



Bayesian Characterisation of Circumbinary Exoplanets with LISA

Michael Katz, Camilla Danielski, Nikolaos Karnesis, Valeriya Korol, Nicola Tamanini, Neil Cornish, Tyson Littenberg

► To cite this version:

Michael Katz, Camilla Danielski, Nikolaos Karnesis, Valeriya Korol, Nicola Tamanini, et al.. Bayesian Characterisation of Circumbinary Exoplanets with LISA. *Monthly Notices of the Royal Astronomical Society*, 2022, 517 (1), pp.697-711. 10.1093/mnras/stac2555 . hal-03677053

HAL Id: hal-03677053




<https://hal.science/hal-03677053>

Submitted on 24 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bayesian characterization of circumbinary sub-stellar objects with *LISA*

Michael L. Katz¹ ,¹★ Camilla Danielski^{2,3} , Nikolaos Karnesis,⁴ Valeriya Korol⁵ , Nicola Tamanini,⁶ Neil J. Cornish⁷ and Tyson B. Littenberg⁸

¹Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Am Mühlenberg 1, D-14476 Potsdam-Golm, Germany

²Instituto de Astrofísica de Andalucía, CSIC, Glorieta de la Astronomía, E-18008, Granada, Spain

³Sorbonne Universités, UPMC Université Paris 6 et CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98 bis bd Arago, 75014 Paris, France

⁴Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece

⁵Institute for Gravitational Wave Astronomy & School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

⁶Laboratoire des 2 Infinis - Toulouse (L2IT-IN2P3), CNRS, UPS, F-31062 Toulouse Cedex 9, France

⁷eXtreme Gravity Institute, Department of Physics, Montana State University, Bozeman, MT 59717, USA

⁸NASA Marshall Space Flight Center, Huntsville, AL 35811, USA

Accepted 2022 August 23. Received 2022 August 3; in original form 2022 May 6

ABSTRACT

The Laser Interferometer Space Antenna (*LISA*) will detect and characterize $\sim 10^4$ Galactic Binaries, consisting predominantly of two white dwarfs (WDs). An interesting prospect within this population is a third object – another WD star, a circumbinary exoplanet (CBP), or a brown dwarf (BD) – in orbit about the inner WD pair. We present the first fully Bayesian detection and posterior analysis of substellar objects with *LISA*, focusing on the characterization of CBPs. We used an optimistic astrophysically motivated catalogue of these CBP third-body sources, including their orbital eccentricity around the inner binary for the first time. We examined Bayesian evidence computations for detectability, as well as the effects on the posterior distributions for both the inner binary parameters and the third-body parameters. We find that the posterior behaviour bifurcates based on whether the third-body period is above or below half the observation time. Additionally, we find that undetectable third-body sources can bias the inner binary parameters whether or not the correct template is used. We used the information retrieved from the study of the CBP population to make an initial conservative prediction for the number of detectable BD systems in the original catalogue. We end with commentary on the predicted qualitative effects on *LISA* global fitting and Galactic Binary population analysis. The procedure used in this work is generic and can be directly applied to other astrophysical effects expected within the Galactic Binary population.

Key words: gravitational waves – white dwarfs – planets and satellites: detection.

1 INTRODUCTION

The Laser Interferometer Space Antenna (*LISA*; Amaro-Seoane et al. 2017), a future space-based gravitational-wave detector set to launch in the mid-2030s, will add new gravitational-wave observations to the high-frequency observations of ground-based gravitational-wave detectors (LVK Collaboration 2018; Abbott et al. 2019). *LISA* will observe in the milliHertz regime of the spectrum, an area where many different types of astrophysical sources are expected to be present (Amaro-Seoane et al. 2022). An important class of *LISA* sources are compact Galactic Binaries (GB): a binary consisting of two compact objects with short-period orbits located within the Milky Way or surrounding close satellites. Typically, GBs will include two white dwarfs (WDs), but some GBs may contain one or more stellar-origin black holes or neutron stars (Amaro-Seoane et al. 2017). There are roughly $\sim 10^7$ GBs in the Milky Way emitting GW in the *LISA* frequency band (e.g. Korol et al. 2022). Of these, $\sim 10^4$

will be individually resolvable with signal-to-noise ratios (S/N) of ~ 10 – 10^3 (e.g. Timpano, Rubbo & Cornish 2006a; Cornish & Robson 2017; Korol et al. 2017; Lamberts et al. 2019; Breivik et al. 2020a). Electromagnetic observations of double white dwarf (DWD) binaries have found $\gtrsim 20$ (Kupfer et al. 2018) individually resolvable GB *LISA* sources based on their electromagnetically determined properties. These systems are the so-called ‘Verification Galactic Binaries’. The remaining unresolved GBs at low S/N will combine into a confusion background that is expected to lie above the low-frequency part of the *LISA* sensitivity curve (e.g. Robson & Cornish 2017).

GBs are quasimonochromatic sources that evolve slowly over the duration of the *LISA* observation. These systems will be long-lived and present in *LISA* data from the beginning to the end of its observation. They emit gravitational waves deep in the inspiralling stage, from tens of thousand to million years prior to merger. The overlap of resolvable and unresolvable GBs, in addition to the rest of the other *LISA* sources, will greatly complicate the analysis of the *LISA* data. Proper modelling of the *LISA* spacecraft orbits and the observed *LISA* instrumental noise will also be necessary for

* E-mail: mikekatz04@gmail.com

accurate analysis, even in the presence of expected non-stationary and non-Gaussian noise effects. The main method suggested to solve this problem is using global fitting techniques that fit the parameters of many astrophysical sources and noise properties simultaneously (Cornish & Crowder 2005). The global fit of GBs specifically has been examined recently in Littenberg et al. (2020). This analysis was performed assuming a basic set of astrophysical GBs. It proved the capabilities of global fit analyses and discussed further development of the method to include more realistic effects of the GB population and *LISA* detector setup.

Many science questions can be addressed with *LISA* observations of GBs on different scales ranging from substellar/stellar to the entire Galaxy (Amaro-Seoane et al. 2022). For example, at the population level, the characteristics of the observed confusion background from GBs and the spatial distribution of resolved ones can help to understand the stellar distribution and properties of the Milky Way Galaxy (Benacquista & Holley-Bockelmann 2006; Adams, Cornish & Littenberg 2012; Korol, Rossi & Barausse 2019; Breivik, Mingarelli & Larson 2020b; Wilhelm et al. 2021; Georgousi et al. 2022). The population of resolvable GBs can also be used to constrain tests of general relativity (GR; Littenberg & Yunes 2019). Resolvable source observations, both for individual sources and the population as a whole, can shed light on the complex astrophysics of multiple-object systems (e.g. Kremer et al. 2017).

Roughly 13 per cent of low-mass stellar systems contain three or more stars (Tokovinin 2014; Fuhrmann et al. 2017). Low-mass binaries with periods of less than three days are expected to exist in hierarchical systems at a fraction of ~ 96 per cent (Tokovinin et al. 2006; Qian et al. 2011). The predicted detection rate of these systems by *LISA* is uncertain; however, it is reasonable to expect these systems to play a large role in the observed GB population. Astrophysical effects, such as Kozai–Lidov oscillations, can cause hardening of the inner binary, driving it to closer separations (Kozai 1962; Lidov 1962). This may enlarge the fraction of observed GB systems containing a third-body (Fabrycky & Tremaine 2007; Thompson 2011; Fang, Thompson & Hirata 2018; Toonen, Perets & Hamers 2018). Within the substellar objects (SSOs) regime, analytical studies by Martin & Triaud (2016) showed that the Kozai–Lidov evolution rarely affects the inner binary in systems hosting a tertiary planet due to its low mass. However, in the brown-dwarf regime, there is a significant change in the Kozai–Lidov behaviour for tight triple systems, where the ratio between the outer and inner binary semimajor axes is less than 15 (Martin & Triaud 2016).

In this paper, we are specifically interested in the observations of hierarchical triple systems where a circumbinary exoplanet (CBP) is in orbit around an inner binary consisting of a DWD. We examine the posterior distributions stemming from a population of these sources (see Section 2). We also examine observations of the population of circumbinary brown dwarfs (BD). As we will discuss, the number of potentially detectable¹ BDs in the population is much higher than CBPs, making a full analysis of the BD systems computationally costly. Therefore, we focus our detailed study on CBP systems and provide general comments on the BD systems in the catalogue, which will be examined in more detail in the future (Danielski et al. in preparation). The cut off between the two object classes was chosen to be 16 M_J (as discussed in Spiegel, Burrows & Milsom 2011), instead of 13 M_J , the deuterium burning limit.

These populations have been examined previously in Robson et al. (2018), Tamanini & Danielski (2019), Danielski et al. (2019), Tamanini et al. (2020), and Kang, Liu & Shao (2021). The statistical Information Matrix² was used to analyse the detection and characterization of CBPs and BDs with *LISA* in Tamanini & Danielski (2019) and Danielski et al. (2019). Recently, the Information Matrix was also deployed to study CBP detection and characterization by *Taiji*, a space-based interferometer comparable in scope to *LISA* (Kang et al. 2021). In Robson et al. (2018), fast frequency-domain waveforms containing the third-body effect were constructed and an initial Bayesian analysis was performed for systems with three stellar-mass objects in a hierarchical triple. A key finding from this work was initial relations comparing different regimes of P_2 , the period of the perturber, compared to T_{obs} , the *LISA* observation time. With $P_2/T_{\text{obs}} \gg 1$, the effect of the perturber is not detectable. Near $P_2/T_{\text{obs}} \lesssim 10$, the companion object can create a detectable Doppler shift in the waveform. As the period of the third body decreases so that $P_2/T_{\text{obs}} \lesssim 1$, the eccentricity and period of the outer orbit can be characterized. While these results relate to systems with a third body that is of similar stellar mass to the inner DWD constituents, we will show that similar relations of P_2 to T_{obs} govern the posterior behaviour of CBP systems. Specifically, there is a clear difference between behaviour above and below $P_2 = T_{\text{obs}}/2$, which represents the Nyquist criterion for sampling of the full third-body orbit.

Here, we expand on this previous work by performing a full Bayesian analysis, including posterior and evidence estimation, on the observability and characterization of individual instances of these triple systems by *LISA*. We will examine how many systems, stemming from a coherent and astrophysically motivated population, produce a waveform where the effect of the third body is statistically significant compared to the base inner-binary waveform. In addition to purely detecting and characterizing these sources, we strive to illuminate issues that may arise from fitting base two-body DWD templates to true source waveforms that include the presence of the third body. The third-body-inclusive waveforms used in this work are also the first to include eccentric third-body orbits in the waveform description when analysing CBP systems in *LISA*.

In Section 2, we describe the population of sources used in this work. Sections 3 and 4 detail the DWD waveform model and Bayesian analysis techniques employed, respectively. The results of our analysis pipeline are given in Section 5. In Section 6, we discuss these results within the current context of DWD analysis with *LISA*. We conclude our remarks in Section 7. Throughout this work, we use geometrical units with $G = c = 1$.

2 POPULATION OF CIRCUMBINARY SUBSTELLAR OBJECTS

In this work we consider hierarchical triple systems consisting of an inner detached DWD binary and a tertiary SSOs. Specifically, we focus on DWD systems only as they will be the most numerous among the other GBs accessible with *LISA* (Amaro-Seoane et al. 2022) and thus more likely to provide triple systems. As a first step, we assemble the DWD population, constituting the inner component of the hierarchical triple system. Next, we inject the SSO population

¹By ‘detectable’ we mean ‘distinguishable’, according to the Bayesian evidence ratio (see Section 4).

²The ‘Information Matrix’ is also referred to as the ‘Fisher Information Matrix’. In order to foster a welcoming community in our field, we have not used this term due to its namesake’s connection with the science of and personal advocacy for Eugenics (e.g. Evans 2020).

into the already formed DWD population i.e. neglecting co-evolution of the stellar binary and the third object.

