



HAL
open science

A Sun-like star orbiting a black hole

Kareem El-Badry, Hans-Walter Rix, Eliot Quataert, Andrew W. Howard,
Howard Isaacson, Jim Fuller, Keith Hawkins, Katelyn Breivik, Kaze W. K.
Wong, Antonio C. Rodriguez, et al.

► **To cite this version:**

Kareem El-Badry, Hans-Walter Rix, Eliot Quataert, Andrew W. Howard, Howard Isaacson, et al..
A Sun-like star orbiting a black hole. Monthly Notices of the Royal Astronomical Society, 2022,
10.1093/mnras/stac3140 . insu-03849364

HAL Id: insu-03849364











<https://insu.hal.science/insu-03849364>

Submitted on 7 Apr 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.

A Sun-like star orbiting a black hole

Kareem El-Badry ^{1,2,3}★ Hans-Walter Rix,³ Eliot Quataert ⁴, Andrew W. Howard,⁵ Howard Isaacson,^{6,7} Jim Fuller ⁵, Keith Hawkins ⁸, Katelyn Breivik,⁹ Kaze W. K. Wong,⁹ Antonio C. Rodriguez,⁵ Charlie Conroy,¹ Sahar Shahaf ¹⁰, Tsevi Mazeh ¹¹, Frédéric Arenou,¹² Kevin B. Burdge ¹³, Dolev Bashi ¹¹, Simchon Faigler,¹¹ Daniel R. Weisz ⁶, Rhys Seeburger ³, Silvia Almada Monter³ and Jennifer Wojno³

¹Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

²Harvard Society of Fellows, 78 Mount Auburn Street, Cambridge, MA 02138, USA

³Max-Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany

⁴Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁵Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

⁶Department of Astronomy, University of California, Berkeley, 501 Campbell Hall #3411, Berkeley, CA 94720, USA

⁷Centre for Astrophysics, University of Southern Queensland, Toowoomba, QLD 4350, Australia

⁸Department of Astronomy, The University of Texas at Austin, 2515 Speedway Boulevard, Austin, TX 78712, USA

⁹Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY, 10010, USA

¹⁰Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel

¹¹School of Physics and Astronomy, Tel Aviv University, Tel Aviv, 6997801, Israel

¹²GEPI, Observatoire de Paris, Université PSL, CNRS, 5 Place Jules Janssen, F-92190 Meudon, France

¹³MIT-Kavli Institute for Astrophysics and Space Research 77 Massachusetts Ave. Cambridge, MA 02139, USA

Accepted 2022 October 27. Received 2022 October 26; in original form 2022 September 14

ABSTRACT

We report discovery of a bright, nearby ($G = 13.8$; $d = 480$ pc) Sun-like star orbiting a dark object. We identified the system as a black hole candidate via its astrometric orbital solution from the *Gaia* mission. Radial velocities validated and refined the *Gaia* solution, and spectroscopy ruled out significant light contributions from another star. Joint modelling of radial velocities and astrometry constrains the companion mass of $M_2 = 9.62 \pm 0.18 M_\odot$. The spectroscopic orbit alone sets a minimum companion mass of $M_2 > 5 M_\odot$; if the companion were a $5 M_\odot$ star, it would be 500 times more luminous than the entire system. These constraints are insensitive to the mass of the luminous star, which appears as a slowly rotating G dwarf ($T_{\text{eff}} = 5850$ K, $\log g = 4.5$, $M = 0.93 M_\odot$), with near-solar metallicity ($[\text{Fe}/\text{H}] = -0.2$) and an unremarkable abundance pattern. We find no plausible astrophysical scenario that can explain the orbit and does not involve a black hole. The orbital period, $P_{\text{orb}} = 185.6$ d, is longer than that of any known stellar-mass black hole binary. The system's modest eccentricity ($e = 0.45$), high metallicity, and thin-disc Galactic orbit suggest that it was born in the Milky Way disc with at most a weak natal kick. How the system formed is uncertain. Common envelope evolution can only produce the system's wide orbit under extreme and likely unphysical assumptions. Formation models involving triples or dynamical assembly in an open cluster may be more promising. This is the nearest known black hole by a factor of 3, and its discovery suggests the existence of a sizable population of dormant black holes in binaries. Future *Gaia* releases will likely facilitate the discovery of dozens more.

Key words: binaries: spectroscopic – stars: black holes.

1 INTRODUCTION

The Milky Way is expected to contain of order 10^8 stellar mass black holes (BHs), an unknown fraction of which are in binaries. The inventory of known and suspected BHs consists of about 20 dynamically confirmed BHs in X-ray binaries, an additional ~ 50 X-ray sources suspected to contain a BH based on their X-ray properties (e.g. McClintock & Remillard 2006; Remillard & McClintock 2006; Corral-Santana et al. 2016), a few X-ray quiet binaries in which a BH

is suspected on dynamical grounds, and an isolated BH candidate discovered via microlensing (Lam et al. 2022; Mroz, Udalski & Gould 2022; Sahu et al. 2022). In X-ray bright systems, a BH accretes material from a close companion through stable Roche lobe overflow or stellar winds. Based on the distance distribution and outburst properties of known BH X-ray binaries, it has been estimated that of order 10^3 such systems exist in the Milky Way (Corral-Santana et al. 2016) – only a tiny fraction of the expected total Galactic BH population.

BH X-ray binaries all have relatively short orbital periods. In Roche lobe overflowing systems with low-mass ($\lesssim 1 M_\odot$) main-sequence donors [low-mass X-ray binaries (LMXBs)], mass transfer

* E-mail: kareem.el-badry@cfa.harvard.edu

only occurs at periods $P_{\text{orb}} \lesssim 1$ d. Systems with massive or evolved donors can overflow their Roche lobes at somewhat longer periods. The longest orbital period for a known BH X-ray binary is 33 d, in GRS 1915+105 (Greiner, Cuby & McCaughrean 2001), where the donor is a red giant. Binary population synthesis models predict that a large fraction of the total BH + normal star binary population may be found in wider binaries (e.g. Breivik, Chatterjee & Larson 2017; Chawla et al. 2022), where there is no significant mass transfer. These longer-period BH binaries are difficult to find because they are not X-ray bright, but they may represent the vast majority of the BH binary population. Searches for dormant BH binaries began even before the identification of the first BHs in X-ray binaries (Guseinov & Zel'dovich 1966; Trimble & Thorne 1969). The intervening decades have witnessed the proliferation of vast spectroscopic and photometric surveys well-suited for finding dormant BHs, but only a few solid candidates have been identified (e.g. Giesers et al. 2018, 2019; Mahy et al. 2022; Shenar et al. 2022a)

Gravitational wave observations have in the last decade also begun to detect large number of binaries containing stellar-mass BHs (e.g. The LIGO Scientific Collaboration 2021). The BHs in these systems likely represent a similarly rare outcome of the binary evolution process to X-ray binaries, but their enormous gravitational-wave luminosities during merger allow them to be detected all throughout the Universe. The formation channels of these merging BHs and their evolutionary relation to local BHs in X-ray binaries are still uncertain.

The *Gaia* mission opens a new window on the Galactic binary population – including, potentially, BHs in binaries – through large-scale astrometric orbit measurements. *Gaia* has been predicted to discover large numbers of BHs in binaries, with the predicted number varying by more than 4 orders of magnitude between different studies (e.g. Breivik et al. 2017; Mashian & Loeb 2017; Andrews, Breivik & Chatterjee 2019; Shao & Li 2019; Wiktorowicz et al. 2020; Chawla et al. 2022; Janssens et al. 2022; Shikauchi, Tanikawa & Kawanaka 2022). The large dispersion in these predictions reflects both inherent uncertainties in binary evolution modelling and different assumptions about the *Gaia* selection function.

The first binary orbital solutions from *Gaia* were recently published in the mission’s third data release (Gaia Collaboration 2022b, a), including about 170 000 astrometric solutions and 190 000 spectroscopic solutions. Early assessments of the BH candidate population in these data sets have been carried out for both spectroscopic solutions (e.g. El-Badry & Rix 2022; Jayasinghe et al. 2022) and astrometric solutions (e.g. Andrews, Taggart & Foley 2022; Shahaf et al. 2022). The DR3 binary sample represents a factor of ~ 100 increase in sample size over all previously published samples of binary orbital solutions, and is thus a promising data set in which to search for rare objects. At the same time, stringent signal-to-noise ratio (SNR) cuts were applied to the sample of orbital solutions that was actually published in *Gaia* DR3. The DR3 binary sample thus represents only a few per cent of what is expected to be achievable in the mission’s data releases DR4 and DR5.

This paper presents detailed follow-up of one astrometric BH binary candidate, which we found to be the most compelling candidate published in DR3. Our follow-up confirms beyond reasonable doubt the object’s nature as binary containing a normal star and at least one dormant BH. The remainder of this paper is organized as follows. Section 2 describes how we identified the source as a promising BH candidate. Section 3 presents the radial velocities (RVs) from archival surveys and our follow-up campaign. In Section 4, we constrain the mass of the unseen companion using the RVs and *Gaia* astrometric solution. Section 5 describes analysis of the luminous star’s spectrum,

including estimates of the atmospheric parameters and abundance pattern, and the non-detection of a luminous companion. Section 6 describes the system’s Galactic orbit, and Section 7 discusses X-ray and radio upper limits. We compare the object to other BHs and BH imposters in Section 8, where we also discuss its evolutionary history. Section 9 discusses constraints on the occurrence rate of wide BH companions to normal stars. Finally, we summarize our results and conclude in Section 10. The appendices provide further details on several aspects of our data and modelling.

2 DISCOVERY

In a search for compact object companions to normal stars with astrometric binary solutions from *Gaia*, we considered all 168 065 sources in the `gaiadr3.nss.two.body.orbit` catalog with purely astrometric (`nss_solution_type = Orbital`) or joint astrometric/spectroscopic (`nss_solution_type = AstroSpectroSB1`) solutions. These solutions describe the ellipse traced on the sky by each source’s G –band light centroid due to binary motion. In the `AstroSpectroSB1` solutions, which are only available for bright ($G \lesssim 13$) sources, RVs and astrometry are fit simultaneously.

Our selection strategy and the candidates we considered are described in Appendix E. In brief, we searched for astrometric solutions with unusually large photocentre ellipses at fixed orbital period, exploiting the facts that (a) massive companions have larger orbits at fixed period due to Kepler’s third law, and (b) dark companions produce larger photocentre wobbles than would luminous companions of the same mass (e.g. van de Kamp 1975). This yielded six initially promising sources. Individual vetting and spectroscopic follow-up showed that in four of the six sources, the astrometric solutions are spurious, making a BH companion unlikely. In one case, the astrometric solution may be correct, but the luminous star is a giant and the orbital period is longer than the *Gaia* DR3 baseline, making the reliability of the solution and the nature of the companion difficult to assess without long-term spectroscopic monitoring. One candidate emerged as very promising, *Gaia* DR3 4373465352415301632. We colloquially refer to the source as *Gaia* BH1.

2.1 Properties of the luminous source

Basic observables of the luminous source are summarized in Fig. 1. It is a bright ($G = 13.77$) solar-type star in Ophiuchus ($\alpha = 17:28:41.1$; $\delta = -00:34:52$). The field (upper left) is moderately but not highly crowded. The nearest companion detected by *Gaia* is 5.1 magnitudes fainter at a distance of 3.9 arcsec. None of the *Gaia*-detected neighbours have parallaxes and proper motions consistent with being bound to the source. In the colour–magnitude diagram, the source appears as a solar-type main-sequence star. The DR3 orbital astrometric solution puts it at a distance of 477 ± 4 pc.

The *Gaia* astrometric solution has a photocentre ellipse with semimajor axis $a_0 = 3.00 \pm 0.22$ mas.¹ Given the constraints on the source’s parallax, this corresponds to a projected photocentre semimajor axis of $a_0/\varpi = 1.44 \pm 0.11$ au. If the photocentre traced the semimajor axis (which it does for a dark companion in the limit of $M_*/M_2 \rightarrow 0$), this would imply a total dynamical mass of

¹Here and elsewhere in the paper, we quantify uncertainties on parameters derived from the *Gaia* solution using Monte Carlo samples from the covariance matrix.

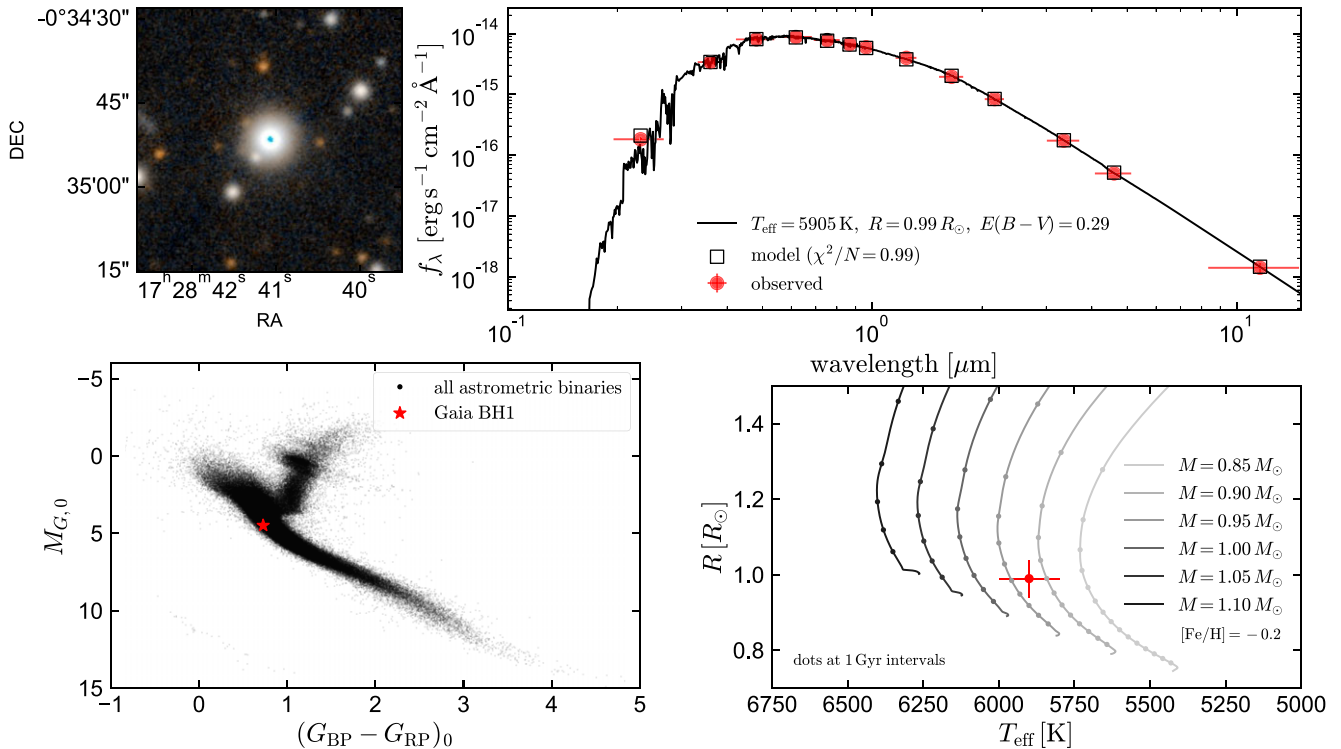


Figure 1. Properties of the luminous star. Upper left shows a 45 arcsec wide z/g PanSTARRS postage stamp centred on the source. Upper right shows the broad-band SED and a single-star fit. Red points show observed photometry; open black squares show the integrated fluxes predicted for the model spectrum (black line). Lower left compares the source to the rest of the *Gaia* astrometric binary sample on the colour–magnitude diagram. Lower right compares the temperature and radius measured from the SED to single-star MIST evolutionary models with $[\text{Fe}/\text{H}] = -0.2$ (as measured from spectroscopy), which suggest a mass of $M_* \approx 0.93 M_\odot$.

$M_{\text{tot}} = 11.5 \pm 2.7 M_\odot$. For larger M_*/M_2 or a luminous secondary, the implied total mass would be larger. If we take the mass of the luminous star to be $M_* = 0.93 \pm 0.05 M_\odot$ (as implied by its temperature and radius; see below), the astrometric mass ratio function (see Shahaf et al. 2019) is $\mathcal{A} = 2.32 \pm 0.17$. This is much larger than the maximum value of $\mathcal{A} \approx 0.6$ that can be achieved for systems with main-sequence components, including hierarchical triples. For a dark companion, this value of \mathcal{A} implies a mass ratio $q = M_2/M_* = 14.2 \pm 2.8$.

We retrieved photometry of the source in the GALEX NUV band (Martin et al. 2005), SDSS u band (Padmanabhan et al. 2008), PanSTARRS $grizy$ bands (Chambers et al. 2016), 2MASS JHK bands (Skrutskie et al. 2006), and WISE $W1 W2 W3$ bands (Wright et al. 2010), and fit the resulting spectral energy distribution (SED) with a single-star model. We set an uncertainty floor of 0.02 mag to allow for photometric calibration issues and imperfect models. We predict bandpass-integrated magnitudes using empirically calibrated model spectral from the BaSeL library (v2.2; Lejeune, Cuisinier & Buser 1997, 1998). We assume a Cardelli, Clayton & Mathis (1989) extinction law with total-to-selective extinction ratio $R_V = 3.1$, and we adopt a prior on the reddening $E(B - V) = 0.30 \pm 0.03$ based on the Green et al. (2019) 3D dust map. We use PYTELLIBS² to interpolate between model SEDs, and PYPHOT³ to calculate synthetic photometry. We then fit the SED using EMCEE (Foreman-Mackey et al. 2013) to sample from the posterior, with the temperature, radius, metallicity, and reddening sampled as free parameters.

The results are shown in the upper-right panel of Fig. 1. A single-star model yields an excellent fit to the data, with $\chi^2/N_{\text{data}} = 0.99$, where N_{data} is the number of photometric points. The inferred temperature and radius correspond to a solar-type star near the main sequence, and evolutionary models then predict a mass of $M_* \approx 0.93 \pm 0.05 M_\odot$. The inferred mass is consistent between MIST (Choi et al. 2016) and PARSEC (Marigo et al. 2017) models within $0.02 M_\odot$. The fact that the source falls near the main sequence provides independent confirmation that the distance inferred from the *Gaia* astrometry is not catastrophically wrong, and suggest that there is no bright companion.

We inspected the ASAS-SN V – and g – band light curves of the source (Kochanek et al. 2017), which contain 3300 photometric epochs over a 10-yr period, with a typical uncertainty of 0.03 mag. This did not reveal any significant periodic or long-term photometric variability. The photometry from ZTF (Bellm et al. 2019), which is more precise but has a shorter baseline, also did not reveal any significant variability.

3 RADIAL VELOCITIES

3.1 Archival data

Gaia BH1 was observed twice by the LAMOST survey (Cui et al. 2012), which obtained low-resolution ($R \approx 1800$) spectra in 2017 and 2019. The LAMOST spectra revealed a main-sequence G star with reported $T_{\text{eff}} = 5863$, $\log g = 4.36$, and $[\text{Fe}/\text{H}] = -0.29$. The two RVs measured by LAMOST were 20.0 ± 4.0 and $8.9 \pm 5.6 \text{ km s}^{-1}$,

²<https://mfouesneau.github.io/pystellibs/>

³<https://mfouesneau.github.io/pyphot/>

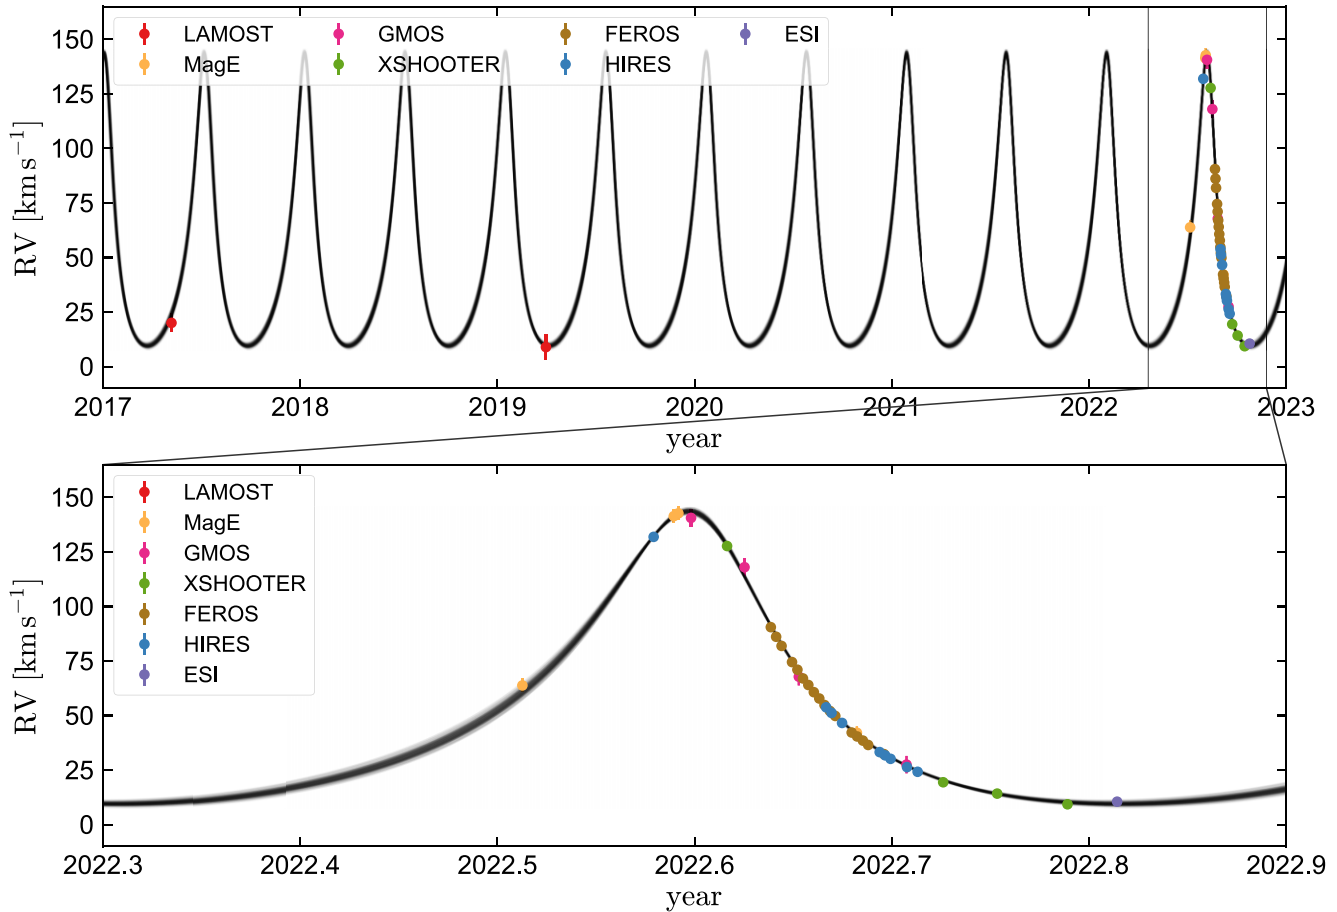


Figure 2. RVs. Points with error bars are measurements; grey lines are draws from the posterior when jointly fitting these RVs and the *Gaia* astrometric constraints. Top panel shows all available RVs, including observations by the LAMOST survey in 2017 and 2019; bottom panel highlights our follow-up in 2022. The best-fitting solution has a period of 186 d, eccentricity 0.45, and RV semi-amplitude of 67 km s^{-1} . Together with the inclination constraint from astrometry, this implies a companion mass of $9.62 \pm 0.18 M_{\odot}$.

both not far from the mean RV of 23.0 ± 2.6 reported in *Gaia* DR3,⁴ and consistent with no RV variation at all. However, the LAMOST observations both happened to occur at times when the astrometric orbit predicted the luminous star would be near apastron (Fig. 2), and thus did not rule out the *Gaia* astrometric solution. Analysis of the *Gaia* scanning law (Appendix B) showed that most of the *Gaia* observations of the source also occurred near apastron. We thus initiated a spectroscopic follow-up campaign.

