



Even a Precessing Clock Is Right Twice per Orbit—The Superperiods of eRO-QPE2 and Challenges for Quasiperiodic Eruption Orbital Models

R. Arcodia^{1,2,22}, G. Miniutti³, J. Chakraborty¹, A. Franchini⁴, M. Giustini³, I. Linial^{5,6,22}, A. Mummery⁷, L. Bertassi^{8,9,10}, M. Bonetti^{8,9}, E. Kara¹, A. Merloni¹¹, A. Motta⁸, G. Ponti^{11,12,13}, E. Quintin¹⁴, R. Soria¹⁵, P. Baldini¹¹, J. Buchner¹¹, M. Dotti^{8,9,10}, P. C. Fragile¹⁶, A. Ingram¹⁷, M. Middleton¹⁸, C. Panagiotou¹, A. Sesana^{8,9,10}, P. Yao¹⁹, A. Rau¹¹, F. M. Vincentelli^{18,20}, M. Guolo²¹, and R. Saxton¹⁴

¹ Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

² Black Hole Initiative at Harvard University, 20 Garden Street, Cambridge, MA 02138, USA

³ Centro de Astrobiología (CAB), CSIC-INTA, Camino Bajo del Castillo s/n, 28692, Villanueva de la Cañada, Madrid, Spain

⁴ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy

⁵ Department of Physics and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

⁶ Center for Cosmology and Particle Physics, Physics Department, New York University, New York, NY 10003, USA

⁷ School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA

⁸ Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

⁹ INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

¹⁰ INAF—Osservatorio Astronomico di Brera, via Brera 20, I-20121 Milano, Italy

¹¹ Max-Planck-Institut für extraterrestrische Physik (MPE), Gießenbachstraße 1, 85748, Garching bei München, Germany

¹² INAF-Osservatorio Astronomico di Brera, Via E. Bianchi 46, I-23807 Merate (LC), Italy

¹³ Como Lake Center for Astrophysics (CLAP), DiSAT, Università degli Studi dell’Insubria, via Valleggio 11, I-22100 Como, Italy

¹⁴ European Space Astronomy Centre (ESAC), European Space Agency (ESA), Madrid, Spain

¹⁵ INAF-Osservatorio Astrofisico di Torino, Strada Osservatorio 20, I-10025 Pino Torinese, Italy

¹⁶ Department of Physics and Astronomy, College of Charleston, 66 George Street, Charleston, SC 29424, USA

¹⁷ School of Mathematics, Statistics, and Physics, Newcastle University, Newcastle upon Tyne NE1 7RU, UK

¹⁸ School of Physics and Astronomy, University of Southampton, University Road, Southampton SO17 1BJ, UK

¹⁹ Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

²⁰ Fluid and Complex Systems Centre, Coventry University, CV1 5FB, UK

²¹ Bloomberg Center for Physics and Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA

Received 2026 February 25; revised 2026 April 3; accepted 2026 April 9; published 2026 May 21

Abstract

We present $O - C$ (“observed minus calculated”) timing analysis of the quasiperiodic eruption (QPE) source eRO-QPE2 with a multimission X-ray campaign including 32 observed eruptions spanning a month (i.e., 325 cycles). For disk instability models, the $O - C$ is consistent with a damped random walk of the recurrence, albeit with highly uncertain parameters. For models with an underlying orbital clock, we infer a period of $P \sim 2.24$ hr and two hierarchical superperiodic modulations with periods of ~ 4.4 days ($\sim 47 P$) and ~ 95 days ($\sim 1000 P$). We found no negative period derivative, with $|\dot{P}| \lesssim 2 \times 10^{-6} \text{ s s}^{-1}$ at 3σ . This limits gravitational-wave decay disfavoring high-eccentricity white dwarfs and high-mass/eccentricity intermediate-mass black holes. For disk-collision models, where the $\dot{P}$ from gas drag and the QPE integrated energy bound the local disk density, a main-sequence star is disfavored as extreme mass-ratio inspiral (EMRI) secondary unless debris streams are present. The correlated odd/even $O - C$ disfavors both disk crossings per orbit being observed. Interpreting the data with one *observed* event per orbit, the short modulation is consistent with apsidal precession for $a \sim 140 R_g$, $e \approx 0.1$, and $M_{\text{BH}} \approx 1.5 \times 10^5 M_\odot$. The longer modulation (much less constrained) is inconsistent with EMRI nodal precession, and disk precession is allowed in a limited parameter volume, while a hierarchical triple with an outer black hole at ~ 0.4 mpc and mass $\sim (0.1\text{--}1) \times M_{\text{BH}}$ is viable. However, more robust EMRI trajectory models yield no reliable solution, possibly due to narrow likelihood peaks in a high-dimensional space with sparse data.

Unified Astronomy Thesaurus concepts: X-ray astronomy (1810); Black holes (162); Active galactic nuclei (16); Extragalactic astronomy (506); X-ray bursts (1814); X-ray transient sources (1852); Timing variation methods (1703); Gravitational waves (678)

Materials only available in the online version of record: machine-readable table

1. Introduction

Quasiperiodic eruptions (QPEs) are soft X-ray flares observed from galactic nuclei evolving on the timescales of hours to days

(G. Miniutti et al. 2019; M. Giustini et al. 2020; R. Arcodia et al. 2021, 2024b, 2025; J. Chakraborty et al. 2021, 2025b; E. Quintin et al. 2023; S. Bykov et al. 2024; M. Nicholl et al. 2024; L. Hernández-García et al. 2025; P. Baldini et al. 2026). When in quiescence, the emission from an accretion disk is detected (M. Nicholl et al. 2024; M. Guolo et al. 2025a, 2025b; T. Wevers et al. 2025), which, at least for some sources, is due to a previous tidal disruption event (TDE; e.g., E. Quintin et al. 2023; S. Bykov et al. 2024; M. Nicholl et al. 2024; J. Chakraborty et al. 2025b; M. Guolo et al. 2025a). The eruptions detected on top of this

²² NASA Einstein Fellow.



continuum show a characteristic spectral evolution with a harder rise compared to the flare decay (R. Arcodia et al. 2022, 2024b; G. Miniutti et al. 2023b; M. Nicholl et al. 2024; J. Chakraborty et al. 2025b; L. Hernández-García et al. 2025). While the average duty cycle is around $\approx 10\%$ – 20% across all sources, their timing properties are somewhat diverse in terms of the scatter around the average QPE peak-to-peak recurrence time (t_{recur}) and flare amplitude.

The presence of an underlying quasiperiodic clock is, however, ubiquitous and reminiscent of an orbital phenomenon, thus inspiring many related ideas. Among the various orbital models,²³ QPEs have been modeled as binary self-lensing (A. Ingram et al. 2021), accretion disk tearing (A. Raj & C. J. Nixon 2021), Lense–Thirring precession in a super-Eddington flow (M. Middleton et al. 2025), partial tidal disruptions (pTDEs) from high-eccentricity white dwarfs (WDs; e.g., A. King 2020; M. Wang et al. 2022; J.-H. Chen et al. 2023) or from mildly eccentric stellar orbits (I. Linial & R. Sari 2023; W. Lu & E. Quataert 2023), quasi-circular pTDEs in a circumbinary disk (D. J. D’Orazio et al. 2007), or disk spiral density waves excited by an orbiter (S. A. Dodd et al. 2025). More prominently, QPEs have been modeled as a stellar-mass object in a low- or moderate-eccentricity orbit crossing the accretion disk around the primary massive black hole (MBH), generally twice per orbit (J. Xian et al. 2021; A. Franchini et al. 2023; I. Linial & B. D. Metzger 2023; H. Tagawa & Z. Haiman 2023; C. Zhou et al. 2024). Many of these interpretations imply that QPEs are the electromagnetic precursor or counterpart of extreme mass-ratio inspirals (EMRIs), which, depending on the precise nature of the orbiter, may be detectable by the recently adopted Laser Interferometer Space Antenna (M. Colpi et al. 2024) and TianQin (J. Luo et al. 2016), or by future μHz gravitational-wave detectors (A. Sesana et al. 2021; see also T. Suzuguchi et al. 2026; Y. Zhan et al. 2026).

The latest developments of these disk collision models suggest that flares may either result from shocked disk gas ejected during the collision (e.g., A. Franchini et al. 2023; I. Linial & B. D. Metzger 2023) or by the interaction between stellar debris streams and the disk (e.g., I. Linial et al. 2025; P. Z. Yao et al. 2025). The latter emission mechanism may solve the tension of the elevated integrated energy per burst seen in some QPE sources, which is higher than what a stellar-mass orbiter could eject piercing through the disk (J. Chakraborty et al. 2025c; W. Guo & R.-F. Shen 2025; A. Mummery 2025). In general, most current QPE disk collision models predict that eruptions would be almost equally observable both during ingress into and egress out of the disk (but see I. Linial et al. 2025), although radiation hydrodynamic simulation shows very asymmetric ejecta and significantly different luminosities (X. Huang et al. 2025; T. Janković et al. 2026; but see A. T.-L. Lam et al. 2025). Having two observable events per orbit implies that the orbital period of the system is traced by $\sim 2t_{\text{recur}}$ (i.e., by the separation between eruptions of the same even/odd parity), which provides an estimate of the orbit size if the black hole mass is known. Thus, this model predicts that the inferred orbit size is smaller than the accretion disk, which is an important consistency check passed at least for the QPE sources AT2019qiz, GSN 069, AT2022upj, and eRO-QPE2

(M. Nicholl et al. 2024; J. Chakraborty et al. 2025b; M. Guolo et al. 2025b; T. Wevers et al. 2025). However, in GSN 069, the accretion disk was estimated to be large enough already in 2014 data when QPEs were not observed (M. Guolo et al. 2025a), thus leaving the observed appearance and disappearance of QPEs as a function of disk flux currently unexplained by QPE disk collision models (but see, e.g., A. Franchini et al. 2023; G. Miniutti et al. 2023a).

In the proposed EMRI orbital models, the observed eruptions would trace the orbital evolution of the system, which may open new avenues for BH parameter inference such as BH mass (M_{BH}) and spin (χ ; as discussed in more detail in J. Chakraborty et al. 2025a; C. Zhou et al. 2025a). Some first attempts have been made using EMRI trajectory models, finding plausible solutions in both single epochs and multiepoch observations of QPE sources (e.g., C. Zhou et al. 2025b). More empirical data analysis tests have also been made to study the quasiperiodicity of QPE sources in the context of EMRI models (e.g., R. Arcodia et al. 2024a; J. Chakraborty et al. 2024; D. Pasham et al. 2024; G. Miniutti et al. 2025). In particular, G. Miniutti et al. (2025) studied time delays of the eruptions in GSN 069 and found that delays for odd and even eruptions are correlated, while a key prediction of EMRI models is that apsidal precession would induce an anticorrelation (by imprinting a negative delay to the ingress collision and a positive one to the egress collision with respect to the observer, and vice versa). Nonetheless, the sparse X-ray data did not allow G. Miniutti et al. (2025) to draw unambiguous conclusions on the timing behavior of the source. In summary, while the EMRI-disk collision models have revealed a promising way forward for the QPE field, no single flavor is able to explain all observables of all QPE sources (e.g., W. Guo & R.-F. Shen 2025; G. Miniutti et al. 2025; A. Mummery 2025).

The goal of this work is to improve on the current state-of-the-art of QPE timing campaigns and on their interpretation with current orbital models. For this, we designed and carried out a multimission X-ray campaign of eRO-QPE2, which has been the most regular QPE source discovered to date in terms of its timing and spectral properties (R. Arcodia et al. 2024a; D. Pasham et al. 2024). We collected data with XMM-Newton, Swift/XRT, NICER, and Einstein Probe, with the main observing baseline being enclosed within four XMM-Newton observations spanning about a month in 2024 July and ancillary data taken more sparsely until 2024 December (see Section 2). We performed $O - C$ (“observed minus calculated”) timing analysis (Section 3) and comparisons against existing orbital models outlining their successes and failures in Sections 4 and 5.

2. X-Ray Analysis and Results

eRO-QPE2 has been observed by XMM-Newton (F. Jansen et al. 2001, hereafter XMM) several times starting from the discovery of QPEs in the system in 2020 (R. Arcodia et al. 2021). This work focuses on a campaign performed by XMM in 2024 June–July (PI: Arcodia) designed to test quantitatively the periodic behavior of eRO-QPE2. This campaign includes four XMM observations spanning about a month of baseline (IDs 0941270101 to -401, taken on 2024 June 28, July 4, July 16, and July 28, respectively), referred to hereafter as “XMM1,” “XMM2,” “XMM3,” and “XMM4,” respectively. In addition, due to XMM4 being incomplete compared to the

²³ But see X. Pan et al. (2023) and K. Kaur et al. (2023), for some examples of accretion disk instability models.




Figure 1. X-ray light curves of the four XMM-Newton epochs of the main campaign, with the elapsed time relative to the start of the first observation ($t_{0,\text{XMM1}}$), which corresponds to MJD = 60489.6419. Red lines and contours show the eruption peak times and 1σ uncertainties fitted with an eruption parametric model. We show the identification number N_{QPE} (see Table B1).

requested exposure, part of it was rescheduled on 2024 December 5 (ID -501, hereafter “XMM5”); thus, more observations were later added in between XMM4 and XMM5. The total exposure (with the pn camera) for XMM1-4 is ~ 134.4 ks, with 7.2 ks more in XMM5. We show XMM1-4 observations with their 19 observed eruptions in Figure 1, and report more details on the processing in Appendix A. For comparison with timing properties, we extracted and fitted XMM quiescence spectra at all epochs following precisely the analysis reported in R. Arcodia et al. (2024a), to which we refer for more details.

In addition to XMM data, we scheduled Target of Opportunity (ToO) observations with Swift-XRT (hereafter XRT; PI: Arcodia), NICER (PI: Arcodia), and Einstein Probe FXT (hereafter EP; PI: Rau), throughout the XMM1-4 campaign and beyond, for a total of ~ 42.8 ks, ~ 81.3 ks, and ~ 11.9 ks, for XRT, NICER, and EP, respectively. We report details on data processing in Appendix A. We show examples of XRT, NICER, and EP data in Figure 2. In general, adding these ancillary datasets results in the detection of 13 eruptions peaks during the XMM1-4 campaign (five with XRT, eight

with NICER), plus an additional six between XMM4 and XMM5 (one with XRT, three with NICER, and two with EP, to add to the two observed in XMM5). Thus, the total number of eruptions observed during the XMM1-4 campaign is 32 (while extending to XMM5 adds eight more). Finally, we note that for all light curves, we performed a barycenter correction with DE405 ephemeris.

2.1. X-Ray Eruptions Arrival Times

We estimated the arrival time of QPEs from their peak, but we have also extracted the start time of QPEs and verified a posteriori that this choice does not impact any of our results (see Appendix B). For XMM light curves, we estimated the peak arrival time fitting each eruption with a parametric double-exponential model, which has proven effective in modeling the fast rise and slow decay of QPE sources (e.g., R. Arcodia et al. 2022, 2024b; J. Chakraborty et al. 2024; L. Hernández-García et al. 2025). For eRO-QPE2, while the bursts may appear symmetric due to their shorter duration and




Figure 2. X-ray light curves of illustrative NICER and XRT eruptions (top panels, bottom-left panel), and the EP observation (bottom-right panel). Different instruments are shown with different symbols: squares for XRT (top left, top right), stars for NICER (top middle, bottom left), triangles for EP (bottom right). As in Figure 1, the QPE arrival time (red line, with shaded regions representing 1σ uncertainties; see Appendix A) and identification numbers are shown (see Table B1). The shaded gray profile is not a fit, but shows a silhouette of the second eruption observed by EP (renormalized arbitrarily) to guide the eye.

lower signal-to-noise compared to other QPE sources, their phase-folded profile shows the same asymmetry (see Figure A.1 in R. Arcodia et al. 2024a). We visually inspected all fits to ensure adequate residuals were obtained for all eruptions, and performed more simulations and tests to confirm that the parametric model choice for the QPE eruption profiles has no impact on our timing solutions (see Appendix B). For EP data, since the eruptions are fully resolved within the ~ 3 ks exposures (Figure 2), we adopt the same fitting method applied to XMM data. For XRT and NICER eruptions, we note that their typical continuous snapshots are ~ 500 – 1000 s. Eruptions in eRO-QPE2 only last ~ 1700 s from rise to decay based on XMM data (R. Arcodia et al. 2024a), and since the eruption-absorbed (i.e., observed) peak flux of $F_{0.5-2.0 \text{ keV}} \sim 4 \times 10^{13} \text{ erg s}^{-1} \text{ cm}^{-2}$ (R. Arcodia et al. 2024a) is only marginally above detectability for both XRT and NICER, the detectable time for XRT and NICER reduces even further. In addition, the eruption recurrence time of ~ 2.3 hr (R. Arcodia et al. 2024a) is also compatible with the orbital gaps of ~ 1.5 – 2.0 hr for XRT and NICER; thus, detecting and identifying an eruption peak is much more challenging. However, the higher flux limit comes to our advantage, as any XRT or NICER snapshot detection would be conveniently associated to being close to the peak (e.g., $\gtrsim 0.02 \text{ c s}^{-1}$ for XRT), while the wings of the eruptions can only be marginally detected (e.g., $\lesssim 0.006 \text{ c s}^{-1}$ for XRT), and the quiescence is not detectable at all (e.g., $\lesssim 0.002 \text{ c s}^{-1}$ for XRT). We report all of the details of our XRT and NICER analyses to estimate the QPE peaks in Appendix B. We show the estimated arrival times for all instruments in Figures 1 and 2 with red vertical lines and related contours for their 1σ uncertainties.