2.1 Inner DWD binaries

We follow a binary population synthesis approach to obtain a representative Galactic DWD population. Specifically, we assemble our mock population based on the DWD evolution model constructed using population synthesis code SEBA (Portegies Zwart & Verbunt 1996; Nelemans et al. 2001; Toonen, Nelemans & Portegies Zwart 2012), publicly available as part of the AMUSE environment (Portegies Zwart et al. 2009). The choice of the fiducial model is motivated by the fact that it yields the DWD space density in a good agreement with that derived from currently available spectroscopically selected samples, and it reproduces the general trend of the observed DWD mass ratio distribution, which peaks at around unity (Toonen et al. 2012, 2017). We refer for a detailed description of this model to Toonen et al. (2012) and we describe below only its main characteristics.

The initial (i.e. zero-age main-sequence) population is assembled with a Monte Carlo technique. The initial binary fraction is assumed to be of 50 per cent (Raghavan et al. 2010; Duchêne & Kraus 2013), and metallicity is set to the Solar value for all binaries. The mass of the primary star – initially more massive of the two – is sampled between 0.95 and 10 M_{\odot} , according to the initial mass function of Kroupa, Tout & Gilmore (1993). The mass of the secondary star is defined to obtain a flat mass ratio distribution between 0 and 1 (e.g. Duchêne & Kraus 2013). Binaries' semimajor axes are sampled from a log-uniform distribution extending up to 10⁶ R_{\odot} (Abt 1983; Raghavan et al. 2010; Duchêne & Kraus 2013), while orbit eccentricities are sampled from a thermal distribution (Heggie 1975).

SEBA evolves the obtained initial population until both stars become WDs, following prescriptions for processes shaping the binary evolution path; these include mass and angular momentum transfer, common envelope evolution, magnetic braking, and gravitational radiation (Portegies Zwart & Verbunt 1996; Toonen et al. 2012, and references therein). It is important to mention that to obtain a close DWD pair emitting GWs in LISA today, binaries typically experience two common envelope phases (e.g. Nelemans et al. 2001). Thus, the assumption about the common envelope evolution is of particular importance and has been found to be the largest source of uncertainty in binary population synthesis models. In our fiducial $\gamma\alpha$ model, the first common envelope phase is typically described by the γ formalism (based on the angular momentum balance equation that allows for both shrinkage and the widening of the orbit), while the second is described by the α formalism (based on the energy balance equation that always leads to the shrinkage of the orbit). This scenario has been derived based on the reconstruction of the evolutionary paths of individually observed DWD binaries (Nelemans et al. 2000; Nelemans & Tout 2005; van der Sluis, Verbunt & Pols 2006).

Next, we distribute DWDs in a Milky Way-like galaxy as in Korol et al. (2019). We adopt a simplified Galactic potential composed of an exponential stellar disc and a spherical central bulge, which contains the majority of stars in the Milky Way. To model the star formation history of the Galaxy, we use the plane-projected star formation rate from a chemo-spectrophotometric model of Boissier & Prantzos (1999), and we assume the age of the Galaxy to be of 13.5 Gyr. The obtained integrated star formation history reproduces well the observed Galactic star formation history inferred from single WDs (Fantin et al. 2019).

Finally, we randomly draw binary initial phases and polarization angles from a uniform distribution, while the binary inclination angle

ι_1 (the '1' indicates the inner binary orbit) is drawn from a uniform distribution in $\cos \iota_1$.

2.2 Pre-selection of detectable DWDs

We obtain a catalogue amounting to $\sim 2.9 \times 10^7$ DWDs emitting in the LISA frequency band. As the next step, following the methodology described in Timpano, Rubbo & Cornish (2006b), Crowder & Cornish (2007), Nissanke et al. (2012), and Karnesis et al. (2021), we estimate the unresolved confusion foreground signal based on our DWD population, which then allows us to pre-select detectable DWDs based on an S/N criterion. Assuming a LISA nominal mission duration of 4 yr (i.e. $T_{\text{obs}} = 4$ yr) and an S/N threshold of 7, we obtain a subset of $\sim 2.5 \times 10^4$ individually resolvable DWDs to be used in the subsequent steps of our analysis.

2.3 Tertiary substellar objects

The characteristics of the population of injected SSOs follow the optimistic scenario reported in Danielski et al. (2019). We adopt this optimistic population to have more systems to study in detail. Consequently, the absolute number of LISA detections that we report here is likely on the high side, and should not be taken as representative of the numbers that LISA will realistically be able to deliver. Furthermore, we note that the occurrence rate of the SSOs was based on the observational WD pollution constraints (i.e. 50 per cent of the DWD population, Koester, Gänsicke & Farihi 2014), and that the optimistic scenario was originally defined to have systems with an outer circular orbit. The detection efficiency of circumbinary objects by LISA will be presented in an upcoming work through a Bayesian study of multiple population scenarios that includes the third-body eccentricity (Danielski et al. in preparation). Here, we focus on the data analysis strategy only.

Following Danielski et al. (2019), we use a log₁₀ uniform distribution for the semimajor axis a : $\log_{10} \mathcal{U}_a$ ($a_{\text{min}} - 200$ au), and a uniform distribution on the third-body mass m_c : \mathcal{U}_{m_c} ($1 M_{\oplus} - 0.08 M_{\odot}$). Novelties with respect to Danielski et al. (2019) are both the different lower bound a_{min} for the distribution of the semimajor axis, and the inclusion of the eccentricity as an extra parameter defining the orbit of the third body. The lower bound a_{min} is now physically motivated and is provided by the stability criterion of Ballantyne et al. (2021) which, assuming a total DWD mass of 2 M_{\odot} separated by 1 R_{\odot} with zero eccentricity, yields $a_{\text{min},1} = 2.39 R_{\odot}$. These values have been chosen as representatives among the highest possible values within our DWD population, to provide a conservative condition for the minimum value of a . This condition holds unless the eccentricity of the SSO is so high that the third body enters the Roche limit of one of the WDs at perihelion, in which case it would be destroyed by tidal gravitational forces. We thus conservatively set the maximum allowed eccentricity $e_{2,\text{max}}$ (the '2' here indicates the outer orbit) using the Roche limit condition (Roche 1849) which, given a , can simply be computed as

$$a(1 - e_{2,\text{max}}) = 2.44 \cdot R_{\oplus} \cdot \left(\frac{\rho_{\text{WD}}}{\rho_{\text{SSO}}} \right)^{1/3} + \frac{R_{\odot}}{2}. \quad (1)$$

Here, we assumed co-planar orbits, a typical WD with radius R_{\oplus} , and where the last addend on the right hand side of the equation accounts for the maximum separation of the two WDs from the center of mass of the DWD system (we set this to $R_{\odot}/2$). We explicitly used R_{\odot} as it roughly corresponds to the maximum separation between a binary WD in our population. This simple choice provides a conservative approach which is valid for all the systems considered in our analysis.

The WD density has been taken to be $\rho_{\text{WD}} = 10^9 \text{ kg m}^{-3}$, while the density of the third-body ρ_{SSO} has been estimated from a linear fit of the data reported in fig. 1 of Hatzes & Rauer (2015), with a flat floor imposed below 1 M_J . Since the eccentricity cannot be lower than zero, the condition in equation (1) at $e_{2,\text{max}} = 0$ sets an effective minimum semimajor axis a that all SSOs must respect, namely $a_{\text{min}, 2}$. For each circumbinary system, the minimum of the semimajor axis distribution a_{min} is thus set as the largest separation value chosen between $a_{\text{min}, 1}$ and $a_{\text{min}, 2}$.

We now prescribe the distribution of eccentricities for the SSO population. Since the SSO population around DWDs is unprobed, and due to a handful of detections of P-type bodies orbiting main sequence binaries, we used the distribution of the eccentricity of single-star planetary systems. Specifically, we combined various distributions found in the literature, as to cover a wide range of SSO-to-binary separations. When the distribution was unknown, a uniform distribution was applied. In detail, covering a global range from a_{min} to 200 au, we set:

(i) $a_{\text{min}} < a \leq 5$ au: uniform distribution $\mathcal{U}_e(0, e_{2,\text{max}})$ for both CBPs and BDs.

(ii) $5 \text{ au} < a \leq 200$ au: beta distribution $f(e_2|\alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} e_2^{\alpha-1} (1-e_2)^{\beta-1}$, where Γ is the gamma function, and where $[\alpha, \beta]$ are [30, 200] and [2.30, 1.65] for CBPs and BDs, respectively. These values were presented in Bowler, Blunt & Nielsen (2020) where the authors studied the eccentricity distribution of directly imaged planets orbiting main sequence, or younger, binaries in the 5–100 au range. Given that the eccentricity increment in stable post-common-envelope systems is very small for separations below 200 au (Veras et al. 2011), we adopted such distributions for our DWDs systems for the whole range up to 200 au.

In both cases the eccentricity values were drawn from the respective distributions with an upper limit $e_{2,\text{max}}$ chosen to avoid the tidal disruption of the SSO, as given by equation (1).

3 GALACTIC BINARY WAVEFORMS

DWDs are comparatively simple systems to model in terms of their gravitational waveforms, which can be attributed to their slow orbital evolution. Adding in effects that occur on slower time-scales than the DWD orbital time-scale, like accretion or multibody systems, creates small changes to the phasing of the base waveform. If this phasing difference is significant enough, the astrophysical effect will be detectable. Here, we will discuss the basics of DWD waveforms and the process of adding the effect of a third body. We point the interested reader to Robson et al. (2018) for more detailed information on the initial construction of these third-body-inclusive waveforms. The waveforms used here are identical to those used in that paper.

The two gravitational-wave polarizations from a DWD are given by

$$h_+(t) = \frac{2\mathcal{M}}{D_L} (\pi f_{\text{gw}}(t))^{2/3} (1 + \cos^2 \iota_1) \cos \Psi_{\text{gw}} \quad (2)$$

and

$$h_\times(t) = -\frac{4\mathcal{M}}{D_L} (\pi f_{\text{gw}}(t))^{2/3} \cos \iota_1 \sin \Psi_{\text{gw}}, \quad (3)$$

where \mathcal{M} is the chirp mass; f_{gw} is the instantaneous gravitational wave frequency; D_L is the luminosity distance; ι_1 is the inclination of the inner binary orbit; and Ψ_{gw} is the gravitational wave phase over time. The terms in equations (2 and 3) are given at the Solar System Barycentre. The phase is related to the integral over frequency: $\Psi_{\text{gw}} = \phi_0 + 2\pi \int^t f_{\text{gw}}(t') dt'$, with ϕ_0 representing an initial arbitrary phase shift.

The instantaneous gravitational-wave frequency is governed by the astrophysical evolution of the source. The time-scale over which the frequency changes for a typical DWD is slow: $f/\dot{f} \sim 10^{-3} \text{ Hz} / 10^{-17} \text{ Hz s}^{-1} \sim 10^6 \text{ yr}$ (Robson et al. 2018). This slow evolution allows for a Taylor expansion in the frequency:

$$f_{\text{gw}}(t) = f_0 + \dot{f}_0 t + \frac{1}{2} \ddot{f}_0 t^2 + \mathcal{O}(t^3). \quad (4)$$

The t^0 term, f_0 , is referred to as the initial frequency or the ‘carrier frequency’ in the following waveform description.

