X-ray variability in the eROSITA era

David Bogensberger

SIGILL. UNIVERSITAT. LUDOVIC. MAXIM.

München 2022

X-ray variability in the eROSITA era

David Bogensberger

Dissertation
der Fakultät für Physik
der Ludwig-Maximilians-Universität
München

vorgelegt von
David Bogensberger
aus Wien

München, den 16.12.2022

Erstgutachter: Prof. Dr. Kirpal Nandra
Zweitgutachter: Prof. Dr. Andreas Burkert
Tag der mündlichen Prüfung: 01.03.2023

Contents

List of Figures

List of Tables

Zusammenfassung

Die helle Röntgenemission von Röntgendoppelsternen und aktiven Galaxienkernen kommt aus der unmittelbaren Umgebung des Ereignishorizonts der Schwarzen Löcher. Weil diese Region verhältnismäßig klein ist, verändert sich die Helligkeit oft innerhalb kurzer Zeit. Die Untersuchung der Röntgen-Variabilität ermöglicht uns, die physikalischen Mechanismen in dieser extremen Region, in der die Physik an ihre Grenzen stößt wird, zu verstehen.

Zu diesem Zweck habe ich die Variabilität tausender Röntgenquellen untersucht, die während der vollständigen Himmelsdurchmusterungen durch *eROSITA* beobachtet wurden. Eine solche Untersuchung wird durch die Eigenschaften der Beobachtungen erschwert, besonders durch die variablen Expositionszeiten. Deshalb habe ich gebräuchliche Methoden an *eROSITA* angepasst, und weiter verbessert. Ich habe Schwellenwerte zur Erkennung signifikanter Variabilität erstellt, und zwar als Funktion der Anzahl der Datenpunkte in der Lichtkurve, und der durchschnittlichen Zählrate der Quelle. Ferner habe ich eine neue Methode entwickelt, um die normalisierte intrinsische Varianz auszurechnen, die deutlich genauer als bisherige Methoden ist. Außerdem habe ich die Größe der systematischen Fehler, sowie das überschüssige Rauschen in Periodogrammen bestimmt.

Für die Variabilitätsanalyse habe ich die Region um den südlichen Ekliptikpol ausgewählt, weil *eROSITA* die Röntgenquellen dort am häufigsten beobachten konnte. Ich habe meine Methoden hier eingesetzt und dadurch 453 eigenständige signifikant variable Quellen entdeckt. Ich habe ihre optischen Spektren beobachtet und als entweder galaktisch oder extragalaktisch eingeordnet. Schließlich habe ich einige variable Quellen mit interessanten und einzigartigen Eigenschaften genauer untersucht.

Außerdem habe ich die faszinierenden und einzigartigen Eigenschaften des Ausbruchs des Röntgendoppelsterns Swift J1658.2-4242 im Jahr 2018 genau analysiert. In diesem Ausbruch habe ich einige Flip-Flops beobachtet, die einen größeren Helligkeitsunterschied zwischen der hellen und der dunklen Phase aufwiesen als in allen bisherigen Beobachtungen. In vielen Röntgendoppelsternen werden quasi-periodische Oszillationen gemessen, die in Typ A, B, und C untergliedert werden. In meinen Beobachtungen habe ich den ersten bisher beobachteten direkten Sprung zwischen Typ A und Typ C feststellen können. Außerdem habe ich herausgefunden, dass die schnellen Flip-Flop Wechsel zwischen den hellen und dunklen Phasen immer nur zu ganzzahligen Vielfachen einer fundamentalen Periode von 2.8 ks geschehen. Eine Analyse der Röntgenspektren hat gezeigt, dass die Helligkeitsschwankungen vor allem durch Veränderungen der Temperatur verursacht werden.

Abstract

The bright X-ray emission from Black Hole Transients (BHTs) and Active Galactic Nuclei (AGNs) originates from the region close to the black hole event horizon. Due to the comparably small size of this region, the observed flux often varies significantly on short timescales. Probing the variability properties of these systems helps us identify and understand the physical mechanisms at work in these extreme environments.