3.2 Spectroscopic follow-up

We obtained follow-up spectra using several instruments. Details about the observing setup, data reduction, and analysis for each instrument are listed in Appendix A. The first follow-up observation yielded an RV of $\approx 64 \text{ km s}^{-1}$, which was clearly different from the LAMOST and *Gaia* RVs and suggested that the *Gaia* astrometric solution might be correct. We obtained 39 spectra over the course

⁴For sources fainter than $G_{\text{RVs}} = 12$, the *Gaia* RV and its uncertainty are calculated from the peak and curvature of the CCF, which is constructed as the average of the CCFs from all individual visits. The *Gaia* RV thus approximately represents the mode of the RVs at all times the source was observed.

of 4 months, as described in Appendix A and summarized in Fig. 2. The follow-up RVs span most of the orbit’s predicted dynamic range and broadly validate the *Gaia* solution.

We measured RVs via cross-correlation with a synthetic template spectrum, which we took from the BOSZ library of Kurucz model spectra (Kurucz 1970, 1979, 1993) computed by Bohlin et al. (2017). The instruments we used have spectral resolutions ranging from $R \approx 4300$ to $R \approx 55\,000$, and in most cases the data have high SNR ($\gtrsim 20$ per pixel; see Table D1). Because we used several different instruments, with different line spread functions and wavelength coverage, the RV uncertainty is dominated by uncertainties in wavelength calibration and zeropoint offsets between different spectrographs. When possible, we performed flexure corrections using sky lines to minimize such offsets, and we used telluric and interstellar absorption lines to verify stability and consistency of the wavelength solutions.

Including calibration uncertainties, the per-epoch RV uncertainties range from 0.1 to 4 km s^{-1} . The RV zero point is set by telluric wavelength calibration of the Keck/High Resolution Echelle Spectrometer (HIRES) spectra, as described by Chubak et al. (2012). This brings the RV zero point to the Nidever et al. (2002) scale within $\sim 0.1 \text{ km s}^{-1}$. Together with the two archival LAMOST RVs, our follow-up RVs are sufficient to fully constrain the orbit, even without including information from the astrometric solution (Section 4.3).




Figure 3. Left-hand panel shows predicted motion of Gaia BH1’s photocentre on the sky over a 6-yr window. Middle and right-hand panels show the predicted orbit with parallax and proper motion removed. Middle panel shows joint constraints from RVs and astrometry; right-hand panel shows constraints from the *Gaia* astrometric solution alone.

4 COMPANION MASS

We explore three different approaches to constraining the orbit and companion mass: (a) simultaneously fitting the RVs and *Gaia* astrometric constraints, (b) using only the *Gaia* astrometric constraints, and (c) fitting only the RVs.

4.1 Joint astrometric + RV orbit fitting

The *Gaia* orbital solution is parametrized as joint constraints on 12 astrometric parameters: the five standard parameters for single-star solutions (`ra`, `dec`, `pmra`, `pmdec`, `parallax`) and seven additional parameters describing the photocentre ellipse (`period`, `t_periastron`, `eccentricity`, `a_thiele_innes`, `b_thiele_innes`, `f_thiele_innes`, `g_thiele_innes`). The Thiele–Innes coefficients describe the ellipse orientation and are transformations of the standard Campbell orbital elements (e.g. Halbwachs et al. 2022). Individual-epoch astrometric measurements are not published in DR3.

We refer to the vector of best-fitting astrometric parameters constrained by *Gaia* as μ_{ast} , and to the corresponding covariance matrix as Σ_{ast} . The latter is constructed from the `corr_vec` column in the `gaiadr3.nss_two_body_orbit` table. We include in our joint fit the five standard astrometric parameters as well as the period, eccentricity, inclination, angle of the ascending node Ω , argument of periastron ω , periastron time, centre-of-mass velocity, luminous star mass M_* , and companion mass M_2 . For each call of the likelihood function, we then predict the corresponding vector of astrometric quantities, θ_{ast} , and corresponding likelihood:

$$\ln L_{\text{ast}} = -\frac{1}{2} (\theta_{\text{ast}} - \mu_{\text{ast}})^{\text{T}} \Sigma_{\text{ast}}^{-1} (\theta_{\text{ast}} - \mu_{\text{ast}}). \quad (1)$$

We neglect terms in the likelihood function that are independent of θ_{ast} . The five single-star parameters for our purposes are nuisance parameters (they do not constrain the companion mass), but we include them in our fit as free parameters in order to properly account for covariances in the parameters constrained by *Gaia*. When predicting the parameters describing the photocentre ellipse, we assume the companion is dark. We predict the Thiele–Innes coefficients using the standard transformations (e.g. Binnendijk 1960).

We additionally predict the RVs of the luminous star at the array of times at which spectra were obtained, t_i . This leads to a RV term in the likelihood,

$$\ln L_{\text{RVs}} = -\frac{1}{2} \sum_i \frac{(\text{RV}_{\text{pred}}(t_i) - \text{RV}_i)^2}{\sigma_{\text{RV},i}^2}, \quad (2)$$

where RV_i and $\text{RV}_{\text{pred}}(t_i)$ are the measured and predicted RVs. The full likelihood is then

$$\ln L = \ln L_{\text{ast}} + \ln L_{\text{RVs}}. \quad (3)$$

A more optimal approach would be to fit the epoch-level astrometric data and the RVs simultaneously, but this will only become possible after epoch-level astrometric data are published in DR4. Since the *Gaia* astrometric fits are essentially pure likelihood constraints (i.e. they are not calculated with informative priors), it is possible to mix individual-epoch RV measurements and the *Gaia* astrometric constraints, without risk of multiplying priors.

We use flat priors on all parameters except M_* , for which we use a normal distribution informed by isochrones, $M_*/M_{\odot} \sim \mathcal{N}(0.93, 0.05)$. We sample from the posterior using EMCEE (Foreman-Mackey et al. 2013) with 64 walkers, taking 3000 steps per walker after a burn-in period of 3000 steps. The results of this fitting are visualized in Figs 2, 3, and C2. We infer an unseen companion mass $M_2 = 9.62 \pm 0.18 M_{\odot}$.

Fig. 2 compares the measured RVs to draws from the posterior. The scatter across posterior samples at fixed time is diagnostic of the uncertainty in the RV solution. Unsurprisingly, uncertainties are larger during phases where there are fewer measured RVs. Overall, the fit is good: we obtain a solution that both matches the observed RVs and predicts an astrometric orbit consistent with the *Gaia* constraints. This did not have to occur, and it speaks to the quality of the astrometric solution. Our constraints from the joint fit are listed in Table 1. Following the *Gaia* convention, the periastron time T_p is reported with respect to epoch 2016.0; i.e. JD 2457389.0.

Fig. 3 shows the predicted astrometric orbit of the photocentre, which we calculate using PYSTROMETRY (Sahlmann 2019). The left-hand panel shows the predicted motion of the source on the plane of the sky over a 6-yr period. The black line (which is actually many narrow lines representing independent posterior samples) shows the total motion due to proper motion, parallax, and orbital motion. The

Table 1. Physical parameters and 1σ uncertainties for both components of Gaia BH1. The inclination, $i \approx 127$ deg, is equivalent to $i \approx 53$ deg for a spectroscopic orbit. We compare constraints on the orbit based on both astrometry and RVs (third block), astrometry alone (fourth block), and RVs alone (fifth block).

Properties of the unresolved source		
RA	α (deg)	262.17120816
Dec.	δ (deg)	-0.58109202
Apparent magnitude	G (mag)	13.77
Parallax	ϖ (mas)	2.09 ± 0.02
Proper motion in RA	μ_α^* (mas yr $^{-1}$)	-7.70 ± 0.020
Proper motion in Dec.	μ_δ (mas yr $^{-1}$)	-25.85 ± 0.027
Tangential velocity	v_\perp (km s $^{-1}$)	61.0 ± 0.5
Extinction	$E(B - V)$ (mag)	0.30 ± 0.03
Parameters of the G star		
Effective temperature	T_{eff} (K)	5850 ± 50
Surface gravity	$\log(g/\text{cm s}^{-2})$	4.55 ± 0.16
Projected rotation velocity	$v \sin i$ (km s $^{-1}$)	< 3.5
Radius	R (R_\odot)	0.99 ± 0.05
Bolometric luminosity	L (L_\odot)	1.06 ± 0.04
Mass	M (M_\odot)	0.93 ± 0.05
Metallicity	(Fe/H)	-0.2 ± 0.05
Abundance pattern	(X/Fe)	Table 2
Parameters of the orbit (astrometry + RVs)		
Orbital period	P_{orb} (d)	185.59 ± 0.05
Photocentre semimajor axis	a_0 (mas)	2.67 ± 0.02
Semimajor axis	a (au)	1.40 ± 0.01
Eccentricity	e	0.451 ± 0.005
Inclination	i (deg)	126.6 ± 0.4
Periastron time	T_p (JD-2457389)	-1.1 ± 0.7
Ascending node angle	Ω (deg)	97.8 ± 1.0
Argument of periastron	ω (deg)	12.8 ± 1.1
BH mass	M_2 (M_\odot)	9.62 ± 0.18
Centre-of-mass RV	γ (km s $^{-1}$)	46.6 ± 0.6
Parameters of the orbit (astrometry only)		
Orbital period	P_{orb} (d)	185.77 ± 0.31
Photocentre semimajor axis	a_0 (mas)	3.00 ± 0.22
Eccentricity	e	0.48 ± 0.07
Inclination	i (deg)	121.2 ± 2.8
Periastron time	T_p (JD-2457389)	-12.0 ± 6.3
Ascending node angle	Ω (deg)	89.6 ± 3.7
Argument of periastron	ω (deg)	-10.3 ± 11.6
Parameters of the orbit (RVs only)		
Orbital period	P_{orb} (d)	185.6 (fixed) 184.8 ± 0.7 (if not fixed)
RV semi-amplitude	K_\star (km s $^{-1}$)	66.7 ± 0.6
Periastron time	T_p (JD-2457389)	2411.8 ± 0.2
Eccentricity	e	0.447 ± 0.005
Centre-of-mass RV	γ (km s $^{-1}$)	47.1 ± 0.6
Argument of periastron	ω (deg)	13.9 ± 1.2
Mass function	f_m (M_\odot)	4.08 ± 0.08

dotted blue line isolates the contribution due to parallax and proper motion alone. The middle and right-hand panels show the predicted orbital motion, with parallax and proper motion removed, based on the joint astrometry + RV fit (middle) or the pure astrometry fit (right; see below). The orbits predicted in the two cases are similar, but the fit with RVs included is better constrained, and has a slightly smaller photocentre semimajor axis.

4.2 Purely astrometric constraints

We also explored fitting the orbit without the constraints from RVs; i.e. simply removing the $\ln L_{\text{RVs}}$ term from equation (3). The results of

this exercise are described in Appendix C, where we also compare the RVs predicted by the astrometry-only solution to the measured RVs. This comparison allows us to assess the reliability of the astrometric solution.

In brief, we find that the measured RVs are consistent with the purely astrometric solution within its uncertainties. The joint RVs + astrometry fit yields a marginally lower companion mass and more face-on orbit than the best-fitting pure astrometry solution, but the astrometric parameter constraints from our joint *Gaia* + RVs fit are consistent with the purely astrometric constraints at the 1.6σ level (Table C1). We also report the purely astrometric constraints in Table 1. The joint RVs + astrometry constraints are much tighter than those from astrometry alone.

4.2.1 Astrometric uncertainties

Although we find no indication that the astrometric solution is unreliable, there is one way in which it is unusual: the uncertainties on the parameters describing the photocentre ellipse are significantly larger than is typical for a source with $G = 13.77$. This is reflected in the uncertainty on the photocentre ellipse semimajor axis, $\sigma_{a_0} = 0.22$ mas. The median uncertainty for sources with $13.6 < G < 13.9$ in the `nss_two_body_orbit` table is only $\sigma_{a_0} = 0.026$ mas, and only 1.4 per cent of sources in that magnitude range have larger σ_{a_0} than Gaia BH1.

The unusually large σ_{a_0} appears to be a result of two factors. First, astrometric solutions with larger a_0 have larger σ_{a_0} (see Appendix G); the median σ_{a_0} for sources in the same magnitude range with $2.5 \leq a_0/\text{mas} \leq 3.5$ is 0.064 mas, and Gaia BH1 is only in the 85th percentile of σ_{a_0} for such sources. Secondly, as we discuss in Appendix B, all *Gaia* observations of the source that contributed to the DR3 solution covered the same $\approx 50\%$ of the orbit, even though they were spread over more than five orbital periods. Despite this, the combination of astrometry and RVs constrains the orbit well, and leads to a tighter constraint on a_0 than is achieved for typical *Gaia* sources at the same apparent magnitude. In any case, a catastrophic problem with the astrometric solution is firmly ruled out by the RVs, which are consistent with the astrometric solution and require a BH-mass companion even without the astrometric constraints.

4.3 Solution based on radial velocities alone

Although the *Gaia* astrometric solution provides strong constraints on the binary's orbit, it is useful to consider constraints based only on the measured RVs, which are insensitive to any possible problems with the *Gaia* solution. To this end, we fit the RVs with a model with the standard six orbital parameters for a single-lined binary. We first search for the maximum-likelihood solution with simulated annealing (e.g. Iglesias-Marzoa, López-Morales & Jesús Arévalo Morales 2015) and then use MCMC to explore the posterior in the vicinity of the maximum-likelihood solution. The dense sampling of RVs in the last 2 months of our observations allows us to robustly constrain the orbit. We first fit the RVs with no period constraint; this yielded a solution with $P_{\text{orb}} = 184.8 \pm 0.7$ d. Since this is consistent with the *Gaia* solution, but the *Gaia* solution is based on a longer time baseline, we then fixed the period to $P_{\text{orb}} = 185.6$ d, as inferred in Section 4.1. The constraints on the RV solution obtained in this way are reported in Table 1, and the best-fitting RV curve is shown in Fig. 4. The reduced χ^2 is 0.54, suggesting that the fit is good and the RV uncertainties are overestimated on average.

Fitting only the RVs yields an RV semi-amplitude and eccentricity consistent with the predictions of the purely astrometric and




Figure 4. Top panel shows the best-fitting orbital solution to the RVs, ignoring the astrometric solution. We show two periods; duplicated datapoints are transparent. Bottom panel shows the resulting constraints on the companion’s mass. Black line assumes a mass of $0.93 M_{\odot}$ for the luminous star as implied by the isochrones; red-dashed line assumes (pathologically) that the star is a massless test particle. Shaded region shows the *Gaia* astrometric inclination constraint, which corresponds to a companion mass of $9.50 \pm 0.20 M_{\odot}$. Even if the *Gaia* astrometric solution were completely wrong and the luminous star had negligible mass, the implied mass of the companion would be $> 4 M_{\odot}$, which is inconsistent with the observed SED and spectra for any luminous companion.

astrometric + RV fits. The main difference is that the inclination is unknown in the pure-RV solution, and the uncertainties on other parameters are somewhat larger. The periastron time inferred in this fit is 13 orbital cycles later than the one reported in the *Gaia* solution, as it corresponds to the periastron passage in 2022 August that is actually covered by our RVs.

The constraint on the companion mass from RVs can be expressed in terms of the spectroscopic mass function,

$$f(M_2)_{\text{spec}} = M_2 \left(\frac{M_2}{M_2 + M_{\star}} \right)^2 \sin^3 i = \frac{P_{\text{orb}} K_{\star}^3}{2\pi G} (1 - e^2)^{3/2}. \quad (4)$$

P_{orb} , K_{\star} , and e are inferred directly from the RVs. Since both the terms $\left(\frac{M_2}{M_2 + M_{\star}} \right)^2$ and $\sin^3 i$ must be less than one, $f(M_2)_{\text{spec}} = 4.08 \pm 0.08$ sets an absolute lower limit on the mass of the unseen companion. This constraint is illustrated in the lower panel of Fig. 4, where we plot the best-fitting orbit and residuals, and the companion mass required to explain the mass function for different inclinations and M_{\star} .

The figure shows that the RVs cannot accommodate a companion mass below $4 M_{\odot}$, independent of the astrometric solution and the mass of the G star. When a plausible mass for the G star is assumed, this limit increases to $5.5 M_{\odot}$. We also show the inclination implied by the *Gaia* orbital solution. Assuming this inclination and $M_{\star} = 0.93 \pm 0.05 M_{\odot}$ leads to a companion mass of $9.50 \pm 0.20 M_{\odot}$, in

good agreement with the value we inferred by simultaneously fitting the RVs and astrometry. The fact that this value is not identical to the $M_2 = 9.62 \pm 0.18 M_{\odot}$ inferred in the joint fit reflects the fact that the joint RV + astrometry fit implies a slightly larger K_{\star} .

4.3.1 Astrometric mass function

To appreciate the additional information provided by the astrometric solution, we can consider the astrometric mass function,

$$f(M_2)_{\text{ast}} = \left(\frac{a_0}{\varpi} \right)^3 \left(\frac{P_{\text{orb}}}{\text{yr}} \right)^{-2} = M_2 \left(\frac{M_2}{M_{\star} + M_2} \right)^2 \quad (5)$$

where the second equality holds only for a dark companion. The quantity after the second equality is equivalent to $f(M_2)_{\text{spec}}$ except for a factor of $\sin^3 i$. In the limit of small uncertainties, $f(M_2)_{\text{ast}}$ sets a dynamical limit on the mass of the companion that is more stringent than the spectroscopic mass function. In particular, $f(M_2)_{\text{ast}}$ is the mass that the orbit would imply if the G star were a massless test particle.

The constraints from astrometry alone yield $f(M_2)_{\text{ast}} = 9.31 \pm 1.17 M_{\odot}$. Those from joint fitting of the RVs and astrometry yield $f(M_2)_{\text{ast}} = 8.00 \pm 0.16 M_{\odot}$. When we adopt $M_{\star} = 0.93 \pm 0.05 M_{\odot}$ and solve for M_2 , the joint constraint yields $M_2 = 9.62 \pm 0.18 M_{\odot}$.

4.3.2 Possibility of underestimated uncertainties

It is worth considering whether the astrometric uncertainties could be underestimated. Given that the *Gaia* single-star astrometric solutions have been shown to have uncertainties underestimated by $\sim 30\%$ near $G = 13.8$ (e.g. El-Badry, Rix & Heintz 2021), and there is modest ($\sim 1.6\sigma$) tension between the joint fit and the pure-astrometry solution (Appendix C) this is not implausible. To explore the possible effects of underestimated astrometric uncertainties on our constraints, we multiplied all the astrometric uncertainties by 2, constructed the covariance matrix assuming the same correlation matrix (this is equivalent to multiplying the covariance matrix by 4), and repeated the joint fit of astrometry and RVs. This yielded a companion mass constraint of $M_2 = 9.54 \pm 0.26 M_{\odot}$ – similar to our fiducial result, and still a rather tight constraint.

5 ANALYSIS OF THE G STAR SPECTRUM

We analysed the spectrum of the G star using several different methods. Most of our analysis used the Keck HIRES spectrum obtained on JD 2459791 using the standard California Planet Survey setup (CPS; Howard et al. 2010). We first fit the spectrum using the empirical `SpecMatch-Emp` code (Yee, Petigura & von Braun 2017), which compares the continuum-normalized rest-frame spectrum to a library of HIRES spectra of FGKM dwarfs and subgiants observed with the standard CPS setup and analysed with traditional spectroscopic methods. It estimates uncertainties based on the dispersion in parameters of objects in the library with spectra most similar to an observed spectrum. `SpecMatch-Emp` yielded an effective temperature $T_{\text{eff}} = 5801 \pm 110$ K and metallicity $[\text{Fe}/\text{H}] = -0.34 \pm 0.09$.

In Fig. 5, we compare the HIRES spectrum of Gaia BH1 to that of its closest match in the library, HD 33636. The parameters of that star, as measured by Valenti & Fischer (2005) from the HIRES spectrum and synthetic spectral models, are $T_{\text{eff}} = 5904$ K, $\log g = 4.43$, $v \sin i = 3.1 \text{ km s}^{-1}$, and $[\text{Fe}/\text{H}] = -0.13$. The region of the spectrum shown is centred on the Mg I b triplet and also contains




Figure 5. Comparison of the HIRES spectrum of Gaia BH1 (gold) and the standard star HD 33636 (blue). Black line shows the difference between the two spectra. Cutout is centred on the Mg I b triplet; other identifiable lines are labelled. The two spectra are very similar, indicating that the surface properties and abundances of the G star are normal for thin-disc stars in the solar neighbourhood.

strong lines of the elements Fe, Ni, Ti, Cr, Cu, Ca, and Y, many of which are labelled. The two spectra are strikingly similar. This simple comparison suggests that the luminous star in Gaia BH1 is a rather unremarkable solar-type star. Quantitative analysis of the spectrum yielded similar conclusions, as described below.

We also fit the spectrum using the HIRES-adapted implementation of *the Cannon* (Ness et al. 2015) that was trained by Rice & Brewer (2020). *The Cannon* uses a library of spectra with known stellar labels (i.e. atmospheric parameters and abundances) to build a spectral model that smoothly interpolates between spectra, predicting the normalized flux at a given wavelength as a polynomial function of labels. This data-driven spectral model is then used to fit for the labels that can best reproduce an observed spectrum through a standard maximum-likelihood method. The *Cannon* fit yielded parameters $T_{\text{eff}} = 5863$ K, $\log g = 4.43$, $v \sin i = 1.34$ km s $^{-1}$, and $[\text{Fe}/\text{H}] = -0.29$, similar to the values inferred by *SpecMatch-Emp*. It also returned constraints on the abundances of 14 metals beside iron. The full set of labels inferred by *the Cannon* can be found in Appendix F.