The slow evolution of f_{gw} also allows for a separation of time-scales in a ‘fast-slow’ decomposition. The fast time-scale is the orbital time-scale: $1/f \lesssim 1 \text{ h}$. The slow time-scale is mostly governed, in addition to the weak effect of the frequency evolution of the source, by the motion of the *LISA* constellation in its heliocentric orbit. Therefore, the slow time-scale is $\sim 1 \text{ yr}$. The large time-scale difference was leveraged in Cornish & Littenberg (2007) to create a fast frequency-domain waveform for DWDs observed by *LISA*. The slow portion builds a sparsely sampled time-domain signal that accounts for the frequency evolution, as well as the projection of the signal on to the time-dependent arms of the *LISA* constellation. It is here where the ecliptic latitude, λ , and ecliptic longitude, β_{sky} , enter the computation. The sparse time-domain signal is then transformed to the frequency domain and convolved with a δ function at the carrier frequency. The convolution produces projections along the *LISA* constellation arms that are combined into time-delay interferometry (TDI; Armstrong, Estabrook & Tinto 1999; Tinto & Armstrong 1999; Estabrook, Tinto & Armstrong 2000; Dhurandhar, Nayak & Vinet 2002; Tinto & Dhurandhar 2005; Baghi et al. 2021; Page & Littenberg 2021; Tinto & Dhurandhar 2021; Vallisneri et al. 2021) observables: X, Y, Z . These three TDI variables are correlated between them, but one could linearly combine them to build the noise orthogonal Vallisneri (2005) and Prince et al. (2002)

$$A = \frac{1}{\sqrt{2}} (Z - X), \quad (5)$$

$$E = \frac{1}{\sqrt{6}} (X - 2Y + Z), \quad (6)$$

$$T = \frac{1}{\sqrt{3}} (X + Y + Z). \quad (7)$$

These three TDI channels – A, E , and T – are the final template waveforms that will go into the Bayesian analysis described in Section 4.

3.1 Inclusion of the third body in the waveform

The waveform, as previously described, includes only the contributions from GR to the shrinking of the two-body DWD orbit. In order to include the effect of the third body on the waveform, we need to only modify Ψ_{gw} . The adjustment must account for the acceleration of the center of mass of the inner binary caused by the third-body perturber. This information is included by red-shifting (or blue-shifting) the gravitational-wave frequency based on the line-of-site velocity, v_{\parallel} , of the inner binary:

$$\Psi_{\text{gw}} = \phi_0 + 2\pi \int^t [1 + v_{\parallel}(t')] f_{\text{gw}}(t') dt'. \quad (8)$$

The line-of-site velocity is given by Robson et al. (2018)

$$v_{\parallel}(t) = \mathcal{A}_2 [\sin(\varphi_2 + \varpi) + e_2 \sin(\varpi)], \quad (9)$$

where φ_2 is the orbital phase of the outer orbit, and e_2 is the eccentricity of the outer orbit. The effective amplitude factor of the




Figure 1. A diagram illustrating the angular description of the third body (m_c) orbiting the inner DWD (m_1). The orbit is described in the ecliptic reference frame. The orbital phase is φ_2 . The three Euler rotation angles from the ecliptic plane are $-\omega_2$, $-\iota_2$, and $-\Omega_2$ (Robson et al. 2018).

perturber, \mathcal{A}_2 , is given by

$$\mathcal{A}_2 = \frac{m_c}{m_2} \sqrt{\frac{m_2}{p_2}} \bar{A}, \quad (10)$$

where m_c is the mass of the perturber, m_2 is the total mass of the system, and p_2 is the semilatus rectum of the outer orbit. \bar{A} and ϖ are based on the orientation of the third-body orbit to the observer, where \bar{A} is an effective amplitude and ϖ is an effective phase. Fig. 1 shows the orientation angles describing the orbit of the perturber involved in the calculation of \bar{A} and ϖ . The three Euler rotation angles from the ecliptic plane are $-\omega_2$, $-\iota_2$, and $-\Omega_2$. With $\theta = \pi/2 - \beta_{\text{sky}}$ and $\phi = \lambda$ as the polar and azimuthal coordinates of the sky location, respectively, we define $C(\theta, \iota_2, \phi, \Omega_2) = \cos \theta \sin \iota_2 + \sin \theta \cos \iota_2 \sin(\phi - \Omega_2)$ and $S(\theta, \phi, \Omega_2) = \sin \theta \cos(\phi - \Omega_2)$. The effective amplitude is then $\bar{A}^2 = C^2 + S^2$ and the effective phase is $\varpi = \omega_2 + \bar{\phi}$, where $\tan \bar{\phi} = \frac{C}{S}$. Due to degeneracies between the parameters needed to determine \mathcal{A}_2 , it is not possible to gain independent information on all of them. Therefore, the sampling algorithm described in Section 4 will compress all of these angles and only sample in \mathcal{A}_2 . For more detailed information on these constructions, see Robson et al. (2018).

The line-of-sight velocity is determined from equation (9) and the inversion of Kepler's equation. The phase contribution from the line-of-sight velocity (v_{\parallel} term in equation 8) is integrated numerically for each system. Within the fast GB waveform generation methodology, the effect of the line-of-sight velocity on the waveform evolves on a time-scale of $P_2 \sim 1$ yr, which means it evolves on a 'slow' time-scale compared to the orbital time-scale of the inner binary. This slow phase change is directly incorporated into the determination of Ψ_{gw} in addition to the frequency evolution of the source and the effect of the LISA orbital motion. With the addition of the phase correction due to the perturber, the acceptable sparsity of the slow time-domain part of the waveform now depends on the value of P_2 . For values of P_2 that are much shorter than ~ 1 yr, the number of samples is increased to properly sample the third-body orbit.

4 BAYESIAN ANALYSIS WITH LISA

Gravitational-wave data analysis is centred around extracting information on a signal, $s(t)$, which is submerged in noise, $n(t)$, with the

full data stream, $d(t)$, consisting of the sum of these two components: $d(t) = s(t) + n(t)$. To perform this extraction, templates, $h(t)$, are designed to match $s(t)$ to the greatest degree possible. In our work, h will be generated directly in the frequency domain with the fast waveform generator discussed in Section 3.

The statistical analysis is governed by Bayes' theorem:

$$p(\bar{\Theta}|d, \Lambda) = \frac{p(d|\bar{\Theta}, \Lambda) p(\bar{\Theta}|\Lambda)}{p(d|\Lambda)}, \quad (11)$$

where $\bar{\Theta}$ is the parametrization of model Λ . In this work, there are two different models: the base template and third-body template. The posterior probability on the parameters of the source is $p(\bar{\Theta}|d, \Lambda)$. The prior density, $p(\bar{\Theta}|\Lambda)$, represents the incorporation of prior knowledge into the posterior distribution on the parameters. The likelihood, $\mathcal{L}(\bar{\Theta}|\Lambda) = p(d|\bar{\Theta}, \Lambda)$, is the gravitational-wave specific quantity that will be computed in the process of determining the posterior distribution. It represents the probability that the observed data stream could be produced by a source with parameters $\bar{\Theta}$ under model assumption Λ . The denominator on the right-hand side of equation (11) is the evidence, which is an integral of the numerator over all of parameter space: $Z(\Lambda) = p(d|\Lambda) = \int_{\bar{\Theta}} \mathcal{L}(\bar{\Theta}|\Lambda) p(\bar{\Theta}|\Lambda) d\bar{\Theta}$. This term is generally intractable to compute exactly in the gravitational-wave case. However, there are approximate methods to compute this term that will be discussed below.

The gravitational-wave likelihood has the shape of the Gaussian distribution on the data stream residual after the template is applied:

$$\ln \mathcal{L} \propto -\frac{1}{2} \langle d - h | d - h \rangle, \quad (12)$$

where $\langle a | b \rangle$ is the noise-weighted inner product between two time-domain datastreams, $a(t)$ and $b(t)$. In the analysis performed for this work, equation (12) is factored: $\langle d - h | d - h \rangle = \langle d | d \rangle + \langle h | h \rangle - 2 \langle d | h \rangle$. From these terms, we get the optimal S/N achievable by template h : $\sqrt{\langle h | h \rangle}$, and the extracted S/N of a template against the data stream: $\langle d | h \rangle / \sqrt{\langle h | h \rangle}$. Assuming the noise is stationary and Gaussian, we can write down the integral that gives the noise-weighted inner product:

$$\langle a | b \rangle = 4 \text{Re} \int_0^\infty \frac{\tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} df, \quad (13)$$

where $\tilde{a}(f)$ is the Fourier transform of $a(t)$: $\tilde{a}(f) = \mathcal{F}\{a(t)\}$; and $S_n(f)$ is the one-sided power spectral density (PSD) in the noise. For the noise PSD, we use the 'SciRDv1' noise curve from the LISA Data Challenges Working Group (LISA Science Study Team 2018). Please note the inner product is really a sum over the inner products in channels A and E . We do not include the T channel in our analysis because it is not sensitive to the signals examined here under the assumed equal-arm length configuration to the orbit of the LISA constellation. In equation (13), the inner product can be maximized over the initial phase parameter by taking the modulus of the complex integral rather than taking just its real value. We will use maximization over initial phase to make a first cut of undetectable systems. This will be discussed further in Section 5.

4.1 Search and parameter estimation with MCMC

Throughout our analysis, we will use Markov Chain Monte Carlo (MCMC) sampling techniques to draw samples from the posterior distribution. Our MCMC sampler (ERYN; Karnesis et al. in preparation) is based on EMCEE (Foreman-Mackey et al. 2013) and uses the parallel tempering scheme from PTMCEE (Vousden, Farr & Mandel

Table 1. ‘Confidence’ levels based on the Bayes Factor, B_{12} (Raftery 1996; Cornish & Littenberg 2007). We consider $2\log B_{12} \geq 5$ to indicate a detectable source.

B_{12}	$2\log B_{12}$	Evidence for model 1
<1	<0	Negative (supports model 2)
1–3	0–2	Not worth more than a bare mention
3–12	2–5	Positive
12–150	5–10	Strong
>150	>10	Very strong

2016). The proposal we use is the Stretch proposal from Goodman & Weare (2010). In Section 5.1, we use these techniques to search for the maximum likelihood of a base template driven only by GR compared against an injection waveform containing the effect of a third body. After locating maximum likelihood values, we use MCMC runs to generate full posterior distributions for computing the evidence via thermodynamic integration. The log of the evidence (relative to the prior) at $T = 1$ is given by (Goggans & Chi 2004; Lartillot & Philippe 2006)

$$\log Z(1) = \log Z(0) + \int_0^1 d\beta \langle \log \mathcal{L} \rangle_\beta, \quad (14)$$

where $\beta = 1/T$ and $\langle \log \mathcal{L} \rangle_\beta$ is the average of the log-likelihood at inverse temperature β . With a normalized prior distribution $\log Z(0) = 0$. In practice, this integral is approximated with a sum across discrete rungs of a temperature ladder between $\beta = 0$ and $\beta = 1$. Then, the Bayes Factor is simply the ratio of the evidence between two models: $B_{12} = Z_1/Z_2$. A Bayes Factor of less (more) than 1 indicates model 2 (1) is favoured. Here, we will compute the log of the Bayes Factor. Table 1 gives basic ‘confidence’ levels for different ranges of B_{12} and $2\log B_{12}$.

The maximum likelihood, posterior distribution, and evidence are determined by using MCMC algorithms with an ensemble of $n_w = 26$ walkers in each of $n_T = 1000$ temperatures (26 000 total walkers). The waveform and likelihood codes are accelerated with Graphics Processing Units (GPU), which allows for a large number of temperatures to be evaluated to better determine the temperature-likelihood profile to be integrated. The target distribution is examined with the cold chain at a temperature of 1. The highest temperature is set to infinity to probe the prior distribution. 100 inverse temperatures ($\beta = 1/T$) are log-spaced between 10^{-10} and 10^{-4} . This set spans the portion of the temperature ladder where the signal is entirely suppressed by the tempering ($S/N/\sqrt{T} \sim 1$). The other 900 inverse temperatures are log-spaced from 10^{-4} to 10^0 to resolve with high density the turnover where the signal is no longer suppressed. This setup also captures in great detail all temperatures where the signal can be found.

All prior distributions on each parameter are treated independently. The amplitude (A) prior is log-uniform between 10^{-24} and 10^{-20} . The initial frequency (f_0) prior is uniform between 0.5 and 20 mHz. The priors on the ϕ_0 and λ are uniform from zero to 2π , while ψ is uniform from zero to π . The inclination angle is uniform in $\cos i_1$. The ecliptic latitude is uniform in $\sin \beta_{\text{sky}}$. The frequency derivative prior is determined separately for each system so that its range spans the derivative values associated with a range of chirp masses at each system’s true initial frequency. The chirp mass range is 10^{-3} to $1.0 M_\odot$, in agreement with the simulated population (cf. Section 2.1).