To this end, we set out to investigate the X-ray variability observed during the *eROSITA* all-sky surveys. Due to the unique characteristics of *eROSITA*, particularly its varying exposure times per observation, various variability methodologies had to be adapted to yield accurate results. Therefore, I determined thresholds for selecting variable sources for an extensive range of number of bins and mean count rates. I also defined a new method of estimating the normalised intrinsic variance of a source that is significantly more accurate than previous methods, especially at low count rates. Additionally, I determined the size of the sampling errors and the excess noise in periodograms.

X-ray sources located close to the South Ecliptic Pole (SEP) were observed most frequently by *eROSITA*. Therefore, this field is the ideal choice for long-term variability analysis. I used the variability methods I defined to identify 453 unique significantly variable sources within 3° of the SEP in the first three all-sky surveys. Next, I distinguished sources as being either likely galactic or likely extragalactic and observed and fitted the optical spectra of thousands of *eROSITA*-selected sources. Finally, I investigated some of the significantly variable sources, identifying unique effects.

I also investigated the intriguing and unique properties of the 2018 outburst of the black hole transient Swift J1658.2-4242. It prominently features several flip-flops, which have the appearance of top-hat functions in a light curve. These flip-flops had a larger flux difference between the dim and bright states than has ever been observed. X-ray binary light curves often feature quasi-periodic oscillations, which are grouped into three types; A, B, and C. For the first time, I identified a direct transition between QPO types A and C that does not involve a type B QPO in between. I found that transitions between flip-flop states only occur at multiples of a fundamental period of 2.8 ks, which provides a significant constraint for models seeking to describe this phenomenon. By fitting the X-ray spectra in the bright and dim states, I found that the 77% difference in source flux is mainly caused by a change in the temperature of the inner accretion disc.

Chapter 1

Introduction

1.1 Black holes across the mass range

1.1.1 Black holes

Black holes (BHs) are the most extreme astrophysical objects. They form when an extremely powerful gravitational field overcomes the repulsive nuclear forces which arise from the Pauli exclusion principle (Pauli, 1925). In such instances, matter collapses further, presumably down to a point, or ring, the so-called singularity (Penrose, 1965). Around this singularity exists a mathematical boundary known as the event horizon. No information can be passed from inside the event horizon to outside observers. All worldlines of all photons emitted from within the event horizon travel invariably towards the singularity.

Position and linear momentum aside, the no-hair-theorem of BHs (Israel, 1967) states that they can be completely described with just three independent parameters; their mass, spin, and charge. The charge is probably essentially zero in most real-world instances and can be mostly ignored from an observational standpoint. Nevertheless, the mass and spin are very relevant parameters for analysing and understanding the zoo of all observed BHs.

Crucially, the size and shape of the event horizon of a BH are entirely defined by its mass and spin. A non-rotating Schwarzschild BH has a spherical event horizon with a radius of:

$$r_S = \frac{2GM}{c^2}, \tag{1.1}$$

where r_S is the so-called Schwarzschild radius (Schwarzschild, 1916), G is the gravitational constant, M is the BH mass, and c is the speed of light. Relatively speaking, this is an exceptionally small size. For instance, a non-rotating 10 $\mathrm{M}_\odot$ BH has a Schwarzschild radius of merely 29.5 km. This should be considered in comparison to the size of some supergiant stars of similar mass, which can have radii of $\approx 10^{11}$ m.

The second key parameter relevant to the study of BHs, is the spin parameter, a, which is described by the equation:

$$a = \frac{Jc}{GM^2}, \tag{1.2}$$

where J is the angular momentum. It has a maximum value, which limits the spin parameter to lie in the interval $0 \leq a < 1$. The spin parameter is often observed close to its upper limit. For instance, Zhao et al. (2021) have observed that the spin parameter of Cygnus X-1 is $a > 0.9985$ at a 3σ level of confidence. In contrast, Gou et al. (2010) found that A0620–00 has a significantly lower spin parameter of merely 0.12 ± 0.19. The observational study of BHs has focused more on measuring their masses than their spins. This is mainly because the spin is more challenging to determine. However, BHs with the same mass but different spins will affect their surroundings in substantially different ways. A rotating Kerr BH can induce frame dragging of the spacetime in its vicinity, known as the Lense-Thirring effect. This does not occur for non-rotating Schwarzschild BHs. The innermost stable circular orbit (ISCO) for matter around a non-rotating BH is located at a radius of:

$$r_{ISCO,a=0} = \frac{6GM}{c^2} = 3r_S. \tag{1.3}$$

In contrast, a BH approaching maximal spin ($a \to 1$) has an ISCO for matter rotating in the same direction, located at a radius approaching:

$$r_{ISCO,a\to 1} = \frac{GM}{c^2} = \frac{1}{2}r_S. \tag{1.4}$$

The smaller the ISCO is, the more gravitational potential energy is released by matter falling into the BH. Orbits with radii smaller than r_{ISCO} are unstable, and matter invariably falls into the BH without requiring any other interaction to reduce its angular momentum. Unlike neutron stars, BHs have no hard surface, so no kinetic energy is released when matter collides with the BH. The amount of gravitational potential energy (E_g) released when a test mass m falls into a BH is quantified by the efficiency η as:

$$\eta = \frac{E_g}{mc^2}. \tag{1.5}$$

For a non-rotating, $a = 0$ BH, $\eta = 0.057$. That is already 15 times greater than the efficiency at which nuclear binding energy is released in the nuclear fusion of deuterium and tritium to helium; $\eta = 0.00375$. In contrast, for matter falling onto a maximally rotating $a \to 1$ Kerr BH, if it spins in the same direction as the BH, $\eta = 0.432$ (Kovács et al., 2011). Therefore, the accretion of matter onto BHs is the single most efficient known mechanism for obtaining energy per unit mass.

BHs sit at the boundary of known physics, at the tension between general relativity and quantum mechanics. Consequently, an accurate treatment of their properties requires the use of both theories. As such, research into their various behaviours is of crucial importance to build toward a better understanding of the physical world.

1.1.2 Stellar-mass black holes

There are two distinct groups of BHs relevant to observational astronomy. They are distinguished by their masses; there are stellar-mass BHs, and supermassive black holes (SMBH).

Stellar-mass BHs typically have masses in the $5-30\ \mathrm{M}_\odot$ range. The lowest mass stellar-mass BH found so far is located in the binary system 2MASS J05215658+4359220 and has been observed to have a mass of merely $3.3^{+2.8}_{-0.7}\ \mathrm{M}_\odot$ (Thompson et al., 2019). Stellar-mass BHs are formed when a massive star, with a mass of $\gtrsim 20\ \mathrm{M}_\odot$, reaches the end of its lifetime, when its core has fused lighter atomic nuclei all the way to iron, at which point no more energy can be obtained through nuclear fusion (Fryer, 1999). The lack of an internal power source whose radiation pressure maintains the temporary stability of the star against the pull of gravity is lost, which causes the core to contract (Oppenheimer and Snyder, 1939). This collapse releases on the order of $\approx 10^{51}$ ergs of gravitational potential energy, which triggers a Type II supernova explosion (Kasen and Woosley, 2009). The collapsed core exceeds the Tolman–Oppenheimer–Volkoff limit (Tolman, 1939; Oppenheimer and Volkoff, 1939) and becomes a BH. Even more massive stars, with masses of $\gtrsim 40-50\ \mathrm{M}_\odot$, are thought to collapse directly to BHs at the end of their lifetimes, without any supernova explosion.

It is estimated that there are about 10^8 stellar-mass BHs in the Milky Way alone, as about 0.1% of stars are sufficiently massive to collapse to BHs at the end of their lifetimes (Shapiro and Teukolsky, 1983). The number of known stellar-mass BHs is, however, substantially smaller.

Most of the known stellar-mass BHs are found in close binary systems with a companion star. This is because it is easiest to detect BHs if they have a companion star. These systems are known as X-ray binaries (XRBs) due to their large X-ray luminosities. However, not all of those contain BHs, as some have a neutron star as their compact object. We currently know of 59 such XRBs (Corral-Santana et al., 2016). It is, however, also estimated that there are a total of ≈ 1300 XRBs in the Milky Way, with the vast majority remaining undetected so far (Corral-Santana et al., 2016).