3. $O - C$ Analysis

We performed the $O - C$ analysis, which compares the residuals between the observed arrival times (“O”) with those computed (“C”) starting from a reference event ($T_{0,\text{est}}$) and assuming a constant period (P_{est} , estimated from available data) as a function of the elapsed number of events N_{QPE} . This method (e.g., C. Sterken 2005) has been widely used for stellar systems and binaries (e.g., P. Zasche et al. 2009; G. Hajdu et al. 2015) and, more recently, also for QPEs (J. Chakraborty et al. 2024, 2026; G. Miniutti et al. 2025) and other repeating nuclear transients (A. V. Payne et al. 2021). In its simplest form, namely, for a strictly periodic system, the $O - C$ residuals are linear with a slope given by the “inaccuracy” of the estimated period P_{est} compared to the true period P , and a constant of the order of a small offset between the estimated start of the reference event and the true one: $O - C = (T_0 - T_{0,\text{est}}) + (P - P_{\text{est}})N_{\text{QPE}}$. Most quasiperiodic systems reveal additional residuals to this linear trend, which, once the noise model is known (see Section 3.1), may be modeled with extra terms. For instance, a quadratic term due to a change in the true period, adding or subtracting $\pm(1/2 \dot{P} P)N_{\text{QPE}}^2$, or sinusoidal modulations, adding $A_{\text{mod}} \sin(2\pi N_{\text{QPE}}/P_{\text{mod}} + \phi_{\text{mod}})$, where P_{mod} is defined in dimensionless units of P . For reference, for our XMM1–4 dataset, the association with N_{QPE} was performed with $T_{0,\text{est}}$ taken from the fitted peak time of the first eruption of XMM1 (i.e., 732.69 s after the start of XMM1, MJD = 60489.6419, which is the origin of the time axis, e.g., Figure 1), and P_{est} taken from the average peak-to-peak QPE recurrence time of the XMM1, XMM2, XMM3, and XMM4 observations (i.e., 8055.71 s). We report more details on the event identification with these ephemerides in Appendix B, which we stress is not ambiguous




Figure 3. $O - C$ data for all eruptions in the XMM1-4 campaign. Different symbols indicate different instruments (circles for XMM, stars for NICER, squares for XRT; see Figures 1 and 2). The bottom subpanel shows a zoom-in of the first consecutive events in XMM1.

due to the XMM1-4 campaign being designed with this analysis in mind. We show the resulting $O - C$ data in Figure 3 (and report the values in Table B1 in Appendix B).

In this work, our main results are reported performing $O - C$ analysis including all eruptions together, thus assuming that the periodic phenomenon driving the $O - C$ delays is traced by each eruption and the separation t_{recur} . The main reason for this choice is that it is the most agnostic, but for completeness, and for reasons that will become clear throughout this work, we also performed $O - C$ analysis by separating odd and even eruptions (i.e., a scenario in which the period is traced by $\sim 2t_{\text{recur}}$; Appendix B.4). We show the same $O - C$ data as in Figure 3, but separating odd and even bursts, in Figure 4. The goal of this plot is to highlight how the delays of odd and even eruptions appear correlated, in that consecutive bursts within a single epoch have compatible delays despite the range covered by all epochs (≈ 1 hr). This suggests that separating odd and even eruptions does not provide more information than analyzing all eruptions together.

3.1. The Source of Noise in the $O - C$

From a first glance at the $O - C$ data in Figure 3, the best-sampled epochs ($N_{\text{QPE}} \sim 0-250$) show that the data span delays between approximately -0.4 and ~ 0 hr, and that XMM4 shows a positive delay significantly offset from the previous epochs, suggesting that on the timescale comparable to (or longer than) our XMM1-4 baseline, there is a large-amplitude component driving the delays. Whether these $O - C$ residuals are due to the deterministic quadratic and sinusoidal

components shown above depends on the nature of the noise in the $O - C$ plot. While the $O - C$ technique is, per se, model independent and only assumes that a given event is quasiperiodic, determining the type of noise is instead somewhat model dependent (see, e.g., Appendix A of J. Chakraborty et al. 2026).

For instance, with a QPE origin from accretion disk instabilities, one may expect that the quasiperiod wobbles in a red-noise fashion. For an underlying orbital clock, one may instead expect the noise to be white, and if that of a given event is fully independent from the others, so too would be the noise in the $O - C$ delay data. However, it has been shown that if instead the jitter in the arrival time of a given event depends on the jitter of the previous ones, this white noise in time units cumulates with the number of events similarly to red noise in $O - C$ delay space (C. Koen 2006). Thus, establishing a noise model for the $O - C$ of eRO-QPE2 implies making assumptions on the QPE origin.

We report in Appendix B our attempts at modeling the $O - C$ data of eRO-QPE2 with a damped random walk (DRW) model, mimicking a red-noise process modulating the QPE recurrence time. An important result is that red-noise can reproduce the $O - C$ data as well as or better than deterministic components, which is interesting on its own, although perhaps unsurprising given that it is a two-parameter model spanning orders of magnitude in its posteriors (see Figure B1 in Appendix B), and considering the relatively sparse dataset (even if it is perhaps the best one yet for QPE sources). Furthermore, the fitted damping timescale τ spans orders of magnitude and mostly values greater than our baseline, thus predicting that on short timescales QPE recurrence times would be highly correlated. However, the autocorrelation coefficient of consecutive recurrence times in the XMM1-4 epoch is compatible with ~ 0 . Perhaps most fundamentally, this uncertain DRW model has little constraining or predictive power for existing QPE models. Thus, we defer further discussion to future work.

In this work, our main focus will be on testing existing QPE orbital models. In this case, we make the assumption that the underlying $O - C$ noise is white, which is equivalent to assuming that the astrophysical uncertainties of each QPE arrival time are independent from those of other events. For disk collision models, each crossing is due to the deterministic orbital clock, and the astrophysical uncertainty comes from the offset between disk crossing and emission, which is independent for each collision. In particular, the mean of this offset can be zero given that $O - C$ delays are scaled from a reference epoch. A proxy for the standard deviation of this zero-mean offset during the XMM1-4 epoch could be obtained by subtracting in quadrature the measurement error of eruption peak times (~ 79 s) from the dispersion in QPE recurrences (~ 82 s), obtaining ~ 23 s. Neglecting this small offset dispersion does not affect our conclusions significantly, given that the measurement error of XMM data is a few times higher, and that of XRT/NICER points is $\gtrsim 10-20$ times higher.

3.2. $O - C$ Fitting with Deterministic Components

We implemented different $O - C$ models by using some or all of the components described at the start of this section. We derived posterior probability distributions and the Bayesian evidence of each fit with the nested sampling algorithm MLFriends (J. Buchner 2019) using the UltraNest package (J. Buchner 2021) v 4.4.0, adopting uniform priors for all




Figure 4. Same as Figure 3, but for even (top) and odd (bottom) eruptions separately, with the opposite parity shown in transparency. In XMM data (i.e., circles), which are the only ones with consecutive events, data points with the opposite parity within a given epoch are covered as they have compatible delays.

parameters. The linear component may be subtracted for more intuitive visualization of the additional components, although we note that it is always fitted in simultaneity with all other parameters.

From a quick glance at Figure 3, it is evident that a linear term alone, namely a model like $O - C = (T_0 - T_{0,\text{est}}) + (P - P_{\text{est}})N_{\text{QPE}}$, is inadequate to describe all of the deviations in the delays. In Appendix B we report details on our model comparison, in which we added model components iteratively stopping at the simplest model that significantly improves the Bayesian evidence by many orders of magnitude. Here, we only show results from the selected best-fit model (`lin+mod1+mod2`). The equation used was:

$$O - C = (T_0 - T_{0,\text{est}}) + (P - P_{\text{est}})N_{\text{QPE}} + A_{\text{mod},1} \sin\left(\frac{2\pi N_{\text{QPE}}}{P_{\text{mod},1}} + \phi_{\text{mod},1}\right) + A_{\text{mod},2} \sin\left(\frac{2\pi N_{\text{QPE}}}{P_{\text{mod},2}} + \phi_{\text{mod},2}\right),$$

which has eight free parameters, namely T_0 , P , $A_{\text{mod},1}$, $P_{\text{mod},1}$, $\phi_{\text{mod},1}$, $A_{\text{mod},2}$, $P_{\text{mod},2}$, and $\phi_{\text{mod},2}$. The prior bounds of the periods of the two modulations, given their obvious degeneracy, are separated at $N \sim 100$. We show this fit in Figure 5 and the related corner plot in Appendix B (Figure B2). The fitted true period of the XMM1-4 epoch is $P = (8064 \pm 7)$ s, or ~ 2.24 hr, which is subtracted from both data and model in Figure 5.




Figure 5. $O - C$ data (as in Figure 3) and best-fit model (solid line for median, shaded contours for 1σ and 3σ), in which the fitted linear trend is subtracted for visualization.

Since the linear component has been subtracted, if eRO-QPE2 were a perfectly periodic system, all data points would lie close to the horizontal line at $O - C = 0$. The two superperiod modulations appear hierarchical, in that one with a smaller amplitude and period is superimposed to another with larger amplitude and period. Both fitted periods are far from the prior bounds chosen to separate them. The faster superperiod is $P_{\text{mod},1} = (47.0 \pm 0.6)P$, which corresponds to $P_{\text{mod},1} \sim 4.4$ days, with an amplitude $A_{\text{mod},1} = 289^{+30}_{-31}$ s. The longer modulation is less constrained, and the formal fit from the XMM1-4 campaign best-fit model yields a 1σ (3σ) lower limit on $P_{\text{mod},2} > 917.6P$ ($> 600.1P$) and $A_{\text{mod},2} = 4425^{+1299}_{-1456}$ s. In Appendix B.3.1, we reported our attempts at using ancillary data taken after XMM4 to assess the possible bound constraint on $P_{\text{mod},2}$. Using two further NICER eruptions $\sim (12\text{--}13)$ days after XMM4 suggests $P_{\text{mod},2}$ is bound around $\sim 1016P$ (~ 95 days).

Based on the best-fit `lin+mod1+mod2` model (and using archival data in support, see Appendix B), the XMM1-4 epoch is consistent with the absence of a strong dissipative term, namely a large period decrease or increase $\dot{P}$. We note that, as extensively discussed in R. Arcodia et al. (2024a) and G. Miniutti et al. (2025), comparing recurrence times across epochs that are far apart in time is not a faithful probe of a true $\dot{P}$, unless all of the precise evolution terms are known. With our $O - C$ baseline, we are able to put a sensitive constraint on the possible period decay. We subtracted a quadratic component, $-(1/2 \dot{P} P)N_{\text{QPE}}^2$, to the best-fit model to infer an upper limit on the period decrease $\dot{P}$ while marginalizing over the other model components describing the evolution of the system. We note that the sign in the $\dot{P}$ model component is switched to negative so that positive values of $\dot{P}$ can be sampled log-uniformly. First, we left all model components

free to vary within the original uniform prior ranges. The fit is of comparable quality, with $\Delta \log Z \sim 1$; thus, the additional dissipative component is not statistically required. In fact, the $\dot{P}$ is bound to the lowest prior limit, which was chosen to be 10^{-8} s s^{-1} given that it corresponds to a few seconds for the $-(1/2 \dot{P} P) N_{\text{QPE}}^2$ term with 325 eruptions (below which it is not meaningful to sample). This fit yields a 3σ upper limit at $\dot{P} < 3.9 \times 10^{-6} \text{ s s}^{-1}$. We also performed a fit allowing the parameters to vary only within the 10th–90th percentile range of the best-fit model, so that $\dot{P}$ is marginalized over the best-fit uncertainties and additional degeneracies are kept under control. The 3σ upper limit of this fit is compatible with the previous, at $\dot{P} < 2.3 \times 10^{-6} \text{ s s}^{-1}$.

Finally, we note that performing the same exact fitting procedure on odd and even eruptions separately (jointly fitting all parameters except constants and phases) obtains consistent results modulo a factor of 2 in the superperiods related to the P (now traced by $\sim 2t_{\text{recur}}$). However, as discussed more at length in Section 5, we found it does not appear to be the correct setup to study QPE sources; thus, we only report this analysis in Appendix B.4 for future reference.

4. Interpretation of the $O - C$ Analysis within the EMRI Framework

One natural way to interpret superperiod modulations within orbital phenomena is with precession terms. Of course, a realistic orbital evolution would not be described by a perfect sine even for a simple source like eRO-QPE2; thus, these tests should be considered a first exploratory data-model comparison. For proper comparisons with the EMRI framework, we note that adopting a single QPE recurrence as the periodic cycle implies interpreting QPEs as either a single disk collision per orbit (e.g., with eccentric orbits and/or disks), or two, but only one being electromagnetically observable (e.g., see the discussions in X. Huang et al. 2025; I. Linial et al. 2025). Sections 4.1 and 4.2 discuss possible origins of the shorter and longer modulations, respectively. The $\dot{P}$ constraint, and its implications for the type of the possible orbiter, is discussed in Section 4.3.

4.1. The Shorter Superperiod Modulation: Apsidal Precession?

One of the possible superorbital periods in an EMRI system is due to GR apsidal precession of the orbiter within the orbital plane:

$$\frac{P_{\text{aps}}}{P_{\text{orb}}} = \frac{(1 - e^2) a}{3 R_g} \sim \frac{a}{3 R_g}, \quad (1)$$

where we ignore higher-order corrections due to BH spin, and where the last term holds for low eccentricities (with a $\lesssim 10\%$ correction up to $e \lesssim 0.3$). This regime seems appropriate for a system like eRO-QPE2 showing rather regular timings and flare amplitudes (R. Arcodia et al. 2024a; D. Pasham et al. 2024) implying relatively consistent collision radii with respect to the MBH.

If the orbital period is known independently, the apsidal precession period P_{aps} only depends on the primary M_{BH} ; thus,

it can be used to constrain it:

$$M_{\text{BH}} = \frac{c^3}{6 \sqrt{3} \pi G} P_{\text{orb}}^{5/2} P_{\text{aps}}^{-3/2}. \quad (2)$$

We first try to identify the shortest of the fitted superorbital modulations with apsidal precession. As the $O - C$ simultaneously constrained $P = P_{\text{orb}}$ and $P_{\text{mod},1} = P_{\text{aps}}$, it provides direct constraints for M_{BH} (see more discussion in recent theoretical works; e.g., J. Chakraborty et al. 2025a and C. Zhou et al. 2025a). Our best-fit model parameters yield $M_{\text{BH}} = (1.53 \pm 0.03) \times 10^5 M_{\odot}$, by scaling the observed periods to rest frame using $z = 0.0175$ (R. Arcodia et al. 2021). Naturally, the quoted uncertainties are only the statistical ones from the joint posterior of P and $P_{\text{mod},1}$, while the systematics are unknown at this stage. Given the simplistic nature of the model components in the $O - C$ analysis and the finite data sampling, they may be large. Thus, we consider this constraint as a rough estimate $M_{\text{BH}} \sim 1.5 \times 10^5 M_{\odot}$. We note that eRO-QPE2 has two estimates in the literature from standard techniques: $M_{\text{BH}} - \sigma$ host-galaxy scaling relations yielded $M_{\text{BH}} \sim 10^5 M_{\odot}$ (with the usual ~ 0.4 – 0.5 dex uncertainties; T. Wevers et al. 2024), while X-ray-to-UV accretion disk modeling yielded $M_{\text{BH}} = 7.9_{-3.9}^{+7.9} \times 10^5 M_{\odot}$ (T. Wevers et al. 2025). Given the simplistic nature of the $O - C$ modeling, the overall agreement with the values from the literature is encouraging (keeping the unknown systematics in mind).

As a further consistency check, we used the modulation amplitude to verify that it is self-consistent with P_{aps} . From Equation (1), the expected semimajor axis a of the orbit is $\sim 141 R_g$. The maximum $O - C$ amplitude associated with light travel time induced by apsidal precession (which is obtained for an EMRI orbital plane close to edge-on) is $\sim a/c$, thus ~ 106 s. Another effect of apsidal precession is that the true anomaly, thus the velocity, changes at the collision point. For low eccentricity, the difference between the velocity at its maximum and minimum (i.e., at pericenter and apocenter) relative to an analogous circular orbit (i.e., the mean value about which apsidal precession would modulate the delays) is of the order of $\sim 2e$. The difference between time delays at these maximum and minimum is thus of the order of $\sim 2e / (2\pi/P) \sim eP/\pi$, which for e within ~ 0.01 – 0.1 would imply delays of ~ 25 – 250 s. Thus, since the observed $A_{\text{mod},1} \sim 289$ s, the system would require an eccentricity of ~ 0.1 to produce the observed time delays via apsidal precession. Given that the scenario discussed here refers to a single observable event per orbit (even if two disk crossing occur), an eccentricity of ~ 0.1 is reasonable for a regular system like eRO-QPE2 (e.g., in terms of maintaining a regular flare amplitude).

4.2. The Origin of the Longer Modulation

As discussed in the previous section and Appendix B, the most conservative estimate for the longer modulation is a lower limit, while our analysis with ancillary data strengthens the constraints to being $\sim 1000 P$ (or ~ 93 days). Regardless, these values can already provide meaningful constraints. Both apsidal and nodal EMRI precessions are bound to delay amplitudes of order $\sim a/c$, far smaller than the ~ 3 – 4 ks observed for the longer modulation. Following the reasoning presented in G. Miniutti et al. (2025), two possible sources of superperiod modulations are disk precession and an outer MBH binary (MBHB), discussed in the subsections below.

For clarity, we note that the statistical requirement of this longer modulation is partially driven by the assumption of underlying white noise, which seemed appropriate for an origin from disk collisions. If future modeling will indicate otherwise, and state that the astrophysical uncertainty on each QPE arrival time depends on the previous events, then a cumulative white-noise mimicking red-noise would be more appropriate (e.g., C. Koen 2006; J. Chakraborty et al. 2026). However, we note that the period jitter scatter would need to be $>3\text{--}4$ times larger than the estimated ~ 23 s in order to reproduce the $O - C$ delays of XMM4 at large $N \sim 325$.

4.2.1. Disk Precession

For the disk precession case, we computed the precession period t_{prec} adopting the rigid-precession framework from A. Franchini et al. (2016). We computed t_{prec} for a range of spin χ of the primary MBH between $\chi = (0.05, 1)$ and slope of the disk surface density profile, $\Sigma \propto \left(\frac{R}{R_g}\right)^{-p_\Sigma}$, within $p_\Sigma = (-1.5, 1.5)$. We adopted as fiducial values a disk aspect ratio of 0.1 and $500R_g$ for the disk outer disk radius.²⁴ We show this 2D (χ, p_Σ) map in the top panel of Figure 6, which shows that t_{prec} decreases for increasing spin and steeper disk density profiles. The fitted putative precession period from the $O - C$ data is $P_{\text{mod},2}$, and we show its median value and 3σ range in cyan. At face value, the top panel of Figure 6 would show that there is ample margin for a disk precession solution to provide the observed $P_{\text{mod},2}$. However, we have also considered the disk alignment problem. As summarized in A. Franchini et al. (2016), alignment would occur on the timescale over which the disk cools and thins (t_{thin} ; e.g., N. Stone & A. Loeb 2012) and also at a rate inversely proportional to the disk viscosity α (t_{align} ; A. Franchini et al. 2016). Both of these timescales can be computed for values of the 2D map in the top panel of Figure 6, assuming a disk aspect ratio and viscosity parameter, and compared to the observed disk lifetime.