The third-body angular parameter ϖ is uniform from zero to 2π . The time of periastron passage (T_2) prior is uniform from zero to P_2 for each specific system. The eccentricity prior was allowed to

vary uniformly from 0 to 0.985. The P_2 prior was also uniform. Its lower bound was 0.1 times the injection value. If $P_2 < T_{\text{obs}}/2$ with $T_{\text{obs}} = 4$ yr, the upper bound was chosen to be 2.2 yr. If the third-body period was between 2 and 3.2 yr, the maximum allowable P_2 value was 32 yr. If the injection value was greater than 3.2 yr, the upper bound on P_2 was set to 10 times the injection value. The choice of prior on A_2 is complicated because it depends on a large number of input parameters, some of which are not accessible in the extraction process. Therefore, we chose an expansive A_2 prior given by the 1D marginalized distribution on A_2 from the injection catalogue. This distribution was determined using KALEPY (Kelley 2021). When the P_2 prior limits were adjusted, the A_2 distribution was also changed to reflect the new period bounds. These prior choices were verified throughout the process to ensure they did not affect the results.

5 CBP AND BD CATALOGUE ANALYSIS

5.1 Locating the maximum likelihood point

The determination of the detectability of the SSOs in the catalogue is done through various steps. The first step is to compute the phase-maximized likelihood ($\ln \mathcal{L}_\phi$) between the third-body injection waveform (d) and the GR-only template waveform (h). This computation is performed at parameters of the base template that are exactly equal to the inner binary parameters of the third-body injection (except for ϕ_0 due to the phase maximization). If $\ln \mathcal{L}_\phi$ for a given source in the population is greater than -2 (and, by definition, less than zero), we consider this third-body undetectable. The value of -2 was chosen conservatively based on initial investigations of the evidence-based detectability study described below. We also eliminate sources from consideration if their S/N is less than 10. These two cuts eliminate 8349/12737 [66 per cent] sources from the original catalogue [6278/10287 (61 per cent) BDs and 2071/2450 (85 per cent) CBPs].

Subsequently, the remaining sources are tested with stochastic sampling via MCMC to further refine the determination of the maximum likelihood ($\ln \mathcal{L}_{\text{max}}$) by allowing all parameters to vary throughout the prior volume. This is a necessary operation because, in general, when the template model does not match the injection model, $\ln \mathcal{L}_{\text{max}}$ will not be found at the exact injection parameters. This can be viewed effectively as a ‘burn-in’ or ‘search’ phase: the code searches for combinations of parameters to best mimic the effect on the waveform created by the third-body orbit. The search is run until it converges to a maximum likelihood value. If this value is between -2 and zero, the associated third-body source is considered undetectable. For many sources in the catalogue, but not all, the search does locate a value for $\ln \mathcal{L}_{\text{max}}$ which is noticeably higher than $\ln \mathcal{L}_\phi$ determined in the first step described above. These sources require a bias against the injection parameters in order for the base template to match the effect of the perturber.

Fig. 2 shows the $\ln \mathcal{L}_{\text{max}}$ value found for each system that passed the first cut with $\ln \mathcal{L}_\phi$ at the injection parameters. These are shown for both CBP and BD systems with applicable cuts for each population shown with dashed lines. The mass of the third body is shown on the horizontal axis and the period of the third body is indicated by colour according to the colour bar. The separation in SSO period P_2 is stark and will be discussed further below. It is also clear that the importance of having a larger third-body mass in creating a strong enough waveform perturbation for detection.

This likelihood maximization process removes 1982/4388 [45 per cent] of the remaining sources [272/379 (72 per cent) CBPs and 1710/4009 (43 per cent) BDs], leaving 107 CBPs and 2299 BDs.




Figure 2. Output from the search for likelihood maximization using MCMC techniques. All sources shown have $\ln \mathcal{L}_\phi < -2$ from the initial template comparison between the base template and the third-body template at the injection parameters. The scatter points show, vertically, the $\ln \mathcal{L}_{\max}$ achieved after letting the parameters for the base template vary across the prior domain. The horizontal axis gives the true value of the perturber’s mass. Our chosen mass separation between CBPs and BDs is displayed with the black dashed line ($m_c = 16 M_J$). For CBP systems ($m_c < 16 M_J$), a conservative $\ln \mathcal{L}_{\max}$ cut of less than -2 (magenta dashed line) was applied to choose systems for accurate evidence estimation due to the computation’s high cost. After applying the accurate evidence estimation to all CBP systems with $\ln \mathcal{L}_{\max} < -2$, a conservative cut of $\ln \mathcal{L}_{\max} < -20$ (orange-dashed line) was determined as a reliable detection cut for determining detectable BD systems.

5.2 Estimating the evidence

While full evidence computations of the large number of remaining BDs would be computationally intensive, the smaller number of remaining CBPs allows us to properly estimate the evidence ratio for each system using thermodynamic integration. For this reason, we chose to focus on the posterior properties of the CBP population. We then address the detections of BDs by determining a conservative maximum log-likelihood difference that encapsulates all detectable systems according to the detailed evidence computations of the CBP sources. This allows us to use the results from the last section to make a conservative detectability cut for the BD systems based on the maximum log-likelihood of the GR-only template.

Fig. 3 shows five CBP sources from top to bottom that aid in clarifying the various properties of the evidence computations. These individual sources and their properties will be discussed shortly. The 1D marginalized histogram of the outer period P_2 is shown for each source in the left column. The centre column compares the average log-likelihood over inverse temperature between the two models. The evidence integral (equation 14) as a function of inverse temperature is shown in the right column.

The behaviour of the average log-likelihood over inverse temperature (centre column) is similar for all systems tested. At low β , where the signal is suppressed, the log-likelihoods are effectively equivalent. In the transition range, where the signal is found, there is some unexpected and interesting behaviour. The signal suppression turnover point generally follows proportionally to the S/N of the source. Given the base template is inherently incorrect and the third-body template has no modelling error compared to the injection, it is expected that the average S/N of a given temperature rung will be higher for the third-body template. As can be seen in the centre column of Fig. 3, this is consistently not the case: the base template finds the signal at a slightly lower inverse temperature compared to the third-body template. We believe this is a manifestation of the five-dimension penalty associated with the addition of the third-body parameters. As the inverse temperature moves above this cross-over point, the average log-likelihood of the third-body template remains less than the base template for a stretch of inverse temperature values. Then, for higher inverse temperatures, the average log-likelihood for the third-body template increases and crosses above the base template. Between these two cross-over points, the evidence integral difference accumulates a negative value, indicating favourability of the base model. This log-evidence-difference penalty is roughly -2 to -4 with shorter periods incurring a more negative penalty, which can be observed in the visual width between the two curves at the turnover point. This penalty can be more negative for some systems with shorter periods. The sources that are detectable have a high enough log-likelihood difference between the templates in the cold chain ($\beta = 1$) that the difference integral accrues enough positive effect on the evidence ratio to overcome the dimensionality penalty.

Beyond the general behaviour of the evidence integral, separate consideration must be given towards the systems with $P_2 < T_{\text{obs}}/2$ versus systems with $P_2 > T_{\text{obs}}/2$, which we will refer to as short- and long-period systems, respectively. This middle point represents the Nyquist criterion for proper sampling of the full third-body orbit: at $P_2 < T_{\text{obs}}/2$ the orbital motion of the perturber is sampled well enough to resolve its frequency (or period), generally preserving a more Gaussian structure on the P_2 1D marginalized posterior if the third body is detectable. The short-period binaries provide a strong indication confirming the accuracy of the evidence computations: short-period perturbers that are *detectable* are represented by a P_2 histogram that is a tight Gaussian dropping off until it reaches 10 or fewer samples in a bin on both sides. The first row in Fig. 3 is an example of a P_2 histogram for a detectable short-period source. The histogram shows the log base 10 of the density to show in detail the samples near the edge of the distribution’s tails. Note that it does not extend throughout the prior range. Additionally, in some of these P_2 distributions, there is a slight skew towards larger periods.

An undetectable short-period third body is shown in the second row of Fig. 3. There is a roughly Gaussian distribution around the true value; however, the Gaussian falls off into a background created by the prior distribution, which is not easily seen with linear density on the vertical axis but can be seen clearly with the log scale. In some examples of undetectable short-period sources, artifacts can also be observed at $\text{yr}/2$ and 1 yr due to confusion related to the orbit of the LISA constellation.

Some detectable sources with periods less than $T_{\text{obs}}/4 = 1 \text{ yr}$ will show a tight Gaussian around the true parameter, with a very small percentage of samples scattered throughout the prior range. This is shown in the third row of Fig. 3. There is a key difference qualitatively indicating this source is detectable compared to sources like those shown in the second row: the small number of samples found away from the true peak are confined to only a small number of other bins