To investigate constraints on the G star’s projected rotation velocity $v \sin i$, we convolved an $R = 300\,000$ Kurucz spectrum from the BOSZ library with a range of rotational broadening kernels and with the $R \approx 55\,000$ HIRES instrumental broadening kernel, which we model as a Gaussian. We find that we can infer a robust upper limit of $v \sin i < 3.5$ km s $^{-1}$: a higher value would broaden the lines more than observed. However, the actual value of $v \sin i$ is not well-constrained because at lower values of $v \sin i$, instrumental broadening dominates, and the adopted value of $v \sin i$ has little effect on the predicted line profiles. For $R = 1 R_{\odot}$, this translates to a rotation period $P_{\text{rot}} \gtrsim 14 \times (\sin i)$ d. Assuming $\sin i \sim 1$ and using

a canonical gyrochronological scaling relation for solar-type stars (e.g. Skumanich 1972; Mamajek & Hillenbrand 2008), this implies an age of $\gtrsim 1$ Gyr. Application of such scaling relations is of course only appropriate if the G star was not previously tidally synchronized and spun down by interaction with the companion.

5.1 Detailed abundances

We fit the HIRES spectrum using the Brussels Automatic Code for Characterizing High accuracy Spectra (BACCHUS; Masseron, Merle & Hawkins 2016) with the same set up as in Hawkins, Lucey & Curtis (2020b). BACCHUS enables us to derive the stellar atmospheric parameters, including the effective temperature (T_{eff}), surface gravity ($\log g$), metallicity ($[\text{Fe}/\text{H}]$), and microturbulent velocity (v_{micro}) by assuming Fe excitation/ionization balance; i.e. the requirement that lines with different excitation potentials all imply the same abundances. For the set up of BACCHUS, we make use of the fifth version of the Gaia–ESO atomic line list (Heiter et al. 2021). Hyperfine structure splitting is included for Sc I, V I Mn I, Co I, Cu I, Ba II, Eu II, La II, Pr II, Nd II, and Sm II (see more details in Heiter et al. 2021). We also include molecular line lists for the following species: CH (Masseron et al. 2014), and CN, NH, OH, MgH, and C $_2$ (T. Masseron, private communication). Finally, we also include the SiH molecular line list from the Kurucz line lists.⁵ Spectral synthesis for BACCHUS is done using the TURBOSPECTRUM (Alvarez & Plez 1998; Plez 2012) code along with the line lists listed above and the MARCS model atmosphere grid (Gustafsson et al. 2008).

⁵<http://kurucz.harvard.edu/linelists/linesmol/>

Table 2. Parameters of the G star inferred by BACCHUS. N_{lines} is the number of independent absorption lines used to infer each element.

Parameter	BACCHUS Constraint	N_{lines}
T_{eff} (K)	5883 ± 27	
$\log(g/\text{cm s}^{-2})$	4.55 ± 0.16	
v_{micro} (km s^{-1})	0.79 ± 0.05	
[Fe/H]	-0.19 ± 0.01	
[Na/Fe]	-0.07 ± 0.15	10
[Mg/Fe]	-0.08 ± 0.14	9
[Al/Fe]	0.06 ± 0.04	2
[Si/Fe]	-0.04 ± 0.06	11
[Ca/Fe]	0.01 ± 0.05	15
[Ti/Fe]	-0.06 ± 0.15	39
[V/Fe]	-0.19 ± 0.21	14
[Sc/Fe]	0.12 ± 0.05	9
[Cr/Fe]	-0.07 ± 0.09	15
[Mn/Fe]	-0.20 ± 0.03	7
[Co/Fe]	-0.22 ± 0.14	6
[Ni/Fe]	-0.06 ± 0.11	17
[Cu/Fe]	-0.24 ± 0.07	5
[Zn/Fe]	-0.24 ± 0.12	3
[Y/Fe]	-0.23 ± 0.11	5
[Zr/Fe]	-0.04 ± 0.07	5
[Ba/Fe]	0.17 ± 0.04	3
[La/Fe]	0.13 ± 0.10	4
[Nd/Fe]	-0.01 ± 0.18	6
[Eu/Fe]	-0.04 ± 0.10	1
A(Li)	2.30 ± 0.1	1

The stellar atmospheric parameters derived from BACCHUS can be found in Table 2.

Once the stellar parameters were determined, the model atmosphere is fixed and individual abundances were derived using BACCHUS’ ‘abund’ module. For each spectral absorption feature, this module creates a set of synthetic spectra, which range between $-0.6 < [X/\text{Fe}] < +0.6$ dex, and performs a χ^2 minimization between the observed and synthetic spectra. The reported atmospheric $[X/\text{Fe}]$ abundances are the median of derived $[X/\text{Fe}]$ across all lines for a given species. The uncertainty in the atmospheric $[X/\text{Fe}]$ is defined as the dispersion of the $[X/\text{H}]$ abundance measured across all lines for a species. If only 1 absorption line is used, we conservatively assume a $[X/\text{Fe}]$ uncertainty of 0.10 dex. For a more detailed discussion of the BACCHUS code, we refer the reader to section 3 of both Hawkins et al. (2020a, 2020b). We ran BACCHUS on the full HIRES spectrum, after merging individual de-blazed orders and performing a preliminary continuum normalization using a polynomial spline. Further normalization is performed by BACCHUS. The resulting stellar parameters and abundances are listed in Table 2.

Overall, the abundance pattern of the G star is typical of the thin disc. When comparing the star’s position in the $[X/\text{Fe}]$ versus $[\text{Fe}/\text{H}]$ plane to the population of stars in the solar neighbourhood with abundances measured by Bensby, Feltzing & Oey (2014), it falls within the $\sim 1\sigma$ observed scatter for most elements, and within 2σ for all. Given that the G star is expected to have a thin convective envelope (making mixing of accreted material into the interior inefficient), this suggests that it did not suffer much pollution from its companion.

The measured lithium abundance, $A(\text{Li}) = 2.3 \pm 0.1$, is also not unusual. For solar-type stars, lithium abundance can be used as an age indicator, because the surface lithium abundance is depleted over time (e.g. Skumanich 1972; Baumann et al. 2010). The age – $A(\text{Li})$

correlation is well-studied for Sun-like stars (e.g. Ramírez et al. 2012; Carlos, Nissen & Meléndez 2016), and for solar twins, $A(\text{Li}) = 2.3$ corresponds to an age of about 1 Gyr. However, $A(\text{Li})$ varies strongly with both T_{eff} and $[\text{Fe}/\text{H}]$ (higher T_{eff} and lower $[\text{Fe}/\text{H}]$ both produce a thinner convective envelope and thus imply an older age at fixed $A(\text{Li})$), so caution should be taken in applying the relation from solar twins to a star somewhat warmer and more metal-poor than the sun. More robustly, the measured abundance allows us to rule out youth: the star falls well below the Hyades (age ≈ 600 Myr) in the $A(\text{Li})$ versus T_{eff} plane (Takeda et al. 2013), so we can confidently rule out ages younger than this. Comparison of the G star’s effective temperature and radius to isochrones (Fig. 1) suggests an age of $\gtrsim 4$ Gyr.

We find no evidence for pollution of the G star’s photosphere by α -elements synthesized during the companion’s death, as has been reported in some BH companions (e.g. Israelian et al. 1999; González Hernández et al. 2011). The barium abundance is slightly enhanced compared to the solar value, with $[\text{Ba}/\text{Fe}] = 0.17 \pm 0.04$. Strong barium enhancement is often interpreted as a result of accretion of the AGB wind of a binary companion (‘barium stars’; McClure & Woodsworth 1990), usually in stars with white dwarf companions. However, the $[\text{Ba}/\text{Fe}]$ we measure in Gaia BH1 is not high enough to qualify it as a barium star and is probably unrelated to the companion. We also do not find significant enrichment of other s -process elements, or of the r -process element europium.

5.1.1 Limits on a luminous companion

As part of the CPS pipeline, the HIRES spectrum of Gaia BH1 was analysed with the Reamatch tool (Kolbl et al. 2015), which searches for a second peak in the cross-correlation function (CCF) after subtracting the best-fitting template. For Gaia BH1, this yielded a null result, with a single narrow peak in the CCF and no evidence of another luminous star. Kolbl et al. (2015) found that for luminous binaries with typical solar-type primaries, reamatch has a $> 90\%$ detection rate for companions that contribute $> 3\%$ of the light in the optical in a single-epoch HIRES spectrum. These limits of course depend on the spectral type and rotation rate of the secondary, and a rapidly rotating secondary would be harder to detect. However, we find that even a secondary that contributes pure continuum cannot contribute more than 10 per cent of the light in the HIRES spectrum, or it becomes impossible to achieve a good match to the observed line depths with SpecMatch-Emp. This can be recognized simply from the depth of the observed absorption lines (Fig. 5): several lines have depths that reach 10 per cent of the continuum flux. There cannot be another source contributing more than $\approx 10\%$ of the light at these wavelengths, or the spectrum of the G star would have to reach negative fluxes in the absorption lines to produced the observed total spectrum.

We also find no evidence of a luminous companion at longer or shorter wavelengths. Spectra from X-shooter provide good SNR over the full optical and NIR wavelength range, from 3100 to 24 000 Å, and are consistent with a single G star. This and the GALEX + WISE photometry allows us to rule out pathological companions that are hot and small or cool and large.

6 GALACTIC ORBIT

To investigate the past trajectory of Gaia BH1, we used its parallax and proper motion from the Gaia astrometric binary solution, as well as the centre-of-mass RV inferred from the spectroscopic fit, as




Figure 6. Galactic orbit of Gaia BH1, calculated backward for 500 Myr from the measured proper motion and centre-of-mass RV. For comparison, we show the Sun’s orbit calculated over the same period. The orbit is typical of a thin-disc star, ruling out a large natal kick or dynamical formation in a globular cluster.

starting points to compute its orbit backward in time for 500 Myr using GALPY (Bovy 2015). We used the Milky Way potential from McMillan (2017). The result is shown in Fig. 6; for comparison, we also show the orbit of the Sun. The orbit appears typical of a few-Gyr old thin-disc star, with excursions above the disc mid-plane limited to ± 250 pc. The orbit also never comes very close to the Galactic centre, and is not aligned with the orbits of any globular clusters (GCs). These conclusions are insensitive to the assumed Galactic potential and uncertainties in the source’s proper motion or centre-of-mass RV.

7 X-RAY AND RADIO UPPER LIMITS

We checked archival all-sky X-ray and radio surveys to place upper limits on the flux from Gaia BH1. We used the second *ROSAT* all-sky survey (2RXS) catalog (Boller et al. 2016) and the Very Large Array sky survey (VLASS) Epoch 1 CIRADA catalog (Gordon et al. 2020; Lacy et al. 2020), respectively. At the position of Gaia BH1, the nearest *ROSAT* source (unassociated with any optical counterpart) is 2RXS J172552.5-003516, located $26'4''.1$ away. The median separation between sources in *ROSAT* and SDSS sources is about 12 arcsec (Agüeros et al. 2009), which would place Gaia BH1 130σ away if a true association. Therefore, assuming no *ROSAT* detection, we place upper limits on the flux of Gaia BH1 to be the limit of the 2RXS catalog, $\sim 10^{-13}$ erg s $^{-1}$ cm $^{-2}$ (Boller et al. 2016). This corresponds to a luminosity upper limit of $L_X \lesssim 3 \times 10^{30}$ erg s $^{-1}$ in the 0.1–2.4 keV range.

In the radio, the nearest VLASS source (also unassociated with an optical counterpart) is VLASS1QLCIR J172823.79-003035.0, located $6'5''.3$ away. The beam size of VLASS is $2''.5$ (Lacy et al. 2020), which would place Gaia BH1 146σ away. Assuming no VLASS detection, we can place upper limits on the flux of Gaia BH1 to be the 3σ limit of a single epoch of the VLASS catalog, ~ 0.36 mJy, which in the 2–4 GHz range (S-band) corresponds to a flux limit of $\lesssim 7 \times 10^{-18}$ erg s $^{-1}$ cm $^{-2}$ (Lacy et al. 2020). This corresponds to a radio luminosity upper limit of $L_R \lesssim 2 \times 10^{26}$ erg s $^{-1}$ in the S-band.

7.1 Could Gaia BH1 be detected with X-ray or radio observations?

We briefly consider whether an X-ray or radio detection is expected. A rough estimate of the expected accretion rate on to the BH can be obtained under the assumption that a spherically symmetric wind

from the G star is accreted at the Hoyle–Lyttleton rate:

$$\begin{aligned} \dot{M}_{\text{acc}} &= \frac{G^2 M_{\text{BH}}^2 \dot{M}_{\text{wind}}}{v_{\text{wind}}^4 a^2} \\ &= 6 \times 10^{-18} M_{\odot} \text{ yr}^{-1} \left(\frac{M_{\text{BH}}}{10 M_{\odot}} \right)^2 \left(\frac{\dot{M}_{\text{wind}}}{10^{-14} M_{\odot} \text{ yr}^{-1}} \right) \times \\ &\quad \left(\frac{v_{\text{wind}}}{600 \text{ km s}^{-1}} \right)^{-4} \left(\frac{a}{1 \text{ au}} \right)^{-2}. \end{aligned} \quad (6)$$

Here, \dot{M}_{wind} is the wind mass-loss rate of the G star and v_{wind} is the wind’s velocity. The predicted \dot{M}_{acc} is very low: less than $10^{-10} \dot{M}_{\text{edd}}$, where $\dot{M}_{\text{edd}} \sim 2 \times 10^{-7} M_{\odot} \text{ yr}^{-1} (M_{\text{BH}} / [10 M_{\odot}])$ is the Eddington rate with 10 per cent efficiency. Assuming accretion on to the BH results in an X-ray luminosity $L_X = \eta \dot{M}_{\text{acc}} c^2$, the expected flux at Earth is

$$\begin{aligned} F_X &= \frac{\eta \dot{M} c^2}{4\pi d^2} \\ &= 1.1 \times 10^{-15} \text{ erg s}^{-1} \text{ cm}^{-2} \left(\frac{\eta}{0.1} \right) \left(\frac{M_{\text{BH}}}{10 M_{\odot}} \right)^2 \\ &\quad \times \left(\frac{\dot{M}_{\text{wind}}}{10^{-14} M_{\odot} \text{ yr}^{-1}} \right) \left(\frac{v_{\text{wind}}}{600 \text{ km s}^{-1}} \right)^{-4} \left(\frac{a}{1 \text{ au}} \right)^{-2} \left(\frac{d}{500 \text{ pc}} \right)^{-2}. \end{aligned} \quad (7)$$

For radiatively efficient accretion with $\eta \sim 0.1$, this is within the flux limits of a deep *Chandra* ACIS observation, which are of order 10^{-16} erg s $^{-1}$ cm $^{-2}$. Unfortunately, models for advection-dominated accretion flows predict very low radiative efficiencies at the relevant accretion rates (e.g. Narayan & Yi 1995; Quataert & Narayan 1999). For example, extrapolating the models of Sharma et al. (2007) yields an expected $\eta \sim 10^{-6}$, corresponding to $F_X \sim 10^{-20}$ erg s $^{-1}$ cm $^{-2}$. We conclude that an X-ray detection would be surprising, though it cannot hurt to look.

The expected radio luminosity can be estimated by extrapolating the radio–X-ray correlation for quiescent BH X-ray binaries (e.g. Merloni, Heinz & di Matteo 2003; Gallo et al. 2006). This predicts a radio flux of order 10 μ Jy at 8 GHz for $\eta = 0.1$ (below the VLASS limit, but detectable with the VLA), or ~ 1 nJ for $\eta = 10^{-6}$ (not detectable in the near future). Uncertainty in the G star’s mass-loss rate (which could plausibly fall between 10^{-14} and $10^{-13} M_{\odot} \text{ yr}^{-1}$) and wind velocity near the BH (plausibly between 300 and 600 km s $^{-1}$) leads to more than an order of magnitude uncertainty in these estimates, in addition to the uncertainty in η .

8 DISCUSSION

8.1 Comparison to recent BH imposters

Many recent dormant BH candidates in binaries have turned out not to be BHs, but rather (in most cases) mass-transfer binaries containing undermassive, overluminous stars in short-lived evolutionary phases. Gaia BH1 is different from these systems in several ways.

First, the evidence for a BH companion does not depend on the mass of the luminous star, as it did in LB-1, HR 6819, and NGC 1850 BH1, which all contained low-mass stripped stars masquerading as more massive main-sequence stars (Bodensteiner et al. 2020; Shenar et al. 2020; El-Badry & Quataert 2021; El-Badry & Burdge 2022). In Gaia BH1, the mass of the luminous star could be zero, and the mass of the companion would still vastly exceed the limit allowed by the SED for normal stellar companion, the Chandrasekhar mass, and the maximum neutron star mass.

Secondly, the BH nature of the companion does not depend on a low assumed inclination, as was for example the case in NGC 2004 #115 (Lennon et al. 2021; El-Badry, Burdge & Mróz 2022a), LB-1, and NGC 1850 BH1. The inclination of Gaia BH1 could be edge-on, and RVs alone would still require the companion to be a BH. In addition, reliable constraints on the inclination are available from the *Gaia* astrometric solution; such constraints have not been available for any previous candidates.

Thirdly, there is no evidence of ongoing or recent mass transfer (i.e. ellipsoidal variability or emission lines from a disc), as there was in all the candidates discussed above except NGC 2004 #115, as well as other recent candidates with red-giant primaries (e.g. Jayasinghe et al. 2021; El-Badry et al. 2022b). The orbital period is long enough, and the luminous star small enough, that there is no plausible evolutionary scenario in which it was recently stripped and either component is in an unusual evolutionary state. A stripped-star scenario is also disfavored by the unremarkable optical spectrum and surface abundances of the G star, and by its measured surface gravity ($\log[g/(\text{cm s}^{-2})] = 4.55 \pm 0.16$), which implies a mass $M_* = 10^{\log g} R^2 / G = 1.27 \pm 0.54 M_\odot$.

Finally, the luminous star in Gaia BH1 has a luminosity of only $1 L_\odot$, which is hundreds to thousands of times fainter than the luminous stars in all the BH imposters discussed above. This makes it vastly harder to hide any luminous companion in the SED: any normal star with a mass exceeding $\approx 0.4 M_\odot$ would easily be detected in the spectrum.

8.2 Nature of the unseen companion

Considering just the RVs and the properties of the G star, it seems incontrovertible that the companion is a $> 5 M_\odot$ dark object. Given the good agreement between the observed RVs and those predicted by the astrometric solution, we are also inclined to trust the inclination constraint from the astrometric solution, and thus the most likely companion mass is $\sim 10 M_\odot$.

The companion has not been detected in X-rays or at radio wavelengths, and it lacks other observables associated with accreting BHs (e.g. outbursts or rapid flickering). An X-ray or radio detection would be unexpected given the weak stellar wind expected from the G star and wide separation. Constraints on the nature of the companion thus come down to probabilistic arguments conditioned on its inferred mass and low luminosity.

The mass of the companion is consistent, for example, with a single $10 M_\odot$ BH, 2 lower-mass BHs, 5 neutron stars, 10 massive white dwarfs, or 200 brown dwarfs. Scenarios involving more than

two objects inside the G star's orbit are difficult to assemble and unlikely to be dynamically stable. As we discuss in Section 8.4, a scenario in which the dark object is a close binary with total mass $\sim 10 M_\odot$ (containing 2 BHs or a BH and a neutron star) may be plausible. We thus proceed under the assumption that the unseen companion is either a single BH or a binary containing two compact objects, at least one of which is a BH.

8.3 Comparison to other BHs

Fig. 7 compares Gaia BH1 to other stellar-mass BHs with known masses and orbits in the literature. Red and blue points show low- and high-mass X-ray binaries, whose parameters we take from Remillard & McClintock (2006) and the BlackCAT catalog of X-ray transients introduced by Corral-Santana et al. (2016). We take the most recent mass estimate for Cyg X-1 from Miller-Jones et al. (2021). We also show the binaries VFTS 243 (in the LMC; Shenar et al. 2022a) and HD 130298 (in the Milky Way; Mahy et al. 2022), which are both single-lined binaries containing $\sim 25 M_\odot$ O stars and unseen companions suspected to be BHs. Finally, we show in magenta two binaries in the GC NGC 3201 discovered with MUSE (Giesers et al. 2018, 2019) with suspected BH companions. For the detached systems besides Gaia BH1, the inclination is unknown, so only a lower limit on the BH mass can be inferred. Other detached BH candidates exist in the literature (e.g. Qian, Liao & Fernández Lajús 2008; Casares et al. 2014; Khokhlov et al. 2018; Thompson et al. 2019; Gomez & Grindlay 2021), but we do not include them because we consider their status as BHs more uncertain. Similarly, there are other X-ray bright binaries proposed to contain BHs that we do not include (e.g. Cyg X-3, SS 433) because a neutron star is not fully ruled out. We exclude IC 10 X-1 because the BH's mass is very uncertain (Laycock, Maccarone & Christodoulou 2015). We note that for many of the X-ray binaries included in the figure, mass estimates across different studies differ by significantly more than the reported uncertainties in individual studies.

Unsurprisingly, the X-ray bright systems are concentrated at short periods, as they involve accretion from the companion. All the systems with low-mass donors (red symbols) and $P_{\text{orb}} \gtrsim 0.5$ days have donors that are somewhat evolved and overflow their Roche lobes. The X-ray bright systems with OB-type companions are at slightly longer periods and are fed by wind accretion from donors that nearly fill their Roche lobes. The detached systems with OB primaries, VFTS 243 and HD 130298, are at only slight longer periods than the HMXBs, and thus might evolve to form systems similar to Cyg X-1.

The orbital period and companion star in Gaia BH1 are most similar to in NGC 3201 #12560, the long-period BH candidate discovered in the GC NGC 3201 with MUSE (Giesers et al. 2018). That system has an orbital period of 167 d, eccentricity 0.61, and spectroscopic mass function of $3.25 M_\odot$, with an $\approx 0.81 M_\odot$ luminous star. Since that system's inclination is unknown, the companion mass could be any value above $4.5 M_\odot$. The important difference, however, is that NGC 3201 #12560 is in a GC, where dynamical capture, exchange, and disruption of binaries are all efficient. Two apparent BHs were discovered in NGC 3201 among only 3553 stars with multi-epoch RVs. This implies a BH companion rate of order 1000 times higher than that in the local Galactic field (Section 9). Such a dramatic enhancement in the rate of BH companions in GCs can be most readily understood if the binaries are formed dynamically through capture or exchange (e.g. Fabian, Pringle & Rees 1975; Rasio, Pfahl & Rappaport 2000; Kremer et al. 2018), and in any case, the long-




Figure 7. Comparison of Gaia BH1 (black) to known BH X-ray binaries. Red and blue symbols correspond to accreting BHs with low- and high-mass companions. Magenta symbols show detached binaries in the globular cluster NGC 3201, and cyan points show detached binaries in which the luminous star is a high-mass ($\gtrsim 20 M_{\odot}$) star. The properties of Gaia BH1 are most similar to the detached spectroscopic binary NGC 3201 #12560 (Giesers et al. 2018). However, that system is in a globular cluster and thus probably formed through a different channel. The closest system from an evolutionary standpoint is perhaps the LMXB GRS 1915+105, which has a red giant donor and was likely detached when it was a main-sequence star.

term survival of a primordial binary with a 167 d orbit in such a dense stellar system is unlikely.