The disk lifetime for eRO-QPE2 is at least ~ 4.3 yr, from 2020 August (R. Arcodia et al. 2021) to 2024 December (this work, XMM5), and to the authors' knowledge, the most recent eruptions observed from eRO-QPE2 are from an EP observation taken on 2025 August 12 (PI: Rau) and by XMM in 2025 December (PI: Guolo), extending the lifetime to $\gtrsim 5.5$ yr. As A. Franchini et al. (2016) concluded, t_{align} is usually faster than t_{thin} , so we focused on that quantity while varying α between 0.01 and 0.1. The dashed black lines in the top panel of Figure 6 correspond to an alignment timescale equal to the lifetime constraint of ~ 5.5 yr; thus, only disk solutions to the left of this line would induce ongoing precession through to today (i.e., $>t_{\text{life}}/t_{\text{prec}} \sim 18$ precession cycles). Comparing the black line upper limits with the observed constraint in cyan shows tension with any $\alpha \gtrsim 0.01$. As the t_{life} constraint increases over time with more detections, it would move the black line lower limits to the left, shrinking the parameter space further. We verified a posteriori that changing the outer disk radius value to $200R_g$ or $800R_g$ would not change the following conclusions




Figure 6. Top panel: parameter space for the precession period of a rigidly precessing compact disk (A. Franchini et al. 2016) as a function of BH spin (χ) and the slope of the disk surface density profile (p_Σ). While the fitted $P_{\text{mod},2}$ (cyan contour) can be reproduced with realistic combinations of χ and p_Σ , the same solutions provide predictions for the alignment timescale, which appears shorter than the current QPE/disk lifetime of ~ 5.5 yr for a disk viscosity $\alpha = 0.1$, and barely compatible with $\alpha = 0.01$. Bottom panel: parameter space for an outer MBH binary (separation a_{bin} , outer black hole mass M_2 , and outer binary inclination) given the fitted $A_{\text{mod},2}$ and $P_{\text{mod},2}$ from the $O - C$ analysis (Figure 5). The EMRI primary black hole mass was fixed to $M_1 \sim 1.5 \times 10^5 M_\odot$ (estimated from the possible apsidal precession constraint), which is shown with a black vertical line. Allowed $M_2 \lesssim M_1$ solutions indicate $a_{\text{bin}} \sim 0.4$ mpc, while no solutions exist for $M_2 > M_1$.

significantly, but only the α value in tension: for a smaller R_{out} , we obtained that $\alpha \gtrsim 0.1$ are not allowed, while for the larger R_{out} , no realistic α is allowed altogether. A more conclusive statement is hindered by the many assumptions and uncertainties in the parameter values (i.e., the disk radial extent, viscosity parameter, and aspect ratio) involved in generating precession and alignment timescales (e.g., see the possible effect of the gas stream in a TDE disk producing a quasi-stable warped configuration; M. Xiang-Gruess et al. 2016).

Finally, we note that the parameter space required to reproduce the amplitude of delays (i.e., $A_{\text{mod},2}$) is also relatively small, as we estimated using the EMRI code QPE-FIT (J. Chakraborty et al. 2025a) that the relative inclination between EMRI orbit and accretion disk has to be in a relatively narrow range between $\approx 15^\circ$ and 25° to imprint time delays as large as the observed $\sim (3\text{--}4)$ ks. Hence, while there is a disk precession solution with a rigidly precessing disk as in A. Franchini et al. (2016), the parameter volume that it requires to satisfy observations is relatively small, and it will

²⁴ An estimate of the disk outer radius has been obtained in T. Wevers et al. (2025) via disk SED fitting, albeit self-consistently with a larger M_{BH} estimate. Once converted to a physical size, it would correspond to $\sim 1760 R_g$ for an $M_{\text{BH}} \sim 1.5 \times 10^5 M_\odot$, and for such extended disks, the rigid-precession framework would cease to be appropriate (A. Franchini et al. 2016), unless disk tearing/breaking occurs out to a sufficiently large radius.

shrink further as the lifetime lower limit increases with future observations. Naturally, there is a clear prediction that the longer precession would disappear due to disk alignment within the next few years.

4.2.2. Hierarchical Triple System

As in the disk precession case, we can use the observed $A_{\text{mod},2}$ and $P_{\text{mod},2}$ to explore the allowed parameter space of a hierarchical triple with an inner EMRI (M_1 and m_2) and outer MBHB (M_1 and M_2 ; e.g., see G. Miniutti et al. 2025). From the superperiod modulation amplitude, we can constrain the outer binary separation,

$$a_{\text{bin}} = \frac{c A_{\text{mod},2} (1 + q)}{\sin i_{\text{out}}}, \quad (3)$$

given a mass ratio $q = M_1/M_2$ and inclination i_{out} with respect to the observer. From the apsidal precession constraint above, the primary of the EMRI system is $M_1 \sim 1.5 \times 10^5 M_\odot$, whereas M_2 is a free parameter. Under the assumption that the longer superperiod modulation is due to an MBHB, from Kepler's third law, we can numerically constrain the allowed i_{out} and M_2 . We find that the observed $A_{\text{mod},2}$ can only be reproduced by an almost face-on binary with $M_2 \sim M_1$, or by a progressively smaller M_2 and more inclined outer orbit, down to an edge-on $M_2 \sim 0.2 M_1$. No solutions can be found for $M_2 > M_1$ that satisfy the observed $A_{\text{mod},2}$ and $P_{\text{mod},2}$. We show these constraints in the bottom panel of Figure 6 as 1σ and 3σ contours, where uncertainties are propagated from sampling the posteriors of $A_{\text{mod},2}$, $P_{\text{mod},2}$, and $M_1 = M_{\text{BH}}$ (where we note the truncated 3σ contours due to the bound $P_{\text{mod},2}$ lower-limit posterior). Interestingly, the only solutions that reproduce $A_{\text{mod},2}$ and $P_{\text{mod},2}$ point toward an outer binary separation $a_{\text{bin}} \sim 0.4$ mpc. Compared to the EMRI binary separation of $\sim 141 R_g \sim 10^{-3}$ mpc, the ratio between outer and inner binary separation is ≈ 400 , thus yielding a stable system (R. A. Mardling & S. J. Aarseth 2001). The related merger timescales of the outer binary solutions are all $\gtrsim 2$ Gyr. We further investigated the stability of this hierarchical triple system from the Kozai–Lidov mechanism (S. Naoz 2016). For the low eccentricities used in all of the above estimates (e.g., $e \sim 0.1$ for the EMRI and $e_{\text{out}} \sim 0.05$ for the outer binary), this effect is much slower than the other limiting factors present in the system (e.g., stellar ablation from collisions, disk dissipation). From the largest to the smallest M_2 in the bottom panel of Figure 6, the Kozai–Lidov timescale at the quadrupole-level approximation (S. Naoz 2016) spans from ~ 2300 to $\sim 19,500$ yr. Finally, it is interesting that the constraint of $M_2 < M_1$ to reproduce $A_{\text{mod},2}$ and $P_{\text{mod},2}$ implies that $M_{\text{tot}} < 2 M_1 \lesssim 3 \times 10^5 M_\odot$, which is compatible with the $M_{\text{BH}}-\sigma$ mass ($M_{\text{tot}} \sim 10^5 M_\odot$; T. Wevers et al. 2024) within the ~ 0.5 dex systematics.

Finally, if the system were part of a wider MBH binary, it would induce Doppler boosting on the disk emission $\Delta F_X/F_X$. Following again the reasoning already applied to GSN 069 data (G. Miniutti et al. 2025), we can write the predicted

Doppler boost as:

$$\frac{\Delta F_X}{F_X} \simeq 2\pi(3 - \alpha_X) \frac{A_{\text{mod},2}}{P_{\text{mod},2}} \sim 0.03, \quad (4)$$

evaluated at the maximum of the phase to yield the semiamplitude, and with a spectral slope $\alpha_X \simeq -9$. Hence, even if the longer modulation were due to an outer binary, it would not be detectable via phase-dependent X-ray flux variability. This is self-consistent with the fact that the observed quiescence flux of eRO-QPE2 has been remarkably constant within uncertainties for the last few years (R. Arcodia et al. 2024a; D. Pasham et al. 2024). We also confirmed this comparing the quiescence flux across the new XMM1-5 epochs, which we found to be compatible within uncertainties.

4.3. Constraints on the Type of Orbiter

Our best-fit model does not require a quadratic dissipative component such as a period increase or decrease. Thus, we estimated an upper limit on the absolute value of the period decrease, which is bound at $\dot{P} < (2 - 3) \times 10^{-6}$ (Section 3).

GW decay. First, we discuss the constraint from our 3σ upper limits on the GW orbital decay. We compute a range of $\dot{P}_{\text{GW}}$ as a function of a range of EMRI secondary mass m_2 and EMRI eccentricity e , given the fitted $M_1 \sim 1.5 \times 10^5 M_\odot$ and $P = 2.23$ hr from the $O-C$, using the P. C. Peters (1964) formula:

$$\dot{P}_{\text{GW}} = -\frac{192\pi}{5} \frac{G^{5/3}}{c^5} \left(\frac{2\pi}{P}\right)^{5/3} \frac{M_1 m_2}{(M_1 + m_2)^{1/3}} \times \frac{1 + \frac{73}{24}e^2 + \frac{37}{96}e^4}{(1 - e^2)^{7/2}}. \quad (5)$$

Figure 7 shows the predicted $\dot{P}_{\text{GW}}$ as a function of m_2 and e , with yellow lines showing our fitted 3σ upper limit constraints. Interestingly, our $O-C$ campaign rules out high eccentricities at all secondary masses including subsolar secondaries. For instance, e is only allowed to be < 0.92 at 3σ for $0.5 M_\odot$, and < 0.94 at 3σ for $0.2 M_\odot$. Effectively, this would disfavor some of the proposed QPE models involving partial disruptions from a white dwarf on an extremely eccentric orbit (A. King 2020; M. Wang et al. 2022; J.-H. Chen et al. 2023), leaving a very limited parameter space. Furthermore, orbiting intermediate-mass BHs (IMBHs; A. T.-L. Lam et al. 2025) would also be disfavored following the yellow curves in Figure 7. To give some values, $m_2 \gtrsim 250 M_\odot$ is ruled out at $e \sim 0.5$, and $m_2 \gtrsim 1130 M_\odot$ at $e \sim 0.1$. Perhaps this latter constraint is not surprising based on previous arguments disfavoring IMBH secondaries based on rates and accretion disk depletion arguments (e.g., A. Franchini et al. 2023; I. Linial & B. D. Metzger 2023; I. Linial et al. 2025; A. Mummery 2025). Nonetheless, our $O-C$ analysis provides an independent empirical constraint for specific (m_2, e) pairs.

Gas drag. We carried out the same exercise for the orbital decay due to hydrodynamical gas drag from the disk collisions (e.g., I. Linial & B. D. Metzger 2023; R. Arcodia et al. 2024a; I. Linial & E. Quataert 2024; C. Zhou et al. 2025b). For a stellar orbiter, this effect is expected to be larger than that from GW emission. Ignoring possible long-term changes due to the change of relative inclination between the EMRI orbit and the disk (R. Arcodia et al. 2024a), the orbital decay of a star of




Figure 7. Map of the absolute value of the predicted period decrease ($\dot{P}$) due to GW emission for a range of EMRI secondary mass (m_2) and eccentricity, given the fitted $M_1 \sim 1.5 \times 10^5 M_\odot$ and $P = 2.23$ hr. The fitted 3σ $\dot{P}$ upper limits from the $O - C$ campaign are shown with yellow lines, highlighting the allowed parameter space. The dashed line is for model parameters fitted within the 10th–90th percentiles of the best-fit model, while the dotted–dashed line is for free parameters within original prior bounds (see the text for details). Their difference shows the (relatively low) impact of parameter degeneracy. Our constraints disfavor high eccentricities at all masses down to subsolar orbiters, and moderate and low eccentricities for IMBH-like masses (e.g., $m_2 \gtrsim 250 M_\odot$ is ruled out at $e \sim 0.5$, and $m_2 \gtrsim 1130 M_\odot$ at $e \sim 0.1$).

mass m_* and radius R_* , which impacts a disk of surface density Σ , is approximately:

$$\begin{aligned} \dot{P}_{\text{drag}} &\approx -\frac{3\pi R_*^2 \Sigma}{m_*} \\ &\approx -2.3 \times 10^{-6} \left(\frac{m_*}{M_\odot}\right)^{-1} \left(\frac{R_*}{R_\odot}\right)^2 \left(\frac{\Sigma}{10^5 \text{ g cm}^{-2}}\right). \end{aligned} \quad (6)$$

Thus, our $\dot{P}$ upper limit may provide simultaneous constraints for the orbiting star and the accretion disk density. We can restrict our analysis to a motivated range for m_* and R_* using both the QPE lifetime (t_{life}), assuming that for a star, the dominating factors for the period decrease are the collisions themselves (P. Z. Yao et al. 2025), and the tidal radius, $r_T \sim R_* (M_{\text{BH}}/m_*)^{1/3}$, which is a lower bound to the collision radius (e.g., T. Wevers et al. 2025). Assuming a main-sequence star ($R_* \propto m_*^{0.8}$), the condition $a > r_T$ imposes $m_* \lesssim 0.72 M_\odot$, using the fitted $P = 2.24$ hr as orbital period. At the same time, there is another constraint from the QPE lifetime $t_{\text{life}} \gtrsim 5.5$ yr. This imposes a condition $t_{\text{life}} \sim t_{\text{recur}} m_*/\Delta m_* \gtrsim 5.5$ yr, where the QPE period is the recurrence time t_{recur} for the $O - C$ with all eruptions, and Δm_* is the mass lost through collisions (I. Linial et al. 2025; P. Z. Yao et al. 2025). This imposes an upper limit on the fractional mass loss $\Delta m_*/m_* \lesssim 4.6 \times 10^{-5}$.

The fractional mass loss can be expressed in proportion to the ratio between the ram pressure experienced by the star during disk crossing and the star’s mean internal pressure (I. Linial & B. D. Metzger 2023; P. Z. Yao et al. 2025). They are, respectively, $p_{\text{ram}} \approx \frac{1}{2} \rho_{\text{mid}} v_{\text{sh}}^2 \approx \rho_{\text{mid}} v_K^2$, where $\rho_{\text{mid}} = \Sigma/(2H_{\text{disk}})$ is the disk mass density, and $v_K \sim \sqrt{GM_{\text{BH}}/a}$ is the Keplerian velocity at the collision radius $\sim a$, and $p_* \sim (Gm_*^2)/(4\pi R_*^4)$. We adopt a range of possible star masses, $0.1 \lesssim (m_*/M_\odot) \lesssim 0.7$, and use the




Figure 8. Map of the absolute value of the predicted period decrease ($\dot{P}$) induced by gas drag through disk collisions onto an orbiting main-sequence star, for a range of accretion disk surface density Σ (at the collision radius) and mass of the secondary star m_* . Additional constraints are: the tidal radius ($a > r_T$, black shaded region); the QPE lifetime ($\Delta m_*/m_* \sim t_{\text{recur}}/t_{\text{life}}$, with $t_{\text{life}} \gtrsim 5.5$ yr; orange line); the QPE integrated energy (yellow) matched to the collision kinetic energy for the star as a “bullet,” for two radiative efficiency values ($\epsilon_{\text{rad}} = 1$ and 0.1 , as labeled); the equivalent constraint for colliding stellar streams (dotted–dashed yellow line, for $\epsilon_{\text{rad}} = 0.1$); the $O - C$ $\dot{P}$ upper limit (red line). There are effectively no (Σ, m_*) solutions for a “bullet” (for reasonable $\epsilon_{\text{rad}} \lesssim 0.1$), while there are some for stellar streams.

fractional mass-loss constraint from the QPE lifetime to restrict possible Σ values for any given m_* . At the collision radius $a \gg R_g$ and for accretion radiative efficiency of η and dimensionless Eddington-normalized accretion rate $\dot{m}$, we can approximate $H_{\text{disk}} \simeq \frac{3}{8\eta} R_g \dot{m}$ (J. Frank et al. 2002). For the collision regime appropriate for a short-period system like eRO-QPE2 (i.e., multiple shock; P. Z. Yao et al. 2025) and noting that $H_{\text{disk}} \approx 0.1 R_\odot$ (thus, $H_{\text{disk}} \lesssim R_*$ for m_* range being studied), from Equation (4) in P. Z. Yao et al. (2025), we obtain:

$$\begin{aligned} \Delta m_*/m_* &\sim 0.03 \left(\frac{H_{\text{disk}}}{R_*}\right) \left(\frac{p_{\text{ram}}}{p_*}\right) \\ &\sim 3 \times 10^{-5} \left(\frac{\Sigma}{10^5 \text{ g cm}^{-2}}\right) \left(\frac{m_*}{0.5 M_\odot}\right)^{2/5}. \end{aligned} \quad (7)$$

Here we assumed a main-sequence star, and given that $p_{\text{ram}} \propto 1/H_{\text{disk}}$, the only free parameters other than R_* and m_* are in the ratio M_{BH}/a through the Keplerian velocity at the collision point, which is essentially the inverse of the semimajor axis in R_g units (which we take from our fitted $O - C$ values, i.e., $a/R_g \sim 141$).

Both $\dot{P}$ and $\Delta m_*/m_*$ upper limits provide a range of allowed m_* and Σ (at the collision radius) for the collision model with a stellar orbiter. We show this parameter space in Figure 8 where the $O - C$ $\dot{P}$ upper limit is shown in red, while the $\Delta m_*/m_* \lesssim 4.6 \times 10^{-5}$ constraint is shown in orange. The latter appears more stringent, although it requires more uncertain assumptions compared to the empirical $O - C$ constraint, so it should be interpreted with caution. Nonetheless, we note that the ablation constraint on Σ will become more and more stringent as more eruptions are observed and the ratio $t_{\text{recur}}/t_{\text{life}}$ decreases.

At the same time, there is a minimum disk density $\Sigma_{\min}$ required to produce the observed integrated energy per flare via disk collisions (A. Mummery 2025). For the case of a stellar “bullet” the relevant cross section to infer the ejected disk mass is given by $2\pi R_*^2$, where the factor of 2 is due to the azimuthal flow of the disk. We note that the range of stellar masses allowed in Figure 8 corresponds to $R_* < 0.77 R_\odot$. The average observed integrated energy is $E_{\text{QPE}} \sim 1.1 \times 10^{45}$ erg (R. Arcodia et al. 2024a), and for an emission process efficiency ε_{rad} bound to be smaller than 1, we expect this to be a lower limit for the kinetic energy in the collision $E_{\text{col}} = 2\pi R_*^2 \Sigma v_K^2 > E_{\text{QPE}}$. This relation provides the lower limit shown with a yellow line in Figure 8, while we show the same relation, but with $\varepsilon_{\text{rad}} = 0.1$, with a dotted line. If the true efficiency is close to 10%, there is virtually no parameter space for colliding main-sequence “bullets.” Indeed, the efficiency is expected to be small ($\varepsilon_{\text{rad}} \ll 1$) if QPEs are powered by a stellar “bullet” crossing the disk (e.g., I. Linial et al. 2025; A. Mummery 2025). Hence, current estimates of $\Sigma_{\max}$ from the $O - C \dot{P}$ and $\Sigma_{\min}$ from the integrated energy would leave no parameter space for a main-sequence star “bullet.”