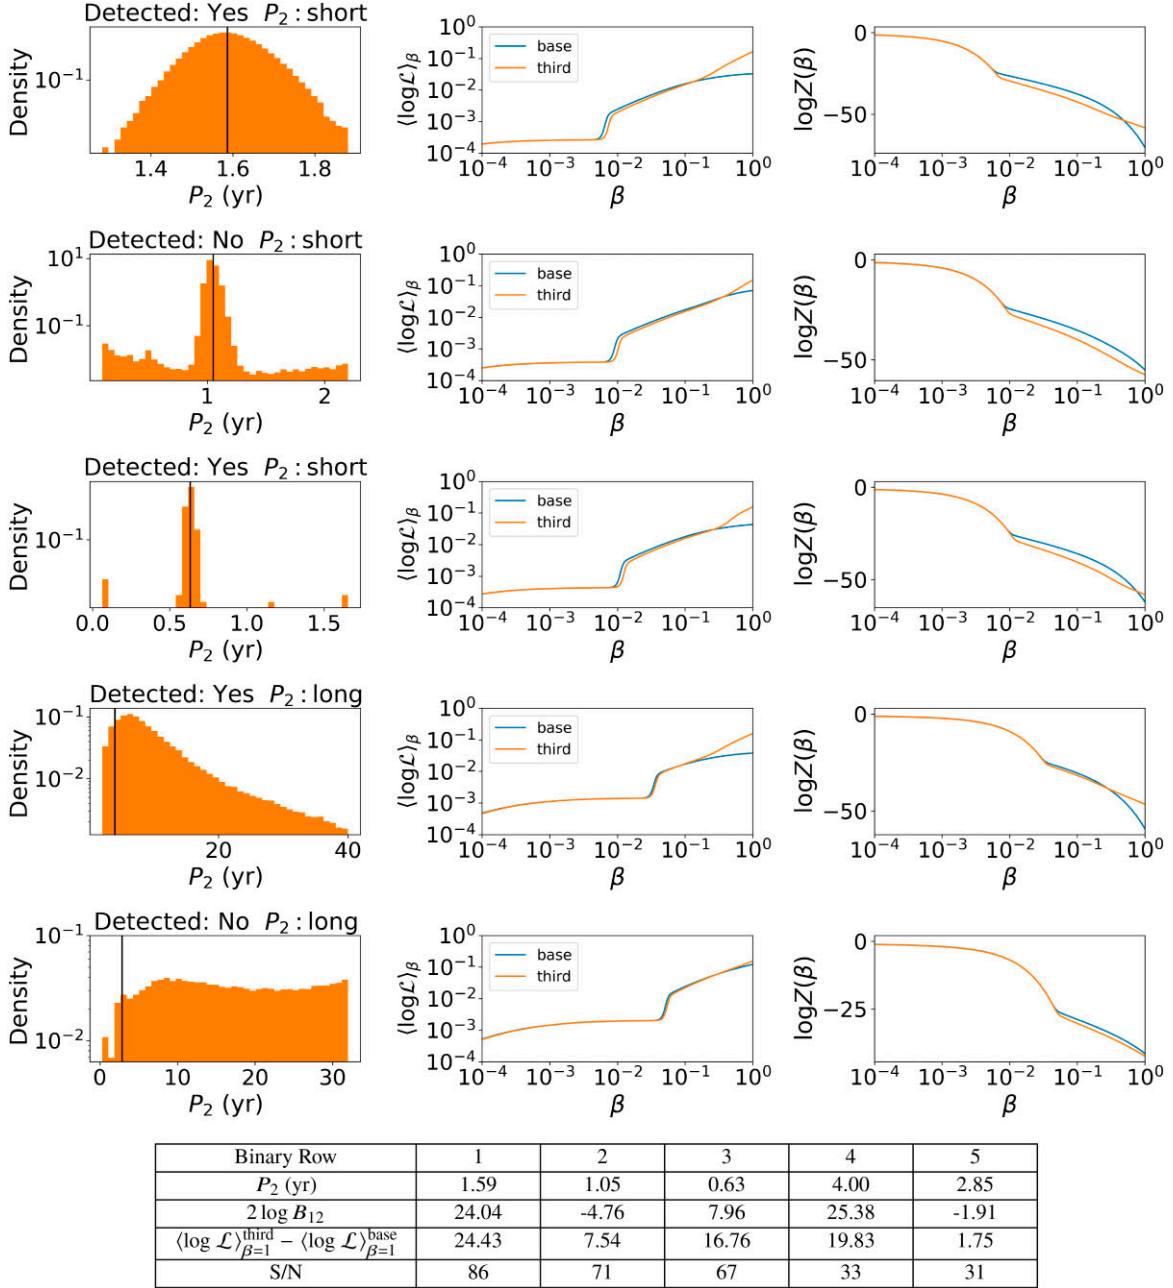


Figure 3. Five examples of results from the accurate evidence estimation analysis. The first column is the 1D marginalized histogram on the third-body orbital period (P_2). Please note the vertical axis is log-scaled. The values of P_2 , $2 \log B_{12}$, difference in the cold-chain ($\beta = 1$) average log-likelihoods between the two templates ($\langle \log \mathcal{L} \rangle_{\beta=1}^{\text{third}} - \langle \log \mathcal{L} \rangle_{\beta=1}^{\text{base}}$), and S/N are given in the table below the figure. The mass and eccentricity (m_c , e_2) of the five systems from top to bottom are (6.11M $_J$, 0.67), (10.87M $_J$, 0.34), (13.71M $_J$, 0.13), (13.01M $_J$, 0.11), and (7.31M $_J$, 0.60), respectively. The detectability of each source is given as (Y)es or (N)o in the titles of left column, followed by whether the third-body period is in the short- or long-period grouping. The second column shows the average log-likelihood at each inverse temperature; it is this curve that is integrated to calculate the evidence. The base template results are shown in blue and the third-body results are shown in orange. The final column illustrates the thermodynamic integral as a function of the inverse temperature. The evidence-related curves are analysed in detail in Section 5.2.

clearly showing there is no continuity in the potential background prior samples. All undetectable sources yield posterior distributions that are non-zero across the prior range.

With $P_2 > T_{\text{obs}}/2$, the ability to resolve the period of the perturber, as well as its Gaussianity in its posterior, is diminished. Representative examples of a detectable and an undetectable long-period source are shown in rows four and five of Fig. 3, respectively. The P_2 histogram of the detectable long-period source is heavily skewed to longer

periods filling the entire prior range. It peaks at longer periods than the true value, which is observed in sources with $P_2 \geq T_{\text{obs}}$. For $T_{\text{obs}}/2 < P_2 \lesssim T_{\text{obs}}$, the distribution peaks at the true value. In long-period cases, where the Nyquist criterion on the third-body orbit is not met, every source is effectively consistent with longer periods until the effect diminishes and the right side of the posterior reaches approximately zero. Undetectable long-period sources (e.g. row five in Fig. 3) show two distinct properties compared to detectable

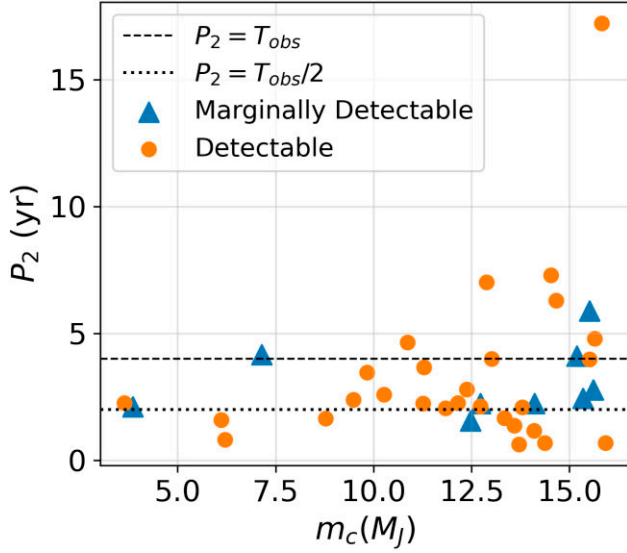


Figure 4. Scatter plot showing the third-body mass (horizontal axis) and period (vertical axis) of detectable (orange dots) and marginally detectable (blue triangles) CBP systems. It must be noted, these are the injection values, not recovered. Horizontal lines correspond to T_{obs} and $T_{\text{obs}}/2$ as indicated in the legend.

sources. The first difference is that the distribution does not tend towards lower counts at longer periods; rather, the histograms tend to turn upward at the end due to the upper boundary on the uniform prior, indicating the signal is not matched well at the large period end of the prior. The second difference, which is specific to sources with $T_{\text{obs}}/2 < P_2 \lesssim T_{\text{obs}}$, is a low-density region appearing at $P_2/2$, which can be seen to the left of the histogram in the fifth row.

The detectable and undetectable short- and long-period events segregate into a few roughly distinct evidence groups. The short-period undetectable binaries are usually found at the lowest evidence values due to the larger penalty at the S/N-temperature turnover point. Long-period undetectable sources populate the evidence spectrum from $-5 \lesssim 2\log B_{12} \lesssim 0$, with marginally detectable long-period sources having evidence values of less than 5 and greater than 0. All of the detectable sources with $2\log B_{12} > 5$ congregate together. The lack of short-period sources found between $-5 \lesssim 2\log B_{12} \lesssim 5$ once again highlights the Nyquist criterion on the third-body orbit: sources with $P_2 < T_{\text{obs}}/2$ are mostly either strongly detectable or strongly undetectable due to the clearer resolution of their orbital frequency.

5.3 Detectable population

We will first examine the detectable CBP population with each source's evidence ratio estimate as the detection metric. We make a conservative cut at $2\log B_{12} > 5$, indicating strong favourability of the third-body template. We find that there are 28/2450 [1 per cent] detectable CBP systems within our modelled catalogue. The planetary mass and orbital periods of the detected and marginally detected ($0 < 2\log B_{12} < 5$) sources are shown in Fig. 4. Most detections are found with $P_2 < T_{\text{obs}}$ with more detections at larger mass, as would be expected since increased mass strengthens the effect of the perturber. The lowest mass detected is $\sim 4 M_J$. We remind the reader that the population detected is a function of the orbital architecture of the injected SSO population, which was built through a single stochastic draw from the various orbital parameter distributions (Section 2.3).

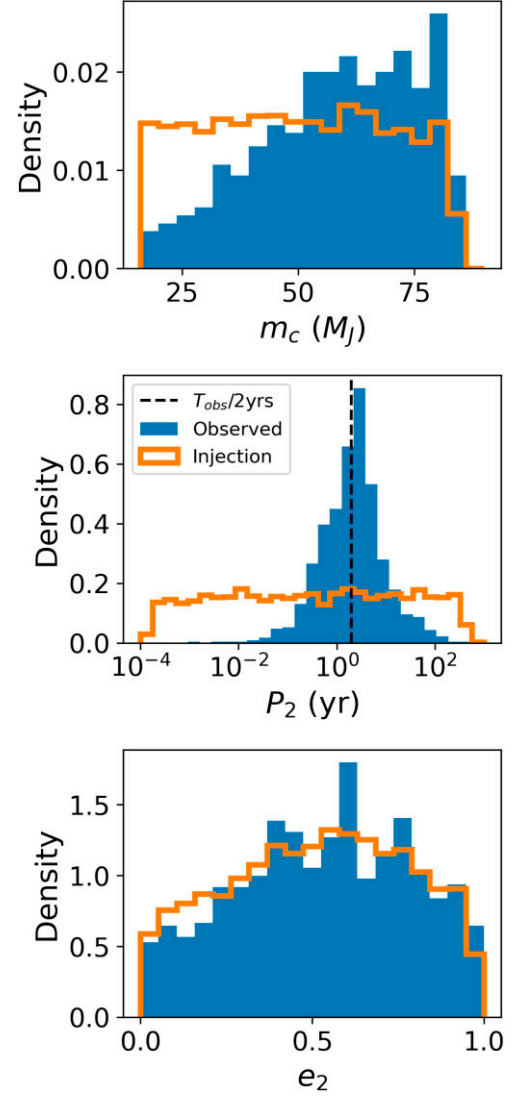


Figure 5. Histograms showing the injection parameters of BD systems around inner DWDs. The mass (m_c), period (P_2), and eccentricity (e_2) of the third body are shown from top to bottom, respectively. The injected population is shown in orange. The detectable population is shown in blue. A future work is needed to understand how these histograms vary over different population models, as well as when examining recovered parameters, rather than injected.

To robustly determine the range of detectable masses, multiple injected populations must be analysed. Such work is currently underway and will be presented in an upcoming paper (Danielski et al. in preparation).

Following the evidence computations on the CBP systems, and given that all CBP sources with $\Delta \ln \mathcal{L} > 20$ had detectable evidence ratios, we decided that, for BD systems, a cut of -20 on the maximum log-likelihood achieved by the base template is sufficient for likely detection. This value was spot-checked with a smaller sample of BD systems. With this conservative cut, 951/10287 [9 per cent] BD systems are detected. This cut probably removes a few hundred detectable sources.

The mass, period, and eccentricity of the detectable BD population is shown in Fig. 5. Please note these are the histograms of the injected values that are detected, not observed values. As expected, the highest number of detections is seen at higher mass and the