Given the orbit of Gaia BH1 in the Galactic disc, formation from an isolated binary seems somewhat more likely than a dynamical formation channel. In this sense, the system is qualitatively different from any other binary shown in Fig. 7. The most closely related system may be the BH + giant star ‘microquasar’ GRS 1915+105 ($P_{\text{orb}} = 33$ d; Castro-Tirado, Brandt & Lund 1992; Greiner et al. 2001), which may have formed through a similar channel to Gaia BH1 but with a significantly closer separation, or through a completely different channel.

8.4 Evolutionary history

The progenitor of the BH very likely had a mass of at least $20 M_{\odot}$. For example, the models of Sukhbold et al. (2016) predict that at solar metallicity, an initial mass of $30 - 50 M_{\odot}$ is required to produce a helium core mass of $\gtrsim 9 M_{\odot}$ (their fig. 19). Such a star would reach a radius of order 10 au as a supergiant if it were allowed to evolve in isolation, which is significantly larger than the present-day separation of the G star and BH. This suggests that the two stars likely interacted prior to the formation of the BH. Given the extreme mass ratio, this interaction is expected to lead to a common envelope (CE) episode, in which the G star was engulfed by the envelope of the massive star, ultimately spiraling in and ejecting some or all of the envelope. Given the short lifetimes of massive stars, the G star would still be contracting toward the main sequence as the BH progenitor completed its evolution.

It is not clear that a $\lesssim 1 M_{\odot}$ star can survive a CE episode with a $\gtrsim 20 M_{\odot}$ companion. Indeed, many calculations in the literature

find that ejection of the massive star’s envelope by a low-mass star is impossible because there is not enough available orbital energy (e.g. Portegies Zwart, Verbunt & Ergma 1997; Podsiadlowski, Rappaport & Han 2003; Justham, Rappaport & Podsiadlowski 2006). The same challenge applies to modelling the formation of BH LMXBs, whose evolutionary channels are still poorly understood. We first review the ‘standard’ formation channel through a CE event, and then discuss possible alternatives.

8.4.1 Common envelope channel

In the CE scenario, the initial separation of the BH progenitor and the G star would have been in the range of 5–15 au (Portegies Zwart et al. 1997). The binary would have emerged from the CE episode as a $\leq 1 M_{\odot}$ star in a close orbit with a $\gtrsim 10 M_{\odot}$ helium core, which might also have retained some of its envelope. The standard $\alpha\lambda$ CE prescription predicts that the ratio of the separations before and after the CE episode is (e.g. Webbink 1984)

$$\frac{a_f}{a_i} = \frac{M_c}{M_1} \left(1 + \frac{2M_e}{\alpha\lambda r_L M_2} \right)^{-1}, \quad (8)$$

where M_c and M_e are the mass of the BH progenitor’s core and envelope, $M_1 = M_c + M_e$ is the total mass of the progenitor, r_L is its Roche lobe radius in units of the semimajor axis, and M_2 is the mass of the G star. α and λ are dimensionless parameters, respectively describing the fraction of the G star’s orbital energy that goes into ejecting the BH progenitor’s envelope, and the binding energy of the BH progenitor’s envelope. If we take $M_c = 10 M_{\odot}$, $M_e = 20 M_{\odot}$, $r_L = 0.65$, $M_2 = 1 M_{\odot}$, $\alpha = 1$, and $\lambda = 0.5$ (a rather optimistic choice of λ , e.g. Klencki et al. 2021), we obtain $a_f/a_i = 0.0025$, such

that for an initial separation of $a_i = 2000 R_\odot$, the final separation is $a_f = 5 R_\odot$. This is just wide enough that the final configuration *might* be stable and avoid a merger, but only for a narrow range of a_i . For lower values of $\alpha\lambda$, there are no initial separations that both lead to Roche lobe overflow of the BH progenitor and avoid a merger (e.g. Portegies Zwart et al. 1997). In any case, the predicted final separation is much closer than the observed $a \approx 300 R_\odot$.

Stripped of its envelope, the helium core is expected to drive a prodigious Wolf–Rayet wind (e.g. Woosley, Langer & Weaver 1995), which will both diminish the final mass of the compact remnant and widen the orbit. Exactly how much mass is lost in this stage is uncertain (e.g. Woosley et al. 1995; Fryer & Kalogera 2001), and attempts to measure the relevant wind mass-loss rates observationally are complicated by factors such as clumping (e.g. Smith 2014; Shenar 2022). In Gaia BH1, a strong wind cannot have persisted for very long, or the helium core mass would have fallen below the required $10 M_\odot$. Eventually, the helium core will collapse to a BH, perhaps accompanied by some additional mass loss and/or an asymmetric natal kick.

The wide orbit of Gaia BH1 and the unremarkable surface abundances of the G star strongly suggest that there has not been significant mass transfer from the G star to the BH. We can thus – in contrast with the situation for typical BH LMXBs (e.g. Podsiadlowski, Rappaport & Pfahl 2002; Pfahl, Rappaport & Podsiadlowski 2003; Justham et al. 2006; Fragos & McClintock 2015) – rule out a scenario in which the luminous star was initially significantly more massive than today and has been stripped down to sub-solar mass by Roche lobe overflow. We can similarly rule out a scenario in which the initial mass of the BH was lower than today, and it grew significantly by accretion.

The orbital eccentricity, $e \approx 0.45$, places a constraint on mass loss and natal kicks during the BH’s formation. If the natal kick was purely due to mass loss (and not an asymmetric stellar death) and the orbit was circular before the BH formed (e.g. Blaauw 1961), then the predicted final eccentricity would be

$$e = \frac{\Delta M_{\text{BH}}}{M_{\text{BH}} + M_\star}, \quad (9)$$

such that mass loss of $\Delta M_{\text{BH}} \approx 4.8 M_\odot$ would be required to fully explain the observed eccentricity. Here, ΔM_{BH} represents any mass lost during the BH progenitor’s death that does not fall into the BH. The net velocity imparted to the centre-of-mass in this case is $v_{\text{sys}} = e \times v_{\text{BH prog}}$, where $v_{\text{BH prog}}$ is the orbital velocity of the BH progenitor at the time of its death. Irrespective of the orbital period at the time, $v_{\text{BH prog}}$ was likely rather low given the extreme mass ratio, and so (symmetric) mass loss alone is unlikely to have resulted in a system velocity larger than $\sim 10 \text{ km s}^{-1}$. Loss of $\gtrsim 4 M_\odot$ from the system at the time of the BH’s formation seems somewhat unlikely in the CE channel because the core should collapse rapidly and most of the envelope should have been removed by the CE episode, but the details of how the CE episode would proceed are uncertain.

On the other hand, if a kick occurred mainly due to asymmetries during the BH progenitor’s death (e.g. asymmetric ejecta or neutrino emission), the final eccentricity depends on the kick velocity and direction, and on the period at the end of the CE episode (e.g. Brandt & Podsiadlowski 1995; Hurley, Tout & Pols 2002). In the limit of no mass loss, the systemic velocity induced by a kick is simply $v_{\text{kick}} \times q/(1+q)$, which in the limit of $q \gg 1$ is close to v_{kick} . Here $q = M_2/M_\star$.

Forming Gaia BH1 via a CE is quite challenging for two reasons: (1) the G star may not have had enough orbital energy to eject the envelope of its much more massive companion (this is generically

a problem for BH LMXB formation models, but is particularly problematic here because the G star cannot have been born with a higher mass than observed today), and (2) if CE ejection was successful, the post-CE separation is expected to be significantly closer ($\sim 5 R_\odot$) than the wide orbit observed today ($\sim 300 R_\odot$) under standard assumptions for CE ejection efficiencies. A fine-tuned natal kick to the BH could in principle widen the orbit, but this would also make the final orbit highly eccentric and impart a large systemic velocity to the binary’s centre of mass, both of which are not observed. For example, assuming an optimistic post-CE separation of $20 R_\odot$ and random kick orientations, we find (e.g. Brandt & Podsiadlowski 1995) that kicks of order 200 km s^{-1} are required to widen the orbit to $a \gtrsim 200 R_\odot$. Such kicks result in a minimum eccentricity of 0.90 and a median of 0.95, both much larger than observed.

Given the seemingly impossible nature of forming Gaia BH1 through the CE channel under standard assumptions, we performed a study of more extreme assumptions for CE and BH natal kicks to determine the required combination of assumptions to produce a binary with the observed properties of Gaia BH1. We follow the methods introduced in Wong et al. (2022) that uses COSMIC, a rapid binary population synthesis code, to explore all possible combinations of Zero Age Main Sequence (ZAMS) properties, CE assumptions, and natal kick assumptions; for an in-depth discussion of how COSMIC evolves binary-star populations, see Breivik et al. (2020).

Instead of selecting a single set of assumptions for CE and natal kicks, we use emcee to sample the posterior of the masses, orbital period, and eccentricity of the binary hosting Gaia BH1 with combinations of the ZAMS mass, orbital period, and eccentricity, as well as the CE ejection efficiency (α) and the BH natal kick strength (v_{kick}) and direction (θ, ϕ) and the mean anomaly of the orbit at the time of core collapse (\mathcal{M}) as model parameters. We place uniform priors on all ZAMS binary parameters as well as the CE and natal kick assumptions. We assign wide limits for the ZAMS primary mass prior between $1 M_\odot$ and $150 M_\odot$ and narrow limits for the ZAMS secondary mass prior between $0.5 M_\odot$ and $2 M_\odot$ since the G star is not expected to have significantly gained or lost mass during the binary evolution. The orbital period and eccentricity priors are limited to 500 and 6000 d and 0 and 1, respectively. We limit our prior on α to be between 0.1 and 20 and our prior on v_{kick} to be between 0 km s^{-1} and 300 km s^{-1} based on the relatively wide and moderately eccentric orbital configuration of Gaia BH1’s binary host. We place uniform priors on the unit sphere for the natal kick angles, θ, ϕ , and all allowed values of the mean anomaly, \mathcal{M} , at core collapse ($0\text{--}360^\circ$). Finally, we fix the metallicity of all simulated binaries to $Z = 0.63 Z_\odot$, where $Z_\odot = 0.02$, consistent with $[\text{Fe}/\text{H}] = -0.2$ under the assumption that the G star follows the solar abundance pattern.

Since the age of the observed G star is not known to great precision, we require the properties of our simulated BH binaries to match the observed properties of the binary hosting Gaia BH1 just after the formation of the BH. This choice does not affect our results because the G star’s presence on the main sequence implies that tides are not expected to alter the binary’s orbit between the formation of the BH and the present day.

We initialize 1024 walkers uniformly over the prior space in ZAMS mass, orbital period, eccentricity, natal kick strength and direction, and CE ejection efficiency. We evolve each walker for 100 000 steps, thin the chains by a factor of 10, and retain the final 2000 steps of each chain to ensure a proper burn-in of our sampling.

We find that the preferred evolutionary pathway is for binary to begin with an orbital period between 1600 and 4800 d depending on the initial eccentricity, which ranges between 0 and 0.5. The

BH progenitor’s ZAMS mass is preferred to be between 65 and 75 M_{\odot} , while the G star’s ZAMS mass is preferred to be between 0.8 and 1.1 M_{\odot} . From this initial configuration, the BH progenitor loses roughly 20 M_{\odot} due to strong stellar winds which widens the binary out to periods of ~ 7000 d. As the BH progenitor evolves off the main sequence and begins core helium burning, its envelope expands and tides work to circularize and shrink the binary back to orbital periods between near ~ 3000 d. Near the end of core helium burning, the BH progenitor fills its Roche lobe, and the mass transfer becomes dynamically unstable and forms a CE due to both the convective envelope of the BH progenitor and the orbital evolution, which drives the binary components close together due to their highly asymmetric mass ratios. The BH progenitor’s envelope is fully stripped in the CE, leaving behind a $\sim 20 M_{\odot}$ Wolf–Rayet star in a ~ 20 day orbit with the G star. The Wolf–Rayet star’s strong wind mass loss widens the orbit to roughly 80 d. Finally the BH progenitor reaches core collapse and explodes which both reduces the BH progenitor mass from $\sim 12 M_{\odot}$ to the mass of Gaia BH1 and imparts a natal kick that increases the eccentricity and orbital period to reach the present day observed values.

Fig. 8 shows the combination of ZAMS parameters and binary evolution assumptions that produce Gaia BH1-like binaries whose masses, orbital period, and eccentricity match Gaia BH1’s present day properties for the 50th and 90th percentiles of each distribution. Due to the moderately wide orbit of Gaia BH1, our models prefer low natal kick strengths with $v_{\text{kick}} < 45 \text{ km s}^{-1}$, which is consistent with Gaia BH1’s orbit lying near the plane of the Galactic disc. There is a slight preference for kicks centred in the plane of the orbit ($\theta \sim 90^{\circ}$), though there is support over nearly the entire allowed range. A correlation between ϕ and v_{kick} exists such that larger kicks prefer azimuthal angles near 90° and 270° . This is due to a higher likelihood of the binary to remain bound for larger natal kicks when the orbital and kick velocities are pointed in opposite directions. We do not show samples for the mean anomaly since there is a near equal preference for all values of \mathcal{M} . This is not unexpected since the binary’s orbit is circularized during the CE, thus fixing a constant orbital speed.

While the ZAMS masses are not correlated with the ZAMS orbital period, ZAMS eccentricity, or natal kick parameters, there is a correlation between the ZAMS primary mass and the CE ejection efficiency. This is expected since an increased primary mass will have a larger envelope which requires a larger, more efficient, α to produce a successful CE ejection. The ZAMS orbital period is also correlated with α ; longer initial periods require lower, less efficient α ’s to shrink the orbit to the proper separation after the CE.

Our models prefer extremely efficient CE ejections with α constrained to always be larger than 5 with a preference for $\alpha = 14$. We remind the reader that α represents the fraction of the G star’s orbital energy that goes into unbinding the BH progenitor’s envelope. Values of α significantly larger than 1 signify that the liberated orbital energy alone is insufficient to eject the envelope, and that some other source of energy is required. While there are large uncertainties in the binding energy of the BH progenitors in our models for which $\lambda \sim 0.5$, it is highly unlikely that such uncertainties will change λ by more than a factor of 10. Furthermore, while it is possible for α values to exceed unity due to additional energy sources like recombination, it is unlikely that such sources will dominate the energy budget (Ivanova, Justham & Podsiadlowski 2015; Ivanova 2018). It is also possible that CE proceeds differently than the standard $\alpha\lambda$ prescription (e.g. Hirai & Mandel 2022), though we leave this for a future study. We thus conclude that Gaia BH1 can be formed through standard binary evolution with a CE only under extreme (and likely unphysical) assumptions about how the CE evolution proceeds.

In this sense, Gaia BH1 is reminiscent of the several known pulsars in wide, moderately eccentric binaries ($P_{\text{orb}} = 80\text{--}400$ d) with companions that appear to be low-mass main-sequence stars (e.g. Phinney & Verbunt 1991; Champion et al. 2008; Parent et al. 2022). These systems are also too wide to form via a CE, but too close to have avoided one. It is possible that the companions are white dwarfs, but their non-zero eccentricities are difficult to explain in this scenario. Perhaps they formed through a similar channel to Gaia BH1. We consider possible alternative formation scenarios below.

8.4.2 Formation from a progenitor that never became a giant

A CE event can be avoided if the BH progenitor never became very large. Models predict that at sufficiently high masses ($M \gtrsim 50 M_{\odot}$, depending on wind prescriptions and rotation rates), stars do not expand to become red supergiants, but instead lose their hydrogen envelopes to winds and outbursts during and shortly after their main-sequence evolution, reaching maximum radii of order $50 R_{\odot}$ (e.g. Humphreys & Davidson 1994; Higgins & Vink 2019). A scenario in which the BH progenitor in Gaia BH1 never became a red supergiant has the attractive feature that it could seem to avoid a CE episode. It requires, however, that the G star formed uncomfortably close to the BH progenitor. The orbit would have expanded by a factor of ~ 5 due to the BH progenitor’s mass loss, as the quantity $a \times (M_1 + M_2)$ is conserved under adiabatic orbit evolution (Jeans 1924). This in turn requires an initial separation of $\lesssim 60 R_{\odot}$ and an initial period of less than 10 d. It seems somewhat unlikely that a pre-main sequence solar-mass star could form and survive so close to a massive star, but observational constraints on the existence of such companions to O stars – which would have luminosity ratios $\gtrsim 10^5$ and angular separations of order 0.1 mas at typical Galactic distances – are extraordinarily difficult to obtain (e.g. Rizzuto et al. 2013; Sana et al. 2014).

A related set of models allow rapidly rotating massive stars to avoid becoming red supergiants due to mixing, such that a chemical gradient never forms and a majority of the star is converted to helium (‘quasi-homogeneous evolution’; e.g. Maeder 1987; Woosley & Heger 2006). Such a scenario could also avoid the problems with the CE channel, because the G star could have been born in an orbit close to the current separation, never interacting with the BH companion. The BH progenitor could have also had a somewhat lower initial mass, perhaps down to 15 M_{\odot} . However, most models predict that such evolution only occurs at low metallicity, because winds remove too much angular momentum at high metallicity (e.g. Yoon, Langer & Norman 2006; Brott et al. 2011). The $[\text{Fe}/\text{H}] = -0.2$ we measure for the G star thus presents a challenge for this scenario.

8.4.3 Formation through three-body dynamics

Mechanisms for forming BH LMXBs have been proposed involving tidal capture of inner binaries in hierarchical triples (e.g. Naoz et al. 2016) or very wide BH + normal star binaries perturbed by encounters with field stars (e.g. Michaely & Perets 2016). In both cases, tidal effects can bring the binary into a close orbit only if a periastron passage occurs within which the star comes within a few stellar radii of the BH. If this occurs, tides can efficiently dissipate orbital energy, locking the star into a close orbit and forming an LMXB. Such a scenario seems unlikely to have operated in Gaia BH1, because the observed periastron distance is too wide for tides to be important.

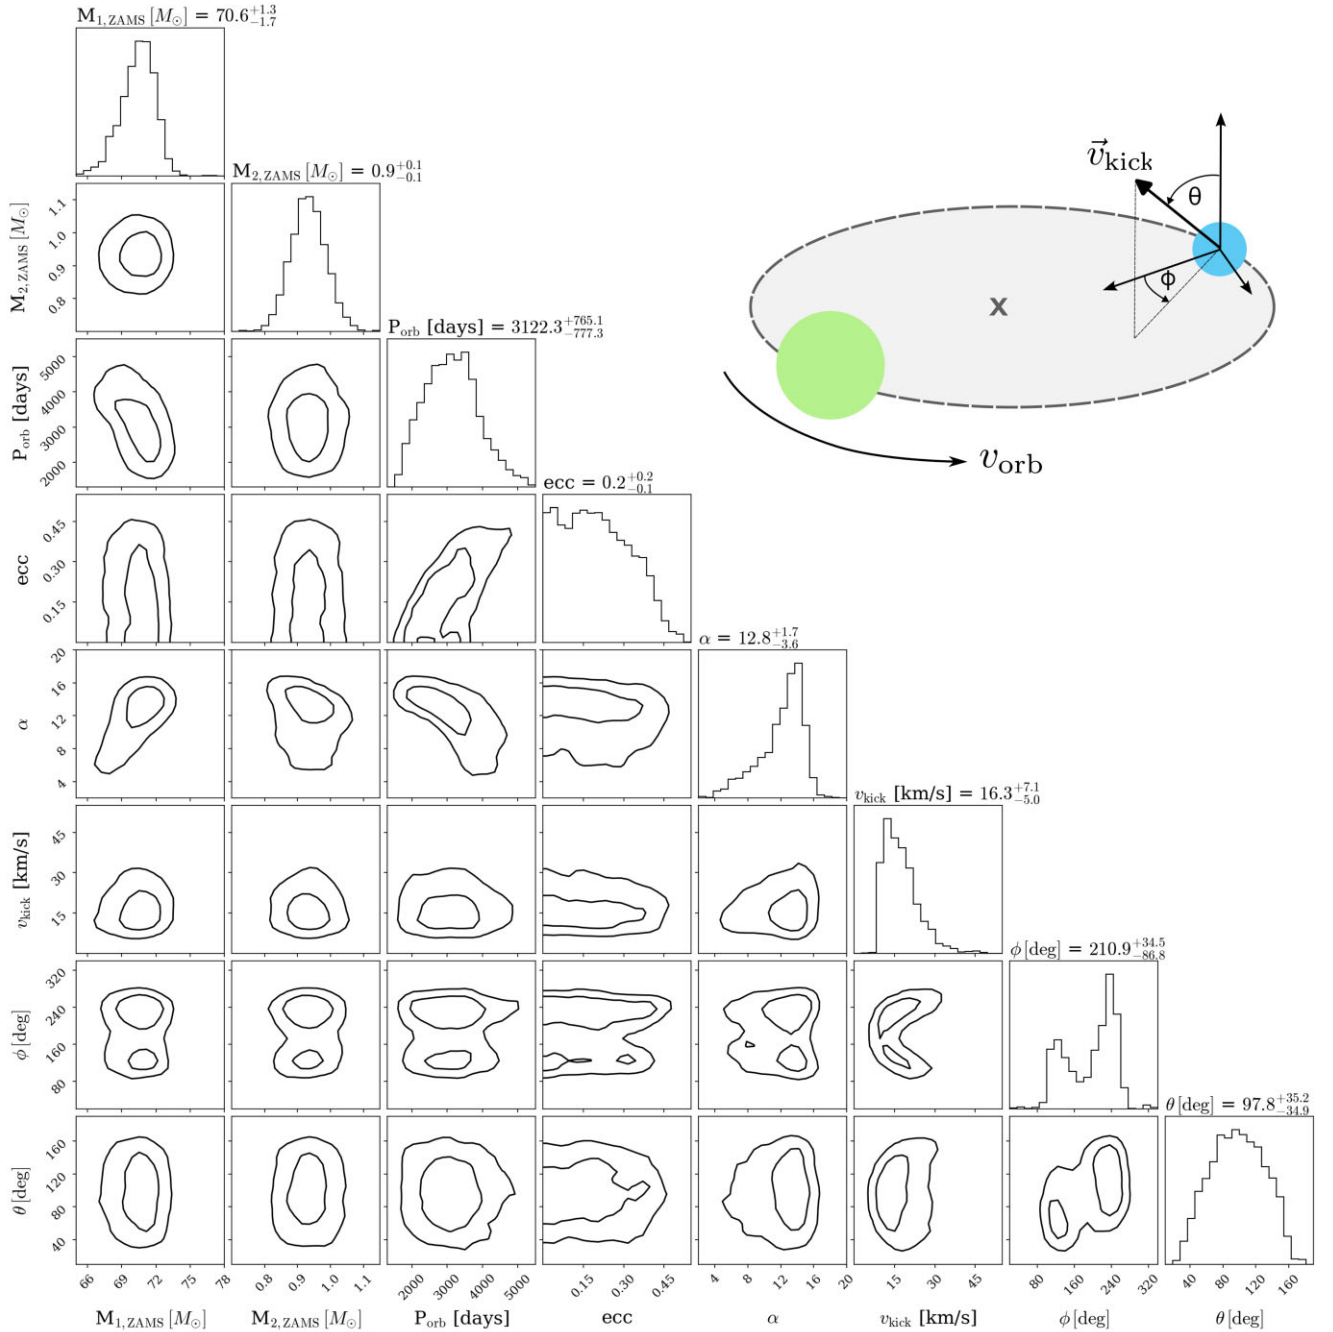


Figure 8. The combination of sampled ZAMS component masses, initial orbital period, and initial eccentricity with uncertain binary evolution model parameters for CE (α) and BH natal kicks (v_{kick} , ϕ , θ) which evolve to produce BHs in binaries with properties similar to Gaia BH1 and its binary host. We do not show distributions for the mean anomaly (\mathcal{M}) because there is not a significant preferred angle nor any appreciable correlations with any other parameters. The contours show the 50th and 90th percentiles of each distribution. The diagram in the upper-right corner illustrates the orbital configuration and natal kick direction for a Wolf–Rayet BH progenitor (blue) at core collapse in a circular orbit with a G star (green). Our models strongly prefer small natal kicks at BH formation and extremely efficient CE ejection efficiency.

