω_0 and $\dot{\omega}_0$ is sufficiently broad to cover the desired range (shown in Figure 10(a) for ω_0, analogous for $\dot{\omega}_0$). This is all we need to compute the likelihood at this grid point, marginalizing over A and ϕ as described below.

A.3. Linear Parameters and Analytic Marginalization

Having chosen ω_0, $\dot{\omega}_0$, σ, and τ, we could also sample the chirp amplitude A, phase ϕ, and mean m_i to directly compute the likelihood of Equation (A2); however, this would be suboptimal. Instead, we take advantage of the Gaussian form of Equation (A2) to analytically marginalize over $\{A, \phi, m_i\}$ for any given choice of $\{\omega_0, \dot{\omega}_0, \sigma, \tau\}$, thus effectively reducing the number of parameters to sample from 7 to 4. This vastly improves sampling efficiency without losing information about the suppressed parameters, which can be restored in postprocessing.

The analytical marginalization is made possible by the factorization properties of the Gaussian (D. W. Hogg et al. 2020), which we can fully leverage by reparameterizing the chirp template in Equation (3). Concretely, we reexpress the chirp amplitude A and phase ϕ in terms of the corresponding *linear quadratures* (e.g., M. Isi 2023),

$$a = A\cos\phi, \;\; b = A\sin\phi, \quad \text{(A11)}$$

so that $A = \sqrt{a^2 + b^2}$ and $\phi = \tan^{-1}(b/a)$. In terms of these quantities, the template of Equation (3) becomes

$$y_{\rm chirp}(t) = a\cos[2\pi f(t)] + b\sin[2\pi f(t)], \quad \text{(A12a)}$$

$$= a\, y_c(t) + b\, y_s(t), \quad \text{(A12b)}$$

with $y_c(t; \omega_0, \dot{\omega}_0) = \cos[2\pi f(t)]$ and $y_s(t) = \sin[2\pi f(t)]$ implicit functions of ω_0 and $\dot{\omega}_0$. Defining also $y_m(t) = 1$, we can

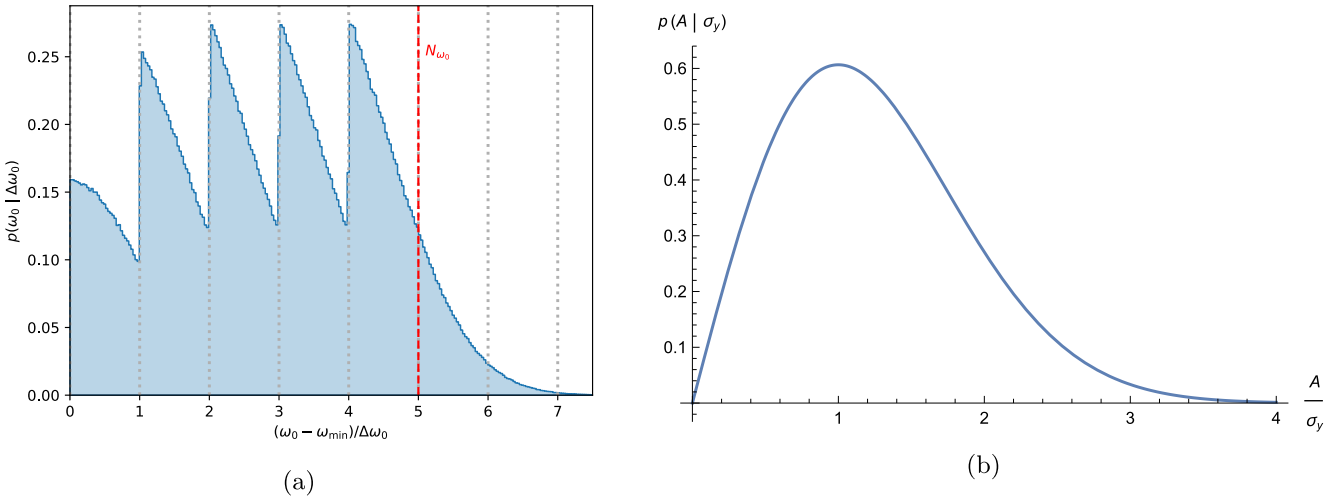


Figure 10. Effective priors for the frequency ω_0 (left) and chirp amplitude A (right). (a) The effective prior on ω_0 results from the sampling strategy associated with Equation (A10), which defines a frequency grid of N_{ω_0} points separated by $\Delta\omega_0$ and starting at values not smaller than $\omega_{\min}$ (for the figure, we set $N_{\omega_0} = 5$); the distribution is broad, with spikes at integer multiples of $\Delta\omega_0$. (b) The effective prior on A is a Rayleigh distribution with scale $\sigma_y = \mathrm{std}(y)$, per Equation (A18). The prior for $\dot{\omega}_0$ is analogous to (a), with a grid of $N_{\dot{\omega}_0}$ points separated by $\Delta\dot{\omega}_0$ and starting at values not smaller than $\dot{\omega}_{\min}$.