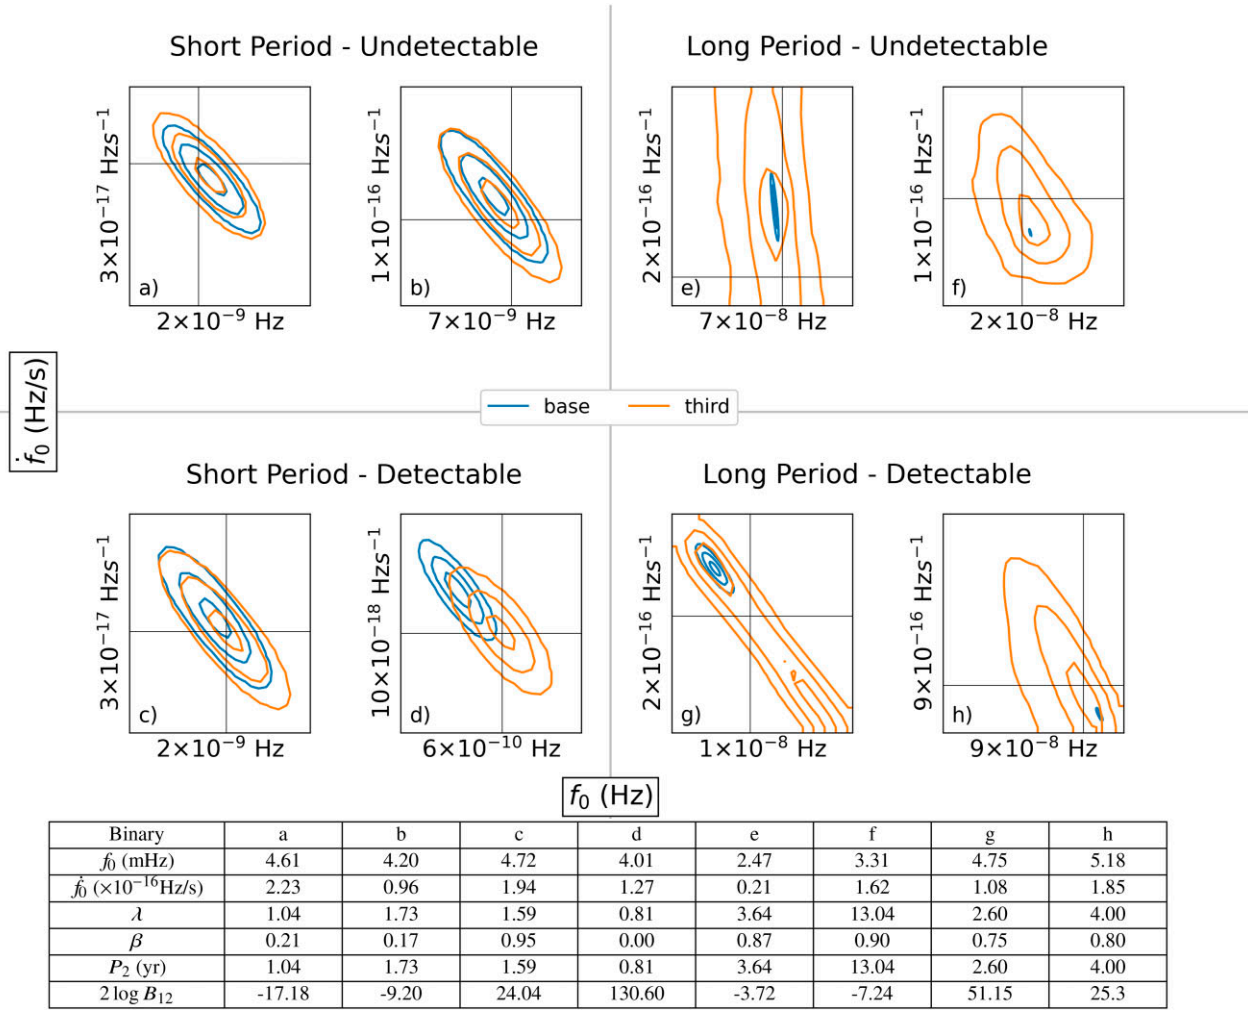


Figure 6. The 2D marginalized posterior distributions (1σ , 2σ , and 3σ contours) in the $f_0 - \dot{f}_0$ plane for eight CBP systems in our catalogue. The distribution for the base GR-only template is shown in blue and the distribution including the third-body effect is shown in orange. The true injection parameters are shown with the horizontal and vertical thin black lines. The injection always consists of the third-body inclusive template. The upper and lower rows display undetectable ($\log B_{12} < 5$) and detectable ($\log B_{12} \geq 5$) systems, respectively. The left two columns are sources with $P_2 < T_{\text{obs}}/2$ ($T_{\text{obs}} = 4$ yr), which represents the Nyquist sampling criterion to completely sample the third-body orbital evolution. The right two columns are third-body objects with periods longer than $T_{\text{obs}}/2$. The pertinent injection parameters and the value of $2 \log B_{12}$ for each source are given in the table. The values given along the horizontal and vertical axis of each plot are the spread in frequency and frequency derivative, respectively. Undetectable sources lead to a third-body posterior distribution that usually centres on the mean of the base template distribution, indicating a bias on recovered parameters will occur whether or not the correct template is used. Detectable short-period sources have a roughly Gaussian character to their posterior distributions in $f_0 - \dot{f}_0$. On the contrary, detectable sources with periods longer than half the observation time have wider and more unpredictable shapes.

period histogram peaks at $P_2 < T_{\text{obs}}$, confirming the trends presented in Tamanini & Danielski (2019) and Danielski et al. (2019). The eccentricity distribution of observed sources does strongly resemble the injected population. Initial checks on BD systems show weaker constraints on the eccentricity (which matches output from the CBP analysis), but a deeper examination is required at a population level. This is a topic for future work (Danielski et al. in preparation).

5.4 CBP posterior analysis

The detailed evidence computations for the CBP systems also produced posterior distributions in the cold chains. Here, we will focus on examining patterns observed in important CBP system parameters across the population by looking at 2D marginalized posterior distributions. Full posterior distributions are available upon request to the authors.

Fig. 6 shows the relation between the frequency and the frequency derivative over eight example systems. The first thing to note is that the maximum log-likelihood locating operation from Section 5.1 preferentially preserves sources with higher frequency and frequency derivative on the inner binary. This is due to the increased inner binary frequency content in the chirping signal that must contend with the doppler shifting from the outer perturber. Higher $f_0 - \dot{f}_0$ sources without any third body tend to have posterior distributions that are multivariate correlated Gaussians. The Gaussian behaviour is still preserved in the posterior distributions on the base GR-only template, even while the distributions are biased away from the true parameters. When examining the bias of the base template, it must be noted that, in general, the base template holds a stronger constraint on the frequency and frequency derivative compared to the third-body template due to the lower dimensionality and lack of the Doppler shifting effect of the perturbing body.

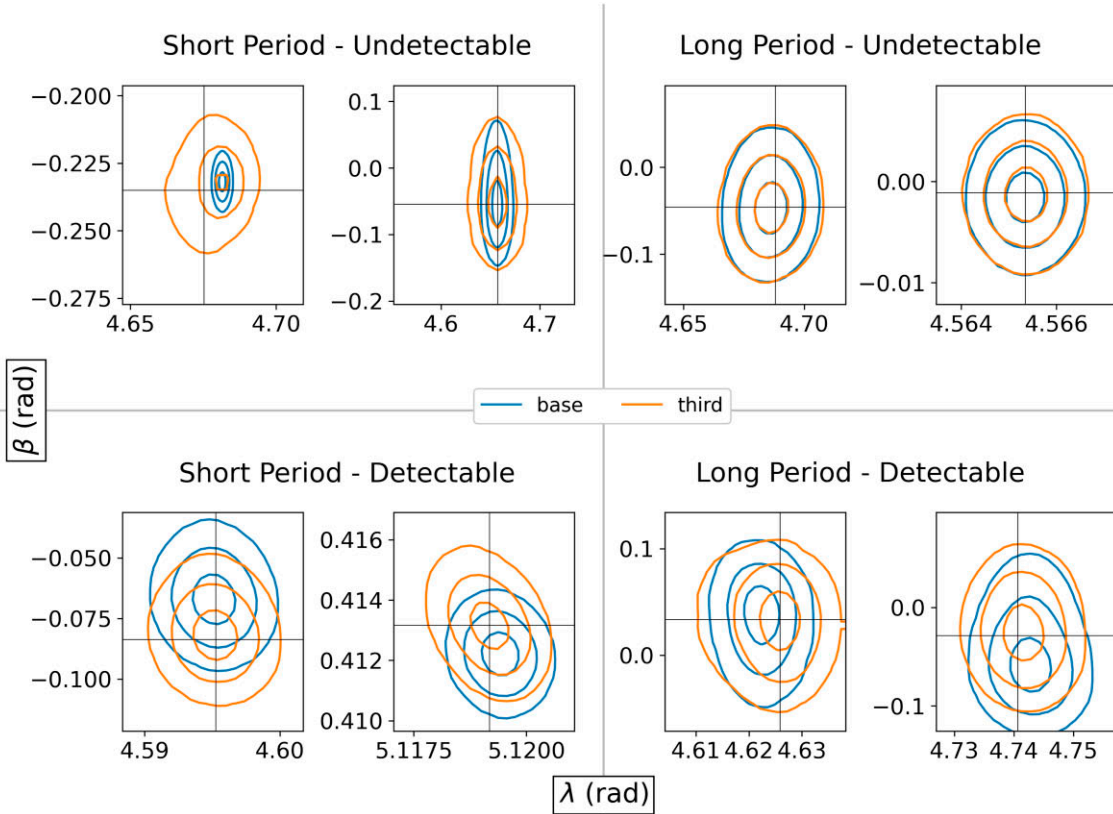


Figure 7. Marginalized posterior distributions in the sky localization. The setup is the same as Fig. 6 with undetectable CBP sources in the top row, detectable sources in the bottom row, short-period ($P_2 < T_{\text{obs}}/2$) systems in the left two columns, and long-period systems ($P_2 > T_{\text{obs}}/2$) in the right two columns. The posteriors for the base template and third-body template are shown in blue and orange, respectively. The true parameters are shown with the thin black vertical and horizontal lines. The values of P_2 , $2\log B_{12}$, and the injected sky locations for each source are given in the table in Fig. 6. All MCMC runs are performed by first injecting the third-body template into the data stream. Undetectable short-period sources centre on the mean of the base-template distribution and can be wider than the base-template distribution due to the confusion of the third-body effect with the Doppler shift caused by the *LISA* motion that aids in probing the ecliptic longitude. The undetectable long-period sources have effectively identical distributions between the two types of templates. Detectable short- and long-period CBP systems usually maintain a similar character between the two different templates, with a bias on the means of the base-template distributions.

The behaviour of the third-body template can once again be separated into groups with short ($P_2 < T_{\text{obs}}/2$) and long ($P_2 > T_{\text{obs}}/2$) periods. The short-period systems (*a*, *b*, *c*, and *d* in Fig. 6) all have multivariate Gaussians for their posteriors. Detectable short-period sources (*c* and *d* in Fig. 6) show a third-body template distribution that is centered on the true point. The base template distribution shows a slight bias from the true point, with the relative magnitude of the bias proportional to the evidence ratio.

Undetectable short-period sources (*a* and *b* in Fig. 6) have distributions that are in general not centered on the true parameters, but centered on, or near, the base template distribution. In these cases, the base template distribution, which inherently exists in a lower dimensional space, pulls the posterior weight away from the higher dimensional third-body template. This can also be directly observed in the maximum likelihoods found in the third-body distributions: even if an MCMC walker was started near the maximum likelihood point with a log-likelihood that was approximately zero, over the course of burn-in the walkers would all drift towards the base template distribution, leaving the maximum log-likelihood in the final MCMC chains to be away from zero and closer to the maximum log-likelihood of the base template distribution. This indicates that, whether we use the correct template model or not, resulting samples will be biased compared to the true parameters.

The long-period sources, which are shown as plots *e*, *f*, *g*, and *h* in Fig. 6, have third-body template behaviour that is no longer Gaussian, once again due to the failure to meet the Nyquist criterion on the sampling of the third-body orbit. The result of this is much wider posteriors with less predictable qualities that are stretched due to the slow doppler shifting caused by the perturber. Undetectable sources (*e* and *f*) tend to peak at the center of the base template distribution. With longer period sources, this is also true for some detectable systems, like the one shown as plot *g* in Fig. 6. This example shows a very unique posterior that has 1σ contours surrounding both the base template distribution and a point that is roughly opposite from the base template distribution across the true point. The true point is contained in the 2σ contour in this case. This once again leads to an issue where either template will create a bias on the extracted parameters, but, in this case, with larger periods, this is even true for detectable third-body systems.

The sky localizations for the same set of sources in Fig. 6 are shown in Fig. 7. For the undetectable short-period sources (*a* and *b*), the third-body template distributions are wider in the ecliptic longitude due to confusion of the Doppler shift from *LISA*'s motion with the Doppler shift from the third-body perturbation. This effect is accompanied by the bias of the third-body distribution towards the centre of the base template distribution, as seen in the undetectable $f_0 - \dot{f}_0$ posteriors. Detectable short-period sources (*c* and *d*) have

roughly the same size posteriors between the two templates with the base template showing some bias from the true parameters. This is also the case for detectable long-period sources (g and h). The long-period systems that are undetectable (e and f) have effectively identical sky distributions between the two templates.

6 DISCUSSION

DWDs containing a companion perturbing object, such as a BD or CBP, lead to many interesting effects on the posterior distributions of these sources. When examining these effects, three different aspects must be considered: (i) the effect on individual source posterior distributions; (ii) the effect on the global fitting of DWD sources in terms of their posterior estimation, as well as the overall convergence of the global fitting algorithm; and (iii) the effects at the population level and the astrophysical implications of these systems. This paper strictly focused on examining (i) in detail, but here we build on our findings to consider expectations for all three important aspects.