8.4.4 Formation in a dense cluster

In dense environments like the cores of GCs, binaries with separations and eccentricities similar to Gaia BH1 can be formed via exchange encounters (e.g. Kremer et al. 2018). Formation in a GC is disfavoured for Gaia BH1 due to its thin-disc-like orbit and high metallicity.

Another possibility is that the Gaia BH1 system was assembled dynamically in an open cluster. Recent work (Fujii & Portegies Zwart 2011; Sana et al. 2017; Ramírez-Tannus et al. 2021) suggests that many massive star binaries are initially formed in long-period orbits, but are subsequently hardened due to dynamical encounters over millions of years. In the case of Gaia BH1, the binary may have formed at longer periods, allowing the BH progenitor to evolve

into a red supergiant and then collapse into a BH within ~ 3 Myr. Dynamical interactions after the BH formation could then shrink the orbital separation to the value observed today. Alternatively, the BH progenitor may have evolved separately from the G star in Gaia BH1. After BH formation, a dynamical exchange within the birth cluster could have allowed it to capture the G star into the current orbit (Banerjee 2018a, b). The binary system could have simultaneously or subsequently been ejected from the cluster at low velocity, or simply remained bound after the cluster dissolved (Schoettler et al. 2019; Dinnbier, Kroupa & Anderson 2022).

Shikauchi et al. (2020) recently investigated the dynamical formation of BH + main-sequence binaries in open clusters and predicted that only a small fraction of such binaries observable by *Gaia* would have formed dynamically. However, their simulations did not include any primordial binaries, and thus did not account for binaries formed through exchange interactions. Further work investigating this formation channel is warranted.

8.4.5 Formation in a hierarchical triple

Given the problems with binary formation models noted above, triple or multibody formation channels may have operated in Gaia BH1. One possibility is that the BH progenitor was originally in a compact binary ($P_{\text{orb}} \lesssim 20$ d) composed of two massive stars. As the stars evolve, they transfer mass to each other, preventing either star from evolving into a red supergiant. A simple possibility is that the inner binary merged during unstable Case B mass transfer, and the merger product never swelled into a red supergiant (Justham, Podsiadlowski & Vink 2014), allowing the G star (originally an outer tertiary) to survive.

Alternatively, the inner binary may have undergone stable mass transfer. Consider a binary with $M_1 \sim 35 M_{\odot}$ and $M_2 \sim 20 M_{\odot}$ in an orbit with $P_{\text{orb}} \sim 4$ d. As the primary evolves, its envelope is stripped during stable Case A mass transfer, producing a helium star which then collapses into a BH with $M_{\text{BH}} \sim 10 M_{\odot}$. During the mass transfer, the orbital period and secondary mass only change slightly. A very similar evolution is thought to produce observed HMXBs such as Cyg X-1 (Qin et al. 2019). The secondary then evolves and has its envelope stripped during a second phase of stable Case A mass transfer, forming a helium star. Depending on the strength of its stellar wind and the core-collapse process, this helium star could collapse into a second BH, or explode and produce a neutron star, which may or may not remain bound. A similar model was proposed by Shenar et al. (2022a) for the binary VFTS 243. At the end of the process, we are left with a BH and possibly a second compact object in a short-period orbit, with a still-bound outer tertiary.

To test this idea, we constructed binary models with MESA (Paxton et al. 2011) using the problem setup described in Fuller & Ma (2019). Assuming non-conservative mass transfer (as expected from the high-mass transfer rates achieved during Case A mass transfer), we find very similar binary evolution histories to those discussed above. For example, if the second phase of mass transfer begins with a $20 M_{\odot}$ donor and a $10 M_{\odot}$ BH in an orbital period of 4 d, it ends with a $5 M_{\odot}$ helium star in a 9.6 d orbit such that the orbit of the G star in Gaia BH1 could plausibly remain stable. As in the channels with single BH progenitors that never became giants, this scenario requires the G star to have formed quite close ($\lesssim 100 R_{\odot}$) from the inner binary. Given this and the expected orbital evolution of the inner binary due to mass transfer and winds, the parameter space of triples that can form a system like Gaia BH1, while not empty, is likely rather small. In the triple scenario, it would also be somewhat

surprising that the only good candidate we identified from *Gaia* DR3 has a (relatively) short period, because the effective search volume is larger for long periods (up to 1000 d; Section 9), and hierarchical triples with long outer periods would also be stable for a wider range of inner binary parameters.

While the triple model may seem unlikely given the rarity of BH HMXBs, it is important to note that the HMXB phase lasts at most $\lesssim 10^6$ yr, while the subsequent quiescent phase can easily last more than 10^{10} yr. If a significant fraction of HMXBs host outer low-mass tertiary stars, this could account for the high space-density of objects like Gaia BH1 relative to HMXBs.

A prediction of this channel is that Gaia BH1 might still harbour a binary compact object, whose orbital motion would induce potentially observable perturbations to the RVs of the outer G star on half the orbital period of the inner binary. Hayashi & Suto (2020) show that the amplitude of these perturbations would scale as $(P_{\text{orb, inner}}/P_{\text{orb, outer}})^{7/3}$; in Gaia BH1, the expected perturbation semi-amplitude ranges from $\sim 2 \text{ m s}^{-1}$ for a 4-d inner period, to $\sim 20 \text{ m s}^{-1}$ for a 10-d inner period, to $\sim 250 \text{ m s}^{-1}$ for a 30-day inner period. Since the G star is bright and slowly rotating, these amplitudes are within the realm of detectability. In addition, an inner binary would cause precession of the G star's orbit, which could be measurable via time-evolution of the orbit's orientation and RVs. The predicted minimum precession period ranges from about 35 yr for a 30-d inner period, to 350 yr for a 4-d inner period (and still longer for shorter inner periods). The RV semi-amplitude of this precession would be large (of order the observed semi-amplitude; 70 km s^{-1}), but a long observing time baseline would be required to detect it.

9 HOW MANY SIMILAR OBJECTS EXIST?

Here, we consider what the discovery of one dormant BH in a binary implies about the broader population. The precision of the conclusions we can draw is of course limited by the Poisson statistics of $N = 1$. We will show, however, that *Gaia* has observed tens of millions of sources around which a BH companion could have been detected, if one existed. This makes possible robust inference about the occurrence rate of BH companions, as long as the selection function of the astrometric catalog can be understood.

As described by Halbwachs et al. (2022), the cuts placed on astrometric solutions published in *Gaia* DR3 were quite conservative, designed to reduce false positives at the expense of completeness. Attempts to fit astrometric orbits were made only for sources satisfying $\text{ruwe} > 1.4$ (indicating a poor fit with a single-star solution) that were observed in at least 12 visibility periods. Sources with evidence of marginally resolved companions or contaminated photometry were likewise discarded. Sources for which a satisfactory fit could be achieved with an accelerating solution were not fit with an orbital solution. After orbital solutions were fit, those which did not satisfy all of the following three conditions were discarded:

$$\frac{\varpi}{\sigma_{\varpi}} > \frac{20000 \text{ days}}{P_{\text{orb}}} \quad (10)$$

$$\frac{a_0}{\sigma_{a_0}} > \frac{158}{\sqrt{P_{\text{orb}}/\text{day}}} \quad (11)$$

$$\sigma_e < 0.079 \ln(P_{\text{orb}}/\text{day}) - 0.244 \quad (12)$$

Here, σ_e is the uncertainty on the orbital eccentricity. As shown by Halbwachs et al. (2022), equation (10) removes a large fraction of otherwise acceptable solutions. After equation (10) is applied, equations (11) and (12) remove a relatively small number of additional solutions.

All these cuts are more aggressive for binaries with shorter periods than those with longer periods. The astrometric solution for Gaia BH1 has $\varpi/\sigma_\varpi = 120$ and $20\,000/P_{\text{orb}} = 107$. That is, it just narrowly passes the cut in equation (10), and would have been excluded if it were 10 per cent more distant. Similarly, it has $a_0/\sigma_{a_0} = 13.6$; if this quantity were smaller than $158/\sqrt{P_{\text{orb}}} = 11.6$, it would have been excluded by equation (11). The fact that Gaia BH1 narrowly escaped the quality cuts applied to the whole catalog has two implications.

First, there are likely other similar systems that could be discovered with only slightly looser cuts. We estimate the number of detectable systems explicitly in Section 9.1. Secondly, BH companions may be more common at shorter periods than at longer periods. *Gaia* DR3 was much more sensitive to binaries with long orbital periods (up to the observing baseline, ≈ 1000 d) than to those with short periods: equation (10) will exclude systems with $\varpi/\sigma_\varpi < 100$ at $P_{\text{orb}} = 200$ d, but only those with $\varpi/\sigma_\varpi < 20$ at $P_{\text{orb}} = 1000$ d. At fixed absolute magnitude, this translates to a factor of ≈ 3 larger distance limit and a factor of ≈ 10 larger survey volume (Section 9.1).

Gaia BH1 has a shorter orbital period than 95 per cent of all objects in DR3 with astrometric solutions. The fact that it appears to be the only credible BH with $P_{\text{orb}} \lesssim 1000$ d, while at longer-periods ($P_{\text{orb}} = 400\text{--}1000$ d), BH companions could have been detected around a much larger number of sources, thus points to a relative paucity of longer-period binaries consisting of a BH and a normal (low-mass) star.

9.1 Effective survey volume of the astrometric binary sample

We now explore around how many sources *Gaia* could have detected a BH companion if one existed. Since ‘detectability’ depends on the astrometric uncertainties, this requires a noise model. We describe such a model in Appendix G.

The actual uncertainties for a given source will depend on factors such as the orbital eccentricity, orientation, and *Gaia* scanning pattern, but we neglect these details. We assume that at $P_{\text{orb}} < 1000$ d, the astrometric uncertainties σ_ϖ and σ_{a_0} depend only on apparent magnitude and a_0 . We constrain the dependence on both parameters empirically using the 170 000 astrometric binary solutions published in DR3. The median astrometric uncertainty for sources with orbital solutions in DR3 ranges from $\sigma_\varpi = 0.020$ mas at $G < 14$, to 0.045 mas at $G = 16$, to 0.08 mas at $G = 17$. The median uncertainty in a_0 is larger by a factor of ≈ 1.5 on average. The astrometric uncertainties also depend on a_0 , increasing by about a factor of 2 from $a_0 = 0.3$ mas to $a_0 = 3$ mas. This reduces the number of BH companions that could be detected, since BHs produce larger a_0 at fixed P_{orb} than luminous companions.

We now assign a hypothetical $10 M_\odot$ dark companion to each source in the *Gaia* catalog and ask whether it could have been detected at a given period. When calculating the expected a_0 , we assign all sources on the main sequence a mass based on their absolute magnitude, and all giants a mass of $1 M_\odot$. We apply the same cuts on apparent magnitude, photometric contamination, and image parameter metrics that were applied to sources in the actual *Gaia* binary catalog (Halbwachs et al. 2022). We calculated the expected uncertainties, σ_ϖ and σ_{a_0} , based on the mean and observed scatter of these quantities among all astrometric solutions with the same G and a_0 , as described in Appendix G. We assume that no binaries with periods within $\pm 15\%$ of 1 yr can be detected.

Fig. 9 shows the number of sources in the `gaia_source` catalog for which a $10 M_\odot$ dark companion would likely pass different quality cuts. The black line corresponds to the detection thresholds actually employed in DR3. As discussed above, there is a strong selection




Figure 9. Number of sources in *Gaia* DR3 in which a $10 M_\odot$ dark companion could be detected. Black line shows the cuts employed in DR3 (equations 10 and 11), which remove most short-period binaries. Dashed red line corresponds to simply requiring $a_0/\sigma_{a_0} > 10$ and $\varpi/\sigma_\varpi > 10$; this results in a much larger searchable volume at short periods. Binaries with periods near 1 yr are excluded, being difficult to detect with astrometry.

bias against shorter orbital periods, which were only accepted at extremely high ϖ/σ_ϖ .

In DR4, *Gaia* is planning to publish epoch-level astrometry for all sources, irrespective of whether they satisfy cuts like equation (10). This will make it possible to identify sources similar to Gaia BH1 at larger distances. The dotted red line in Fig. 9 shows the effective search volume for the looser – but still relatively conservative – threshold of a minimum astrometric SNR of 10. Using this threshold instead of equations (10) and (11) would increase the search volume for $10 M_\odot$ companions at a period of 186 d by a factor of 20. Gains at shorter periods would be even more dramatic. Of course, removing these cuts is likely to result in a large number of spurious solutions – this is why they were applied in the first place. Inspection of the epoch-level astrometric measurements, coupled with spectroscopic follow-up of the most promising candidates, will be necessary and sufficient to eliminate spurious astrometric BH candidates.

In addition, the factor of two improvement in observing baseline between DR3 and DR4 will make DR4 have more precise astrometry and be sensitive to orbits with periods as large as 2000 d, for which the search volume is also larger. While extrapolating future detection rates from the current search criteria and one detection is risky, we conclude that there is good reason to expect dozens more astrometric BHs from *Gaia* DR4. The number could also be higher if the occurrence rate of (wide) BH companions is significantly higher among massive stars than among solar-type stars, as is expected theoretically (e.g. Langer et al. 2020; Janssens et al. 2022). Few massive stars received orbital solutions in DR3, because most are too distant to pass the stringent parallax uncertainty cut of equation (10).

One can infer a rough lower limit on the number of objects similar to Gaia BH1 that exist in the Milky Way. Fig. 9 shows that the effective search volume in 186 d is about 2.5×10^6 sources for which a $10 M_\odot$ companion would be detectable with DR3 cuts. Translating this into a constraint on the occurrence rate of BH companions in a given period range requires an assumed BH companion period distribution, which is obviously very uncertain. If we assume that a fraction f of all low-mass stars have a $10 M_\odot$ BH companion with $100 < P_{\text{orb}}/d < 300$ and the period distribution is log-uniform in this range ($dN/d\log P_{\text{orb}} = \text{const}$), then we find that 1 detection implies $f \approx 4 \times 10^{-7}$; i.e. an occurrence rate of order 1 in 2.5 million for BH

companions to low-mass stars with $100 < P_{\text{orb}}/d < 300$. This is a constraint on the average incidence across a range of stellar masses, sampled in a way that is different from the mass function. A large majority of the sources in this search volume is main-sequence stars with masses between 0.5 and $2 M_{\odot}$. Multiplying by $\sim 10^{11}$ low-mass stars, this would translate to $\sim 40\,000$ similar systems in the Milky Way – a factor of ~ 40 larger than the expected number of BH X-ray binaries (Corral-Santana et al. 2016), but still less than 0.1 per cent of the expected total number of BHs.

We can compare the occurrence rate of Gaia BH1-like systems to the number of stars formed that are massive enough to become BHs. For a Kroupa (2001) IMF, roughly two out of every 1000 stars formed has $M > 20 M_{\odot}$. Thus, if 1 out of every 5000 stars with $M > 20 M_{\odot}$ has a solar-type companion in a Gaia BH1-like orbit when it becomes a BH, this could produce the population implied by Gaia BH1, with 1 in 2.5 million low-mass stars having a BH companion with $100 < P_{\text{orb}}/d < 300$. The ‘1 in 5000’ factor could reflect a dearth of extreme mass ratio companions to massive stars, and/or the fact that a majority of such companions do not evolve to become intact BH + normal star binaries at these separations.

The lack of a solid detection at $P_{\text{orb}} = 400\text{--}1000$ d implies a lower BH companion incidence rate there. Assuming a log-uniform period distribution as above, the probability of detecting at least one BH companion in this period range would be $\gtrsim 50\%$ for $f > 4 \times 10^{-8}$; i.e. if the BH companion occurrence rate were more than 1 in 25 million in this period range. As above, this limit depends sensitively on the assumed period distribution.

Limits at shorter periods are weaker, given the implicit period bias of the selection cuts. For example, the number of sources for which a $10 M_{\odot}$ companion in a 60-d orbit would be detected astrometrically in DR3 is only 0.3 million. At such short periods, the (lack of) detections in *Gaia* SB1 (spectroscopic) solutions becomes a more stringent limit. DR3 contains 5.9 million sources satisfying `rv_method_used = 1` (meaning that RVs were measured from individual epochs) and `rv_nb_transits > 10` (enough epochs that an orbital solution could plausibly be inferred). About 2.5 million of these are near the main sequence. Orbital solutions with $K_{*} > 250 \text{ km s}^{-1}$ were removed from DR3 (Gaia Collaboration 2022a), but this will only remove BH companions to solar-type stars at $P_{\text{orb}} \lesssim 3$ d. This suggests that BHs could have been detected as companions to about 2.5 million sources at periods ranging from 3 to 1000 d, greatly exceeding the astrometric search volume at short periods. A non-detection of BH companions in the DR3 SB1 sample would seem to rule out a large increase in the BH companion incidence rate towards short periods. However, it is still unclear from the published information how SB1 solutions were processed and what cuts were ultimately applied to them. It is possible, for example, that close, tidally synchronized companions to BHs would have had too large $v \sin i$ for RVs to be reliably measured by the *Gaia* RVS pipeline.

9.2 Prospects for detecting similar systems with spectroscopy and photometry

Time-domain spectroscopic surveys provide another promising avenue for detecting more BH + normal star binaries like Gaia BH1. Indeed, Gaia BH1 was observed twice by the LAMOST survey; the source would have been easily recognizable as a promising BH candidate if it had been observed ~ 10 times with suitable cadence. *Gaia* will eventually provide >20 epochs of spectroscopy for most bright, ‘normal’ stars with $G \lesssim 12.5$. The calculations above suggest that at most a few BHs similar to Gaia BH1 will satisfy $G \lesssim 12.5$.

Many-epoch surveys from the ground are needed to reach targets that are fainter or harder to measure RVs for (e.g. hot and rapidly rotating stars). Most current wide-field surveys do not obtain enough epochs per target to constrain the orbital period, but more targeted surveys focused on binarity do (e.g. Giesers et al. 2019; Wang et al. 2021; Shenar et al. 2022b). We note that, except in cases where orbital periods can be constrained with photometry or astrometry, spectroscopic surveys that obtain $\gtrsim 10$ epochs for a modest number of targets are more useful for characterizing binaries than those obtaining 1–3 spectra for a larger number of targets.

Besides astrometry and spectroscopy, another avenue for detecting Gaia BH1-like binaries is through photometry. Ellipsoidal variation is not currently detectable (the expected variability amplitude is $\approx 5 \times 10^{-7}$; Morris & Naftilan 1993) but will become detectable when the G star becomes a red giant. In the meantime, Gaia BH1-like systems could be detected through self-lensing (e.g. Masuda & Hotokezaka 2019). If a binary like Gaia BH1 were viewed edge-on, the G star would brighten by $\approx 5\%$ each time it moved behind the BH. The typical duration of this brightening would be 3–4 h; i.e. a duty cycle of about 10^{-3} . Self-lensing would only be detectable for the $\approx 0.3\%$ of binaries with sufficiently edge-on inclinations to be eclipsing, but given the large expected magnification amplitude, a self-lensing search could extend to significantly fainter targets than are accessible with spectroscopic surveys. Given that several self-lensing binaries containing white dwarfs were detected by *Kepler* with longer periods and much smaller amplitudes than expected for a system like Gaia BH1 (Kruse & Agol 2014; Kawahara et al. 2018), future missions like *Earth 2.0* (Ge et al. 2022) may detect rarer systems containing BHs.

9.3 Other Gaia astrometric candidates

A few other recent works have searched *Gaia* DR3 for dormant BH + normal star binaries. Andrews et al. (2022) analysed the astrometric binary sample and identified 24 candidates for having BH or neutron star companions. They excluded Gaia BH1 from their sample as probably spurious because its orbital period is close to $3\times$ the *Gaia* precession period. Our follow-up has shown that the *Gaia* solution was nevertheless robust.

The candidates Andrews et al. (2022) do consider all have rather low masses for BHs. Only one, *Gaia* DR3 6328149636482597888, has an inferred companion mass above $2 M_{\odot}$ at the 2σ level. We have begun spectroscopic monitoring of this source, for which Andrews et al. (2022) reported $M_2 = 2.71^{+1.50}_{-0.36} M_{\odot}$ (2σ uncertainties), and we find it to be a metal-poor turn-off star ($[\text{Fe}/\text{H}] \approx -1.5$) on a halo-like orbit. Andrews et al. assumed a mass $M_{*} = 1.21 \pm 0.2$ (2σ uncertainties) for the luminous star, but we find this to be an overestimate. Fitting the SED, we find $T_{\text{eff}} = 6150 \pm 100$ K and $R = 1.58 \pm 0.05 R_{\odot}$, which for $[\text{Fe}/\text{H}] = -1.5$ MIST models, we find to correspond to $M_{*} = 0.78 \pm 0.05 M_{\odot}$. Adopting this mass for the luminous star and fitting the *Gaia* astrometric solution, we find $M_2 = 2.25^{+0.61}_{-0.26} M_{\odot}$ (2σ uncertainties), consistent with a neutron star. The source’s RVs thus far show a linear trend consistent with the astrometric solution. This is promising, but the *Gaia* orbital period is 736 d, so RV monitoring over an extended period is required to confirm the reliability of the *Gaia* solution, and more analysis is required to determine the nature of the companion. A BH seems unlikely.