write the full template as

$$m_i + y_{\mathrm{chirp}}(t; a, b) = m_i\, y_m(t) + a\, y_c(t) + b\, y_s(t). \quad \text{(A13)}$$

This reparameterization is advantageous because the a, b and m_i parameters enter the waveform as linear coefficients to produce combinations of the basis functions $y_c(t)$, $y_s(t)$ and $y_m(t)$, which will allow us to analytically marginalize out $\{m_i,\ a,\ b\}$ conditioned on $\{\omega_0, \dot{\omega}_0, \sigma, \tau\}$. To affect the analytical marginalization, we choose Gaussian priors on a, b, and m_i such that

$$a \sim \mathcal{N}(0, \sigma_y^2), \quad \text{(A14a)}$$

$$b \sim \mathcal{N}(0, \sigma_y^2), \quad \text{(A14b)}$$

$$m_i \sim \mathcal{N}(\mu_y, \sigma_y^2), \quad \text{(A14c)}$$

where, as above, $\sigma_y = \mathrm{std}(y)$ is the standard deviation of the observed lightcurve and $\mu_y = \mathrm{mean}(y)$ is its mean. With this choice of prior $p(a, b, m_i)$, we can compute the marginalized likelihood

$$p(y|\omega_0, \dot{\omega}_0, \sigma, \tau) = \int p(y|a, b, m_i, \omega_0, \dot{\omega}_0, \sigma, \tau) \times p(a, b, m_i)\, \mathrm{d}a\, \mathrm{d}b\, \mathrm{d}m_i, \quad \text{(A15)}$$

analytically as explained in D. W. Hogg et al. (2020), i.e., without carrying out the integral explicitly.

Using this framework, the HMC-Gibbs sampler produces a set of samples in $\{\omega_0, \dot{\omega}_0, \sigma, \tau\}$, drawn from the marginal posterior

$$p(\omega_0, \dot{\omega}_0, \sigma, \tau|y) = \int p(\omega_0, \dot{\omega}_0, \sigma, \tau, a, b, m_i|y)\, \mathrm{d}a\, \mathrm{d}b\, \mathrm{d}m_i. \quad \text{(A16)}$$

We restore the a, b, m_i parameters in postprocessing by sampling from the conditional posterior $p(a, b, m_i|\omega_0, \dot{\omega}_0, \sigma, \tau, y)$, which is a Gaussian that can be computed analytically for each sampled point in $\{\omega_0, \dot{\omega}_0, \sigma, \tau\}$ (D. W. Hogg et al. 2020). This uses the fact that

$$p(\omega_0, \dot{\omega}_0, \sigma, \tau, a, b, m_i|y) = p(a, b, m_i|\omega_0, \dot{\omega}_0, \sigma, \tau, y) \times p(a, b, m_i|y)/p(\omega_0, \dot{\omega}_0, \sigma, \tau|y) \quad \text{(A17)}$$

by Bayes theorem.

The analytical marginalization requires setting a Gaussian prior on the linear parameters a, b, and m_i, as in Equation (A14). It can be shown that this is equivalent to a uniform prior on ϕ, but not on A (e.g., M. Isi 2023); rather, the induced prior on A is a Rayleigh distribution with scale σ_y, i.e.,

$$p(A|\sigma_y) = \frac{A}{\sigma_y^2} \exp\left(-\frac{A^2}{2\sigma_y^2}\right), \quad \text{(A18)}$$

as we illustrate in Figure 10(b). Although this distribution is not flat, we find that it is sufficiently broad so as to not bias our inference.

Appendix B
Inference of the Chirp

We demonstrate our ability to measure chirp as a function of binary time to merger, total mass, and orbital period in Figure 11. The z-score of $\dot{f}_0$ drastically drops for longer orbital periods (and consequently for larger t_m), further emphasizing on the importance of observing many cycles to achieve a reliable measurement of chirp.