6.1 Individual source analysis

For an individual source, the inclusion of the third-body effects are extreme: the posteriors become much more complicated, especially for sources with long third-body periods. At a basic level, this means the time necessary to converge to a final posterior distribution can be greatly increased. It was routinely seen throughout this work that the extreme difference (5–10 \times) in the time necessary for the two different templates to converge to their final posterior states. Additionally, important quantities such as the frequency, frequency derivative, and the sky localization of these sources can become biased, even when using the correct template if the third-body source is undetectable.

When working to determine the chirp mass and distance to a system, it is required that the frequency derivative be constrained. As previously mentioned, including third-body perturbers tends to select the higher frequency and frequency derivative sources from the catalogue for having a noticeable effect on the waveform phasing. This effect on the phasing means the frequency derivative of these sources is generally constrained. Biases on the frequency and its derivative can lead to a biased determination of the chirp mass and distance to the system. This is important because the ‘chirping’ sources with measurable frequency derivative are considered the most astrophysically interesting: these systems contain more information to help draw conclusions on any underlying astrophysical effects on the inner binary.

A bias on the sky localization may affect electromagnetic measurements of these sources depending on how large the bias is and the specific characteristics of the EM observatory. The biases observed on the sky location of the CBP third-body sources are not large, indicating this should not be a major issue. If the sources are roughly located where predicted, a telescope may locate them. This may help to further constrain the other parameters by refining the priors on its sky localization and any other parameters than can roughly be determined from the electromagnetic measurements, such as the frequency, frequency derivative (if observed long enough), and inclination.

These individual source aspects can be similarly expected for the much larger number of BD systems. The BD systems contain more mass in the perturber, indicating their effective biases are likely to be stronger. However, systems with stronger effects will lead to less confusion over detectability, potentially causing issues in a relatively

smaller number of systems compared to the CBP population. This will have to be investigated in more detail in future work.

6.2 Astrophysical population analysis

We expect undetectable CBP third-body systems to have little impact on the overall population modelling tests of two-body DWD populations because the number of sources is low compared to the two-body DWD population and the biases are small relative to the full prior domain over which the population tests will be performed. We also expect this to be the case with any detectable sources as these should be inherently separated out from the population of two-body systems. BD perturbers will likely behave similarly with their stronger effect, ensuring the posterior estimates will favour the third-body system.

More work will be needed to understand the modelling of the third-body population and its residual astrophysical uncertainties it leaves behind in the two-body population. Here, we have analysed an optimistic astrophysical scenario with the underlying astrophysical catalogue, allowing us to study the largest diversity of potential sources. More pessimistic catalogues will deal with an overall smaller effect on the underlying population, decreasing the overall impact on parameter estimation and population analysis.

This work has taken an important step towards third-body population analysis with multiple catalogues by building a pipeline for analysing these large and diverse populations in the most computationally efficient manner. Analysing the entire population from the beginning with full evidence estimates would be computationally intractable and wasteful. This pipeline, which proceeds through the stages described in Section 4, minimizes the necessary computational resources while ensuring that no detectable sources are missed. The undetectable sources do need to be examined though in some detail to confirm their expected minimal effect on the two-body DWD population analysis.

Techniques learned over the course of developing, checking, and performing the evidence computations are also ready to be run on a larger population, which will make this extendable to the full ‘potentially detectable’ BD third-body catalogue. It can also be used directly for other types of DWD astrophysical models to determine the detectability of different effects compared to GR-only-driven inspirals.

6.3 Effects on the global fit

Global fitting is already challenging even when using basic GB models (e.g. Littenberg et al. 2020). In context of the kind of time evolving analysis described in Littenberg et al. (2020), where the solution is built up as new data arrive, triple systems are likely to be first picked up using the standard two-body templates, with the effects of the third body becoming discernible as more data are accumulated. The inclusion of all of the different potential DWD astrophysical effects in the global fit will have to be gathered in steps over time as different effects are suggested and their associated waveforms produced. Each effect will have to be studied at the individual source and population level to further understand the prevalence of each model and its effect on the global fit.

The third-body perturber is clearly a complicated effect and will add confusion and computational burden to the global fitting effort. The overall effect, as well as the effect from the CBP population only, will have to be analysed with more population models in the future to better understand the global fitting of these systems. With that said, if the actual *LISA* data are representative of a more

optimistic population, there may be a strong effect on the global fitting procedure.

As discussed above, the parameters extracted for the undetectable third-body sources may be biased, but it is unlikely that the global fit will choose the third-body template over the base template with regularity while sampling in this case. Similarly, the third-body template is expected to be favoured with regularity in the global fit for detectable sources. Systems that are near the detectability threshold may cause more confusion because the algorithm adopted for the global fit will have to operate ‘switching’ between these templates as the sampling proceeds. This will be easier for systems with shorter periods where the $f_0 - \dot{f}_0$ remains Gaussian for the third-body template. An improvement would be to define new MCMC proposals that would aid the transition between the two templates, which is usually a nuanced task.

6.4 A note on an astrophysical \ddot{f} template

Another template parametrization that includes general astrophysical effects is a template with \ddot{f} as a free parameter rather than the value determined from GR. This template is useful for a subclass of astrophysical effects that chirp smoothly enough to be well-approximated by a quadratic in frequency. Astrophysical tests where this template may be useful include constant acceleration due to a third-body (Bonvin et al. 2017; Inayoshi et al. 2017; Randall & Xianyu 2019; Wong, Baibhav & Berti 2019; Tamanini et al. 2020); matter accretion and dynamical friction (Kremer et al. 2017; Caputo et al. 2020; Cardoso & Maselli 2020; Sberna, Toubiana & Miller 2021); and tests of gravity (Damour, Gibbons & Taylor 1988; Barausse, Yunes & Chamberlain 2016). In cases where the effect cannot be approximated as a quadratic, the \ddot{f} template will at worst fit as well as the base template due to its one-higher degree of freedom.

Any (detectable) shorter period third-body source where the orbit is even close to reaching its full angular range will not be well-approximated with a quadratic. In our detectable CBP population, for all sources but the longest period source ($P_2 \sim 17$ yr), the \ddot{f} template does not perform much differently than the base template. For the longest period source, the \ddot{f} template performs marginally better, but this difference is highly dependant on exactly where in its orbit the third body is when the observation is performed. For BD systems with larger masses, longer periods will be observable with a slower change in the Doppler shift. In some initial investigations, we have observed some of these instances showing the \ddot{f} templates fit the third-body effect equally well to the actual third-body template. Further study will be needed to show if the $f_0 - \dot{f}_0 - \ddot{f}_0$ triplets are consistent with any other astrophysical effects. If they are not consistent with other effects, this template can also act as a means of detectability of the third-body effect; if they are consistent with other effects, it will be harder to be certain about the longer-period BD systems. Analysis with this generalizable template will also be quite nuanced with its evidence being computed in a dimensionality (9) much more similar to the base template (8) than the third-body template (13). Confusion between the \ddot{f} template and other specific model templates may also lead to more delay in the convergence of the global fit algorithm if they have similar matches against the data.

7 CONCLUSION

An interesting potential source for *LISA* is a GB with a third body in orbit around it. In this work, we produced the first fully Bayesian analysis on the detection of and posterior estimation for a population of circumbinary SSOs, orbiting an inner DWD pair. We employed

parallel-tempered MCMC techniques to generate posterior distributions and estimate the Bayesian evidence ratio via thermodynamic integration. This analysis also provided the first examination of these substellar mass sources with eccentricity included in their template parametrization.

The posterior behaviour of these unique sources is highly dependent on two aspects: the detectability of the source and the period of the third-body orbit. The period values generally bifurcate into two categories above (long period) and below (short period) $P_2 = T_{\text{obs}}/2$, which represents the Nyquist criterion associated with sampling the full orbit of the perturbing object. Detectable sources ($2\log B_{12} \gtrsim 5$) with short periods remain fairly Gaussian while long-period systems have more unpredictable distributions due to the inadequate sampling of the third-body orbit. Undetectable sources have similar behaviours in terms of remaining Gaussian, but the third-body posterior distributions tend to centre on the mean of the posterior of the base GR-only template, indicating biased parameter measurements whether or not the true template is used.

These SSO sources came from an optimistic catalogue to provide a larger number of potential sources to examine since we were concentrating on the evidence-based detection and posterior estimation of these systems. The catalogue, while optimistic, was useful because it allowed us to consider a more realistic population that contained a large number of systems generated within a large-dimensional parameter space. We stress again that absolute numbers of detections reported here should not be taken as representative of the expected performance of *LISA*, but at most as an upper optimistic estimate. Another investigation taking into account the variability of the underlying population of CBPs and BDs orbiting *LISA* DWDs is underway to address astrophysical and population-related questions.

The computational cost of this study was considerable. Two routes were taken to ensure tractability. First, the waveform code was reformed for GPU computing capability, providing a large acceleration in the necessary likelihood computations. Secondly, the process of determining detectability through evidence ratio estimates and forming full posterior distributions was strategically reserved for only those sources with the potential of third-body detection, as determined through simpler and less time-consuming waveform matching and maximum likelihood estimation techniques. This allowed us to consider ~ 107 CBPs systems ($m_c \leq 16 M_J$) in detail, 28 of which were detectable. Using these systems as a baseline, we found that this optimistic population catalogue also contained a conservative estimate of 954 detectable BDs ($m_c > 16 M_J$).

We note that the qualitative properties of the detected populations found in our study are consistent with those reported in Tamanini & Danielski (2019) and Danielski et al. (2019); however, a direct comparison of the *LISA* detection efficiency was not possible given the new updated features we included in the injected SSO population.

To further analyse the posterior question related to triple GB systems, our study must be expanded with BDs to the posterior level and then in general to other larger objects that may participate in similar triple systems with an inner DWD pair. For these larger mass systems, it will be particularly interesting to re-examine the longer period regime.

Once the initial posterior estimation is understood, studies of these sources across populations will be needed to inform their potential range of prevalence and all of the effects that come with that quality. The first study of this type will be on CBP systems obtained from different scenarios, as this will be the most computationally reasonable and builds entirely from tools designed here. With both posterior distribution and population studies completed, everything

will have to be combined and implemented within the full *LISA* global fit framework.

The overall procedure presented in this work helps provide an initial sense of source detectability, as determined by full Bayesian methods, and potential effects on the overall global fit when including GBs with generic astrophysical prescriptions. This process lays the groundwork for expanding to new and different astrophysical models related to these binary sources, providing a road map to ensure the success in the global fitting of generic GB systems with *LISA*.

ACKNOWLEDGEMENTS

The authors thank Stas Babak and Antoine Petiteau for helpful initial discussions on this project and Silvia Toonen for providing us with the DWD binary population synthesis model. MLK thanks Jonathan Gair and Lorenzo Speri for helpful discussions. CD thanks Dimitri Veras for the helpful discussions. VK acknowledges support from the Netherlands Research Council NWO (Rubicon 019.183EN.015 grant). NK acknowledges the support from the GR-PRODEX 2019 funding program (PEA 4000132310). This research was supported in part through the computational resources and staff contributions provided for the Quest/Grail high-performance computing facility at Northwestern University. NT acknowledges support from the French space agency CNES in the framework of *LISA*. CD acknowledges financial support from the State Agency for Research of the Spanish MCIU through the ‘Center of Excellence Severo Ochoa’ award to the Instituto de Astrofísica de Andalucía (SEV-2017-0709), and the Group project Ref. PID2019-110689RB-I00/AEI/10.13039/501100011033. This work was supported by the COST action CA16104 ‘Gravitational waves, black holes, and fundamental physics’ (GWverse) through a Short-Term Scientific Mission (STSM) awarded to CD (ID 48268) and VK (ID 48272). COST actions are supported by COST (European Cooperation in Science and Technology). NJC appreciates the support of the NASA LISA Foundation Science grant 80NSSC19K0320. This project has received financial support from the CNRS through the MITI interdisciplinary programs. This paper also employed the use of SCIPY (Jones et al. 2001) and MATPLOTLIB (Hunter 2007).