Among the lower-mass candidates identified by Andrews et al., which they propose to be neutron stars, a majority are significantly blueward of the main-sequence. This could indicate that (a) they have unresolved white dwarf companions, (b) they have low metallicities, or (c) the astrometric solutions are spurious. Our spectroscopic

follow-up of these objects has shown that a majority do not have unusually low metallicities. Spurious astrometric solutions or white dwarf companions both remain plausible explanations. The candidates in their sample that are not bluerward of the main-sequence are all sufficiently luminous ($L \gtrsim 1 L_{\odot}$) that white dwarf companions would likely escape detection. Among the neutron-star candidates, we find that when reliable mass estimates for the luminous stars are adopted from SED fitting, we in several cases infer companion masses significantly below the Chandrasekhar limit.

Shahaf et al. (2022) also recently searched for compact object companions among the *Gaia* DR3 astrometric solutions. They identified eight sources with inferred companion masses above $2 M_{\odot}$, including Gaia BH1. Three of the remaining seven sources, including all the sources with inferred M_2 above $3 M_{\odot}$, are ruled out by our spectroscopic follow-up, and likely have spurious astrometric solutions. The remaining candidates in their sample have sufficiently long periods that the astrometric solutions could not yet be tested robustly. Their sample also contains a few dozen neutron star candidates, most with periods ranging from 400 to 1000 d.

One neutron star candidate, *Gaia* DR3 5136025521527939072, was presented by Gaia Collaboration (2022a): it has an *AstroSpectroSB1* solution that was validated by their RV follow-up, and also by our follow-up. However, we find that this source is *also* metal-poor and on a halo-like orbit. Fitting the SED and spectra, we find $[\text{Fe}/\text{H}] = -1.3$, $T_{\text{eff}} = 6400 \text{ K}$, $R = 1.17 R_{\odot}$, and $M_{\star} = 0.77 \pm 0.05 M_{\odot}$, lower than the $M_{\star} = 1.2 M_{\odot}$ assumed by Gaia Collaboration (2022a). Adopting this M_{\star} and fitting the *Gaia* solution, we obtain $M_2 = 1.23 \pm 0.06 M_{\odot}$, which is lower than the $M_2 = 1.5 M_{\odot}$ inferred by Gaia Collaboration (2022a), and consistent with a white dwarf. This object and *Gaia* DR3 6328149636482597888 highlight another potential contaminant for astrometric compact object searches: moderately evolved, low-metallicity stars can be easily mistaken for higher-mass stars near the zero-age main sequence.

Further analysis of the Andrews et al. (2022) and Shahaf et al. (2022) candidates is warranted, as some could indeed host neutron star or very low-mass BH companions. RV follow-up is essential for determining the reliability of the astrometric solutions. As our follow-up has already shown (Appendix E), the probability of encountering a spurious solution (or simply one that is in the tails of the noise distribution) is high when one selects objects in sparsely populated regions of parameter space. For companions with masses consistent with neutron stars, other models (e.g. two white dwarfs) are also harder to rule out.

We conclude that *Gaia* DR3 astrometric solutions contain *one* unambiguous dormant BH, although we cannot exclude the possibility that up to a few low-mass ($\lesssim 3 M_{\odot}$) BHs also exist in the data set. For BHs above $5 M_{\odot}$, Gaia BH1 is the only solid candidate in the astrometric catalog among sources near the main-sequence (see Appendix E; there is one source with a red giant primary that *may* host a BH companion, but RV followup over a few-year period is required to vet it). Among *Gaia* DR3 spectroscopic binaries, there are almost certainly no solid BH companions to main-sequence stars (El-Badry & Rix 2022; Jayasinghe et al. 2022), but some neutron star candidates may turn out to be robust. BH companions to giants might still exist in the catalog, but are more difficult to vet.

We note that if any of the neutron star or BH candidates discussed above turn out to be robust, their evolutionary histories will be a puzzle. Like Gaia BH1, they have orbits that are too wide to be understood as a result of a CE event, but too close to accommodate a red supergiant. In addition, their wide orbits could not survive the typical $\sim 250 \text{ km s}^{-1}$ natal kicks expected for neutron stars. The

existence of Gaia BH1 and these other candidates might thus point to a different dominant channel for forming moderately wide compact object + normal star binaries, such as the hierarchical triple scenario discussed in Section 8.4.5.

While this manuscript was under review, a preprint analysing the same system was posted by Chakrabarti et al. (2022). Their conclusions are broadly consistent with ours. The best-fitting BH mass they report, $M_2 = 11.9_{-1.6}^{+2.0} M_{\odot}$, is marginally higher than the value we found, $M_2 = 9.62 \pm 0.18 M_{\odot}$. We suspect that the difference in our results comes mainly from the fact that their RVs do not include any measurements near periastron, and thus do not constrain the RV amplitude well. This also leads to larger uncertainties in the BH mass. Their best-fitting solution predicts an RV $\approx 20 \text{ km s}^{-1}$ larger than our measurements near periastron, so we consider their constraints less reliable than those reported here. Nevertheless, it is encouraging that an independent analysis, based on independent RV data, reached similar conclusions about the nature of the Gaia BH1 system.

10 SUMMARY AND CONCLUSIONS

We have presented the discovery of Gaia BH1, a new binary system consisting of a $\sim 10 M_{\odot}$ dark object orbited by an otherwise unremarkable Sun-like star (Fig. 1). The system was first identified as a promising BH candidate based on its *Gaia* astrometric solution. Our follow-up spectroscopic observations validated the astrometric solution and allowed us to place tight dynamical constraints on the mass of the dark object (Fig. 2), which are summarized in Table 1. Even without consideration of the astrometric solution, the RVs imply a companion mass that can only be attributed to a BH (Fig. 4; the 5σ lower limit on the spectroscopic mass function is $f(M_2)_{\text{spec}} > 3.68 M_{\odot}$). Scenarios not involving a BH are firmly ruled out by the object's large mass and stringent limits on the light contributions of a luminous companion.

Gaia BH1 is different in several ways from other known BHs in binaries. It is nearby ($d = 480 \text{ pc}$) and bright ($G = 13.8$). Together with the lack of contamination from an accretion disc and the fact that the luminous star is cool and slowly rotating, this makes it possible to study the system in greater detail in the optical than is possible for BH X-ray binaries. For example, we measure chemical abundances for 22 elements in the G star's photosphere (Table 2) – and find them entirely unremarkable. The fact that the orbit is measured with both astrometry and RVs allows for tight constraints on the BH mass, with no inclination degeneracy. Several high-precision RV instruments are capable of measuring RVs for the G star with $\sim 1 \text{ m s}^{-1}$ precision.

The system's properties are difficult to explain with evolutionary models for isolated binaries. The maximum radius of the BH progenitor is expected to have been much larger than the G star's current orbit, leading to an episode of CE evolution in which the G star spiraled in through the BH progenitor's envelope and ultimately ejected it. However, a $1 M_{\odot}$ star is unlikely to have had enough orbital energy to eject the $\sim 20 M_{\odot}$ envelope of a red supergiant, and would be in danger of merging with the core under standard CE assumptions. This is a generic problem for formation models of BH LMXBs. For Gaia BH1, there is an addition challenge: the orbit today is much wider than expected after a CE episode, which would be predicted to leave the G star just a few solar radii from the BH progenitor. One cannot invoke strong natal kicks to the BH to widen the orbit again in this scenario, as this would produce a highly eccentric orbit and high space velocity.

The system's evolution may be better-explained in models in which the G star was initially a wide tertiary companion to a close binary

containing two massive stars. In this case, interactions between the two stars could have prevented either one from expanding to become a red supergiant, such that the G star could have formed in an orbit similar to its current orbit (but somewhat tighter) and remained there ever since. In this case, the dark object could be a binary composed of two lower-mass BHs. High-precision RV follow-up offers the tantalizing possibility of testing this scenario. Alternatively, the binary may have formed through an exchange interaction. Its thin-disc-like Galactic orbit, modest eccentricity, and near-solar metallicity suggest that it did not form dynamically in a GC, but dynamical formation in an open cluster that has since dissolved is more plausible.

Gaia BH1 is about three times nearer to Earth than the next-closest BH known, the LMXB A0620-00 (McClintock & Remillard 1986), whose *Gaia* parallax implies a distance of 1.44 ± 0.25 kpc. This suggests that wide BH binaries like Gaia BH1, while being harder to detect, are significantly more common than BHs in close binaries with ongoing accretion. Given the conservative cuts applied to astrometric binary solutions published in *Gaia* DR3, it is likely that DR4 and DR5 will identify dozens of similar systems. Gaia BH1 will likely remain one of the nearest and brightest BHs at periods $P_{\text{orb}} \lesssim 1000$ d, but even nearer BH binaries may be discovered in longer-period orbits, which could not yet be characterized with DR3 data.

It is both interesting and irksome that the first unambiguous BH discovered by *Gaia* seems to be somewhat of an oddball, and not easily accommodated by standard binary evolution models. One interpretation is that – among solar-type stars, which make up the bulk of the population to which *Gaia* is sensitive – CE evolution leads to the merger of most would-be BH + normal star binaries. In this case, only systems formed through alternative channels will survive. It is clear that BH X-ray binaries – which today are more compact versions of systems like Gaia BH1 – only narrowly escape a merger if they form through a CE, and represent a rare outcome of binary evolution. BH + normal star binaries with $P_{\text{orb}} \gtrsim 10$ yr can avoid a CE, but (a) *Gaia* is not yet sensitive to them, and (b) they may be disrupted by natal kicks. BH + massive star binaries are more likely to both survive CE and undergo stable mass transfer, but the number of massive stars accessible with *Gaia* astrometric orbits is thus far modest. In any case, population demographics of the detached BH + normal star binary population will become clearer in the next decade, as more systems are discovered.

Because of its long orbital period and close distance, the angular size of Gaia BH1's orbit on the sky is more than 10 times larger than any other known BH binary. This makes the system an excellent target for interferometric follow-up. Observations at X-ray and radio wavelengths are also warranted. Although the accretion rate of the G star's wind by the BH is expected to be low and any detectable emission would be faint (Section 7), Gaia BH1 provides a unique opportunity to test imperfectly understood models of accretion flows at very low accretion rates.

In the future, the G star will become a red giant. After an X-ray bright period during which it transfers mass to the BH – perhaps appearing as a wider version of the red giant + BH binary GRS 1915+105 – the system will terminate its evolution as a wide BH + white dwarf binary. If the system turns out to host an inner BH binary, this evolution may be interrupted by its merger, with the resulting recoil potentially disrupting the outer binary.

ACKNOWLEDGEMENTS

We thank Morgan Macleod, E. Sterl Phinney, Rohan Naidu, Lieke van Son, Mathieu Renzo, and Tom Maccarone for helpful discussions, and Vedant Chandra for help with data reduction. We are

grateful to Yuri Beletsky, Sam Kim, Angela Hempel, and Régis Lachaume for observing help. HWR acknowledges the European Research Council for the ERC Advanced Grant [101054731]. This research made use of PYSTROMETRY, an open source PYTHON package for astrometry time-series analysis (Sahlmann 2019). This work made use of ASTROPY,⁶ a community-developed core PYTHON package and an ecosystem of tools and resources for astronomy (Astropy Collaboration 2022).

This project was developed in part at the Gaia Fête, held at the Flatiron Institute's Center for Computational Astrophysics in June 2022, and in part at the Gaia Hike, held at the University of British Columbia in June 2022. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

Based on observations obtained at the international Gemini Observatory, a program of NSF's NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation on behalf of the Gemini Observatory partnership: the National Science Foundation (United States), National Research Council (Canada), Agencia Nacional de Investigación y Desarrollo (Chile), Ministerio de Ciencia, Tecnología e Innovación (Argentina), Ministério da Ciência, Tecnologia, Inovações e Comunicações (Brazil), and Korea Astronomy and Space Science Institute (Republic of Korea).

DATA AVAILABILITY

Data used in this study are available upon request from the corresponding author.

REFERENCES

- Agüeros M. A. et al., 2009, *ApJS*, 181, 444
 Alvarez R., Plez B., 1998, *A&A*, 330, 1109
 Andrews J. J., Breivik K., Chatterjee S., 2019, *ApJ*, 886, 68
 Andrews J. J., Taggart K., Foley R., 2022, preprint ([arXiv:2207.00680](https://arxiv.org/abs/2207.00680))
 Astropy Collaboration, 2022, *ApJ*, 935, 167
 Banerjee S., 2018a, *MNRAS*, 473, 909
 Banerjee S., 2018b, *MNRAS*, 481, 5123
 Baumann P., Ramírez I., Meléndez J., Asplund M., Lind K., 2010, *A&A*, 519, A87
 Bellm E. C. et al., 2019, *PASP*, 131, 018002
 Bensby T., Feltzing S., Oey M. S., 2014, *A&A*, 562, A71
 Binnendijk L., 1960, Philadelphia. University of Pennsylvania Press
 Blaauw A., 1961, *Bull. Astron. Inst. Neth.*, 15, 265
 Bodensteiner J. et al., 2020, *A&A*, 641, A43
 Bohlin R. C., Mészáros S., Fleming S. W., Gordon K. D., Koekemoer A. M., Kovács J., 2017, *AJ*, 153, 234
 Boller T., Freyberg M. J., Trümper J., Haberl F., Voges W., Nandra K., 2016, *A&A*, 588, A103
 Bovy J., 2015, *ApJS*, 216, 29
 Brandt N., Podsiadlowski P., 1995, *MNRAS*, 274, 461
 Breivik K., Chatterjee S., Larson S. L., 2017, *ApJ*, 850, L13
 Breivik K. et al., 2020, *ApJ*, 898, 71
 Brott I. et al., 2011, *A&A*, 530, A115
 Butler R. P., Marcy G. W., Williams E., McCarthy C., Dosanji P., Vogt S. S., 1996, *PASP*, 108, 500
 Cardelli J. A., Clayton G. C., Mathis J. S., 1989, *ApJ*, 345, 245

⁶<http://www.astropy.org>

- Carlos M., Nissen P. E., Meléndez J., 2016, *A&A*, 587, A100
- Casares J., Negueruela I., Ribó M., Ribas I., Paredes J. M., Herrero A., Simón-Díaz S., 2014, *Nature*, 505, 378
- Castro-Tirado A. J., Brandt S., Lund N., 1992, *IAU Circ.*, 5590, 2
- Chakrabarti S. et al., 2022, preprint ([arXiv:2210.05003](https://arxiv.org/abs/2210.05003))
- Chambers K. C. et al., 2016, preprint ([arXiv:1612.05560](https://arxiv.org/abs/1612.05560))
- Champion D. J. et al., 2008, *Science*, 320, 1309
- Chawla C., Chatterjee S., Breivik K., Moorthy C. K., Andrews J. J., Sanderson R. E., 2022, *ApJ*, 931, 107
- Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, *ApJ*, 823, 102
- Chubak C., Marcy G., Fischer D. A., Howard A. W., Isaacson H., Johnson J. A., Wright J. T., 2012, preprint ([arXiv:1207.6212](https://arxiv.org/abs/1207.6212))
- Corral-Santana J. M., Casares J., Muñoz-Darias T., Bauer F. E., Martínez-Pais I. G., Russell D. M., 2016, *A&A*, 587, A61
- Cui X.-Q. et al., 2012, *Res. Astron. Astrophys.*, 12, 1197
- Dinnbier F., Kroupa P., Anderson R. I., 2022, *A&A*, 660, A61
- El-Badry K., Burdge K. B., 2022, *MNRAS*, 511, 24
- El-Badry K., Quataert E., 2021, *MNRAS*, 502, 3436
- El-Badry K., Rix H.-W., 2022, *MNRAS*, 515, 1266
- El-Badry K., Rix H.-W., Heintz T. M., 2021, *MNRAS*, 506, 2269
- El-Badry K., Burdge K. B., Mróz P., 2022a, *MNRAS*, 511, 3089
- El-Badry K., Seeburger R., Jayasinghe T., Rix H.-W., Almada S., Conroy C., Price-Whelan A. M., Burdge K., 2022b, *MNRAS*, 512, 5620
- Fabian A. C., Pringle J. E., Rees M. J., 1975, *MNRAS*, 172, 15P
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Fragos T., McClintock J. E., 2015, *ApJ*, 800, 17
- Fryer C. L., Kalogera V., 2001, *ApJ*, 554, 548
- Fujii M. S., Portegies Zwart S., 2011, *Science*, 334, 1380
- Fuller J., Ma L., 2019, *ApJ*, 881, L1
- Gaia Collaboration, 2022a, preprint ([arXiv:2206.05595](https://arxiv.org/abs/2206.05595))
- Gaia Collaboration, 2022b, preprint ([arXiv:2208.00211](https://arxiv.org/abs/2208.00211))
- Gallo E., Fender R. P., Miller-Jones J. C. A., Merloni A., Jonker P. G., Heinz S., Maccarone T. J., van der Klis M., 2006, *MNRAS*, 370, 1351
- Ge J. et al., 2022, preprint ([arXiv:2206.06693](https://arxiv.org/abs/2206.06693))
- Giesers B. et al., 2018, *MNRAS*, 475, L15
- Giesers B. et al., 2019, *A&A*, 632, A3
- Gomez S., Grindlay J. E., 2021, *ApJ*, 913, 48
- González Hernández J. I., Casares J., Rebolo R., Israelian G., Filippenko A. V., Chornock R., 2011, *ApJ*, 738, 95
- Gordon Y. A. et al., 2020, *Res. Notes Am. Astron. Soc.*, 4, 175
- Green G. M., Schlafly E., Zucker C., Speagle J. S., Finkbeiner D., 2019, *ApJ*, 887, 93
- Greiner J., Cuby J. G., McCaughrean M. J., 2001, *Nature*, 414, 522
- Guseinov O. K., Zel'dovich Y. B., 1966, *Sov. Astro.*, 10, 251
- Gustafsson B., Edvardsson B., Eriksson K., Jørgensen U. G., Nordlund Å., Plez B., 2008, *A&A*, 486, 951
- Halbwachs J.-L. et al., 2022, preprint ([arXiv:2206.05726](https://arxiv.org/abs/2206.05726))
- Hawkins K. et al., 2020a, *MNRAS*, 492, 1164
- Hawkins K., Lucey M., Curtis J., 2020b, *MNRAS*, 496, 2422
- Hayashi T., Suto Y., 2020, *ApJ*, 897, 29
- Heiter U. et al., 2021, *A&A*, 645, A106
- Higgins E. R., Vink J. S., 2019, *A&A*, 622, A50
- Hirai R., Mandel I., 2022, *ApJL*, 937, L42
- Hook I. M., Jørgensen I., Allington-Smith J. R., Davies R. L., Metcalfe N., Murowinski R. G., Crampton D., 2004, *PASP*, 116, 425
- Howard A. W. et al., 2010, *ApJ*, 721, 1467
- Humphreys R. M., Davidson K., 1994, *PASP*, 106, 1025
- Hurley J. R., Tout C. A., Pols O. R., 2002, *MNRAS*, 329, 897
- Iglesias-Marzoa R., López-Morales M., Jesús Arévalo Morales M., 2015, *PASP*, 127, 567
- Israelian G., Rebolo R., Basri G., Casares J., Martín E. L., 1999, *Nature*, 401, 142
- Ivanova N., 2018, *ApJ*, 858, L24
- Ivanova N., Justham S., Podsiadlowski P., 2015, *MNRAS*, 447, 2181
- Janssens S. et al., 2022, *A&A*, 658, A129
- Jayasinghe T. et al., 2021, *MNRAS*, 504, 2577
- Jayasinghe T., Rowan D. M., Thompson T. A., Kochanek C. S., Stanek K. Z., 2022, preprint ([arXiv:2207.05086](https://arxiv.org/abs/2207.05086))
- Jeans J. H., 1924, *MNRAS*, 85, 2
- Justham S., Rappaport S., Podsiadlowski P., 2006, *MNRAS*, 366, 1415
- Justham S., Podsiadlowski P., Vink J. S., 2014, *ApJ*, 796, 121
- Kaufer A., Stahl O., Tubbesing S., Nørregaard P., Avila G., Francois P., Pasquini L., Pizzella A., 1999, *The Messenger*, 95, 8
- Kawahara H., Masuda K., MacLeod M., Latham D. W., Bieryla A., Benomar O., 2018, *AJ*, 155, 144
- Khokhlov S. A. et al., 2018, *ApJ*, 856, 158
- Klencki J., Nelemans G., Istrate A. G., Chruslinska M., 2021, *A&A*, 645, A54
- Kochanek C. S. et al., 2017, *PASP*, 129, 104502
- Kolbl R., Marcy G. W., Isaacson H., Howard A. W., 2015, *AJ*, 149, 18
- Kremer K., Ye C. S., Chatterjee S., Rodriguez C. L., Rasio F. A., 2018, *ApJ*, 855, L15
- Kroupa P., 2001, *MNRAS*, 322, 231
- Kruse E., Agol E., 2014, *Science*, 344, 275
- Kurucz R. L., 1970, *SAO Spec. Rep.*, 309
- Kurucz R. L., 1979, *ApJS*, 40, 1
- Kurucz R. L., 1993, *Kurucz CD-ROM. Smithsonian Astrophysical Observatory, Cambridge, MA*, p. |c1993
- Lacy M. et al., 2020, *PASP*, 132, 035001
- Lam C. Y. et al., 2022, *ApJ*, 933, L23
- Langer N. et al., 2020, *A&A*, 638, A39
- Laycock S. G. T., Maccarone T. J., Christodoulou D. M., 2015, *MNRAS*, 452, L31
- Lejeune T., Cuisinier F., Buser R., 1997, *A&AS*, 125, 229
- Lejeune T., Cuisinier F., Buser R., 1998, *A&AS*, 130, 65
- Lennon D. J., Dufton P. L., Villaseñor J. I., Evans C. J., Langer N., Saxton R., Monageng I. M., Toonen S., 2022, *A&A*, 665, A180
- Maeder A., 1987, *A&A*, 178, 159
- Mahy L. et al., 2022, *A&A*, 664, A159
- Mamajek E. E., Hillenbrand L. A., 2008, *ApJ*, 687, 1264
- Marigo P. et al., 2017, *ApJ*, 835, 77
- Marshall J. L. et al., 2008, in McLean I. S., Casali M. M., eds, *Proc. SPIE Conf. Ser.*, Vol. 7014, *Ground-based and Airborne Instrumentation for Astronomy II. SPIE, Bellingham*, p. 701454
- Martin D. C. et al., 2005, *ApJ*, 619, L1
- Mashian N., Loeb A., 2017, *MNRAS*, 470, 2611
- Masseron T. et al., 2014, *A&A*, 571, A47
- Masseron T., Merle T., Hawkins K., 2016, *Astrophysics Source Code Library*, record ascl:1605.004
- Masuda K., Hotokezaka K., 2019, *ApJ*, 883, 169
- McClintock J. E., Remillard R. A., 1986, *ApJ*, 308, 110
- McClintock J. E., Remillard R. A., 2006, in Lewin W., van der Klis M., eds, *Compact Stellar X-Ray Sources. Cambridge Astrophysics Series*, No. 39. Cambridge Univ. Press, Cambridge, p. 157
- McClure R. D., Woodsworth A. W., 1990, *ApJ*, 352, 709
- McMillan P. J., 2017, *MNRAS*, 465, 76
- Merloni A., Heinz S., di Matteo T., 2003, *MNRAS*, 345, 1057
- Michael E., Perets H. B., 2016, *MNRAS*, 458, 4188
- Miller-Jones J. C. A. et al., 2021, *Science*, 371, 1046
- Morris S. L., Naftilan S. A., 1993, *ApJ*, 419, 344
- Mroz P., Udalski A., Gould A., 2022, *ApJL*, 937, L24
- Naoz S., Fragos T., Geller A., Stephan A. P., Rasio F. A., 2016, *ApJ*, 822, L24
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Ness M., Hogg D. W., Rix H. W., Ho A. Y. Q., Zasowski G., 2015, *ApJ*, 808, 16
- Nidever D. L., Marcy G. W., Butler R. P., Fischer D. A., Vogt S. S., 2002, *ApJS*, 141, 503
- Padmanabhan N. et al., 2008, *ApJ*, 674, 1217
- Parent E. et al., 2022, *ApJ*, 924, 135
- Paxton B., Bildsten L., Dotter A., Herwig F., Lesaffre P., Timmes F., 2011, *ApJS*, 192, 3
- Penoyre Z., Belokurov V., Evans N. W., 2022, *MNRAS*, 513, 5270
- Pfahl E., Rappaport S., Podsiadlowski P., 2003, *ApJ*, 597, 1036