Many BHs may be in wide binaries, in which there is a negligible mass transfer from the companion to the BH. X-ray emission in such systems comes exclusively from the stellar corona rather than from an accretion disc around the BH. Detecting BHs in wide binaries is much more challenging. Their presence can, however, be inferred via the motion of the companion star. The orbit of the companion star is also used to constrain the mass of the BH. For instance, Liu et al. (2019) observed a sinusoidal variation in the radial velocity of a star, which was caused by a companion with a mass of $68^{+11}_{-13}\ \mathrm{M}_\odot$, that did not emit any light. A dark, compact object with that mass can only be a BH. Yamaguchi et al. (2018) also argue that BHs in such systems could be detected via precise astrometric observations of the orbit of stars around massive, dark, compact objects.

A breakthrough in BH astrophysics was achieved with the advent of gravitational wave detection capabilities (LIGO Scientific Collaboration and Virgo Collaboration, 2016). The inspiral ahead of the merger of two massive, compact objects produces powerful gravitational waves, which can be detected over cosmological distances, allowing us to detect stellar-mass BHs in distant galaxies. Furthermore, the shape of the gravitational wave signal enabled the masses of the merging, and the merged BHs to be determined. LIGO and Virgo have already detected a total of 90 gravitational wave events (The LIGO Scientific Collaboration et al., 2021), which were mainly produced by the inspiral of binary BHs.

All of these methods can only detect stellar-mass BHs located in a binary system with a star, neutron star, or BH. However, it is possible that many stellar-mass BHs do not have any companion (Shapiro and Teukolsky, 1983). During the red giant phase of the stellar evolution of massive stars, they commonly grow to envelop any companion star, which often leads to a merger (Ivanova et al., 2013) before the possible collapse to a BH. Asymmetries in supernova explosions also create natal kicks (Janka, 2013; Belczynski et al., 2016), which can eject the newly formed BH from the stellar system at high velocities.

These isolated BHs are exceedingly difficult to detect. They are tiny, having radii of merely a few km, and do not emit any light of their own. So far, the only known way in which such isolated BHs can be found is via the microlensing method (Karolinski and Zhu, 2020). When a BH passes into the line of sight of a distant bright source, gravitational lensing from the BH causes the distant source to be magnified. This is manifested as a brief, sharp increase in the measured flux following a standard temporal evolution. By determining the mass of the lens and identifying whether it emits any light of its own, isolated BHs can be found. So far, there has only been one unambiguous detection of an isolated BH with microlensing (Sahu et al., 2022).

1.1.3 Black hole transients

In XRBs, the separation of the two components is sufficiently small, such that the companion star exceeds its Roche lobe at least some of the time, resulting in mass transfer from the star to the BH (Savonije, 1979). Due to the small size of the BH and the excess angular momentum of the transferred matter, it cannot fall directly onto the BH. It instead forms an accretion disc around it, in which friction transfers angular momentum outwards and increases the temperature of matter falling into ever tighter orbits around the BH (Abramowicz and Fragile, 2013). In this process, accreting matter loses gravitational potential energy, corresponding to a significant fraction of its rest mass energy (η). This energy release heats up the accretion disc to immense temperatures on the order of 10^7 K (Shakura and Sunyaev, 1973). The accretion disc radiates as a black body, having a peak in its emitted energy distribution of several hundred eV, firmly within the X-ray energy range. Additionally, Compton up-scattering in the corona (Rees et al., 1982; Narayan et al., 1996) increases the photon energy even further, producing a non-thermal power law energy distribution. These are the physical mechanisms responsible for the large X-ray luminosities observed in XRBs, giving them their name. XRBs also generate powerful radio jets perpendicular to the plane of the accretion disc (Fender et al., 2009). The activity of the jets is very strongly related to the activity in the accretion disc itself, and these two components evolve together (Fender, 2010).

XRBs can either be persistent sources that are bright all the time or transient sources that have long periods of quiescence between outbursts, in which they shine brightly in X-rays (Belloni et al., 2011). Transient XRBs with a BH as their compact object are called black hole transients (BHTs).