Stellar streams have been invoked to reproduce the high integrated energy of the long-period QPE sources (I. Linial et al. 2025; A. Mummery 2025; P. Z. Yao et al. 2025). Before our $\Sigma_{\max}$ constraint from the $O - C \dot{P}$, there was no reason to exclude a canonical “bullet” collision in eRO-QPE2 based on energetics. Our Figure 8 now suggests otherwise, and we investigate whether stellar streams would relax the tension. The model in I. Linial et al. (2025) indicates that for a system like eRO-QPE2, the stream would penetrate the disk with the effect of increasing the amount of disk mass contributing to the emission, thus significantly lowering $\Sigma_{\min}$. As a stream is, by definition, much more extended than the star, the gas drag $\dot{P}$ map would be equivalent to that shown in Figure 8, but the integrated-energy constraint would be far less stringent. We show the related curve with a dotted-dashed yellow line for $\varepsilon_{\text{rad}} = 0.1$, noting that the radiative efficiency for streams may be larger than in the “bullet” case (I. Linial et al. 2025; A. Mummery 2025). The lower $\Sigma_{\min}$ is due to a larger disk area interacting with the orbiter, increased to $\sim 4r_H^2 \tilde{l}$, where $r_H = a \left(\frac{m_*}{3M_{\text{BH}}} \right)^{1/3}$ is the Hill’s radius and for a stream elongation $\tilde{l} \sim 28$ appropriate for a system like eRO-QPE2 (I. Linial et al. 2025). Hence, the allowed parameter space is now somewhat larger, although we note that it is still limited to $\Sigma \approx (0.5-1) \times 10^5 \text{ g cm}^{-2}$, but neither much larger nor much smaller, and for $m_*/M_\odot \approx (0.4-0.7)$.

We estimated whether realistic accretion disks in a system like eRO-QPE2 can satisfy this requirement assuming a surface density profile $\Sigma(R, t) = \Sigma_0(t) \left(\frac{R}{R_g} \right)^{-p_\Sigma}$. We assumed the disk was created by a TDE of a star with mass m_*^{TDE} and circularization radius twice the tidal radius $2r_T^{\text{TDE}}$. We follow this approach given the mounting evidence connecting QPE accretion flows with previous TDEs (E. Quintin et al. 2023; S. Bykov et al. 2024; M. Nicholl et al. 2024; J. Chakraborty et al. 2025b; M. Guolo et al. 2025a; L. Hernández-García et al. 2025). While, to date, no evidence of a precursor TDE has been found for eRO-QPE2 in particular, we note that adopting active galactic nucleus (AGN)-like steady-state disks would produce higher Σ values compared to an expanding TDE disk with a fixed mass budget. Thus, any parameter space excluded by the TDE disk

profile would be ruled out for steady-state disks too. We constrain the surface density of a TDE disk at the collision radius $R = a = 141 R_g$ by setting mass and angular momentum conservation, $M_{\text{disk}} \leq f^{\text{TDE}} m_*^{\text{TDE}}$ and $J_{\text{disk}} = f^{\text{TDE}} J_*^{\text{TDE}}$, and keeping the smallest of the two Σ_0 estimates, namely between:

$$\Sigma_0(t) = \frac{M_{\text{disk},0} (2 - p_\Sigma)}{2\pi R_g^2} \times \left[\left(\frac{R_{\text{out}}(t)}{R_g} \right)^{2-p_\Sigma} - \left(\frac{R_{\text{in}}}{R_g} \right)^{2-p_\Sigma} \right]^{-1} \quad (8)$$

(which follows from mass conservation), and

$$\Sigma_0(t) = \frac{M_{\text{disk},0} (5 - 2p_\Sigma)}{4\pi R_g^2} \left(\frac{2r_T^{\text{TDE}}}{R_g} \right)^{1/2} \times \left[\left(\frac{R_{\text{out}}(t)}{R_g} \right)^{\frac{5}{2}-p_\Sigma} - \left(\frac{R_{\text{in}}}{R_g} \right)^{\frac{5}{2}-p_\Sigma} \right]^{-1}, \quad (9)$$

which follows from angular momentum conservation. We evaluated these profiles for a range of $f^{\text{TDE}} = (0.01, 0.5)$, $m_*^{\text{TDE}}/M_\odot = (0.5, 4)$ (i.e., combined spanning $M_{\text{disk},0}/M_\odot = 0.005 - 2$), spin $\chi = (-0.998, 0.998)$, which defines the inner radius R_{in} , slopes $p_\Sigma = (0, 1.5)$, and $R_{\text{out}}/R_g = (200, 1000)$. Disk solutions in the desired range $\Sigma \sim (0.5-1) \times 10^5 \text{ g cm}^{-2}$ exist for most p_Σ and χ , but with low $M_{\text{disk},0}$ and/or high R_{out} . In particular, around $R_{\text{out}} \approx 300 R_g$, the allowed $M_{\text{disk},0}$ range is $f_{\text{mass}} m_*^{\text{TDE}}/M_\odot \lesssim 0.006$, while around $R_{\text{out}} \approx 900 R_g$, the 1σ range is $0.01 \lesssim f_{\text{mass}} m_*^{\text{TDE}}/M_\odot \lesssim 0.07$.

As an independent constraint on disk values, for $R_{\text{out}} \gg a$, the local accretion rate equals $\dot{m} \approx L_Q^{\text{bol}}/\eta c^2$, as quasi steady-state is established in the inner disk, where $L_Q^{\text{bol}} \approx 10^{43} \text{ erg s}^{-1}$ is the bolometric quiescent luminosity of eRO-QPE2 (R. Arcodia et al. 2024a; T. Wevers et al. 2025). We can therefore estimate the local disk density as $\Sigma(a) = L_Q/(3\pi\eta c^2 \nu)$, for an α -disk with viscosity $\nu = \alpha \sqrt{GM_{\text{BH}} a} (H_{\text{disk}}/a)^2$. Assuming the fiducial values used throughout this Section for L_Q , M_{BH} , H_{disk}/a , we obtain:

$$\Sigma(a) \sim 1.2 \times 10^6 \text{ g cm}^{-2} \left(\frac{\alpha}{0.1} \right)^{-1} \left(\frac{H_{\text{disk}}/a}{3.4 \times 10^{-3}} \right)^{-2}; \quad (10)$$

thus, we have a disk density ~ 10 times higher than the constraint from gas drag, namely $\Sigma \sim (0.5-1) \times 10^5 \text{ g cm}^{-2}$ (Figure 8). In turn, for self-consistency, this may provide some constraints on α and H_{disk} .

Stripped-envelope stars. We remind the reader that the range of disk density solutions shown in Figure 8 only exists for stellar streams in the case of main-sequence stars. This is in part due to the gas drag $\dot{P}$ upper limit being too low compared to the $\Sigma_{\min}$ required by “bullet” collisions. Thus, we investigated whether there exists a larger parameter space for the “bullet” in the case of a different stellar mass-radius relation, in particular with an orbiter with comparably smaller R_* . Interestingly, this would also relax the tension between collision distances inferred too close to the tidal radius or mass-transfer radius (or even within) for some QPE sources (e.g., A. Franchini et al. 2023; S. A. Dodd et al. 2025; W. Guo & R.-F. Shen 2025; I. Linial et al. 2025). A putative stellar-




Figure 9. Same as Figure 8, but for stripped-envelope stars, namely for comparatively smaller R_* .

EMRI with $R_* \lesssim R_\odot$ for $m_* \gtrsim M_\odot$ is akin to stripped-envelope stars during their phase of core helium fusion (e.g., Y. Göteborg et al. 2018). We fitted a functional form for the mass–radius relation to the models reported in Y. Göteborg et al. (2018) and obtained $\log\left(\frac{R_*}{R_\odot}\right) = 0.6 \log\left(\frac{m_*}{M_\odot}\right) - 0.5$. With this, we produced an analog of Figure 8 for stripped-envelope stars, for a range $m_*/M_\odot = (0.3-3)$ and $R_*/R_\odot = (0.15-0.70)$ (see Figure 9). Orbiters with lower R_*/m_* have a range of solutions even for the “bullet” scenario for $\epsilon_{\text{rad}} \lesssim 0.1$, with allowed disk densities $\Sigma \sim (0.5-1) \times 10^6 \text{ g cm}^{-2}$. This is due to a $\Sigma_{\text{max}} \propto R_*^{-3}$ dependency in the ablation (“QPE lifetime”) term from Equation (7), and $\Sigma_{\text{min}} \propto R_*^{-2}$ in the “QPE energy” lower limit. In particular, around $R_{\text{out}} \approx 300 R_g$, the allowed $1\sigma M_{\text{disk},0}$ is $0.02 \lesssim f_{\text{mass}} m_*^{\text{TDE}}/M_\odot \lesssim 0.03$, while around $R_{\text{out}} \approx 900 R_g$, this range is $0.1 \lesssim f_{\text{mass}} m_*^{\text{TDE}}/M_\odot \lesssim 0.3$. Naturally, similar to the main-sequence case, stellar streams from a stripped-envelope star would yield an even lower Σ_{min} constraint (e.g., see the dotted–dashed line in Figure 9 for $\epsilon_{\text{rad}} = 0.1$), allowing for even smaller $M_{\text{disk},0}$. However, as stripped-envelope stars are much rarer, we note that if a similar requirement were also to be found for other QPE sources, there would be tension with the QPE rates (R. Arcodia et al. 2024c), unless it is the more abundant main-sequence stars that evolve to R_*/m_* values akin to stripped-envelope stars through the disk collisions themselves.

Stellar BH orbiters. The $\dot{P}$ upper limit from the $O - C$ analysis is formally consistent with a stellar BH orbiter (A. Franchini et al. 2023). The Σ_{min} requirement would imply that the influence radius of the BH orbiter is $R_{\text{Bondi},*} \approx R_*$, so that the integrated energy is reproduced, and in fact, A. Franchini et al. (2023) inferred that $m_* = 40-100 M_\odot$ and low enough relative velocities (i.e., for grazing prograde collisions) may provide the required cross section. This means that for the sake of adapting the gas drag equation (Equation (6)) to a BH orbiter with $R_{\text{Bondi},*} \approx R_*$, the only relevant dependence is the orbiter’s mass $\dot{P}_{\text{drag}} \propto \Sigma_{\text{max}}/m_*$, thus allowing the maximum disk density at the collision point to be $\sim 40-100\times$ larger than the main-sequence star case in Figure 8. A BH orbiter at these masses is allowed at moderate and low eccentricity (e.g., $e \lesssim 0.6$; see Figure 7) for a system like eRO-QPE2. Hence, a stellar BH solution exists for a

system like eRO-QPE2. In general, while EMRI rates are highly uncertain, recent work by C. M. Allievi et al. (2026) has shown that if grazing angles between the secondary stellar-mass BH and the disk are required (e.g., see A. Franchini et al. 2023), the stellar-mass BH rates are 3 orders of magnitude lower than the observed QPE rates (R. Arcodia et al. 2024c). In addition, a cross section defined by their Bondi radii is unable to reproduce the integrated energy of the longer-period QPE sources (e.g., I. Linial & B. D. Metzger 2023; I. Linial et al. 2025; A. Mummery 2025). Thus, if QPEs can be triggered by a variety of orbiters, stellar-mass BHs could still be invoked for short-period (and low integrated energy) sources like eRO-QPE2, or if the inclination constraint suppressing the rates is relaxed (C. M. Allievi et al. 2026): for instance, if orbiter–disk interactions from past TDEs or AGN phases and two-body scattering play a significant role in preferentially dragging the BH-EMRI toward low inclinations with the primary BH’s equatorial plane (e.g., N. Jiang & Z. Pan 2025; Y. Zeng & Z. Pan 2026). However, relaxing the rate suppression does not affect the integrated-energy problem for the long-period QPEs (e.g., A. Mummery 2025). Finally, the low relative velocity between the accretion flow and the BH orbiter required to satisfy the collision cross-sectional constraint poses an additional challenge, in that shock velocities $\ll 0.1c$ seem to produce flare temperatures that are below the observed $\sim 100 \text{ eV}$, as shown in I. Vurm et al. (2025).

5. Fitting with EMRI Trajectory Models

Section 4 provides constraints for EMRI models from the empirical $O - C$ analysis. To test EMRI trajectory models directly, we used two different codes, namely the timing analysis code QPE-FIT²⁵ v. 0.1.11 (J. Chakraborty et al. 2025a), while the other is based on A. Franchini et al. (2023), and hereafter named FB23 (but we note it will be published separately by A. Motta et al. 2026, in preparation). Both codes perform Bayesian parameter inference for EMRI and SMBH properties using QPE timings under the assumption that they are due to orbiter–disk collisions around an SMBH. The main difference is that while in QPE-FIT disk precession is disjoint from MBH parameters, in FB23 it is linked to the BH spin. More details on these models and our fits are reported in Appendix C.

In summary, we found no reliable solution to the XMM1-4 and XMM1-5 eRO-QPE2 data. While the EMRI code runs did converge to some solutions, the overall inability of modeling both short-term and long-term delays in QPE arrival times, and the high sensitivity of fitted values to both the number of data points used (e.g., censoring some epochs) and the prior volume, led us to consider all solutions as unreliable. We report more details on priors and fitted parameters in Appendix C, and report here only the useful lessons learned for QPE modeling. In particular, we note that obtaining good residuals and reasonable parameters with converged posteriors is not a reliable confirmation that a given adopted EMRI model setup is correct. Thus, we advise the community to use EMRI codes together with more empirical approaches like the $O - C$ analysis performed here, when possible.

²⁵ Publicly available at <https://github.com/joheenc/QPE-FIT>.

5.1. Two Observed Events per Orbit Are Disfavored

Section 3 reports analysis of $O - C$ data where all eruptions are analyzed; thus, P is traced by t_{recur} . In parallel, we performed analogous analysis and tests for the case of odd and even eruptions separately, where P is traced by $2t_{\text{recur}}$ (see Appendix B.4). This analysis is more appropriate to compare with models predicting two observable events per orbit, such as most current disk collision models for QPEs (e.g., A. Franchini et al. 2023; I. Linial & B. D. Metzger 2023; C. Zhou et al. 2024). As discussed extensively in G. Miniutti et al. (2025) for GSN 069, if QPEs are observed for both ingresses and egresses in the disk, a key prediction of apsidal precession is that odd and even data would be observed in antiphase on this timescale. However, from our $O - C$ analysis, we found that for both superperiod modulations (fitted at $\sim 23P$ and $\sim 500P$ in this case), odd and even eruptions are in phase (e.g., Figures B8 and B9 in Appendix B.4). This is a fundamental inconsistency with current collision models with two observable events per orbit, strongly disfavoring such scenarios. The only way to overcome this tension would be to have an additional faster modulation with comparable or larger amplitude, that correlates odd and even data at the apsidal precession timescale. However, while this was still a possibility for GSN 069 data (G. Miniutti et al. 2025), for our eRO-QPE2 dataset it is not possible to conceal a modulation on the timescales between P and $P_{\text{aps}} \sim 23P$ (Figure B8). As a matter of fact, it would need to be of the order of $\sim 2\text{--}5P$ in order to be sufficiently faster than apsidal and of the same amplitude ($\gtrsim 300$ s), which not only is unlikely to be a physical precession term, but is also not observed in the individual XMM epochs. For instance, XMM1 is the observation with the most consecutive eruptions (eight; see the bottom panel of Figure 3), and it may show, at best, a hint of residuals at the ~ 150 s level.

However, given that a solution with EMRI trajectory codes with two observed crossings per orbit had been found before for eRO-QPE2 (e.g., C. Zhou et al. 2025b; see more details in Section 6), we tried to fit our data with the EMRI trajectory models described above, assuming the standard two observable events per orbit. While both QPE-FIT and FB23 models find a solution with good residuals and reasonable parameters, after deeper analysis, they were found to be merely a numerical artifact. The most useful diagnostic was to generate the mock $O - C$ from the best-fit posteriors of both QPE-FIT and FB23 models and to overlay them to the $O - C$ data. Quite worryingly, the model solutions fitted the higher signal-to-noise XMM data at the crossing of the anticorrelated odd and even predictions from the models (see Appendix C and Figure C1). This corresponds to a precession phase in which the semimajor axis is aligned to the disk plane (still with an inclined orbital plane), and collisions happen at pericenter and apocenter. Clearly, it is unrealistic that all of the XMM observations, and only the XMM observations, were taken during this relatively rare orbital phase. In the attempt to quantify the absurdity of this result, we shifted the best-fit model keeping the XMM observations fixed to quantify the probability of obtaining this alignment by chance. Effectively, we computed the difference between odd and even mock $O - C$ curves from QPE-FIT (top panel of Figure C1) and slid this array by 62 integers toward lower N_{QPE} (i.e., up to ~ 6 days). We obtained that in none of these 62 configurations, namely in no other configuration other than the one we have, we would obtain XMM1-3 datasets at epochs with such a

small difference between odd and even curves (i.e., at the crossings). The result is the same changing any observation between XMM1, XMM2, and XMM3 by $\sim \pm 1$ day.

Given that these unphysical solutions comes in addition to the absence of the predicted anticorrelation between odd and even eruptions, we conclude that QPE EMRI models in which two disk crossings occur, and both are electromagnetically observed, are disfavored.

5.2. One Observed Event per Orbit: No Reliable Solution Is Found

Based on our $O - C$ analysis, in order for QPE EMRI models to be correct they would require that only the ingress or egress induces an observed event. Naturally, it is also possible that one event per orbit occurs in the first place, for instance for elevated orbital eccentricity or disk eccentricity. However, for a regular source like eRO-QPE2 in terms of recurrence time and eruptions amplitude, we do not expect a large diversity in collision distances or high eccentricity. Thus, we adopt the former scenario for further tests, also considering that most current disk collision simulations show rather asymmetric forward and backwards ejecta (X. Huang et al. 2025; T. Jankovič et al. 2026). We adopted QPE-FIT to carry these out, the fastest of the two algorithms used in this work. We report more details on the fitted parameters and comparisons with the observed $O - C$ delays in Appendix C and again report here the main lessons learned. We attempted several fits varying the number of data points (e.g., using XMM1-4 data only, or extending out to XMM5) and the prior volume (see Table C1 in Appendix C), and, worryingly, obtained inconsistent solutions. This remains true after exploring different volume sampling algorithms within Ultraneest. Hence, for caution, we do not consider any of these solutions as conclusive and robust. We suspect that this is due to the application of a model with a large, multi-dimensional, and highly correlated parameter volume and relatively narrow likelihood peaks. Even a relatively transformational campaign such the XMM1-4 dataset presented in this work is too sparse in comparison, let alone the much sparser archival data of other QPE sources.