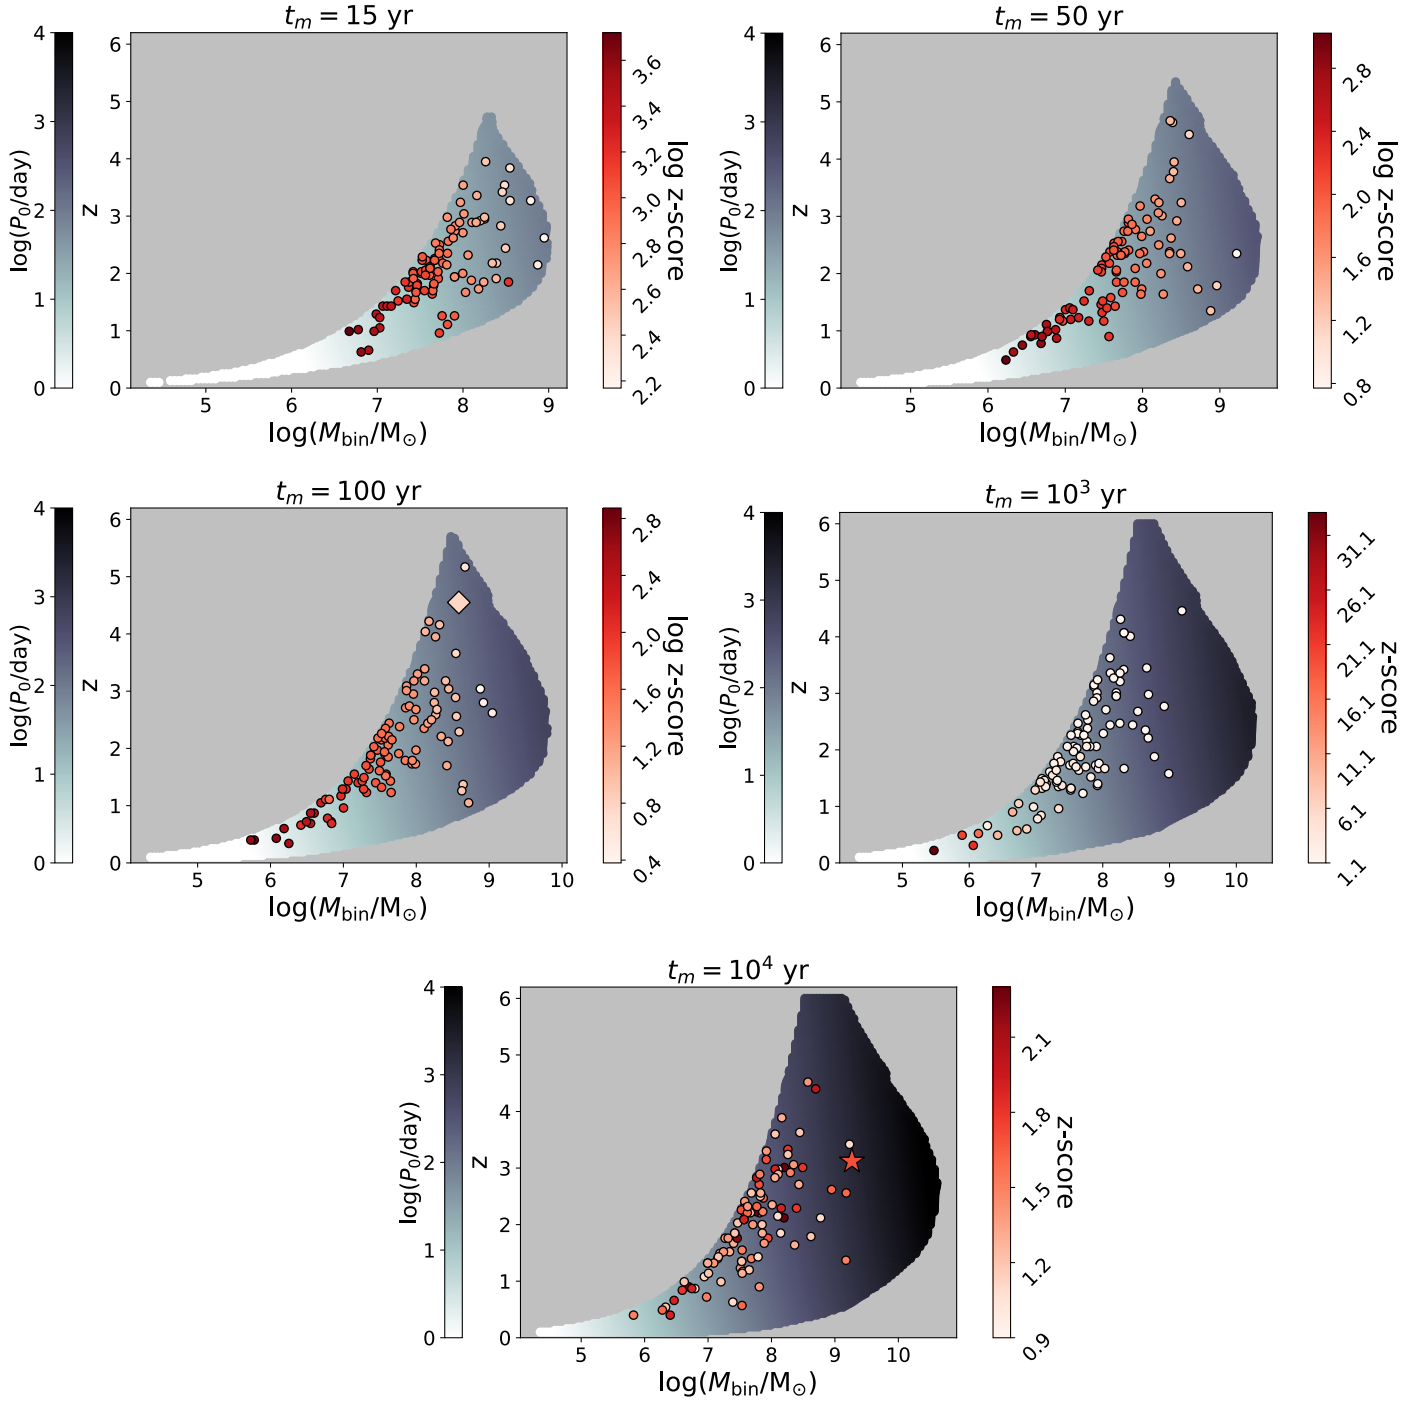


Figure 11. We show the chirp ($\dot{f}_0$) posterior z-scores for the same binaries presented in Figure 4. Notice the difference in ranges and scales on the z-score (red) color gradients. The diamond and the star are the two wide binaries that we highlight in Section 3.1.

Appendix C
Runtime

Ultimately, we will apply our analysis to systematically search for compact MBH binaries in the LSST, beginning in 2025–2026. In Figure 12, we show the runtime of our Bayesian algorithm on the fiducial 100 binaries with $t_m = 50$ yr. Most lightcurves are analyzed under 10 minutes, with a maximum runtime of ∼1 day; the distribution in Figure 12 indicates a median runtime of 5.6 minutes, and the central 68% credible interval between ∼1.1 and ∼51 minutes. This is significantly more efficient than other Bayesian algorithms developed to identify periodicity in quasar lightcurves, which typically take at least a few hours per lightcurve (e.g., C. A. Witt