DATA AVAILABILITY

The waveform codes used in this work are from the publicly available GBGPU package Katz (2022). This package is based on the original fast GB code from Cornish & Littenberg (2007) and the third-body-inclusive fast GB code used in Robson et al. (2018). The maintained and documented version of the new code can be found on GITHUB (<https://github.com/mikekatz04/GBGPU>). The specific code version used in this paper is from an older iteration of GBGPU that can be accessed through the above GITHUB repository here. The code makes use of CUPY (Okuta et al. 2017), NUMPY (van der Walt, Colbert & Varoquaux 2011), CYTHON (Behnel et al. 2011), and a special CYTHON wrapper for CUDA, from McGibbon & Zhao (2012), to parallelize the waveform and likelihood computations for GPUs and then wrap them into PYTHON so they are accessible to the other MCMC and analysis codes used here. The parallelization model involves batching many waveform and likelihood computations together at fixed wall time, therefore, decreasing the per-likelihood cost substantially. During the MCMC runs, typically thousands of likelihoods are computed simultaneously, reducing the per-likelihood computational cost to $\lesssim 5\mu\text{s}$ on an NVIDIA A100 GPU. Please note the codes are written for use on both CPUs and GPUs. For reference, the general per-likelihood computation time on a CPU is ~ 1 ms.

The code for our MCMC sampler (ERYN, Karnesis et al. in preparation) is available upon request to the authors and will eventually be made public.

The search, parameter estimation, and evidence computations required the storage of a large amount of data. These data are also available by request to the authors.

REFERENCES

- Abbott B. P. et al., 2019, *Phys. Rev. X*, 9, 031040
 Abt H. A., 1983, *ARA&A*, 21, 343
 Adams M. R., Cornish N. J., Littenberg T. B., 2012, *Phys. Rev. D*, 86, 124032
 Amaro-Seoane P. et al., 2017, preprint ([arXiv:1702.00786](https://arxiv.org/abs/1702.00786))
 Amaro-Seoane P. et al., 2022, preprint ([arXiv:2203.06016](https://arxiv.org/abs/2203.06016))
 Armstrong J., Estabrook F., Tinto M., 1999, *ApJ*, 527, 814
 Baghi Q., Thorpe J. I., Slutsky J., Baker J., 2021, *Phys. Rev. D*, 103, 042006
 Ballantyne H. A. et al., 2021, *MNRAS*, 507, 4507
 Barausse E., Yunes N., Chamberlain K., 2016, *Phys. Rev. Lett.*, 116, 241104
 Behnel S., Bradshaw R., Citro C., Dalcin L., Seljebotn D. S., Smith K., 2011, *Comput. Sci. Eng.*, 13, 31
 Benacquista M., Holley-Bockelmann K., 2006, *ApJ*, 645, 589
 Boissier S., Prantzos N., 1999, *MNRAS*, 307, 857
 Bonvin C., Caprini C., Sturani R., Tamanini N., 2017, *Phys. Rev. D*, 95, 044029
 Bowler B. P., Blunt S. C., Nielsen E. L., 2020, *ApJ*, 159, 63
 Breivik K. et al., 2020a, *AJ*, 898, 71
 Breivik K., Mingarelli C. M. F., Larson S. L., 2020b, *ApJ*, 901, 4
 Caputo A., Sberna L., Toubiana A., Babak S., Barausse E., Marsat S., Pani P., 2020, *ApJ*, 892, 90
 Cardoso V., Maselli A., 2020, *A&A*, 644, A147
 Cornish N. J., Crowder J., 2005, *Phys. Rev. D*, 72, 043005
 Cornish N. J., Littenberg T. B., 2007, *Phys. Rev. D*, 76, 083006
 Cornish N., Robson T., 2017, *J. Phys. Conf. Ser.*, 840, 012024
 Crowder J., Cornish N. J., 2007, *Phys. Rev. D*, 75, 043008
 Damour T., Gibbons G. W., Taylor J. H., 1988, *Phys. Rev. Lett.*, 61, 1151
 Danielski C., Korol V., Tamanini N., Rossi E. M., 2019, *A&A*, 632, A113
 Dhurandhar S. V., Nayak K. R., Vinet J.-Y., 2002, *Phys. Rev. D*, 65, 102002
 Duchêne G., Kraus A., 2013, *ARA&A*, 51, 269
 Estabrook F. B., Tinto M., Armstrong J. W., 2000, *Phys. Rev. D*, 62, 42002
 Evans R. J., 2020, RA Fisher and the science of hatred. <https://www.newstestman.com/uncategorized/2020/07/ra-fisher-and-science-hatred>
 Fabrycky D., Tremaine S., 2007, *ApJ*, 669, 1298
 Fang X., Thompson T. A., Hirata C. M., 2018, *MNRAS*, 476, 4234
 Fantin N. J. et al., 2019, *ApJ*, 887, 148
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
 Fuhrmann K., Chini R., Kaderhandt L., Chen Z., 2017, *ApJ*, 836, 139
 Georgousi M., Karnesis N., Korol V., Pieroni M., Stergioulas N., 2022, preprint ([arXiv:2204.07349](https://arxiv.org/abs/2204.07349))
 Goggans P. M., Chi Y., 2004, *AIP Conf. Proc.*, 707, 59
 Goodman J., Weare J., 2010, *Commun. Appl. Math. Comput. Sci.*, 5, 65
 Hatzes A. P., Rauer H., 2015, *ApJ*, 810, L25
 Heggie D. C., 1975, *MNRAS*, 173, 729
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Inayoshi K., Tamanini N., Caprini C., Haiman Z., 2017, *Phys. Rev. D*, 96, 063014
 Jones E. et al., 2001 SciPy: Open source scientific tools for Python. <http://www.scipy.org/>
 Kang Y., Liu C., Shao L., 2021, *AJ*, 162, 247
 Karnesis N., Babak S., Pieroni M., Cornish N., Littenberg T., 2021, *Phys. Rev. D*, 104, 043019
 Katz M. L., 2022, [mikekatz04/GBGPU: First official public release!](https://arxiv.org/abs/2203.06016)
 Kelley L. Z., 2021, *J. Open Sour. Softw.*, 6, 2784
 Koester D., Gänsicke B. T., Farihi J., 2014, *A&A*, 566, A34
 Korol V., Rossi E. M., Groot P. J., Nelemans G., Toonen S., Brown A. G. A., 2017, *MNRAS*, 470, 1894
 Korol V., Rossi E. M., Barausse E., 2019, *MNRAS*, 483, 5518

- Korol V., Hallakoun N., Toonen S., Karnesis N., 2022, *MNRAS*, 511, 5936
- Kozai Y., 1962, *AJ*, 67, 591
- Kremer K., Breivik K., Larson S. L., Kalogera V., 2017, *ApJ*, 846, 95
- Kroupa P., Tout C. A., Gilmore G., 1993, *MNRAS*, 262, 545
- Kupfer T. et al., 2018, *MNRAS*, 480, 302
- Lamberts A., Blunt S., Littenberg T. B., Garrison-Kimmel S., Kupfer T., Sanderson R. E., 2019, *MNRAS*, 490, 5888
- Lartillot N., Philippe H., 2006, *Syst. Biol.*, 55, 195
- Lidov M. L., 1962, *Planet. Space Sci.*, 9, 719
- LISA Science Study Team, 2018, LISA Science Requirements Document
- Littenberg T. B., Yunes N., 2019, *Class. Quantum Gravity*, 36, 095017
- Littenberg T., Cornish N., Lackeos K., Robson T., 2020, *Phys. Rev. D*, 101, 123021
- LVK Collaboration, 2018, *Living Rev. Relativ.*, 21, 3
- Martin D. V., Triaud A. H. M. J., 2016, *MNRAS*, 455, L46
- McGibbon R., Zhao Y., 2012, *npcuda-example*
- Nelemans G., Tout C. A., 2005, *MNRAS*, 356, 753
- Nelemans G., Verbunt F., Yungelson L. R., Portegies Zwart S. F., 2000, *A&A*, 360, 1011
- Nelemans G., Yungelson L. R., Portegies Zwart S. F., Verbunt F., 2001, *A&A*, 365, 491
- Nissanke S., Vallisneri M., Nelemans G., Prince T. A., 2012, *ApJ*, 758, 131
- Okuta R., Unno Y., Nishino D., Hido S., Crissman, 2017
- Page J., Littenberg T. B., 2021, *Phys. Rev. D*, 104, 084037
- Portegies Zwart S. F., Verbunt F., 1996, *A&A*, 309, 179
- Portegies Zwart S. et al., 2009, *New Astron.*, 14, 369
- Prince T. A., Tinto M., Larson S. L., Armstrong J. W., 2002, *Phys. Rev. D*, 66, 122002
- Qian S. B., Liu L., Zhu L. Y., He J. J., Yang Y. G., Bernasconi L., 2011, *AJ*, 141, 151
- Raftery A. E., 1996, *Practical Markov Chain Monte Carlo*. Chapman and Hall, London
- Raghavan D. et al., 2010, *ApJS*, 190, 1
- Randall L., Xianyu Z.-Z., 2019, *ApJ*, 878, 75
- Robson T., Cornish N., 2017, *Class. Quantum Gravity*, 34, 244002
- Robson T., Cornish N. J., Tamanini N., Toonen S., 2018, *Phys. Rev. D*, 98, 064012
- Roche É., 1849, *Mém. Sec. Sci. (Montpellier: Académie des sciences et lettres de Montpellier)*, 1, 243
- Sberna L., Toubiana A., Miller M. C., 2021, *ApJ*, 908, 1
- Spiegel D. S., Burrows A., Milsom J. A., 2011, *ApJ*, 727, 57
- Tamanini N., Danielski C., 2019, *Nature Astron.*, 3, 858
- Tamanini N., Klein A., Bonvin C., Barausse E., Caprini C., 2020, *Phys. Rev. D*, 101, 063002
- Thompson T. A., 2011, *ApJ*, 741, 82
- Timpano S. E., Rubbo L. J., Cornish N. J., 2006a, *Phys. Rev. D*, 73, 122001
- Timpano S. E., Rubbo L. J., Cornish N. J., 2006b, *Phys. Rev. D*, 73, 122001
- Tinto M., Armstrong J. W., 1999, *Phys. Rev. D*, 59, 102003
- Tinto M., Dhurandhar S. V., 2005, *Living Rev. Relativ.*, 8, 4
- Tinto M., Dhurandhar S. V., 2021, *Living Rev. Relativ.*, 24, 1
- Tokovinin A., 2014, *AJ*, 147, 87
- Tokovinin A., Thomas S., Sterzik M., Udry S., 2006, *A&A*, 450, 681
- Toonen S., Nelemans G., Portegies Zwart S., 2012, *A&A*, 546, A70
- Toonen S., Hollands M., Gänsicke B. T., Boekholt T., 2017, *A&A*, 602, A16
- Toonen S., Perets H. B., Hamers A. S., 2018, *A&A*, 610, A22
- Vallisneri M., 2005, *Phys. Rev. D*, 71, 22001
- Vallisneri M., Bayle J.-B., Babak S., Petiteau A., 2021, *Phys. Rev. D*, 103, 082001
- van der Sluys M. V., Verbunt F., Pols O. R., 2006, *A&A*, 460, 209
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Veras D., Wyatt M. C., Mustill A. J., Bonsor A., Eldridge J. J., 2011, *MNRAS*, 417, 2104
- Vousden W., Farr W., Mandel I., 2016, *MNRAS*, 455, 1919
- Wilhelm M. J. C., Korol V., Rossi E. M., D’Onghia E., 2021, *MNRAS*, 500, 4958
- Wong K. W. K., Baibhav V., Berti E., 2019, *MNRAS*, 488, 5665

This paper has been typeset from a \LaTeX file prepared by the author.