- Phinney E. S., Verbunt F., 1991, *MNRAS*, 248, 21P
- Plez B., 2012, Astrophysics Source Code Library, record ascl:1205.004
- Podsiadlowski P., Rappaport S., Pfahl E. D., 2002, *ApJ*, 565, 1107
- Podsiadlowski P., Rappaport S., Han Z., 2003, *MNRAS*, 341, 385
- Portegies Zwart S. F., Verbunt F., Ergma E., 1997, *A&A*, 321, 207
- Prochaska J. X., Hennawi J. F., Westfall K. B., Cooke R. J., Wang F., Hsu T., Davies F. B., Farina E. P., 2020, *J. Open Source Softw.*, 5, 2308
- Qian S. B., Liao W. P., Fernández Lajús E., 2008, *ApJ*, 687, 466
- Qin Y., Marchant P., Fragos T., Meynet G., Kalogera V., 2019, *ApJ*, 870, L18
- Quataert E., Narayan R., 1999, *ApJ*, 520, 298
- Raghavan D. et al., 2010, *ApJS*, 190, 1
- Ramírez I., Fish J. R., Lambert D. L., Allende Prieto C., 2012, *ApJ*, 756, 46
- Ramírez-Tannus M. C. et al., 2021, *A&A*, 645, L10
- Rasio F. A., Pfahl E. D., Rappaport S., 2000, *ApJ*, 532, L47
- Remillard R. A., McClintock J. E., 2006, *ARA&A*, 44, 49
- Rice M., Brewer J. M., 2020, *ApJ*, 898, 119
- Rizzuto A. C. et al., 2013, *MNRAS*, 436, 1694
- Sahlmann J., 2019, Johannes-Sahlmann/pystrometry: Release for Zenodo. Zenodo, available at <https://zenodo.org/record/3515526#.Y3RClnZBzIU>
- Sahu K. C. et al., 2022, *ApJ*, 933, 83
- Sana H. et al., 2014, *ApJS*, 215, 15
- Sana H., Ramírez-Tannus M. C., de Koter A., Kaper L., Tramper F., Bik A., 2017, *A&A*, 599, L9
- Schoettler C., Parker R. J., Arnold B., Grimmett L. P., de Bruijne J., Wright N. J., 2019, *MNRAS*, 487, 4615
- Shahaf S., Mazeh T., Faigler S., Holl B., 2019, *MNRAS*, 487, 5610
- Shahaf S., Bashi D., Mazeh T., Faigler S., Arenou F., El-Badry K., Rix H.-W., 2022, preprint ([arXiv:2209.00828](https://arxiv.org/abs/2209.00828))
- Shao Y., Li X.-D., 2019, *ApJ*, 885, 151
- Sharma P., Quataert E., Hammett G. W., Stone J. M., 2007, *ApJ*, 667, 714
- Sheinis A. I., Bolte M., Epps H. W., Kibrick R. I., Miller J. S., Radovan M. V., Bigelow B. C., Sutin B. M., 2002, *PASP*, 114, 851
- Shenar T., 2022, preprint ([arXiv:2208.02614](https://arxiv.org/abs/2208.02614))
- Shenar T. et al., 2020, *A&A*, 639, L6
- Shenar T. et al., 2022a, *Nature Astron.*, 6, 1085
- Shenar T. et al., 2022b, *A&A*, 665, A148
- Shikauchi M., Kumamoto J., Tanikawa A., Fujii M. S., 2020, *PASJ*, 72, 45
- Shikauchi M., Tanikawa A., Kawanaka N., 2022, *ApJ*, 928, 13
- Skrutskie M. F. et al., 2006, *AJ*, 131, 1163
- Skumanich A., 1972, *ApJ*, 171, 565
- Smith N., 2014, *ARA&A*, 52, 487
- Sukhbold T., Ertl T., Woosley S. E., Brown J. M., Janka H. T., 2016, *ApJ*, 821, 38
- Takeda Y., Honda S., Ohnishi T., Ohkubo M., Hirata R., Sadakane K., 2013, *PASJ*, 65, 53
- Tanikawa A., Hattori K., Kawanaka N., Kinugawa T., Shikauchi M., Tsuna D., 2022, preprint ([arXiv:2209.05632](https://arxiv.org/abs/2209.05632))
- The LIGO Scientific Collaboration, 2021, preprint ([arXiv:2111.03606](https://arxiv.org/abs/2111.03606))
- Thompson T. A. et al., 2019, *Science*, 366, 637
- Trimble V. L., Thorne K. S., 1969, *ApJ*, 156, 1013
- Valenti J. A., Fischer D. A., 2005, *ApJS*, 159, 141
- van de Kamp P., 1975, *ARA&A*, 13, 295
- Vernet J. et al., 2011, *A&A*, 536, A105
- Vogt S. S. et al., 1994, in Crawford D. L., Craine E. R., eds, Proc. SPIE Conf. Ser., Vol. 2198, Instrumentation in Astronomy VIII. SPIE, Bellingham, p. 362
- Wang S. et al., 2021, *Res. Astron. Astrophys.*, 21, 292
- Webbink R. F., 1984, *ApJ*, 277, 355
- Wiktorowicz G., Lu Y., Wyrzykowski Ł., Zhang H., Liu J., Justham S., Belczynski K., 2020, *ApJ*, 905, 134
- Wong K. W. K., Breivik K., Farr W. M., Luger R., 2022, preprint ([arXiv:2206.04062](https://arxiv.org/abs/2206.04062))
- Woosley S. E., Heger A., 2006, *ApJ*, 637, 914
- Woosley S. E., Langer N., Weaver T. A., 1995, *ApJ*, 448, 315
- Wright E. L. et al., 2010, *AJ*, 140, 1868
- Yee S. W., Petigura E. A., von Braun K., 2017, *ApJ*, 836, 77
- Yoon S. C., Langer N., Norman C., 2006, *A&A*, 460, 199

APPENDIX A: OBSERVATIONS AND DATA REDUCTION

A1 MagE

We observed Gaia BH1 four times using the Magellan Echellette spectrograph (MagE; Marshall et al. 2008) on the *6.5m Magellan Clay telescope* at Las Campanas Observatory. All observations were carried out with the 0.7 arcsec slit and exposure times ranging from 300 to 600 s. This yielded a typical SNR of 100 per pixel at 5500 Å, spectral resolution $R \approx 5400$, and wavelength coverage of 3500–11 000 Å. To minimize drift in the wavelength solution due to flexure, we obtained a ThAr arc on-sky immediately after each science exposure.

We reduced the spectra using PYPEIT (Prochaska et al. 2020), which performs bias and flat-field correction, cosmic ray removal, wavelength calibration, sky subtraction, extraction of 1d spectra, and heliocentric RV corrections. To assess the stability of the wavelength solution, we observed several standard stars, measured their RVs using the same procedure we used for the science targets, and compared the results to their measured *Gaia* DR3 RVs. We found that the MagE-measured RVs were $1.5 \pm 2 \text{ km s}^{-1}$ larger than those measured by *Gaia*, so we adopted a conservative uncertainty of 3 km s^{-1} for all the MagE RVs.

A2 GMOS

We observed Gaia BH1 four times using the Gemini Multi-Object Spectrograph (GMOS; Hook et al. 2004) on the *8.1m Gemini-North telescope* on Mauna Kea (program GN-2022B-DD-202). We used the R831_G5302 grating with a 0.5 arcsec slit and central wavelength of 5200 Å, leading to wavelength coverage from 4000 to 6400 Å and resolution $R \approx 4400$. We obtained a CuAr arc on-sky immediately after each science exposure. We used 342 s exposures, yielding a typical SNR of 90 per pixel at 6000 Å.

We reduced the data using PYPEIT, which required construction of a new template for the R831 grating at blue wavelengths. Cross-correlation with a template yielded formal RV uncertainties of order 1 km s^{-1} , but we found variations of up to 10 km s^{-1} between RVs measured at the blue and red ends of the spectrum. We adopted a RV uncertainty of 4 km s^{-1} for the GMOS RVs, which was also validated by observations of Gaia BH1 with higher-resolution instruments that occurred nearly concurrently with some of the GMOS observations.

A3 X-shooter

We observed Gaia BH1 four times using the X-shooter spectrograph (Vernet et al. 2011) on the *8.2m UT3 telescope* at the VLT on Cerro Paranal (program 2109.D-5047). We used the 0.5, 0.4, and 0.4 arcsec slits on the UVB, VIS, and NIR arms. This setup yielded a resolution of $R \approx 9700$ in the UVB arm, $R \approx 18 400$ in the VIS arm, and $R \approx 11 600$ in the NIR arm, with near-continuous wavelength coverage from 3100 to 24 000 Å. We reduced the spectra using PYPEIT.

We generally measured RVs using the NIR spectra, whose wavelength solution is calculated from telluric OH lines in the science spectra, yielding a maximally stable wavelength solution. For the first epoch, we inadvertently used a 0.9 arcsec slit in the NIR arm, yielding a lower spectral resolution. We measured the RV from the VIS arm for this epoch. We adopt an uncertainty of 1 km s^{-1} on all the X-shooter RVs.

A4 HIRES

We observed Gaia BH1 nine times using the HIRES (Vogt et al. 1994) on the *10m Keck I telescope* on Maunakea. The data were obtained and reduced using the standard CPS setup (Howard et al. 2010), including use of the C2 decker (0.86×14 arcsec), which yields spectra with $R \approx 55\,000$ and wavelength coverage over most of $3700\text{--}8000\text{ \AA}$, with 2 gaps $100\text{--}200\text{ \AA}$ wide. We used 600 s exposures, yielding SNR of order 40 per pixel at 6000 \AA . The CPS reduction includes sky-subtraction using the long C2 decker.

As part of the CPS pipeline, absolute RVs are measured using the method described by Chubak et al. (2012): to correct for instrumental shifts in the wavelength zero point, the telluric A and B molecular oxygen absorption bands of each star are cross-correlated with those of RV standard stars, whose adopted RVs are matched to the scale of Nidever et al. (2002) and other RV standard catalogs. The resulting RVs are robust to 0.1 km s^{-1} . We use these RVs throughout our analysis. Most of our observations were performed through an iodine cell, making it possible to measure higher-precision RVs from these data in the future, after a higher-SNR template spectrum has been obtained and the method of Butler et al. (1996) can be applied.

A5 FEROS

We observed Gaia BH1 17 times with the Fiberfed Extended Range Optical Spectrograph (FEROS; Kaufer et al. 1999) on the *2.2m ESO/MPG telescope* at La Silla Observatory. We used 2×2 binning in order to reduce read-out noise (allowing us to measure RVs down to $G \approx 16$); this resulted in an effective resolution $R \approx 37\,000$ measured from sky lines. We took 2400 ssures and achieved a typical SNR of 20 per pixel. We reduced the data with the standard ESO MIDAS pipeline, which performs bias-subtraction, flat-fielding, wavelength calibration, and sky subtraction, and merges the spectra from different orders.

On some nights, the wavelength solution drifted from the one established by the afternoon arcs by up to a few km s^{-1} . We corrected this by cross-correlating an arc spectrum (which is taken during each exposure with a second fiber) with a reference arc spectrum taken on the first night of our observations. Based on observations of an RV standard with the same setup, we adopted and uncertainty of 0.3 km s^{-1} for all the FEROS observations, except in one case where poor conditions yielded a formal uncertainty of 0.5 km s^{-1} . Near-concurrent FEROS and HIRES observations of the source allowed us to verify that the two RV zero points are consistent within this tolerance.

A6 ESI

We observed Gaia BH1 one time with the Echellette Spectrograph and Imager (ESI; Sheinis et al. 2002) on the *10m Keck II telescope* on Maunakea. We used a 300 s exposure with the 0.3 arcsec slit, yielding a resolution $R \approx 12\,000$ and $\text{SNR} = 60$, with useful wavelength coverage of $3900\text{--}10\,000\text{ \AA}$. We reduced the data using the MAuna Kea Echelle Extraction pipeline, which performs bias-subtraction, flat-fielding, wavelength calibration, and sky subtraction, and we refined the wavelength solution using sky lines. Based on RV measurements of RV standards on the same night, we adopted an uncertainty of 1.5 km s^{-1} .

APPENDIX B: GAIA OBSERVATIONS

As discussed in Section 4.2.1, the astrometric uncertainties for the parameters describing the orbit of Gaia BH1 are unusually large for the star's apparent magnitude. Here we show that this is likely a consequence of the *Gaia* scanning law.

Epoch-level astrometric data are not published in DR3, but one can determine when a source was observed using the *Gaia* observation

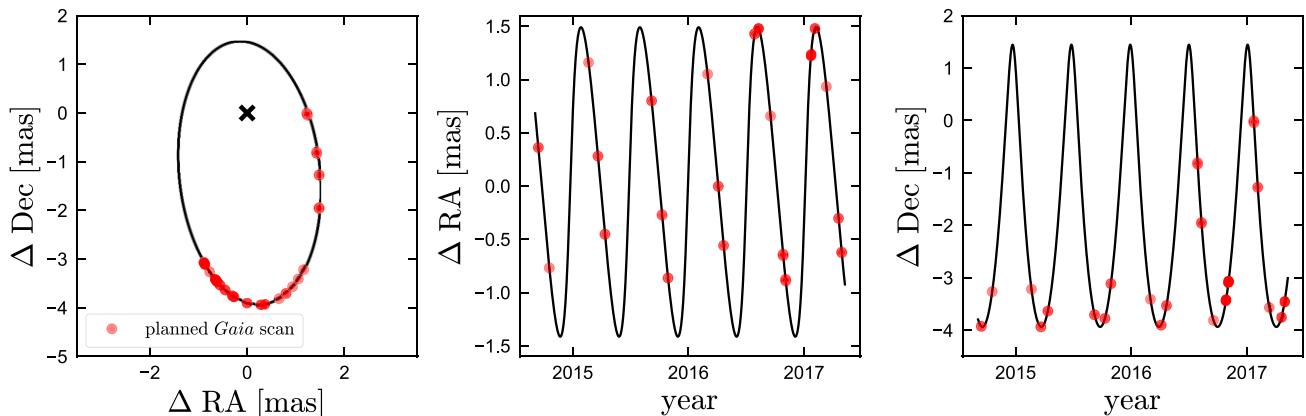


Figure B1. Predicted observation times of Gaia BH1 from the GOST. Black line shows the best-fitting photocentre orbit. Red points show the predicted photocentre positions at the times when GOST predicts *Gaia* would have observed the source. Note that we do not have access to the actual measured ΔRA and ΔDec values; only to the predicted scan times. GOST predicts that *Gaia* observed the source in 21 visibility periods which all fall on one side of the orbit. This may be a result of the fact that the orbital period is a near-integer multiple of *Gaia*'s 63-d precession period, and almost exactly half a year.




Figure C1. Comparison of measured RVs (same as Fig. 2) to predictions from the *Gaia* astrometric solution alone. The centre-of-mass RV (which is not constrained by astrometry) is set to the value inferred from RVs, but the cyan lines otherwise do not know about the measured RVs. The uncertainty in the predictions is larger than when RVs are included in the fit, but the measured RVs are generally consistent with the predictions, validating the *Gaia* solution.

scheduling tool (GOST).⁷ Given a source position, GOST returns a list of observation times when the scanning law predicts that a source will transit across the *Gaia* focal plane. It is not guaranteed that *Gaia* actually obtains data on each source at the predicted times, as the tool does not account for gaps between CCDs and issues that cause temporary gaps in the data stream, such as micrometeor impacts. However, for *Gaia* BH1, GOST predicts 21 visibility periods (i.e. groups of observations separated from other groups by at least 4 d), and the *gaia_source* table reports that 21 visibility periods were indeed used in the astrometric solution. This suggests that the scan times predict by GOST are likely a good approximation of reality.

We visualize these times in Fig. B1. It is clear that although *Gaia* DR3 covered more than five orbital periods of *Gaia* BH1, the phase coverage is far from uniform, and only about half of the orbit in phase was covered by astrometric data. This may be a consequence of the fact that the orbital period is a near-integer multiple of the *Gaia* precession period, and close to half a year. Our RV follow-up shows that despite this limitation, *Gaia* was able to measure an astrometric orbit that was basically correct. However, the incomplete phase coverage make it unsurprising that the uncertainties in the astrometric parameters are larger than usual for a source of this brightness.

GOST predicts that the north side of the orbit will remain poorly sampled even in *Gaia* DR4. We note that a few strategically timed observations with *HST*, *JWST*, GRAVITY+, or MICADO could fill in this part of the orbit and yield improved astrometric precision if analysed together with the epoch-level *Gaia* astrometric data that will become available in DR4.

APPENDIX C: COMPARISON OF THE ASTROMETRIC AND RV SOLUTIONS

As discussed in Section 4, the tightest constraints on the *Gaia* BH1 orbit can be obtained through joint fitting of the RVs and astrometric constraints. However, it is also useful to compare the RVs predicted by the *Gaia* solution alone to the measured RVs, in order to assess whether or not the astrometric solution is consistent with the RVs. Fig. C1 compares the measured RVs to predictions based on the *Gaia* astrometric solution, which are generated from posterior draws of the fit without RVs. Because the astrometric solution does not constrain the centre-of-mass velocity, we set it to 45 km s^{-1} (consistent with the measured value) in all cases.

The measured RVs are basically consistent with the *Gaia*-only solution, in the sense that some posterior draws go through all the RVs. Fig. C2 and Table C1 also show that all inferred parameters are consistent within $\lesssim 2\sigma$. This suggests that the *Gaia* solution and

⁷<https://gaia.esac.esa.int/gost/>

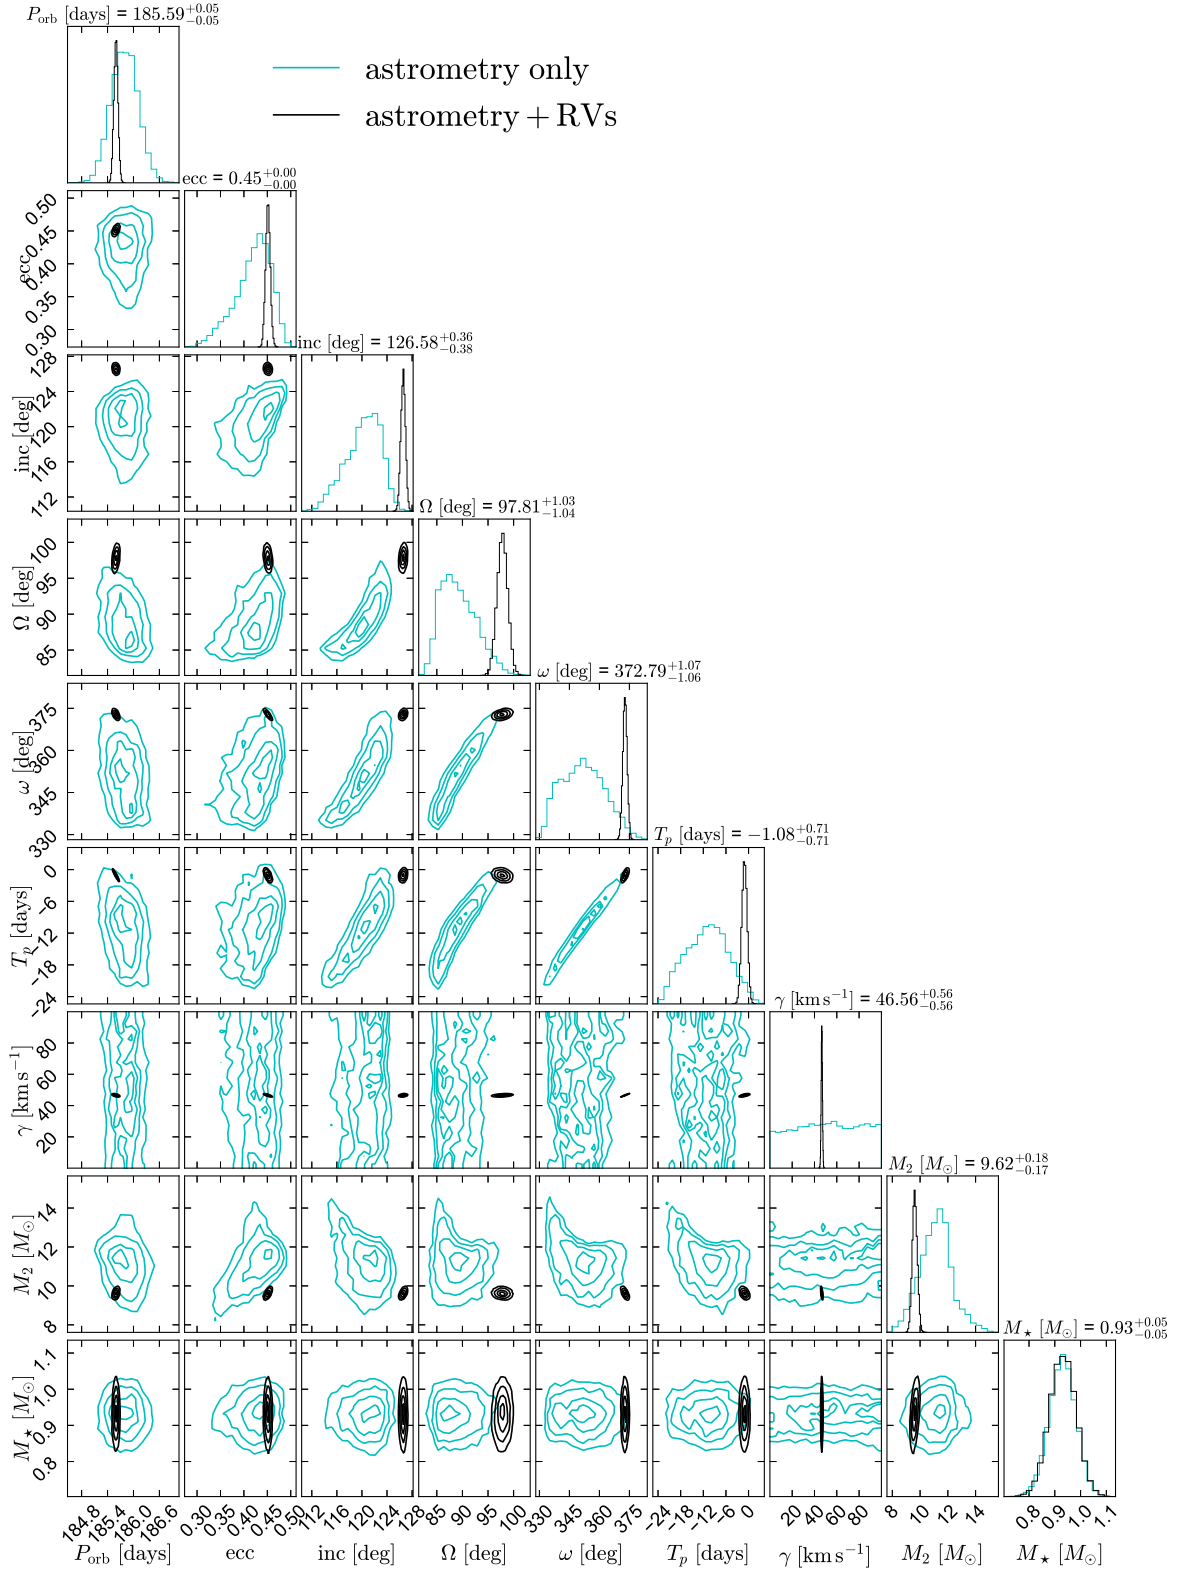


Figure C2. Comparison between constraints obtained from joint fitting of RVs and astrometry (black) and constraints from astrometry alone (cyan). The two sets of constraints are generally consistent, but those with RVs are tighter. Numbers on the diagonals reflect the joint astrometry + RVs constraints. Both sets of constraints are listed in Table 1.