BHTs spend most of their time in a quiescent state, in which their X-ray and radio emission is negligible, and there is no mass transfer from companion to BH. However, typ-

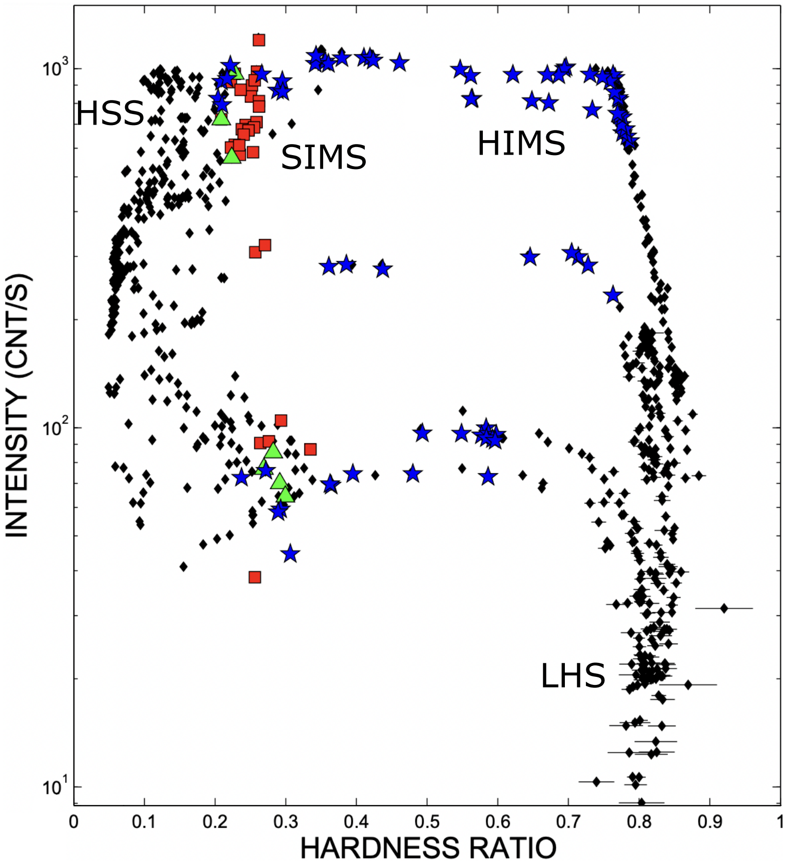

Figure 1.1: The HID of the 2002, 2004, 2007, and 2010 outbursts of the low-mass XRB GX 339-4. Adapted from Belloni et al. (2011).

ically once every few years, the companion star expands beyond its Roche lobe, triggering an outburst of the BHT. During an outburst, the accretion rate onto the BH and the X-ray source flux increase by many orders of magnitude (see e.g. Chen et al., 1997; Remillard and McClintock, 2006). Outbursts typically last for a few months, with the BHT passing through several well-defined states (Belloni et al., 2011). These are distinguished and classified via their X-ray intensity, spectral, and timing properties (Homan and Belloni, 2005). These different states can be visualised on a hardness intensity diagram (HID), which shows the X-ray intensity against the flux ratio between a hard and a soft X-ray energy band. They typically evolve through a hysteresis 'q'-shaped path traversed in an anti-clockwise fashion (Fender et al., 2004). Fig. 1.1 shows the HID of multiple outbursts of the XRB GX 339-4, with labels at the location of specific states (Belloni et al., 2011).

At the start and end of the outburst, the system is located in the lower right region of the HID. It has a low luminosity and a hard X-ray spectrum, which is dominated by a power law component produced via Compton up-scattering of soft thermal X-ray photons from the accretion disc by high energy electrons in the corona (Gilfanov, 2010). In the standard truncated disc model (Esin et al., 1997; Poutanen et al., 1997), the geometrically thin, optically thick accretion disc is terminated at a large radius, so that the disc black body component only makes a small contribution to the X-ray spectrum. The system also features a steady, compact jet (see e.g. Corbel et al., 2001). This is known as the low hard state (LHS).

As the accretion rate and luminosity increase, the truncation radius of the accretion disc starts to decrease. This causes the disc emission to become stronger and the X-ray spectrum to be characterised by a combination of a multi-temperature black body and a power law component. This part of the outburst is classified into two types: the hard intermediate and soft intermediate states (HIMS and SIMS, respectively), which can, for instance, be distinguished by their variability properties. The transition from the HIMS to the SIMS is marked by the steady compact jet switching off, being quenched by more than 3.5 orders of magnitude (Russell et al., 2019) after a short-lived, bright and rapidly flaring transient jet is launched (Fender et al., 2004). Sometimes, a BHT temporarily reaches an anomalous state (AS; Motta et al., 2012), which branches off from the HID at the SIMS towards even higher luminosities and greater X-ray hardness.