et al. 2022). With our current average runtime, we can perform our analysis on 100 million quasars using ∼16 million CPU hours. In future work, we will work on optimizing the algorithm to speed up the search on big datasets. Additionally, the search space for binary parameters, including periodicity and chirp, is much wider than what we consider in this work, which will make the runtime longer. On the other hand, lightcurves with "obvious" periodicity can be identified with a certain "preselection" (faster) approach, like

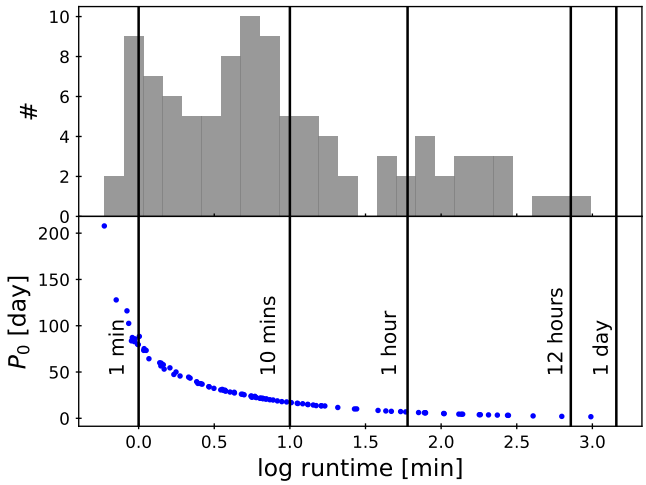


Figure 12. The runtime of our Bayesian inference model for each lightcurve in the fiducial binary population (top right panel in Figure 4). The log-runtime is presented in the histogram (top panel), which shows that most of our lightcurves are analyzed within 10 minutes. We also show runtime versus orbital periods (bottom panel), and we find that the runtime is shorter for binaries with longer periods.