Table C1. Comparison of the *Gaia* astrometric constraints on *Gaia* BH1's photocentre ellipse to constraints from the joint RVs + astrometry fit. All parameters are consistent at the 1.6σ level.

Parameter	units	Astrometry-only constraint	Joint RVs+astrometry constraint	How many σ discrepant?
a_thiele_innes	mas	-0.26 ± 0.17	-0.01 ± 0.05	1.1
b_thiele_innes	mas	2.93 ± 0.18	2.62 ± 0.02	1.6
f_thiele_innes	mas	1.52 ± 0.16	1.61 ± 0.02	0.5
g_thiele_innes	mas	0.53 ± 0.55	-0.36 ± 0.05	1.5
period	days	185.77 ± 0.31	185.59 ± 0.05	0.5
eccentricity		0.49 ± 0.07	0.45 ± 0.01	0.5
t_periastron	days	-12.0 ± 6.3	-1.1 ± 0.7	1.6

its uncertainties are reliable, despite the incomplete orbit coverage discussed in Appendix B.

APPENDIX D: RADIAL VELOCITIES

RVs are listed in Table D1. All times are reported as mid-exposure times in the heliocentric frame (HJD UTC). RVs are also in the

Table D1. RVs.

HJD UTC	RV (km s^{-1})	Instrument	Resolution R	SNR
2457881.2899	20.0 ± 4.0	LAMOST	1800	6
2458574.3663	8.9 ± 5.6	LAMOST	1800	55
2459767.6226	63.8 ± 3	MagE	5400	110
2459791.9186	131.90 ± 0.1	HIRES	55 000	40
2459795.6461	141.4 ± 3	MagE	5400	100
2459796.4995	142.7 ± 3	MagE	5400	100
2459798.8399	140.6 ± 4	GMOS	4300	80
2459805.5101	127.7 ± 1.0	X-shooter	18 400	55
2459808.7388	118.0 ± 4	GMOS	4300	75
2459813.6045	90.5 ± 0.3	FEROS	37 000	20
2459814.5874	86.1 ± 0.3	FEROS	37 000	15
2459815.5927	81.9 ± 0.3	FEROS	37 000	19
2459817.5278	74.5 ± 0.3	FEROS	37 000	20
2459818.5266	71.0 ± 0.3	FEROS	37 000	20
2459818.7870	67.8 ± 4	GMOS	4300	85
2459819.5543	67.0 ± 0.3	FEROS	37 000	20
2459820.5465	64.0 ± 0.3	FEROS	37 000	21
2459821.5669	60.7 ± 0.3	FEROS	37 000	22
2459822.5745	57.8 ± 0.3	FEROS	37 000	22
2459823.5422	54.8 ± 0.3	FEROS	37 000	21
2459823.8525	53.76 ± 0.1	HIRES	55 000	40
2459824.5305	52.1 ± 0.3	FEROS	37 000	18
2459824.8516	51.18 ± 0.1	HIRES	55 000	40
2459825.5361	49.8 ± 0.3	FEROS	37 000	20
2459826.7920	46.59 ± 0.1	HIRES	55 000	40
2459828.5677	42.2 ± 0.3	FEROS	37 000	20
2459829.5373	42.1 ± 3	MagE	5400	80
2459829.5768	40.4 ± 0.3	FEROS	37 000	17
2459830.6452	38.5 ± 0.3	FEROS	37 000	10
2459831.6223	36.5 ± 0.5	FEROS	37 000	5
2459833.7523	33.23 ± 0.1	HIRES	55 000	40
2459834.5509	32.3 ± 0.3	FEROS	37 000	11
2459834.7691	31.74 ± 0.1	HIRES	55 000	40
2459835.7678	30.14 ± 0.1	HIRES	55 000	40
2459838.8082	26.35 ± 0.1	HIRES	55 000	40
2459838.7208	27.5 ± 4	GMOS	4300	78
2459840.7729	24.20 ± 0.1	HIRES	55 000	40
2459845.5069	19.4 ± 1.0	X-shooter	11 600	45
2459855.5012	14.2 ± 1.0	X-shooter	11 600	45
2459868.5128	9.3 ± 1.0	X-shooter	11 600	40
2459877.6978	10.5 ± 1.5	ESI	12 000	60

heliocentric frame and their uncertainties include uncertainties in wavelength calibration. The reported SNR is at the wavelength range we used for RV measurement, which in most cases was near 5400 Å.

APPENDIX E: OTHER CANDIDATES

Fig. E1 illustrates our strategy for selecting BH candidates among the *Gaia* astrometric solutions. We show that binaries containing a BH and a low-mass main-sequence star are expected to have much larger photocentre wobble than binaries containing two luminous stars.

The left-hand panel shows a simulation. Here, we begin with all sources in the `gaia_source` catalog with `parallax_over_error > 50`. We then attached a simulated binary population to these sources, drawing periods from the canonical lognormal distribution for solar-type stars (Raghavan et al. 2010) and mass ratios from a uniform distribution. For each simulated binary, we then calculate the expected photocentre semimajor axis and its uncertainty, and we retain only the binaries that satisfy the cuts imposed on the binary solutions published in DR3. We also generate a population of hierarchical triples, drawing both the inner and outer mass ratios from uniform distributions. We calculate the flux ratios using the G -band mass–magnitude relation tabulated by Janssens et al. (2022). Finally, we generate a population of BH + normal star and neutron star + normal binaries, by assigning $10 M_{\odot}$ and $1.4 M_{\odot}$ dark companions to a random subset of the luminous sources, again drawing periods from the canonical lognormal. This is *not* meant to represent a realistic distribution of orbital periods, masses, or mass ratios, but only to show the expected photocentre orbit size at a given period.

In the right-hand panel, we show the same observables for all objects in DR3 with astrometric solutions (Orbital and AstroSpectroSB1 solutions). As expected, most solutions fall in the region of parameter space that we expect to be populated by luminous binaries and triples. For each binary, we calculate \tilde{M} , a quantity which is related to, but not equal to, the total dynamical mass:

$$\tilde{M} = \frac{4\pi^2}{P_{\text{orb}}^2 G} \left(\frac{a_0}{\varpi} \right)^3. \quad (\text{E1})$$

For a binary with a luminous star and a dark companion, $\frac{a_0}{\varpi} = a \frac{q}{1+q}$, where $q = M_2/M_*$ and a is the semimajor axis. Thus, $\tilde{M} = M_{\text{tot}} \left(\frac{q}{1+q} \right)^3$. For BH companions to low-mass stars, $q/(1+q) \approx 1$, and thus \tilde{M} is only slightly less than M_{tot} . On the other hand, for a binary with two luminous stars, $\frac{a_0}{\varpi} = a \delta_{q\ell}$, where

$$\delta_{q\ell} = \left| \frac{q/(1+\ell) - \ell/(1+\ell)}{(1+q)} \right|, \quad (\text{E2})$$




Figure E1. Parameter space within which we selected BH candidates. Left-hand panel shows a simulation, in which we model a population of luminous (main-sequence) binaries, triples, and compact object + normal star binaries subject to the cuts imposed on binary solutions published in DR3. The relative number of objects in the three populations is chosen arbitrarily. Binaries containing dark companions fall at larger separations at fixed period than binaries containing two (or three) luminous stars. Right-hand panel shows the same quantities for binaries published in DR3. Red star shows Gaia BH1; black points show other initially promising candidates, which we discuss in the text.

Table E1. BH candidates highlighted in Fig. E1. GoF is the `goodness_of_fit` parameter, for which large values indicate a potentially problematic solution (see text for details). `rv_amplitude_robust` is the min-to-max range in the individual-epoch *Gaia* RV measurements (only available for bright sources); ‘expected RV amplitude’ is the min-to-max RV amplitude predicted by the astrometric solution. Sources are discussed individually in the text.

Gaia DR3 source ID	G (mag)	\tilde{M} (M_{\odot})	P_{orb} (d)	$(a_0 \times d)$ (au)	GoF	<code>rv_amplitude_robust</code> (km s^{-1})	expected RV amplitude (km s^{-1})	Verdict
3640889032890567040	9.2	122 ± 47	1076 ± 12	10.2 ± 1.3	10.3	12.9	674 ± 154	✗
4467000291193143808	15.5	119 ± 71	1647 ± 520	13.4 ± 2.2	4.9		358 ± 182	✗
4373465352415301632	13.8	11.5 ± 2.5	185.8 ± 0.3	1.44 ± 0.10	0.3		165 ± 20	✓
6281177228434199296	11.3	10.8 ± 1.6	153.9 ± 0.4	1.24 ± 0.06	8.0	20.1	171 ± 10	✗
5870569352746779008	12.3	6.7 ± 0.5	1352 ± 45	4.52 ± 0.13	3.1	37.0	50 ± 2	⚠
3664684869697065984	11.6	3.5 ± 2.6	1220 ± 233	3.35 ± 1.13	3.6	18.3	54 ± 14	✗

and $\ell = F_2/F_1$ is the ratio of the two stars’ G -band fluxes (e.g. Penoyre, Belokurov & Evans 2022). For binaries with two main-sequence stars, $\delta_{q\ell} \lesssim 0.35$ (e.g. Shahaf et al. 2019), and thus in general $\tilde{M} \lesssim 0.05M_{\text{tot}}$. \tilde{M} is thus a good quantity with which to select massive dark companions to low-mass stars, though more sophisticated methods exist that are sensitive to lower-mass BHs and BH companions to massive stars (e.g. Shahaf et al. 2019, 2022).

With this in mind, we selected objects from *Gaia* DR3 with $\tilde{M} > 3M_{\odot}$. This yielded six objects, which are highlighted in the right-hand panel of Fig. E1 and enumerated in Table E1. Sources brighter than $G_{\text{RVs}} = 12$ have multi-epoch RVs measured (but not yet published) by *Gaia*. The min-to-max range of these RVs is published as `rv_amplitude_robust`. Comparing this to the expected full RV amplitude implied by the *Gaia* solution provides a useful diagnostic of whether the solution is likely to be correct.

Each solution is accompanied by a `goodness_of_fit` value. As discussed by Halbwachs et al. (2022), this quantity is formally expected to follow a normal distribution, $\mathcal{N}(0, 1)$, with values significantly larger than 1 indicative of a problematic solution. However, this may not hold in practice, since the uncertainties in epoch-level measurements can be underestimated. We find that the distribution of `goodness_of_fit` values displays a sharp discontinuity at $G = 13$. Fainter values approximately follow a $\mathcal{N}(0, 1.5)$ distribution, such that a value much larger than 1 is indeed unusual. However,

for brighter solutions, the distribution peaks at ≈ 3.5 , and 13 per cent of solutions of have `goodness_of_fit` > 10 . We thus expect significantly larger `goodness_of_fit` values to be (potentially) consistent with a good solution for bright sources.

We comment briefly on each source:

3640889032890567040: Poor GoF and implausibly high companion mass. Ruled out by low `rv_amplitude_robust`.

4467000291193143808: Period of 1647 d is significantly longer than *Gaia* DR3 baseline. Implausibly high companion mass; poor GoF for the source magnitude.

4373465352415301632: The subject of this paper.

6281177228434199296: Ruled out by our RV follow-up. Low `rv_amplitude_robust` also makes the astrometric solution implausible.

5870569352746779008: Most plausible candidate beside Gaia BH1, but still uncertain. P_{orb} is longer than the *Gaia* DR3 observing baseline. The GoF is significantly larger than 1, but is not unusual for sources in the relevant range of G magnitude. The CMD position is consistent with a lower giant, or possibly a red clump star. The inclination is relatively low (35 ± 1 deg), so the nature of the companion hinges on the correctness of the astrometric solution. That is, if the *Gaia* solution is correct, RVs alone will only yield a spectroscopic mass function near $1M_{\odot}$, which is not high enough to unambiguously rule out a luminous companion to a red giant. The

reliability of the astrometric solution will be easier to assess in DR4, when epoch-level astrometric data are published and a full orbit has been covered. The measured `rv_amplitude_robust` is not too far from the expected peak-to-peak amplitude. We have initiated a spectroscopic follow-up campaign, on which we will report in future work. Thus far, the RVs are consistent with the `AstroSpectroSB1` solution. Further RV monitoring is warranted and necessary. This source was recently investigated in detail by Tanikawa et al. (2022), who found it to be the only credible BH candidate in *Gaia* DR3 with both astrometric and spectroscopic *Gaia* data, but did not carry out follow-up. If the solution turns out to be credible, the source will present a similar puzzle to *Gaia* BH1: the BH progenitor would likely not have fit inside the current orbit, but the orbit is much too wide to be a result of CE evolution. The current orbital separation is such that if the luminous star is a core helium-burning red clump star, mass transfer via winds and/or stable Roche lobe overflow would have occurred prior to the helium flash.

3664684869697065984: Ruled out by APOGEE RVs.

For a $10 M_{\odot}$ BH companion, the cut of $\tilde{M} > 3 M_{\odot}$ would be satisfied for luminous stars with $M_{\star} \lesssim 8 M_{\odot}$. For higher-mass luminous stars, the AMRF, \mathcal{A} statistic designed by Shahaf et al. (2019) is more sensitive. We verified that no objects in the `gaia_dr3_nss_two_body_orbit` catalog with photometric masses above $5 M_{\odot}$ have $\mathcal{A} > 0.6$, as would be required to identify a dark companion with high confidence.

APPENDIX F: HIRES CANNON LABELS

The labels inferred by the *Cannon* fit to the HIRES spectrum of *Gaia* BH1 are listed in Table F1. The method is not expected to yield reliable formal uncertainties on labels, because at high

Table F1. Parameters of the G star, from the HIRES cannon. The cross-validation error is taken from Rice & Brewer (2020) and represents an estimate of the uncertainty in each label.

Parameter	Best fit	Cross-validation error
T_{eff} (K)	5863	56
$\log(g/\text{cm s}^{-2})$	4.43	0.09
$v \sin i$ (km s^{-1})	1.3	0.87
[C/H]	-0.24	0.05
[N/H]	-0.37	0.08
[O/H]	-0.10	0.07
[Na/H]	-0.33	0.05
[Al/H]	-0.28	0.04
[Mg/H]	-0.23	0.04
[Si/H]	-0.26	0.03
[Ca/H]	-0.26	0.03
[Ti/H]	-0.21	0.04
[V/H]	-0.22	0.06
[Cr/H]	-0.30	0.04
[Mn/H]	-0.50	0.05
[Fe/H]	-0.29	0.03
[Ni/H]	-0.34	0.04
[Y/H]	-0.31	0.08

SNR the errors are dominated by systematics in the spectral model rather than photon noise. We thus report the cross-validation error achieved by Rice & Brewer (2020) for each label when testing the model, which provides a rough lower limit on the uncertainties in each label.

The atmospheric parameters and abundances reported by *the Cannon* are quite similar to those we find with BACCHUS (Section 5.1), which we adopt as our fiducial values. Note that *the Cannon* reports [X/H], while for the BACCHUS fit we reported $[X/\text{Fe}] = [X/\text{H}] - [\text{Fe}/\text{H}]$.

APPENDIX G: DEPENDENCE OF GAIA ASTROMETRIC UNCERTAINTIES ON OBSERVABLES

In Fig. G1, we show the mean uncertainty in a_0 and ϖ for DR3 astrometric binary solutions as a function of apparent magnitude and a_0 . We only include binaries with best-fitting $P_{\text{orb}} < 1000$ d, which is roughly the time baseline of observations used to calculate DR3 solutions. The figure shows that typical astrometric uncertainties – especially the uncertainty in a_0 – depend both on apparent magnitude and on a_0 . At fixed apparent magnitude, binaries with $a_0 \approx 3$ mas have uncertainties in a_0 and ϖ that are respectively about $2.5 \times$ and $1.7 \times$ larger than sources with $a_0 \approx 0.3$ mas.

It is not immediately obvious whether this trend reflects an actual dependence of the astrometric uncertainties on a_0 among all solutions that were initially fit, or is mainly a result of the astrometric SNR cuts imposed on the published solutions (equations 10–12). Fig. G2 shows that removal of solutions with larger uncertainties and short periods likely has a significant effect on trends with a_0 : there is a ‘wedge’ of missing solutions with small a_0 and larger uncertainties.

To determine whether an inherent dependence of the astrometric uncertainties on a_0 (or another parameter that scales with it) is required to explain the trends in Fig. G1, we performed the following experiment. Using the simulation setup described in Appendix E, we begin with all sources in the `gaia_source` catalog with `parallax_over_error > 10` and create a simulated binary population that we then subject to the DR3 SNR cuts. To predict the uncertainties in a_0 and ϖ for each simulated binary, we first tabulate the mean and standard deviation of σ_{a_0} and σ_{ϖ} as a function of apparent *G*-band magnitude, *multiply this mean by an a_0 -dependent scaling factor*, and then draw uncertainties for each source from a Gaussian with the appropriate mean and standard deviation. Finally, we compare the simulated distribution of σ_{a_0} and σ_{ϖ} for binaries surviving the quality cuts to the observed distribution shown in Fig. G1.

We parametrize the scaling factor as $f = a \log(a_0/\text{mas}) + b$. We find that when $a = 0$ and $b = 1$ (i.e. no dependence of uncertainties on a_0), the predicted correlation between a_0 and the astrometric uncertainties is weaker than observed. We can best reproduce the observed distribution for σ_{a_0} when $a = 1.3$ and $b = 1.4$; for σ_{ϖ} , we find $a = 0.64$ and $b = 1.2$. We adopt these values when predicting the expected astrometric uncertainties of a simulated binary population in Section 9.

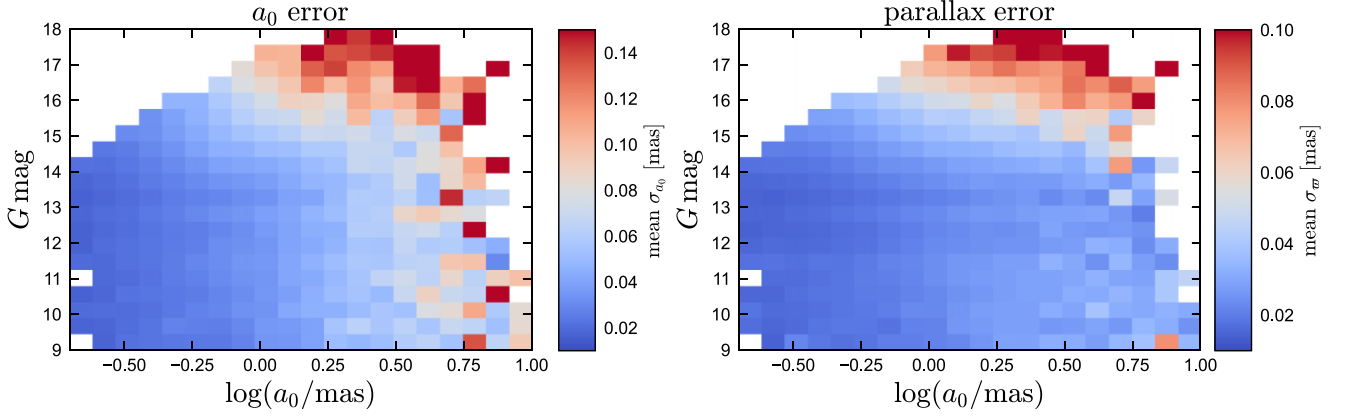


Figure G1. Mean uncertainty in photocentre semimajor axis (a_0 ; left) and parallax (right) for all *Gaia* DR3 astrometric solutions with $P_{\text{orb}} < 1000$ d. Fainter sources and sources with large a_0 have larger astrometric uncertainties. The correlation between a_0 and the astrometric uncertainties appears to be only partially a result of period-dependent quality cuts imposed on sources in DR3 (see text). We incorporate an explicit a_0 -dependence when predicting the astrometric uncertainties of a hypothetical binary observed by *Gaia*.




Figure G2. Uncertainties in parallax (top) and photocentre semimajor axis (a_0 ; bottom) for all sources with astrometric solutions published in DR3. Left column shows uncertainties as a function of a_0 ; right column shows them as a function of P_{orb} . Red star shows Gaia BH1. A correlation is apparent between the astrometric uncertainties and a_0 . This is in part because sources with small a_0 and large uncertainties were removed from the catalog through period-dependent quality cuts (equations 10–12). More structure is apparent as a function of P_{orb} : binaries with periods close to 0.5 and 1 yr, as well as harmonics of the 63 d precession period, have larger uncertainties. Uncertainties in a_0 also increase at $P_{\text{orb}} > 700$ d, likely due to incomplete phase coverage.

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