The truncation radius continues to decrease as the luminosity starts to drop slightly. When the accretion disc extends down to the ISCO, the spectrum is dominated by black body emission from the accretion disc, whereas the Comptonised power law component becomes relatively insignificant. This part of the outburst is characterised as the high soft state (HSS), during which no jet is launched (Fender et al., 1999). Eventually, after the source has faded enough, the hardness increases again, and the BHT returns to the LHS, via the SIMS and HIMS, before finally returning to quiescence. This hardening in the intermediate states occurs at a lower luminosity than the initial softening and features the re-establishment of the compact jet over a few weeks (Russell et al., 2014)). Multiple transitions between the HIMS and the SIMS usually occur in both of these horizontal regions of the HID.

The above describes the standard picture of the evolution of a BHT in an outburst, but it is by no means universal. For example, some outbursts fail to reach some of the softer states before returning to the LHS and then to quiescence (Belloni and Motta, 2016).

1.1.4 Supermassive black holes and active galactic nuclei

The second prominent class of BHs at the other end of the mass scale are the SMBHs, with masses typically found in the range $10^6 - 10^{10}$ $\mathrm{M}_\odot$. The most massive BH known so far is Ton 618, with a mass of 6.6×10^{10} $\mathrm{M}_\odot$ (Shemmer et al., 2004). SMBHs are found at the centre of almost every large galaxy. Their formation mechanism is still unknown, with several possibilities being debated. However, high redshift observations of SMBHs indicate that they must have formed very early, within the first 10^9 years after the Big Bang (Fan, 2006). Even though SMBHs only make up a minuscule portion of the entire mass of the galaxies they reside in, they can have profound effects on their host galaxies. The BH mass was observed to strongly correlate with both the stellar velocity dispersion (Ferrarese and Merritt, 2000) and the total mass (Magorrian et al., 1998) of the bulges of their host galaxies. These correlations between the BH mass and the galaxy properties are thought to be caused by AGN feedback comprising wind outflows and jets from the SMBH system (Kormendy and Ho, 2013). By inverting the correlations, it is possible to estimate the BH mass from a measurement of the stellar velocity dispersion or the mass of the bulge of a galaxy.

SMBHs are thought to grow considerably in mass over cosmic time but at a highly non-uniform rate. They undergo periods of very high accretion rates that are estimated to last for $\approx 10^{7-9}$ years, separated by long periods of quiescence when their accretion rate is negligible (Marconi et al., 2004). In periods with a high accretion rate onto SMBHs, the accreting matter releases so much gravitational potential energy in the process that it can cause the centre of the galaxy to outshine the rest of the galaxy across the electromagnetic spectrum (Masoura et al., 2021). These sources have been observed with a variety of different properties in different wavelength bands and are grouped into the broad category of Active Galactic Nuclei (AGNs). AGNs are the brightest persistent sources in the universe, having a power of up to several 10^{48} erg s^{-1} (Padovani et al., 2017).

AGNs were more common in the past than they are in the universe now; the luminosity density of AGNs was observed to be greatest at $z = 1.2 \pm 0.1$ (Aird et al., 2010). Most of the growth of SMBHs through cosmic time occurred at $z > 1$ (Aird et al., 2010), presumably via several periods of AGN activity. Due to their large luminosities, AGNs can also be detected at high redshifts. In recent years, several AGNs with $z > 6$ have been identified (see e.g. Venemans et al., 2015; Wolf et al., 2021). This constrains the time of formation of the SMBHs, indicating that they must have formed in the very early universe.