6. Comparisons with the Previous Literature on eRO-QPE2

eRO-QPE2 was observed five times after discovery by XMM-Newton, and its long-term spectral and timing properties have been reported in R. Arcodia et al. (2024a) and D. Pasham et al. (2024). The flux of the quiescent accretion disk and the eruptions peak has been found to be consistent within uncertainties across the epochs. Performing the same analysis on the XMM1-5 epochs, we confirm this trend with our new data.

We compared the observed recurrence time t_{recur} from our new data with the literature values in Figure 10. In this work, we used the same method to estimate QPE peak times as in R. Arcodia et al. 2024a, namely an asymmetric parametric model (see Appendix A). In addition, we added two archival observations (IDs 0932590101 and 0932590201, taken on 2023 December 20 and 2024 February 4; PI: Wevers), presented in D. Pasham et al. (2024), but reprocessed and analyzed here following the same procedure as for the rest of the data. The individual consecutive recurrence times are shown in gray in




Figure 10. Long-term evolution of the recurrence time (t_{recur}), here a proxy of the period of eRO-QPE2, relative in time to the XMM1-4 campaign (green data). Archival 2020–2023 data are shown in yellow (R. Arcodia et al. 2024a; D. Pasham et al. 2024), and EP and XMM5 data taken after the XMM1-4 campaign (in addition to recent observations from M. Guolo et al. 2026, in preparation) are shown in brown. Filled colors correspond to the per-epoch mean and standard deviation of t_{recur} , while individual measurements are shown in gray underneath. The green contour shows the 3σ uncertainty predictions, extrapolated forward and backward in time, from the best-fit $O - C$ model fitted against the XMM1-4 data alone. The gray contours show the $\text{lin}+\text{pdot}+\text{mod}$ fit, discarded based on the significant underestimate of the archival data t_{recur} . The right panel is a zoomed-in region corresponding to the dashed gray box of the main panel.

Figure 10, with the average and standard deviation shown in yellow for archival data, green for the XMM1-4 data, and brown for ancillary data taken after the XMM1-4 campaign and reported here for the first time. The right panel of Figure 10 is a zoomed-in region of the newly reported observations. In dark green, we show the extrapolated $O - C$ best-fit model (see Section 3), and in light green, we show the effect of adding a $\dot{P}$ upper limit term ($\dot{P} < 2 \times 10^{-6}$). Notably, the span of the best-fit $O - C$ model fitted at the XMM1-4 epoch (green contours) is consistent within uncertainties with the QPE recurrence time observed after the second epoch (i.e., from 2022 June onward; R. Arcodia et al. 2024a; D. Pasham et al. 2024). The recurrence time in 2020 (first epoch in Figure 10) is confirmed to be somewhat unique, in that t_{recur} was comparably higher than all other epochs, and it was the only epoch with a significant long-short behavior (R. Arcodia et al. 2021, 2024a). Clearly, the $O - C$ solution found at the XMM1-4 epoch (green data and model) does not grasp the full 2020–2025 evolution of eRO-QPE2, as it is expected given that the $O - C$ did not include those data. While the $\dot{P}$ upper limit model relaxes the tension, it is not sufficient to explain the first archival data of eRO-QPE2. A possible explanation for the long-term behavior of eRO-QPE2 on the timescale of years may still be a moderate (consistent with our $\dot{P} < 2 \times 10^{-6}$ upper limit at the 2024 July epoch) dissipative term that is epoch-dependent, for instance, the inclination-dependent gas drag period decrease suggested in R. Arcodia et al. (2024a). A slightly higher $\dot{P}$ term in the past compared to the 2024 July upper limit would qualitatively reproduce the long-term behavior of eRO-QPE2; however, this requires further modeling, which is beyond the scope of this work. In particular, we note that current QPE disk collision theories do not address the large intra-observation variation in recurrence time (highlighted by the spread of the gray points), which may be due to the currently unmodeled delay between the time of disk crossing and the onset of the QPE emission.

Finally, we compare our analysis in this work with previous attempts inferring orbital parameters from eRO-QPE2 timing data. In J. Xian et al. (2025), the authors focused on reproducing the possible t_{recur} decay of eRO-QPE2 (e.g., Figure 10), assuming it is a detection of a period decay. However, comparing average t_{recur} values inferred at different epochs is not a precise trace of the orbital period (e.g., see the discussion in R. Arcodia et al. 2024a;

G. Miniutti et al. 2025) and, in fact, our $O - C$ analysis does not detect a negative $\dot{P}$, excluding it at an absolute value of $\lesssim 2 \times 10^{-6} \text{ s s}^{-1}$. Thus, the solutions found by J. Xian et al. (2025) for eRO-QPE2 are not applicable here. Recently, C. Zhou et al. (2025b) fitted the archival 2020–2023 observations with their EMRI trajectory code by modeling both ingresses into and egresses out of the disk as observed eruptions, and reported solutions with semimajor axis $\approx 500 R_g$ and, consequently, black hole masses around $\approx 10^{4.7-4.8} M_\odot$. There are several reasons why, similarly to our attempts here (Section 5), we do not consider this numerical solution to be fully reliable. First, it was obtained fitting two observable events per orbit, which is a setup strongly disfavored by the $O - C$ odd/even correlation and the data-model comparison shown in Figure C1 obtained with two independent EMRI models. Furthermore, the posterior of the semimajor axis spans $\sim 100\text{--}900 R_g$ within 1σ in C. Zhou et al. (2025b), thus suggesting the difficulty of finding a single solution for the entire period 2020–2023. This is perhaps unsurprising given that, in this work, two independent EMRI models failed at finding a reliable solution on a much better sampled dataset (Appendix C). Finally, in C. Zhou et al. (2025b), a random noise component is added as a free parameter to account for the unknown delay between disk crossing and emission. While this component is likely present and surely required, as discussed above in this work as well, in their fit it is allowed to reach 2 ks in eRO-QPE2, and fitted around ~ 200 s (reaching 300 s at 1σ) and even ~ 700 s in some model runs. However, the $O - C$ delays at these amplitude and short timescales were found to be inconsistent with being dominated by a random distribution, and instead more compatible with a deterministic component (Appendix B.3.1 and Figure B6).

7. Summary and Conclusions

In this work, we reported our timing analysis of eRO-QPE2 with a multimission X-ray campaign. We performed $O - C$ analysis (Section 3), where the only model assumption comes from building the noise model. We found that for models like accretion disk instabilities and red-noise processes, the $O - C$ delays of eRO-QPE2 are reproduced entirely by a QPE recurrence time following a DRW, although the fitted

parameters are highly uncertain and lack predictive power (Appendix B). Instead, assuming an underlying periodic clock, and testing the $O - C$ delays against white noise, we found that eRO-QPE2 data are well described by a periodic source and a sum of two hierarchical superperiod modulations.

The latter scenario provides numerous constraints for current QPE models. For QPEs as colliding EMRIs, $O - C$ data of eRO-QPE2 are inconsistent with, and disfavor, two observable events per orbit. Our analysis shows that odd and even eruptions are in phase at all of the epochs spanned by our baseline (Figures 4, B8, and B9). This result confirms what was found for GSN 069 by G. Miniutti et al. (2025). Furthermore, no reliable solution could be found for models with two eruptions per orbit, even with robust EMRI trajectory models (Section 5 and Figure C1). Hence, EMRI-disk collision models are only viable if only one event per orbit is observed, be it whether only one crossing per orbit occurs, or whether only one of the two crossings is observed. Interpreting the $O - C$ results in this light obtains the following constraints (Section 4):

1. The fitted period is $P \sim 8064 \text{ s} \sim 2.24 \text{ hr}$. The shorter superperiod modulation has a period of ~ 4.4 days (or $\sim 47 P$) and an amplitude of $\sim 289 \text{ s}$ (see Figure B2). For the longer modulation, the most conservative estimate would be to consider it a lower limit, namely $\gtrsim 918 P$ at 1σ ($\gtrsim 600 P$ at 3σ). However, by adding more data to the $O - C$, it appears constrained at ≈ 95 days (or $\approx 1014 P$), with an amplitude $\sim 3.5 \text{ ks}$.
2. There is no evidence of a strong dissipative term, such as a period increase or decrease, and we obtain a 3σ limit on the absolute value of a period decrease at $|\dot{P}| \lesssim 2 \times 10^{-6} \text{ s s}^{-1}$. Comparing this limit with the predicted orbital decay due to GW emission rules out IMBH orbiters down to moderate and low eccentricity (we rule out $\gtrsim 250 M_\odot$ at $e \sim 0.5$, and $\gtrsim 1130 M_\odot$ at $e \sim 0.1$), and high-eccentricity orbiters at all masses including subsolar (Figure 7). Thus, our campaign excludes most of the required mass and eccentricity parameter space for models with WDs on a very high-eccentricity orbit (e.g., A. King 2020; M. Wang et al. 2022; J.-H. Chen et al. 2023).
3. In the orbiter–disk collisions framework, the shorter superperiod modulation over ~ 4.4 days can be interpreted as apsidal precession, which, through Equation (1) and (2), implies semimajor axis $a \sim 140 R_g$ and $M_{\text{BH}} \sim 1.5 \times 10^5 M_\odot$, and the amplitude of the modulation can be reproduced for $e \sim 0.1$.
4. The longer modulation over ≈ 95 days cannot be attributed to EMRI nodal precession, as the observed amplitude of (3–4) ks is far larger than the maximum time delays from EMRI nodal for the most favorable inclination (edge-on disk with low EMRI orbit inclination), which are of the order of $\sim a/c \approx 100 \text{ s}$.
5. While a possible solution for the longer modulation exists for disk precession (for a given spin, disk surface density profile, radial extent, aspect ratio, and viscosity parameter), the same solutions may induce alignment of the disk (e.g., see A. Franchini et al. 2016) on a timescale shorter than the current disk lifetime lower limit of $\sim 5.5 \text{ yr}$ (see the top panel of Figure 6). Thus, the parameter space for disk precession is narrow, and it will only shrink further over time if more observations will increase the QPE lifetime lower limit.

6. A possible solution for the longer modulation exists with a hierarchical triple system, as both amplitude and period are reproduced by an outer binary at a separation of $\sim 0.4 \text{ mpc}$ (thus, a factor ~ 400 larger than the inner EMRI separation, implying a stable system) at any inclination and tertiary masses M_2 within a factor $\sim (0.1-1)$ of the EMRI primary ($M_{\text{BH}} \sim 1.5 \times 10^5$). No solutions exist with $M_2 > M_1$ (see the bottom panel of Figure 6).
7. We investigated the parameter space allowed for the EMRI secondary mass m_* and disk density Σ , in the case of a stellar orbiter and orbital decay due to hydrodynamical gas drag (e.g., I. Linial & B. D. Metzger 2023; P. Z. Yao et al. 2025). As the $O - C \dot{P}$ implies a maximum disk density, and the QPE integrated energy implies a minimum disk density, the parameter space for disk collisions depends on the cross section of the object. Our constraints (see Figures 8 and 9) indicate that main-sequence stars colliding as “bullets” are disfavored. Instead, main-sequence stars in which the colliding area is determined by stellar streams, and/or stars with smaller radii compared to main sequence (e.g., akin to stripped-envelope stars) are allowed. A stellar-mass BH orbiter is also still viable for a system like eRO-QPE2.
8. We also made use of EMRI trajectory models as a consistency check to vet our $O - C$ interpretation with one event per orbit (Section 5), although no fully reliable solution was found. While all model setups converged to a solution, none were able to recover both short-term and long-term timing data. Perhaps more worryingly, we found that best-fit solutions changed based on both the number of data points used (e.g., censoring some epochs) and the prior volume, which led us to consider any of them unreliable (see Appendix C for more details). We suspect that this is due to the application of a model with a large, multidimensional, and highly correlated parameter volume and relatively narrow likelihood peaks. Hence, we advise one to use EMRI codes together with more empirical approaches like the $O - C$ analysis performed here, when possible, and encourage caution in the interpretation of any solution. It is also likely that additional physics is required in the current EMRI-disk collision models.

The results found in this work for eRO-QPE2 can be scaled to the full QPE population with some educated assumptions. For instance, short-period systems are considered less pathological than the longer-period QPEs with higher integrated energy per burst (e.g., A. Mummery 2025; P. Z. Yao et al. 2025). In the latter case, it will be crucial to understand the nature of the emission and the delay between disk crossing and the start or peak of the X-ray emission. For the QPE sources with more complex timing, we note that if the superperiod modulations are not hierarchical, but have comparable periods and/or amplitudes, the expected timing behavior would be more chaotic (see the simulated plots in J. Chakraborty et al. 2025a). Moreover, while a mechanism like apsidal precession is always expected, its amplitude is a function of a and e ; thus, it might not be always detectable. In eRO-QPE2, the longer modulation is the most uncertain component, although from current data, it is currently impossible to reconcile with EMRI nodal precession, hard to reconcile with disk precession, while a hierarchical triple

system is possible. Clearly, not all QPE sources have to share the same configuration; thus, a longer superperiod does not have to be ubiquitous. The $\dot{P}$ upper limit we obtained with this $O - C$ analysis is currently the only reliable data-driven orbital decay constraint for QPE sources. A result that can be generalized is the exclusion of high-eccentricity orbiters at all masses (Figure 7), which formally disfavors partial disruptions of highly eccentric WDs. Similarly, the constraints shown in Figure 8 for stellar orbiters, obtained from the tidal radius and the QPE lifetime bounds, can be estimated for all QPE sources. Hence, this work thus offers some legacy constraints from all of the current QPE models, which are required to be modified given the current inconsistencies. Naturally, this work also opens the way for entirely new paradigms that can be tested with this multimission dataset.

Acknowledgments

R.A. is deeply grateful to Z. Arzoumanian, K. Gendreau, and B. Cenko for the positive response to the NICER and Swift/XRT ToOs that significantly enriched this work, and the flexibility in maneuvering the complex multimission scheduling.

R.A. was supported by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51499.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. G.M. acknowledges support from grant No. PID2023-147338NB-C21 funded by Spanish MICIU/AEI/10.13039/501100011033 and ERDF/EU. This research benefited from interactions at workshops funded by the Gordon and Betty Moore Foundation through grant GBMF5076. A.F. acknowledges financial support from the Unione europea—Next Generation EU, Missione 4 Componente 1CUP G43C24002290001. A.M. acknowledges support from the Ambrose Monell Foundation, the W.M. Keck Foundation and the John N. Bahcall Fellowship Fund at the Institute for Advanced Study. I.L. is supported by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51581.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555, and acknowledges support from a Rothschild Fellowship and The Gruber Foundation, as well as Simons Investigator grant 827103. M.G. is funded by Spanish MICIU/AEI/10.13039/501100011033 and ERDF/EU grant PID2023-147338NB-C21. G.P. acknowledges financial support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program HotMilk (grant agreement No. 865637) and from the Framework per l’Attrazione e il Rafforzamento delle Eccellenze (FARE) per la ricerca in Italia (R20L5S39T9). E.Q. acknowledges support through the European Space Agency (ESA) Research Fellowship Programme in Space Science. A.S. acknowledges the financial support provided under the European Union’s H2020 ERC Consolidator Grant “Binary Massive Black Hole Astrophysics” (B Massive, Grant Agreement: 818691) and Advanced Grant “PINGU” (Grant Agreement: 101142079). M.B. acknowledges support from the Italian Ministry for Universities and Research (MUR) program “Dipartimenti di Eccellenza 2023-2027,” within the framework of the activities of the Centro Bicocca di Cosmologia Quantitativa (BiCoQ).

Software: astropy (Astropy Collaboration et al. 2018), swifttools.ukssdc for light curves (P. A. Evans et al. 2007, 2009).

Appendix A Details on X-Ray Data Processing and Analysis

Here, we report mode details on the X-ray processing and analysis of the multimission dataset used in this work.

A.1. X-Ray Data Processing

XMM data were processed using SAS v. 20.0.0 and HEASoft v. 6.31. Products were extracted from source (and background) regions selecting a circle of $40''$ centered on eRO-QPE2 (in a nearby source-free $40''$ circle). For all epochs, XMM light curves were extracted in the 0.2–2.0 keV range. All good time intervals were kept, and since eruptions are soft and background flares are hard, eruptions can still be confidently detected in 0.2–2.0 keV range (see also R. Arcodia et al. 2024a). We verified a posteriori that the arrival time of the contaminated eruptions is not offset in the presence of background flares, but only more uncertain, by comparing the 0.2–2.0 keV with an even softer 0.2–1.0 keV extraction. In any case, flares affect only part of XMM1 and XMM3, and out of 21 eruptions detected by XMM in total, only five have occurred during a background flare, and only one is significantly affected. For XMM5, the pn camera only recorded half of the first eruption in the exposure, and we used MOS cameras to estimate its peak time.

For XRT, we processed data with the data analysis tool (P. A. Evans et al. 2009), extracting a light curve between 0.3 and 2.0 keV and initially adopting a per-snapshot binning (more details on the analysis below). NICER data were processed using HEASoft v6.33 and NICERDAS v12. Data products were reduced using nicerl2 with no overshoot or undershoot screening, including both orbit day- and nighttime data, and autoscreening disabled. We then grouped the event files into 200 s good-time intervals (GTIs; using `nimake-time` and `niextract-events`), barycenter-corrected the events using `barycorr`, and extracted spectra for each GTI using `nicerl3-spect` to follow a time-resolved spectroscopy approach for reliable source light-curve extraction. Each spectrum was first fit in a broadband range (0.25–10.0 keV for orbit night/0.38–10.0 keV for orbit day) with the SCORPEON semiempirical spectral background model, leaving solar wind charge exchange line normalizations left free to vary. We then added an absorbed blackbody model (`tbabs×ztbabs×bbody`) to assess whether the addition of a source component significantly improved the fit. We consider a source “detection” any GTI in which the addition of the source improved the background-only fit by $\Delta C\text{-stat} \geq 25$, a cutoff determined empirically as a suitable boundary between robust detections and background noise. We refer the reader to J. Chakraborty et al. (2024) for further details. For EP, data were processed using standard recipes with the FXT Data Analysis Software Package (FXTDAS) v1.10. Source light curves were extracted from barycenter-corrected event files between 0.2 and 5 keV with the Heasoft XSELECT task, using circular regions with $60''$ radii. The light curves from the two telescope modules (FXT-A and FXT-B) were then binned with 200 s bins and summed.

A.2. QPE Arrival Times and Uncertainties

Here, we report more details on the estimate of QPE peak times, and the impact of some of the assumptions and analysis techniques on the final results of this work.