the LS periodogram, to narrow down the parameter space for the Bayesian analysis. Other parameters can undergo similar pre-search steps using simpler statistics. We will design the preselection methods for each parameter for a given observation in future work.

ORCID iDs

Chengcheng Xin https://orcid.org/0000-0003-3106-8182
Maximiliano Isi https://orcid.org/0000-0001-8830-8672
Will M. Farr https://orcid.org/0000-0003-1540-8562
Zoltán Haiman https://orcid.org/0000-0003-3633-5403

References

Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023a, ApJL, 951, L8
Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023b, ApJL, 952, L37
Alexander, D. M., & Hickox, R. C. 2012, NewAR, 56, 93
Barret, D., Decourchelle, A., Fabian, A., et al. 2020, AN, 341, 224
Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, Natur, 287, 307
Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, J. Mach. Learn. Res., 20, 28
Burke-Spolaor, S., Blecha, L., Bogdanovic, T., et al. 2018, arXiv:1808.04368
Charisi, M., Bartos, I., Haiman, Z., et al. 2016, MNRAS, 463, 2145
Charisi, M., Taylor, S. R., Runnoe, J., Bogdanovic, T., & Trump, J. R. 2022, MNRAS, 510, 5929
Chen, Y.-C., Liu, X., Liao, W.-T., et al. 2020, MNRAS, 21, 1
Cocchiararo, F., Franchini, A., Lupi, A., & Sesana, A. 2024, A&A, 691, 1
Colpi, M., Danzmann, K., Hewitson, M., et al. 2024, arXiv:2402.07571
Cutler, C., & Flanagan, E. E. 1994, PhRvD, 49, 2658
Davelaar, J., & Haiman, Z. 2022a, PhRvL, 128, 191101
Davelaar, J., & Haiman, Z. 2022b, PhRvD, 105, 103010
Davis, M. C., Grace, K. E., Trump, J. R., et al. 2024, ApJ, 965, 34
Dittmann, A. J., Ryan, G., & Miller, M. C. 2023, ApJL, 949, L30
D'Orazio, D. J., & Di Stefano, R. 2018, MNRAS, 474, 2975
D'Orazio, D. J., Duffell, P. C., & Tiede, C. 2024, AJ, 977, 11
D'Orazio, D. J., Haiman, Z., & Schiminovich, D. 2015, Natur, 525, 351
Dotti, M., Bonetti, M., Rigamonti, F., et al. 2023, MNRAS, 518, 4172
Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, ApJL, 875, L1
Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019b, ApJL, 875, L2
Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2015, MNRAS: Letters, 447, L80
Foreman-Mackey, D. 2018, RNAAS, 2, 31
Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, AJ, 154, 220
Foustoul, V., Webb, N. A., Mignon-Risse, R., et al. 2025, A&A, 699, 19A55
Garg, M., Franchini, A., Lupi, A., Bonetti, M., & Mayer, L. 2025, ApJ, 993, 145
Goulding, A. D., Greene, J. E., Bezanson, R., et al. 2018, PASJ, 70, S37
Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015a, MNRAS, 453, 1562
Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015b, Natur, 518, 74
Guo, H., Liu, X., Tayyaba, Z., & Liao, W. T. 2020, MNRAS, 492, 2910
Haiman, Z. 2017, PhRvD, 2, 023004
Haiman, Z., Kocsis, B., & Menou, K. 2009, ApJ, 700, 1952
Hogg, D. W., Price-Whelan, A. M., & Leistedt, B. 2020, arXiv:2005.14199
Hopkins, P. F., & Hernquist, L. 2009, ApJ, 698, 1550
Huijse, P., Davelaar, J., De Ridder, J., Jannsen, N., & Aerts, C. 2025, arXiv:2505.16884
Ingram, A., Motta, S. E., Aigrain, S., & Karastergiou, A. 2021, MNRAS, 503, 1703
Isi, M. 2023, CQGra, 40, 203001