In the AGN unification model (Antonucci, 1993), various AGN classes, such as radio galaxies, Seyfert galaxies, quasars, blazars, and BL Lacs are all described to be caused by accretion onto a SMBH in the centre of the galaxy, featuring the same basic structure, and geometry. The difference between these different classes of sources is explained as being due to a difference in the viewing angle relative to the plane of the accretion disc. In this model, the structure of an AGN in the region close to the event horizon of the SMBH consists of a geometrically thin, optically thick accretion disc, which is typically described by the Shakura-Sunyaev α-disc model (Shakura and Sunyaev, 1973). A corona and two jets perpendicular to the accretion disc are also located close to the SMBH. The broad line region is elevated slightly above and below the accretion disc, which produces strongly Doppler-broadened emission lines due to the high velocities of orbits at these radii (Peterson, 2006). Further out, there is a geometrically thick dusty torus, which can block the view of the inner parts of the structure if viewed side-on (Nenkova et al., 2008). Finally, there is a narrow line region, which generates significantly narrower Doppler-broadened emission lines because it is located further away from the centre of the BH (Heckman et al., 1981). However, there has been some recent criticism of the AGN unification model, arguing that it is a too simplified description and cannot account for changing look AGNs, which are observed to change from type 1 to type 2, or vice-versa (Spinoglio and Fernández-Ontiveros, 2021).

A recent major development in this field is the ability to image SMBHs. The Event Horizon Telescope (*EHT*; Event Horizon Telescope Collaboration, 2019) combines observations from radio telescopes around the world, utilising very long baseline interferometry to obtain an angular resolution of merely 24 μarcsec (Event Horizon Telescope Collaboration, 2022). The *EHT* has now imaged both the SMBH at the centre of giant elliptical galaxy M87 (Event Horizon Telescope Collaboration, 2019), and Sagittarius A*, the SMBH at the centre of the Milky Way (Event Horizon Telescope Collaboration, 2022). Fig. 1.2 depicts the images obtained by the EHT of these two SMBHs.