A.2.1. The QPE Flare Parametric Model

We note that adopting a given parametric model shape over another only impacts arrival times as an offset term; thus, it has little impact on our conclusions since both $O - C$ techniques and EMRI trajectory codes express peak times relative to an arbitrary start time. To test this quantitatively, we fitted all eruptions also with a symmetric Gaussian model. We found that the average offsets (with their average uncertainty) between double-exponential and Gaussian parametric models are compatible across various epochs (-151 ± 87 s, -168 ± 47 s, -134 ± 38 s, and -132 ± 44 s, respectively, for XMM1-4), and across eruptions of a single epoch as well, which are all found within a ~ 24 s standard deviation (using XMM2 as example). Hence, we confirm the model choice for the QPE eruption profiles has no impact on our timing solutions, and we adopt the double-exponential model, since it generally performs better than a symmetric profile.

A.2.2. Extracting QPE Peak Times from XRT and NICER Data

For XRT, we note that the online tool, while yielding a series of detections around $\gtrsim 0.02 \text{ c s}^{-1}$, also has other more ambiguous cases, including very low count rates and/or errors larger than the count rate value. To treat this dataset more robustly and systematically, we performed aperture photometry in all of the XRT snapshots computing the no-source binomial probability P_b (following the method described in R. Arcodia et al. 2024d). The source aperture adopted was $20''$, while background counts were extracted in an annulus with an inner and outer radius of $80''$ and $320''$, respectively. We start the selection of XRT peaks by retaining the 10 snapshots with $P_b < 0.003$, five of which are within the XMM1-XMM4 campaign. For these, we re-extracted XRT products with different selections of time bins ranging within ~ 300 – 500 s to investigate whether smaller time bins would indicate any intra-snapshot variability and, thus, a more confident presence of a peak. All five snapshots within the XMM1-4, once decomposed in smaller time bins, indicate the presence of the eruptions peak either in a central time bin, or an edge (e.g., Figure 2). Out of the remaining five snapshots with $P_b < 0.003$, which are outside the XMM1-4 campaign (which is not used for the $O - C$ analysis), we only retain one that shows clear evidence of a detected peak.

NICER shares similar orbital constraints, but our processing method adopting spectral analysis in the single good time interval does not suffer from retention of spurious detections as much as the automated XRT tool. Thus, QPE arrival times can be immediately identified. For both XRT and NICER, we cannot fit eruption profiles due to the coarser constraint that these data provide (e.g., Figure 2), and instead assign peak times and errors as follows: if more than one immediately consecutive detection is present (i.e., with no orbital gaps in between), the peak time is assigned to the brightest and, otherwise, to the single data point. To justify the latter choice, we remind the reader that these detected count rates are compatible with expected values based on the peak fluxes observed by XMM and that eruptions are only detectable for ~ 600 – 800 s by XRT and NICER; thus, these detected good time intervals must encapsulate the actual QPE peak. For peak time uncertainties, we proceed as follows: if the peak is somewhat resolved (i.e., if three or more consecutive bins are detected, and the central one is clearly

the brightest), we assign an uncertainty equal to half the time bin size of the light curve, which is 200 s for NICER and typically 300 s for XRT (e.g., see the top-right and bottom-left panels in Figure 2 for examples); if the peak is only partially resolved (i.e., if two or more consecutive bins are detected, but the edge one is clearly the brightest), we assign an uncertainty equal to the full time bin size of the light curve (e.g., see the top-left panel in Figure 2 for an example); if the eruption is only observed with a single data point, the uncertainty on the peak time is 439 s, which is half the total near-peak duration of the eruptions based on XMM data, and the full error span includes the time for which the source is detectable by XRT and NICER (e.g., see the top-middle panel in Figure 2 as an example). In one case ($N_{\text{QPE}} = 231$ observed by XRT), we increase this to 1000 s (which fully includes how long the eruptions last, while still providing useful constraints) given the slightly lower count rate compared to the expected peak, and the lack of neighboring good time intervals.

A.2.3. The Energy Range

We tested the impact of extracting light curves with slightly different low-energy bounds across missions (i.e., 0.2 keV for XMM, 0.3 keV for XRT, 0.4 keV for NICER). We extracted XMM data of XMM2, as an example, between 0.4 and 2.0 keV and compared with the 0.2–2.0 keV light curves. We perform the same analysis described above for XMM data and estimated that the difference between peak times for XMM2 eruption has a mean offset of ~ 13 s, with a standard deviation of ~ 43 s, and it is thus consistent with zero. We also note that this negligible offset is much smaller than the typical uncertainties of XRT and NICER arrival times, which are in the range ~ 300 – 1000 s, depending on the eruption (see Table B1). Hence, this effect is negligible and does not impact our conclusions.

A.2.4. QPE Peak Time versus Start Time

We note that taking the peak of the QPEs as representative to study the quasiperiodicity of the source is the most sensible, and agnostic, choice to make in this work. However, for completeness, we tested whether adopting the start of the eruptions (as done for some EMRI model testing; e.g., C. Zhou et al. 2025a) would imprint any significant differences in the $O - C$ analysis. In essence, we do not expect a significant impact, as QPE durations across epochs were found to be consistent within uncertainties in long-term data of eRO-QPE2 (R. Arcodia et al. 2024a; D. Pasham et al. 2024). To confirm this, we took as start time of the XMM eruptions the time at which the flux reaches $1/e^3$ of the peak flux. We computed the difference between start and peak times and found that it is a constant offset, within uncertainties, around ≈ 500 s. As mentioned above, offsets do not have any impact in the $O - C$ analysis, as all times are relative to an arbitrary start time.

Appendix B More Details on the $O - C$ Analysis

B.1. Event Identification

For our XMM1-4 dataset, the association with N_{QPE} was performed with $T_{0,\text{est}}$ taken from the fitted peak time of the first

eruption of XMM1 (i.e., 732.69 s after the start of XMM1, MJD = 60489.6419, which is the origin of the time axis; e.g., Figure 1), and P_{est} taken from the average peak-to-peak QPE recurrence time of the XMM1, XMM2, XMM3, and XMM4 observations (i.e., 8055.71 s). Identification of an observed eruption with N_{QPE} was thus performed by associating a computed (“C”) event with the closest observed one (“O”), and we report all identification numbers in Table B1, and show the $O - C$ data in Figure 3. Data points from different missions are highlighted with symbols that follow those in Figure 1 and 2, namely circles for XMM, squares for XRT, and stars for NICER. A crucial requirement for this technique to be used, and its results interpreted correctly, is that each event can be associated with an integer N_{QPE} with no ambiguity. All of the eruptions of our dedicated XMM1-4 campaign can be associated unambiguously, as it was designed with this goal in mind. In practice, this is because gaps in between observations are small enough that the cumulative uncertainty of arrival times does not become larger than a small fraction of the QPE recurrence: the predicted “C” is either clearly overlapping in time with an observed eruption, or it is close enough (i.e., much smaller than half of a QPE cycle). More quantitatively, we note that in both Table B1 and Figure 3, most absolute values in the $O - C$ lie around $\approx(0.0-0.2)$ hr, with only a couple of data points around ~ 0.4 hr and only XMM4 around ~ 0.6 hr, all much smaller than the ~ 2.3 hr recurrence.

B.2. Modeling the $O - C$ Data with Red Noise

In the main text, we reported results obtained for models in which an underlying clock is present (i.e., orbital models), and the underlying noise process is assumed to be white. Here, we present an alternative analysis more appropriate for models in which the QPE event is predicted to be regulated and dominated by red-noise processes (e.g., disk instabilities or accretion related processes).

We model the stochastic contribution by the so-called DRW through Gaussian processes (GPs), which allow for the description of the stochastic process through a covariance function and assume the data are drawn from a multivariate Gaussian distribution around a mean function. We refer to L. Bertassi et al. (2025) for more details. In short, the GP log-likelihood L is given by:

$$\log L = -\frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}})\Sigma^{-1}(\mathbf{y} - \bar{\mathbf{y}})^T - \frac{1}{2}|\Sigma| - \frac{n}{2}\log(2\pi), \quad (\text{B1})$$

with $\mathbf{y}$ being the set of observations, $\bar{\mathbf{y}}$ being the mean function of the process, and Σ being the covariance matrix. For the DRW, the covariance function elements are:

$$\Sigma_{i,j} = k(x_i, x_j) = \sigma^2 e^{-\left(\frac{|x_i - x_j|}{\tau}\right)}, \quad (\text{B2})$$

where k is the covariance function, x_i, x_j are two observation times of the time series, σ sets the variance scale of the stochastic process, and τ is the damping timescale of the process beyond which the data become nearly uncorrelated. The GP mean function is taken to be either zero, or a deterministic component consisting of a linear trend (with coefficients a_{lin} and b_{lin}) plus a sinusoidal term (with parameters A, P, ϕ). We used `celerite` to compute the GP likelihood and `raynest` as nested sampling

Table B1

QPE Arrival Times (“O” for observed, “C” for Computed Using an Estimated Period) as Seconds Elapsed from the Start of the XMM1-4 Campaign ($t_{0,\text{XMM1}}$), Which Corresponds to MJD = 60489.6419

N_{QPE}	O (s)	O_{err} (s)	C (s)	Instr.
0	732.69	35.87	732.69	xmm
1	8739.45	167.72	8788.40	xmm
2	16830.53	39.51	16844.11	xmm
3	24781.40	37.43	24899.82	xmm
4	33033.59	35.66	32955.53	xmm
5	41185.92	29.17	41011.24	xmm
6	49251.53	52.84	49066.95	xmm
7	57195.27	20.94	57122.66	xmm
22	177196.76	439.30	177958.30	xrt
29	234335.95	439.30	234348.27	xrt
34	273254.13	439.30	274626.82	nicer
36	289634.68	300.00	290738.24	xrt
45	362394.13	439.30	363239.62	nicer
65	523958.62	31.83	524353.81	xmm
66	531892.37	27.13	532409.52	xmm
67	539934.91	27.59	540465.23	xmm
68	548024.16	27.16	548520.94	xmm
69	556009.37	25.62	556576.65	xmm
121	975273.71	439.30	975473.55	nicer
182	1466300.33	150.00	1466871.84	xrt
198	1595954.82	32.10	1595763.19	xmm
199	1603974.79	20.27	1603818.90	xmm
200	1611967.02	21.89	1611874.61	xmm
201	1620077.86	24.75	1619930.32	xmm
202	1627782.69	439.30	1627986.03	nicer
204	1644134.58	200.00	1644097.45	nicer
213	1716576.54	100.00	1716598.84	nicer
222	1788839.48	439.30	1789100.22	nicer
229	1845185.55	200.00	1845490.19	nicer
230	1852571.47	1000.00	1853545.90	xrt
323	2604787.86	29.48	2602726.89	xmm
324	2612822.99	21.27	2610782.60	xmm

Note. Some eruptions are shown in Figures 1 and 2.

(This table is available in machine-readable form in the [online article](#).)

algorithm. The priors assumed for the three models are shown in Table B2.

We adopted the noise-only model as null and obtained a DRW amplitude $\log \sigma^2 = -0.66^{+0.94}_{-0.54}$ and damping timescale $\log \tau = 2.76^{+0.94}_{-0.61}$. We show in Figure B1 the best-fit model predictions for the $O - C$ and the related corner plot. While successful in reproducing the data, the posteriors and model predictions are highly uncertain. Using the nonzero GP mean function to infer the presence of deterministic signals reveals a worse fit statistically, as the null model is preferred with $\Delta \log Z \sim 4$. However, we note that simulations would be required to assess the significance of this model comparison, which is beyond the scope of this work. Furthermore, we note that the fitted τ spans orders of magnitude and mostly values greater than the total observing baseline, thus predicting that, on short timescales, QPE recurrence times would be highly correlated. To confirm this, we computed a lag1 autocorrelation coefficient with the 15 QPE recurrence times observed by XMM during the XMM1-4 campaign. However, they appear compatible with zero, thus showing no correlation of consecutive recurrence times. While we note that the number of data points for this autocorrelation is limited, this observable

Table B2
Prior Ranges for the Red-noise Model Testing

Parameter	DRW	DRW + mod
a_{lin}	...	$\mathcal{U}(-3, 3)$
b_{lin}	...	$\mathcal{U}(-3, 3)$
$\log A$	...	$\mathcal{U}(-2, 2)$
$\log P$	...	$\mathcal{U}(-1, 4)$
ϕ	...	$\mathcal{U}(0, 2\pi)$
$\log \tau$	$\mathcal{U}(-1, 5)$	$\mathcal{U}(-1, 5)$
$\log \sigma^2$	$\mathcal{U}(-2, 2)$	$\mathcal{U}(-2, 2)$

appears in tension with the fitted τ , thus raising further concerns for the DRW-only model. More tests on much better sampled campaigns are required.

B.3. $O - C$ with Deterministic Components: Model Comparison and Selection

Here, we refer to model comparison and best-fit model selection with deterministic components against white noise. In Section 3, we mainly presented the adopted best-fit model ($\text{lin}+\text{mod}_1+\text{mod}_2$). In support of the full $O - C$ equation and the discussion on the best-fit parameters reported in Section 3 (see, e.g., Figure 5), we show here the best-fit corner plot in Figure B2. As described in Appendix B.3.1, to investigate the possible constraint on the longer modulation, we added two further eruptions identified by extrapolating the best-fit model forward in time. Figure B3 shows the resulting $O - C$ data and model, which confirm the presence of a modulation on timescales longer than our baseline. Here, we also expand on the preliminary models used to select the above as the best-fit model. We first fitted the $O - C$ data with a $\text{lin}+\text{mod}$ model, namely a single modulation obtained fitting the following equation: $O - C = (T_0 - T_{0,\text{est}}) + (P - P_{\text{est}})N_{\text{QPE}} + A_{\text{mod}} \sin(2\pi N_{\text{QPE}}/P_{\text{mod}} + \phi_{\text{mod}})$, which has five free parameters (T_0 , P , A_{mod} , P_{mod} , and $\phi_{\text{mod},1}$). The fit yielded poor residuals (see Figure B4), as the model is not able to capture both the negative delays of XMM1-3, spanning around $O - C \approx -0.4$ hr, and the positive delay of XMM4. This warranted adding more components.

We performed two additional fits, one adding a positive dissipative term, $+(1/2 \dot{P} P)N_{\text{QPE}}^2$ ($\text{lin}+\text{pdot}+\text{mod}$ model, with positive $\dot{P}$), and one adding a second modulation ($\text{lin}+\text{mod}_1+\text{mod}_2$). Both perform significantly better than the $\text{lin}+\text{mod}$, with a difference between the logarithmic Bayesian evidence ($\log Z$) being $\Delta \log Z = 128$ and $\Delta \log Z = 126$, for the $\text{lin}+\text{pdot}+\text{mod}$ and the $\text{lin}+\text{mod}_1+\text{mod}_2$ model, respectively, in comparison with the $\text{lin}+\text{mod}$. We show the $\text{lin}+\text{pdot}+\text{mod}$ model in Figure B5. While these two models appear equally good at describing the XMM1-4 data, we note that the $\text{lin}+\text{pdot}+\text{mod}$ model requires a strong positive dissipative term ($\dot{P} \approx 1.5 \times 10^5 \text{ ss}^{-1}$), which predicts an increasing QPE recurrence time over time. In turn, it predicts that the source had a smaller QPE recurrence time (t_{recur}) in the past, namely ≈ 2.1 hr a year before (in 2023 June) and ≈ 2.0 hr 2 yr before (in 2022 June), and so on. Therefore, this model can be immediately discarded based on archival 2020–2023 data of eRO-QPE2, as from 2022 June onward, t_{recur} has been observed in the ~ 2.25 – 2.30 hr range (R. Arcodia et al. 2024a; D. Pasham et al. 2024), and at even larger values in




Figure B1. Top panel: same as Figure 5, but with the damped random walk model with no deterministic components. Bottom panel: related corner plot.

the 2020 discovery light curve (R. Arcodia et al. 2021). We discussed the long-term data of eRO-QPE2 in more detail in Section 6, and we refer the reader to Figure 10, where the gray contours show the significant discrepancy between the $\text{lin}+\text{pdot}+\text{mod}$ model and the archival data. Thus, the $\text{lin}+\text{pdot}+\text{mod}$ is discarded as a viable description for the $O - C$ data of eRO-QPE2 thanks to the archival observations, and we instead focused on the $\text{lin}+\text{mod}_1+\text{mod}_2$ model (shown in the top panel of Figure 5). Given that it significantly improved the $\text{lin}+\text{mod}$ fit and it also obtains good residuals (bottom panel of Figure 5) with most data consistent within 1σ uncertainties, we refrained from adding any more components to avoid overfitting and adopt it as best-fit model. We note that these results, while using simplistic parametric models, are weakly sensitive to the prior volume adopted: varying amplitude prior bounds between tens to tens of thousands of seconds does not impact the fitted values; phases span the maximum bounds $(0, 2\pi)$; the bounds of the two modulation periods are separated at $N \sim 100$ (chosen visually from the $O - C$ data), and while this imposes a hierarchy, the fitted values are significantly smaller and larger than this separation bound, for the shorter and longer superperiod, respectively. All fitted parameters quoted throughout this work have statistical uncertainties marginalized over all of the model parameters. The systematic uncertainties may be large and are unknown at this stage, and a future monitoring campaign with more observed eruptions is needed.




Figure B2. Corner plot of the best-fit model (lin+mod₁+mod₂), shown in Figure 5.

B.3.1. The Robustness of the Superperiod Modulations

Here, we evaluated the robustness of the two superperiod components fitted by the adopted best-fit model with the deterministic $O - C$ components tested against white noise. First, we tested the short-term deterministic sinusoidal component analyzing the residual $O - C$ data after subtracting the linear component and the longer modulation from the best-fit model. We note that these residuals only impact the following simulations by providing a realistic standard deviation about the longer modulation; thus, it is irrelevant that they were obtained from a fit that includes the shorter modulation too. As a matter of fact, this subtraction approach is the most accurate description of where $O - C$ data would lie

about a longer modulation, compared to the poor fit with a single sinusoidal component (Figure B4). We randomly extracted 100 fake datasets of $O - C$ residuals using the observed standard deviation (~ 0.11 hr), assuming they are randomly distributed as white noise. We added noise to each fake dataset according to the actual measurement errors (i.e., to account for the more uncertain XRT and NICER data points). For each draw, we computed an error-weighted kernel density estimate (KDE) and thus obtained a distribution of 100 random KDEs. To compare these simulated random residuals with the true observed ones, we obtained 100 realizations of the observed dataset by perturbing each measurement according to its error bar (again to account for the high uncertainties in XRT and NICER data) and computed an analogous KDE distribution. We show this




Figure B3. Same as Figure B4, but for the “Best + Aug. data” model. Adding two further eruptions confirms the trend of the longer modulation inferred by the XMM1-4 data alone.