Ivezic, Z., & MacLeod, C. L. 2013, IAUS, 304, 395
Jiang, N., Yang, H., Wang, T., et al. 2022, arXiv:2201.11633
Kauffmann, G., & Haehnelt, M. 2000, MNRAS, 311, 576
Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, ApJ, 698, 895
Kis-Tóth, Á., Haiman, Z., & Frei, Z. 2025, CQGra, 42, 19
Kochanek, C. S., Eisenstein, D. J., Cool, R. J., et al. 2012, ApJS, 200, 20
Kollmeier, J. A., Onken, C. A., Kochanek, C. S., et al. 2006, ApJ, 648, 128
Krauth, L. M., Davelaar, J., Haiman, Z., et al. 2023, MNRAS, 526, 5441
Krauth, L. M., Davelaar, J., Haiman, Z., et al. 2024, PhRvD, 109, 103014
Kulkarni, G., Worseck, G., & Hennawi, J. F. 2019, MNRAS, 31, 1035
Kulkarni, S. R., Harrison, F. A., Grefenstette, B. W., et al. 2021, arXiv:2111.15608
Li, K., Bogdanović, T., & Ballantyne, D. R. 2020, ApJ, 896, 113
Li, K., Bogdanović, T., Ballantyne, D. R., & Bonetti, M. 2022, ApJ, 933, 18
Lomb, N. R. 1976, Ap&SS, 39, 447
LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
LSST Science CollaborationRubin Observatory Project 2023, rubin-sim: Rubin Observatory Simulation Framework v2.4.0, GitHub, https://github.com/lsst/rubin_sim/tree/v2.4.0
MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, ApJ, 721, 1014
Martini, P. 2004, arXiv:astro-ph/0304009
Mushotzky, R. F., Edelson, R., Baumgartner, W., & Gandhi, P. 2011, ApJL, 743, 6
Nandra, K., Barret, D., Barcons, X., et al. 2013, arXiv:1306.2307
O'Neill, D., Tiede, C., D'Orazio, D. J., Haiman, Z., & Macfadyen, A. 2025, ApJ, 993, 206
Peters, P. C., & Mathews, J. 1963, PhRv, 131, 435
Park, K., Xin, C., Davelaar, J., & Haiman, Z. 2025, PhRvD, 111, 063011
Phan, D., Pradhan, N., & Jankowiak, M. 2019, arXiv:1912.11554
Rasmussen, C. E., & Williams, C. K. I. 2006, Gaussian Processes for Machine Learning (MIT Press)
Robnik, J., Bayer, A. E., Charisi, M., et al. 2024, MNRAS, 534, 1609
Saade, M. L., Brightman, M., Stern, D., et al. 2024, ApJ, 966, 104
Scargle, J. D. 1982, ApJ, 263, 835
Shvartzvald, Y., Waxman, E., Gal-Yam, A., et al. 2024, ApJ, 964, 74
Siwek, M., Weinberger, R., & Hernquist, L. 2023, MNRAS, 522, 2707
Smith, K. L., Mushotzky, R. F., Boyd, P. T., et al. 2018, ApJ, 857, 141
Tang, Y., Haiman, Z., & MacFadyen, A. 2018, MNRAS, 476, 2249
Tiede, C., & Orazio, D. J. D. 2024, MNRAS, 527, 6021
Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference Manual (CreateSpace)
Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, MNRAS, 461, 3145
Wagg, T., Broekgaarden, F., & Gültekin, K. 2024, TomWagg/software-citation-station v1.2, Zenodo, doi:10.5281/zenodo.13225824
Wagg, T., & Broekgaarden, F. S. 2024, arXiv:2406.04405
Westernacher-Schneider, J. R., Zrake, J., MacFadyen, A., & Haiman, Z. 2022, PhRvD, 106, 103010
Witt, C. A., Charisi, M., Taylor, S. R., & Burke-Spolaor, S. 2022, ApJ, 936, 89
Xin, C., Charisi, M., Haiman, Z., et al. 2020, MNRAS, 496, 1683
Xin, C., & Haiman, Z. 2021, MNRAS, 506, 2408
Xin, C., & Haiman, Z. 2024, MNRAS, 533, 3164
Yu, W., Richards, G. T., Vogeley, M. S., Moreno, J., & Graham, M. J. 2022, ApJ, 936, 132
Zrake, J., Tiede, C., MacFadyen, A., & Haiman, Z. 2021, ApJL, 909, L13