Figure B4. Top panel: $O - C$ data fitted by the `lin+mod` model (the median is shown with a solid line, and 1σ and 3σ are shown as shaded contours), which obtains poor residuals (shown in the bottom panel).

comparison in the left panel of Figure B6. The observed residuals appear inconsistent with the simulated random white-noise distribution, in that they show a much smaller spread and a multi-peaked structure that is, quite intriguingly, peaking around the minimum, maximum, and zero, consistent with what sinusoidal residuals would show (we show the simulated sinusoidal distribution in brown in Figure B6).

Regarding instead the longer modulation (the `mod2` component), we note that while our XMM1-4 campaign is long enough to infer its presence in the data, it is not long




Figure B5. Same as Figure B4, but for the `lin+pdot+mod` model. While it is a good fit to the XMM1-4 campaign, and while it is not distinguishable from the best-fit model (Figure 5), the presence of a strong dissipative term is discarded based on the significant inconsistency with archival data.

enough to probe its full cycle. In our best-fit model (see the corner plot in Figure B2), the posterior of $P_{\text{mod},2}$ is bound at the maximum value allowed (1300), and we deemed it unnecessary to extend it any farther, given that our campaign probes up to $N_{\text{QPE}} \sim 325$ (and we note a degeneracy with the parameter T_0 that should not be bigger than $\sim P/2$). The parameter $A_{\text{mod},2}$ is also naturally uncertain and spans a few thousand seconds. While it appears a somewhat constrained posterior, one must consider the degeneracy with $P_{\text{mod},2}$ and its upper bound. We note that NICER, XRT, and EP have observed further eruptions after XMM4. They are not part of the main campaign because their N_{QPE} could not be unambiguously identified from the start of the $O - C$ analysis (i.e., using $T_{0,\text{est}}$ and P_{est}). However, using the XMM1-4 best-fit model to identify further QPEs, we were able to add two eruptions observed by NICER around ~ 12.15 days and ~ 13.84 days after XMM4 (identified with $N_{\text{QPE}} = 454$ and 472 , respectively). They were identified with updated $T_{0,\text{est}}^{\text{ext}} = 3095.114$ s and $P_{\text{est}}^{\text{ext}} = 8064.298$ s from the XMM1-4 best-fit model. Their O (O_{err}) values, following the same units and starting point as in Table B1, are 3662927.964 s (± 439.3 s) and 3808453.612 s (± 439.3 s). We note that using the original $T_{0,\text{est}}$ and P_{est} , these two events would have been found halfway between $N_{\text{QPE}} = 454$ – 455 and 472 – 473 , respectively, thus still compatible with the identification using the refined $O - C$ best-fit period. We stress that contrary to the XMM1-4 campaign, this association is model dependent (as it benefits from the fitted modulations to identify N_{QPE} unambiguously), and we only intend to use it for a consistency check on the longer modulation: effectively, whether the best-fit model is a good predictor for the following several days or




Figure B6. Distribution of $O - C$ residuals after subtracting the linear trend and the longer sinusoidal modulation (e.g., Figure 5). The observed residuals are shown in light green and are obtained drawing from the observed uncertainties. Simulated datasets drawn from a random distribution centered at zero, with standard deviation ~ 0.01 hr (which is the one of the observed residuals) are shown in gray.




Figure B7. Posteriors of modulation parameters of the $O - C$ fit. The best-fit model ($\text{lin} + \text{mod}_1 + \text{mod}_2$) is highlighted in green, with other models in addition to interpret the robustness of the best-fit results. The dotted-dashed brown line shows the same best-fit model applied to two further eruptions added after the XMM1-4 campaign.

whether the new solution including the two further bursts changes dramatically. We fit this new dataset and obtain posteriors consistent with the best-fit model for all parameters, within uncertainties. Thanks to the extension of the baseline to 472 events, we also obtain a tighter constraint on the mod_2 component. Figure B7 shows the posteriors of $A_{\text{mod},2}$ and $P_{\text{mod},2}$ with our best-fit model together with the new dataset extending to two further eruptions (“Best + Aug. data”), and the posterior from the analysis of odd and even separately (see

Appendix B.4, corrected with a factor of 2), for a consistency check. Interestingly, the posterior of $P_{\text{mod},2}$ appears bound at $1016^{+119}_{-148} P$, and $A_{\text{mod},2}$ consequently improves to a more precise constraint ($A_{\text{mod},2} = 3521^{+803}_{-938}$ s). The fact that with two further eruptions added in the “Best + Aug. data” the longer modulation is consistent suggests that it is a good description of the timing of the eruptions (we show the related $O - C$ data-model comparison in Appendix B and Figure B3). We note that identifying these eruptions with the best-fit model may only bias the identification toward having additional $O - C$ data within $\pm P/2 \sim 1.1$ hr around the best-fit model (we refer to Figure 5 for visualization), but this does not bias them toward being very close to the prediction as opposed to anywhere else in the $\sim P \sim 2.2$ hr interval around it. This is a useful consistency check that confirms that the longer modulation, if it is the most correct description of the data, is not a spurious artifact driven by the XMM4 eruptions, but it is likely a robust feature in the data.

B.4. $O - C$ Fit with Odd and Even Eruptions Separately

The $O - C$ analysis performed on odd and even bursts separately traces the workflow reported in Section 3. The burst identification is the same as reported in Table B1. The $O - C$ event number for odd eruptions is obtained rescaling N_{QPE} from Table B1 to $N_{\text{QPE,odd}} = (N_{\text{QPE}} - 1)/2$, and those for even eruptions is $N_{\text{QPE,even}} = N_{\text{QPE}}/2$. $T_{0,\text{est}}$ and P_{est} are the same, with the exception of the period now being double that of the fit with all eruptions, namely 16111.42 s. The number of free parameters in all of the fits changes as follows. Most parameters are fitted jointly (e.g., P , $\dot{P}$, A_{mod} , P_{mod} , if present), but constants and phases are not.

Similarly to the case of all eruptions, for the case of deterministic components, the $\text{lin} + \text{mod}$ is inadequate to reproduce the $O - C$ dataset (Figure 4), and both $\text{lin} + \text{pdot} + \text{mod}$ and $\text{lin} + \text{mod}_1 + \text{mod}_2$ equally improve the fit, with $\Delta \log Z = 113$ and $\Delta \log Z = 114$ with respect to the $\text{lin} + \text{mod}$ model, respectively. However, the former is discarded, as it severely underpredicts the QPE recurrence time of archival 2020–2023 data (see Figure 10 for the case of all eruptions). Thus, we adopted the latter as a best-fit model, and we show the $O - C$ model in Figure B8 and related corner plot in Figure B9. The fitted parameters are overall compatible with the fit with all eruptions, modulo a factor of 2 in the periods; for instance, the fitted period at the XMM1-4 epoch is $P = 16127^{+14}_{-11}$ s, or ~ 4.48 hr, which is subtracted from both data and model in Figure B8. Compatible superperiod modulations (modulo the factor of 2) are fitted for both the mod_1 and mod_2 components, namely $P_{\text{mod},1} = (23.4 \pm 0.2) P$ for the faster superperiod, and $\approx 454 P$ at 1σ for the slower. The major result of the fit with odd and even datasets, in relation to their possible interpretation as EMRIs, is that both fitted superperiod modulations are consistent with being in phase. This can be seen by eye in Figure B8 and from the independent phase values in the corner plot (Figure B9). More quantitatively, the phase differences between odd and even $O - C$ datasets are $\Delta \phi_{\text{mod},1} = 0.05 \pm 0.06$ and $\Delta \phi_{\text{mod},2} = 0.02 \pm 0.02$, respectively. Finally, we added a negative quadratic component, $-(1/2 \dot{P} P) N_{\text{QPE}}^2$, to the best fit to infer an upper limit on $\dot{P}$ (sampling positive values for it). We obtained an upper limit on $\dot{P} < 3 \times 10^{-6} \text{ s s}^{-1}$, allowing




Figure B8. Same as Figure 5, but for even (left panels) and odd (right panels) eruptions fitted as separate datasets. The parameters P , $A_{\text{mod},1}$, $P_{\text{mod},1}$, $A_{\text{mod},2}$, and $P_{\text{mod},2}$ are fitted in common, while each dataset has its own constant and phase terms. Odd and even eruptions are found in phase on both short and long timescales. The related corner plot is shown in Figure B9.




Figure B9. Same as Figure B2, but for the fit with odd and even eruptions separately. The best-fit $O - C$ model is shown in Figure B8.




Figure B10. Same as Figure 6, but for odd and even eruptions separately, showing that consistent conclusions for the longer modulation term can be obtained.

parameters to vary only within the 10th–90th percentile range of the best-fit model.

Given the fundamental inconsistency unveiled by the correlated odd and even eruptions, the body text of this article shows primarily analysis and plots from the case of all eruptions fitted together, appropriate for comparing with single observed events per orbit. However, here we report constraints for EMRI models from the parameters inferred in the fit with separated odd and even eruptions, in case future versions of the QPE = EMRI models will require this setup (appropriate, for instance, if only one parity of QPEs is tracing the crossings). As the ratio between the superperiods parameters and the orbital period is now halved ($P_{\text{mod},1} \sim 23.3 P$), Equation (2) now yields $M_{\text{BH}} = (8.7 \pm 0.1) \times 10^5 M_{\odot}$. The amplitude of the modulation is $A_{\text{mod},1} \sim 318$ s, thus compatible with that of the fit with all eruptions; however, the estimated $\sim a/c$ time delay for orbit-size delays is larger (it is a fewer gravitational radii, but an R_g is now a larger physical size for a larger M_{BH}), namely ~ 297 s. This value can be thus achieved also for very low eccentricities. For the longer modulation, EMRI nodal is still an $\sim a/c$ effect and, thus, incompatible with the observed $A_{\text{mod},2}$. We repeated the tests that led to Figure 6 with the values inferred for the odd/even case, and generated Figure B10. Disk precession is disfavored for with the longer modulation since the parameter space for its solution would imply a disk precession alignment timescale shorter than the QPE lifetime in eRO-QPE2 (top panel of Figure B10). For the MBH binary case, no solutions exist for an $M_2 > M_1$ for the observed $A_{\text{mod},2}$ and $P_{\text{mod},2}$, and binaries can exist at $M_2 < M_1$




Figure B11. Same as Figures 7 (top) and 8 (bottom), but for odd and even eruptions separately.

for separations $a_{\text{bin}} \sim 0.75$ mpc, which is a factor ~ 260 larger than the inner EMRI separation $\sim 70 R_g$, which is thus a stable system. However, the total mass, between a factor of 1 and 2 of M_1 , is incompatible with the $M_{\text{BH}} - \sigma$ estimate of $\approx 10^5 M_{\odot}$ (T. Wevers et al. 2024).

Finally, the $\dot{P}$ upper limit for the case of odd/even eruptions is $< 3 \times 10^{-6}$, which would still rule out maximum eccentricity at all secondary masses, and most IMBH solutions (top panel of Figure B11). In the bottom panel of Figure B11, we show the constraints that are obtained from the $O - C \dot{P}$ upper limit and the QPE lifetime in terms of the mass of a star and the density of the accretion disk if the orbital evolution is dominated by gas drag. Compared to Figure 8, a main-sequence star “bullet” has some allowed parameter space even for $\epsilon \lesssim 0.1$, due to the lower $\Sigma_{\text{min}} R = a = 70 R_g$. For $R_{\text{out}} \approx 300 R_g$, appropriate given the estimate from T. Wevers et al. (2025), the range $\Sigma = (0.5 - 1) \times 10^5 \text{ g cm}^{-2}$ can be obtained with a range $0.07 \lesssim M_{\text{disk},0}/M_{\odot} \lesssim 0.17$. Similarly to the case for all eruptions, stellar streams would allow for larger cross sections, thus lower Σ_{min} , thus lower $M_{\text{disk},0}$. Finally, we note that there is a minor difference in some of the slopes as the radius–mass relation breaks to ~ 0.6 for supersolar masses (compared to ~ 0.8 for subsolar).

Appendix C

More Details on the EMRI Trajectory Models

The first code we used is the publicly available QPE timing analysis code `QPE-FIT`²⁶ v. 0.1.11, which performs Bayesian parameter inference for EMRI and SMBH properties using QPE timings under the assumption that they are due to orbiter-disk collisions around an SMBH (J. Chakraborty et al. 2025a). To do so, it assumes a forward model combining an EMRI orbital trajectory $\mathbf{r}(t) = (x(t), y(t), z(t))$ with a planar accretion disk having a rigidly precessing normal vector:

$$\hat{\mathbf{d}}(t) = \begin{bmatrix} \sin \theta_{\text{disk}} \cos(2\pi t/T_{\text{disk}} + \phi_{\text{disk},0}) \\ \sin \theta_{\text{disk}} \sin(2\pi t/T_{\text{disk}} + \phi_{\text{disk},0}) \\ \cos \theta_{\text{disk}} \end{bmatrix}, \quad (\text{C1})$$

where θ_{disk} is the angle between the disk normal vector and the SMBH spin axis, T_{disk} is the disk nodal precession period, and $\phi_{\text{disk},0}$ is an arbitrary initial phase. In general, $\mathbf{r}(t)$ depends on three orbital parameters (semimajor axis a , eccentricity e , and inclination i); three initial phase parameters ($\phi_{r,0}$, $\phi_{\theta,0}$, and $\phi_{\phi,0}$); and two SMBH parameters that influence the spacetime metric and relativistic timing effects (spin χ_* and mass $\log M_*$). Disk crossings occur when the condition $\mathbf{r}(t) \cdot \hat{\mathbf{d}}(t) = 0$ is satisfied. After computing the disk crossings in the orbiter's frame, their timings are further modified by the Shapiro time delay, a function of M_* , and the geometric light travel-time differences, which requires introducing an additional parameter for the observer viewing angle, θ_{obs} . The trajectory $\mathbf{r}(t)$ is relatively expensive to compute, as solutions to the geodesic equations in Kerr spacetime are known analytically only in terms of elliptic integrals (M. van de Meent 2020), which must be solved numerically. Hence, `QPE-FIT` uses the post-Newtonian orbital solutions for bound Kerr geodesics (R. Fujita & W. Hikida 2009), along with GPU-acceleration to simultaneously process large parameter batches in parallel, making the inference problem tractable.

For a given parameter vector $\xi = \{a, e, i, \phi_{r,0}, \phi_{\theta,0}, \phi_{\phi,0}, \chi_*, \log M_*, \theta_{\text{obs}}, \theta_{\text{disk}}, T_{\text{disk}}, \phi_{\text{disk},0}\}$, `QPE-FIT` first computes $\mathbf{r}(t)$ and $\hat{\mathbf{d}}(t)$, then numerically finds the zeros of $\mathbf{r}(t) \cdot \hat{\mathbf{d}}(t)$, accounts for relativistic/geometric time delays to find the model timings $t_{\text{model},i}$, then combines with the observed QPE timings $\{t_{\text{data},i}\}$ and errors $\{\sigma_{\text{data},i}\}$ to estimate the model likelihood via the chi-squared error:

$$\log \mathcal{L}(\xi, \{t_{\text{model},i}, \sigma_{\text{model},i}\}) \propto -\frac{1}{2} \sum_j \sum_i \left(\frac{t_{\text{data},i} - t_{\text{model},i}}{\sigma_{\text{data},i}} \right)^2, \quad (\text{C2})$$

where j indexes over the N_{obs} total observing windows, and i indexes over the N_{QPE} eruptions within each window. Using this log-likelihood function, `QPE-FIT` performs Bayesian inference to derive posterior probability distributions and the evidence using the nested sampling Monte Carlo algorithm `MLFriends` (J. Buchner 2019, 2021) using the `UltraNest` Python package (J. Buchner 2021).

In parallel, we performed an independent fit using an EMRI trajectory model based on the work of A. Franchini et al.

Table C1

Prior Ranges for Various `QPE-FIT` Runs with Two/One Flares per Orbit (`QF-2coll`/`QF-1coll`, Respectively)

Parameter	Two-flare	One-flare (wide)	One-flare (narrow)
a/R_g	$\mathcal{U}(50, 1000)$	$\mathcal{U}(50, 1000)$	$\mathcal{U}(136, 187)$
e	$\mathcal{U}(0, 0.1)$	$\mathcal{U}(0, 0.5)$	$\mathcal{U}(0, 0.5)$
i	$\mathcal{U}(0^\circ, 180^\circ)$	$\mathcal{U}(0^\circ, 180^\circ)$	$\mathcal{U}(0^\circ, 180^\circ)$
$\phi_{r,0}$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$
$\phi_{\theta,0}$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$
$\phi_{\phi,0}$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$
χ_*	$\mathcal{U}(0, 0.998)$	$\mathcal{U}(0, 0.998)$	$\mathcal{U}(0, 0.998)$
$\log(M_*/M_\odot)$	$\mathcal{U}(4.5, 6.5)$	$\mathcal{U}(4.5, 6.5)$	$\mathcal{U}(5, 5.2)$
θ_{obs}	$\mathcal{U}(0, \pi)$	$\mathcal{U}(0, \pi)$	$\mathcal{U}(0, \pi)$
θ_{disk}	$\mathcal{U}(0^\circ, 20^\circ)$	$\mathcal{U}(0^\circ, 20^\circ)$	$\mathcal{U}(0^\circ, 20^\circ)$
$T_{\text{disk}}/P_{\text{orb}}$	$\mathcal{U}(200, 1000)$	$\mathcal{U}(100, 2000)$	$\mathcal{U}(1000, 2000)$
$\phi_{\text{disk},0}$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$	$\mathcal{U}(0, 2\pi)$

Table C2

Prior Ranges for the `FB23` Runs with Two Flares per Orbit (`FB23-2coll`)

Parameter	Two-flare
a/R_g	$\mathcal{U}(50, 300)$
e	$\mathcal{U}(0, 0.9)$
i_{orb}	$\mathcal{U}(0^\circ, 90^\circ)$
χ_*	$\mathcal{U}(0, 1)$
$\log(M_*/M_\odot)$	$\mathcal{U}(3, 8)$
θ_{obs}	$\mathcal{U}(0, \pi)$
i_{disk}	$\mathcal{U}(0^\circ, 90^\circ)$
$\Delta t/h$	$\mathcal{U}(0, 500)$

(2023), thus named hereafter `FB23`, but we note that it will be published in A. Motta et al. (2026, in preparation). Similarly to the `QPE-FIT` code, we compute the EMRI trajectory $\mathbf{r}(t)$ and the time-dependent position of the accretion disk $\mathbf{d}(t)$, identify the impact times numerically by solving $\mathbf{r}(t) \cdot \mathbf{d}(t) = 0$, add the time delays, and use the predicted arrival times to evaluate the likelihood function. The system is described by eight free parameters: mass of the MBH $\log M_*$, semimajor axis a , eccentricity e , orbital inclination i_{orb} , spin χ_* , disk inclination i_{disk} , polar angle of the observer θ_{obs} , and a time offset Δt . The latter encodes the global information on the initial conditions of the system, including the observer's azimuthal angle, the true anomaly, the longitude of the ascending node, and the argument of the periastron. The accretion disk is assumed to be planar, circular, and rigidly precessing, with a power-law density profile $\Sigma(R)$ (A. Franchini et al. 2016). It extends from the innermost stable circular orbit out to $R_{\text{out}} = 300R_g$, and its inclination with respect to the plane perpendicular to the spin remains constant in time, as alignment effects are neglected. Under these assumptions and for a small misalignment angle, the disk precession frequency is computed as the angular-momentum-weighted average of the Lense–Thirring precession frequency Ω_{LT} across the disk. Ω_{LT} at a given disk radius R around an MBH of spin χ_* reads:

$$\Omega_{\text{LT}}(R, \chi_*) = \frac{c^3}{2GM_*} \frac{4\chi_* \left(\frac{R}{R_g}\right)^{-3/2} - 3\chi_*^2 \left(\frac{R}{R_g}\right)^{-2}}{\chi_* + \left(\frac{R}{R_g}\right)^{3/2}},$$

²⁶ Available at <https://github.com/joheenc/QPE-FIT>.




Figure C1. $O - C$ data for odd and even eruptions (from Figure 4) overlaid with the mock $O - C$ delays from EMRI trajectory best-fit models (QPE-FIT run QF-2c011 at the top, the analogous fit with the FB23 model at the bottom, FB23-2c011) assuming both odd and even eruptions trace disk crossing. In both cases, the EMRI model generates anticorrelated odd and even delays on the apsidal precession timescale as expected (even if data are correlated). The fits circumvent this inconsistency by fitting the XMM data (circles), which are the most and those with the most accurate uncertainties, at the crossing between odd and even modes of the modeled $O - C$. Thus, these solutions are not reliable.

where $R_g = GM/c^2$ is the MBH gravitational radius. The angular momentum surface density is given by $L(R) = \Sigma(R)\Omega(R)R^2$, with $\Omega(R)$ being the Keplerian orbital frequency of the gas.

The main difference with QPE-FIT is that the precession period is linked to the MBH spin and the disk properties, providing a stronger constraint to the spin value during the parameter inference. Following A. Franchini et al. (2023), the EMRI trajectory is evolved by integrating the equations of motion within the post-Newtonian framework up to 3.5PN order (A. Bohé et al. 2013; L. Blanchet 2014), thereby accounting for dissipative effects due to gravitational-wave emission. The equations are numerically integrated over the full temporal baseline of the dataset, and once the crossing times are determined, the code selects the events closest to the observed ones to evaluate the model likelihood. Bayesian inference is performed using the nested sampling algorithm Nessai (M. J. Williams et al. 2021), employing 1000 live points and parallelization over 48 CPU cores. Uniform priors are adopted for all parameters, and the maximum-likelihood estimate is taken as the best-fit solution.

C.1. Model Runs

When fitting the QPE timings, we first assumed a model in which two observable eruptions are generated by the




Figure C2. $O - C$ data overlaid with the mock $O - C$ delays from (QPE-FIT). Three separate runs for the one-impact QF-1c011 flavor are shown, as described by the legend. The best-fit models from the 6 month XMM1-5 timing data (purple and green lines) do not reproduce the short-term variability shown by XRT and NICER data, with either wide or narrow priors (see Table C1). The best-fit model of the XMM1-4 dataset (orange), while reproducing relatively well the $O - C$ delays with narrow priors, does not reproduce the censored data point XMM4 (not in the plot).

ascending/descending passages of the orbiter through the accretion disk, as is usually invoked in EMRI models (e.g., J. Xian et al. 2021; A. Franchini et al. 2023; I. Linial & B. D. Metzger 2023; C. Zhou et al. 2024). We performed this sampling run with both QPE-FIT and FB23 models, and name the runs QF-2c011 and FB23-2c011, respectively. The priors of this sampling run are given in the second column of Table C1 for QPE-FIT, and the second column of Table C2 for FB23-2c011. We applied the faster QPE-FIT to the full XMM1-5 dataset, and the fit converged to a solution with parameters $a \sim 91R_g$, $e \sim 0.04$, and $\log M \sim 5.8$. For the FB23-2c011 run, we applied this setup to the XMM1-4 dataset given the much longer run times in integrating over the full temporal baseline in the FB23 model. The fit converged to a solution with parameters $a \sim 205R_g$, $e \sim 0.03$, and $\log M \sim 5.2$. Notably, they are two incompatible solutions, but we note that the semimajor axis agrees within a factor of 2 in physical sizes. Each of them is at face value reasonable for a system like eRO-QPE2 when compared to the $O - C$ analysis performed on odd/even eruptions separately (Appendix B.4) and M_{BH} values in the literature (T. Wevers et al. 2024, 2025). To investigate more in detail their comparison with the empirical $O - C$ data, we overplotted the predicted $O - C$ curves from the best-fit model posteriors to the observed $O - C$ delays in Figure C1. Both solutions have in common that XMM data are fitted at the crossing points between odd and even predicted curves, which is what ultimately led us to consider them unreliable. We refer to the body text (Section 5) for more details on this.

Given that our analysis disfavored EMRI models with two observable collisions per orbit, we have provided a possible interpretation of our $O - C$ results assuming one event per orbit in Section 4. Here, we report out attempts to fit timing data of eRO-QPE2 with QPE-FIT, the fastest of the two algorithms used here, to test whether more robust EMRI trajectory models would find a compatible solution (QF-1c011). In summary, no solution stands the test of either data censoring (varying the application to XMM1-4 data or XMM1-5 data) nor significant prior volume changes. Hence, no fully reliable solution is found for the case of one collision

per orbit either. We report four main setups, although many more tests have been done to test the algorithm robustness to data censoring and prior volume.

We first attempted to find solutions with a wide prior range, choosing bounds on the semimajor axis, eccentricity, SMBH mass, and disk precession period that were agnostic to the empirical fits (second column of Table C1). These fits terminated at a $\log Z = -399 \pm 0.5$, converging to parameters $a \approx (352 \pm 3)R_g$, $e \approx 0.34 \pm 0.01$, $\log M_* \approx 4.6 \pm 0.01$, and $T_{\text{disk}} \approx (1386 \pm 106)P_{\text{orb}}$. It is immediately apparent that this fit does not recover the same apsidal period as the empirical fits, instead showing a shortest oscillation timescale of $a(1 - e^2)/(3R_g) \approx 104P_{\text{orb}}$. More quantitatively, this solution ignores the short-term timing variation shown by intra-XMM observations in the XMM1-4 epoch, mainly shown by NICER and XRT data. Perhaps more worryingly, once only the XMM1-4 data are fitted (thus censoring any data post XMM4), the solution changes. To investigate whether more stability to data censoring would be obtained with a narrow parameters volume, we restricted the prior bounds to values compatible with the $O - C$ solution (third column of Table C1). A solution is found with $\log Z = -321 \pm 0.4$, with $\Delta \log Z = 78$ compared to the runs with the wider bounds, indicating that the multimodality and extremely small posterior/prior volume ratio of this problem presents a challenge for even robust statistical inference algorithms such as nested sampling. The main fitted parameters are $a = (141 \pm 0.5)R_g$, $e = 0.12 \pm 0.01$, $\log M_{\text{BH}}/M_\odot = 5.194 \pm 0.002$, and $P_{\text{disk}}/P = 1364^{+53}_{-211}$. While they appear in remarkable agreement with the $O - C$ parameters (corner plot in Figure B2), once the predicted $O - C$ delays are superimposed to the observed one (similarly to the previous test done with the two-impact scenario; Figure C1), again, the medium-term variability highlighted by XRT and NICER data is largely ignored. Again, quite worryingly, censoring data post-XMM4 changes the fitted parameters to $a \sim 184R_g$, $e \sim 0.16$, and $\log M_{\text{BH}}/M_\odot \sim 5.0$. We show the superposition between mock and observed $O - C$ delays from the narrow prior runs in Figure C2, as an example. Similar inconsistencies are obtained for the wider prior volume. Finally, we note that we explored different volume sampling algorithms within the *Ultraneest* code with no change.

ORCID iDs

R. Arcodia  <https://orcid.org/0000-0003-4054-7978>
 G. Miniutti  <https://orcid.org/0000-0003-0707-4531>
 J. Chakraborty  <https://orcid.org/0000-0002-0568-6000>
 A. Franchini  <https://orcid.org/0000-0002-8400-0969>
 M. Giustini  <https://orcid.org/0000-0002-1329-658X>
 I. Linial  <https://orcid.org/0000-0002-8304-1988>
 M. Bonetti  <https://orcid.org/0000-0001-7889-6810>
 E. Kara  <https://orcid.org/0000-0003-0172-0854>
 A. Merloni  <https://orcid.org/0000-0002-0761-0130>
 A. Motta  <https://orcid.org/0009-0008-6234-6062>
 G. Ponti  <https://orcid.org/0000-0003-0293-3608>
 R. Soria  <https://orcid.org/0000-0002-4622-796X>
 J. Buchner  <https://orcid.org/0000-0003-0426-6634>
 P. C. Fragile  <https://orcid.org/0000-0002-5786-186X>
 A. Ingram  <https://orcid.org/0000-0002-5311-9078>
 M. Middleton  <https://orcid.org/0000-0002-8183-2970>
 C. Panagiotou  <https://orcid.org/0009-0001-9034-6261>
 P. Yao  <https://orcid.org/0000-0003-3024-7218>
 A. Rau  <https://orcid.org/0000-0001-5990-6243>

M. Guolo  <https://orcid.org/0000-0002-5063-0751>
 R. Saxton  <https://orcid.org/0000-0002-4912-2477>

References

- Allievi, C. M., Broggi, L., Sesana, A., & Bonetti, M. 2026, arXiv:2603.02302
 Arcodia, R., Baldini, P., Merloni, A., et al. 2025, *ApJ*, **989**, L3
 Arcodia, R., Merloni, A., Nandra, K., et al. 2021, *Natur*, **592**, 704
 Arcodia, R., Miniutti, G., Ponti, G., et al. 2022, *A&A*, **662**, A49
 Arcodia, R., Linial, I., Miniutti, G., et al. 2024a, *A&A*, **690**, A80
 Arcodia, R., Liu, Z., Merloni, A., et al. 2024b, *A&A*, **684**, A64
 Arcodia, R., Merloni, A., Buchner, J., et al. 2024c, *A&A*, **684**, L14
 Arcodia, R., Merloni, A., Comparat, J., et al. 2024d, *A&A*, **681**, A97
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Baldini, P., Rau, A., Merloni, A., et al. 2026, *A&A*, **706**, L15
 Bertassi, L., Charisi, M., Buscicchio, R., et al. 2025, arXiv:2512.13688
 Blanchet, L. 2014, *LRR*, **17**, 2
 Bohé, A., Marsat, S., & Blanchet, L. 2013, *CQGra*, **30**, 135009
 Buchner, J. 2019, *PASP*, **131**, 108005
 Buchner, J. 2021, *JOSS*, **6**, 3001
 Bykov, S., Gilfanov, M., Sunyaev, R., & Medvedev, P. 2025, *MNRAS*, **540**, 30
 Chakraborty, J., Arcodia, R., Kara, E., et al. 2024, *ApJ*, **965**, 12
 Chakraborty, J., Drummond, L. V., Bonetti, M., et al. 2025a, *ApJ*, **992**, 120
 Chakraborty, J., Kara, E., Arcodia, R., et al. 2025b, *ApJL*, **983**, L39
 Chakraborty, J., Kara, E., Masterson, M., et al. 2021, *ApJL*, **921**, L40
 Chakraborty, J., Kosec, P., Kara, E., et al. 2025c, *ApJ*, **984**, 124
 Chakraborty, J., Rappaport, S. A., Arcodia, R., et al. 2026, *ApJL*, **1001**, L6
 Chen, J.-H., Shen, R.-F., & Liu, S.-F. 2023, *ApJ*, **947**, 32
 Colpi, M., Danzmann, K., Hewitson, M., et al. 2024, arXiv:2402.07571
 Dodd, S. A., Huang, X., Davis, S. W., & Ramirez-Ruiz, E. 2025, 10.48550/arXiv.2506.19900
 D’Orazio, D. J., Tiede, C., Zwick, L., Hayasaki, K., & Mayer, L. 2025, *ApJ*, **994**, 112
 Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2007, *A&A*, **469**, 379
 Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2009, *MNRAS*, **397**, 1177
 Franchini, A., Bonetti, M., Lupi, A., et al. 2023, *A&A*, **675**, A100
 Franchini, A., Lodato, G., & Facchini, S. 2016, *MNRAS*, **455**, 1946
 Frank, J., King, A., & Raine, D. J. 2002, *Accretion Power in Astrophysics* (3rd ed.; Cambridge Univ. Press)
 Fujita, R., & Hikida, W. 2009, *CQGra*, **26**, 135002
 Giustini, M., Miniutti, G., & Saxton, R. D. 2020, *A&A*, **636**, L2
 Götzberg, Y., de Mink, S. E., Groh, J. H., et al. 2018, *A&A*, **615**, A78
 Guo, W., & Shen, R.-F. 2026, *ApJ*, **998**, 78
 Guolo, M., Mummery, A., Ingram, A., et al. 2025a, *ApJ*, **992**, 114
 Guolo, M., Mummery, A., Wevers, T., et al. 2025b, *ApJ*, **985**, 146
 Hajdu, G., Catelan, M., Jurcsik, J., et al. 2015, *MNRAS*, **449**, L113
 Hernández-García, L., Chakraborty, J., Sánchez-Sáez, P., et al. 2025, *NatAs*, **9**, 895
 Huang, X., Linial, I., & Jiang, Y.-F. 2025, *ApJ*, **993**, 186
 Ingram, A., Motta, S. E., Aigrain, S., & Karastergiou, A. 2021, *MNRAS*, **503**, 1703
 Janković, T., Bonnerot, C., Karpov, S., & Jurca, A. 2026, arXiv:2602.02656
 Jansen, F., Lumb, D., Altieri, B., et al. 2001, *A&A*, **365**, L1
 Jiang, N., & Pan, Z. 2025, *ApJL*, **983**, L18
 Kaur, K., Stone, N. C., & Gilbaum, S. 2023, *MNRAS*, **524**, 1269
 King, A. 2020, *MNRAS*, **493**, L120
 Koen, C. 2006, *MNRAS*, **365**, 489
 Lam, A. T.-L., Shibata, M., Kawaguchi, K., & Pelle, J. 2025, *PhRvD*, **112**, 083006
 Linial, I., & Metzger, B. D. 2023, *ApJ*, **957**, 34
 Linial, I., Metzger, B. D., & Quataert, E. 2025, *ApJ*, **991**, 147
 Linial, I., & Quataert, E. 2024, *MNRAS*, **527**, 4317
 Linial, I., & Sari, R. 2023, *ApJ*, **945**, 86
 Lu, W., & Quataert, E. 2023, *MNRAS*, **524**, 6247
 Luo, J., Chen, L.-S., Duan, H.-Z., et al. 2016, *CQGra*, **33**, 035010
 Mardling, R. A., & Aarseth, S. J. 2001, *MNRAS*, **321**, 398
 Middleton, M., Gúrpide, A., Kwan, T. M., et al. 2025, *MNRAS*, **537**, 1688
 Miniutti, G., Franchini, A., Bonetti, M., et al. 2025, *A&A*, **693**, A179
 Miniutti, G., Giustini, M., Arcodia, R., et al. 2023a, *A&A*, **674**, L1
 Miniutti, G., Giustini, M., Arcodia, R., et al. 2023b, *A&A*, **670**, A93
 Miniutti, G., Saxton, R. D., Giustini, M., et al. 2019, *Natur*, **573**, 381
 Mummery, A. 2025, arXiv:2504.21456

- Naoz, S. 2016, *ARA&A*, **54**, 441
- Nicholl, M., Pasham, D. R., Mummery, A., et al. 2024, *Natur*, **634**, 804
- Pan, X., Li, S.-L., & Cao, X. 2023, *ApJ*, **952**, 32
- Pasham, D., Kejriwal, S., Coughlin, E., et al. 2024, arXiv:2411.00289
- Payne, A. V., Shappee, B. J., Hinkle, J. T., et al. 2021, *ApJ*, **910**, 125
- Peters, P. C. 1964, *Phys. Rev.*, **136**, B1224
- Quintin, E., Webb, N. A., Guillot, S., et al. 2023, *A&A*, **675**, A152
- Raj, A., & Nixon, C. J. 2021, *ApJ*, **909**, 82
- Sesana, A., Korsakova, N., Arca Sedda, M., et al. 2021, *ExA*, **51**, 1333
- Sterken, C. 2005, ASPC, **335**, 3
- Stone, N., & Loeb, A. 2012, *PhRvL*, **108**, 061302
- Suzuguchi, T., Omiya, H., & Takeda, H. 2026, *PASJ*, **78**, 185
- Tagawa, H., & Haiman, Z. 2023, *MNRAS*, **526**, 69
- van de Meent, M. 2020, *CQGra*, **37**, 145007
- Vurm, I., Linial, I., & Metzger, B. D. 2025, *ApJ*, **983**, 40
- Wang, M., Yin, J., Ma, Y., & Wu, Q. 2022, *ApJ*, **933**, 225
- Wevers, T., French, K. D., Zabludoff, A. I., et al. 2024, *ApJL*, **970**, L23
- Wevers, T., Guolo, M., Lockwood, S., et al. 2025, *ApJL*, **980**, L1
- Williams, M. J., Veitch, J., & Messenger, C. 2021, *PhRvD*, **103**, 103006
- Xian, J., Zhang, F., Dou, L., & Chen, Z. 2025, *ApJ*, **987**, 171
- Xian, J., Zhang, F., Dou, L., He, J., & Shu, X. 2021, *ApJL*, **921**, L32
- Xiang-Gruess, M., Ivanov, P. B., & Papaloizou, J. C. B. 2016, *MNRAS*, **463**, 2242
- Yao, P. Z., Quataert, E., Jiang, Y.-F., Lu, W., & White, C. J. 2025, *ApJ*, **978**, 91
- Zasche, P., Liakos, A., Niarchos, P., et al. 2009, *NewA*, **14**, 121
- Zeng, Y., & Pan, Z. 2026, arXiv:2601.11925
- Zhan, Y., Wang, D., Yi, S.-X., & Wang, F.-Y. 2026, *ApJ*, **997**, 134
- Zhou, C., Huang, L., Guo, K., Li, Y.-P., & Pan, Z. 2024, *PhRvD*, **109**, 103031
- Zhou, C., Pan, Z., Jiang, N., & Zhao, W. 2025a, *MNRAS*, **543**, 1816
- Zhou, C., Zeng, Y., & Pan, Z. 2025b, *ApJ*, **985**, 242