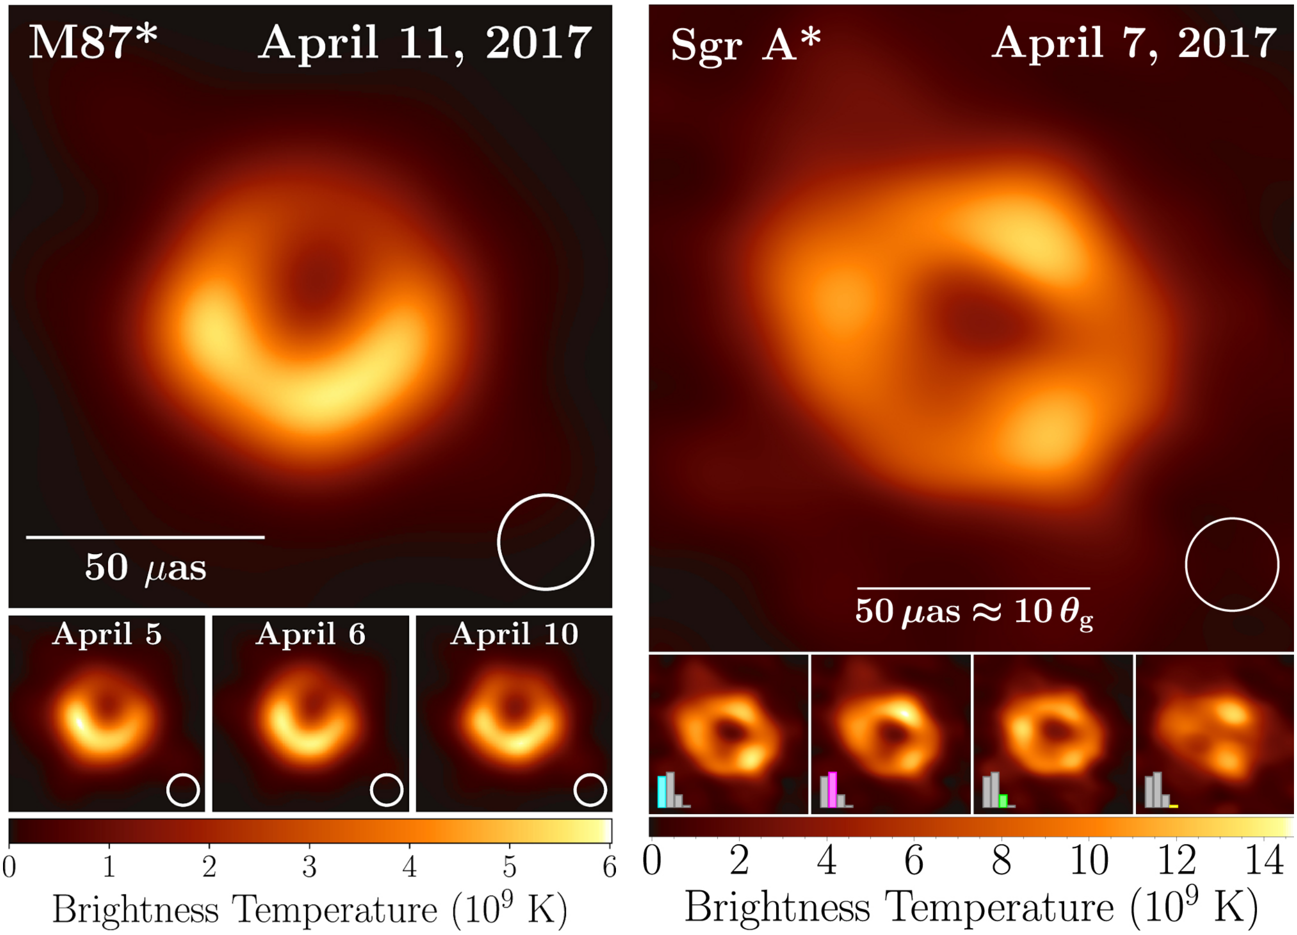

Figure 1.2: The images of the SMBH at the centre of galaxy M87 (left, credit: Event Horizon Telescope Collaboration, 2019), and Sagittarius A* at the centre of the Milky Way, as observed by the EHT (right, credit: Event Horizon Telescope Collaboration, 2022).

There is a large mass gap between the stellar-mass BHs and the SMBHs, indicating that they must have been formed by different mechanisms. BHs that lie in this mass gap, having masses in the range $10^2 - 10^5$ $M_\odot$, are labelled as intermediate-mass black holes (IMBHs). Most of the IMBHs identified so far are detected at the upper limit of this mass range and are associated with low luminosity AGNs (Chilingarian et al., 2018). However, only one IMBH is known at the lower end of the mass range. It has a mass of 142 $M_\odot$ and was detected to have been formed from the merger of two very massive stellar-mass BHs (LIGO Scientific Collaboration and Virgo Collaboration, 2020). Currently, the IMBHs that fall into the two categories of low luminosity AGNs, or stellar BH mergers, are still separated by more than 2 orders of magnitude in mass. This could indicate that the known IMBHs are merely the extreme mass limits of stellar-mass BHs and SMBHs, which fall slightly outside of the standard range of masses considered for these two types of BHs. This further supports the idea that SMBHs formed from a massive seed rather than from countless mergers of stellar-mass BHs.

Although the mass and size of BHTs and AGNs are vastly different, they feature many similarities. They are both powered by accretion onto a BH, feature an accretion disc, corona, and radio jets, and are defined mainly by irregular, comparatively short periods of high luminosity between long periods of quiescence. Additionally, it was found that the accretion process, the timescales of changes in the accretion process, and the types of variability involved could be matched when accounting for the different masses of the two classes of systems (McHardy et al., 2006). This similarity means that research into the properties of BHTs may also be helpful for understanding AGNs and vice versa.

1.2 Black hole variability properties

1.2.1 Shortest variability timescale

The time it takes an astronomical object to change its brightness indicates the size of its light-emitting region. For continuously variable objects, the information about a change in brightness has to propagate within the source to take effect. Therefore, a preliminary estimate of the upper bound of the size of a continuously variable source can be made by equating the shortest timescale of observed changes to half of the light-crossing timescale:

$$r_{max} \approx 2c\tau. \tag{1.6}$$

In this equation, r_{max} is the maximum size that the variable object can have, c is the speed of light, and τ is the shortest timescale over which continuous brightness changes could be detected. The factor of 2 is accurate, if the source is spherical or cylindrical with a much smaller height than radius, and if the information of a change in brightness is propagated radially from the centre through a lossless medium with a refractive index of 1. It is, therefore, a theoretical upper limit. Equation 1.6 also does not apply if there are multiple emission regions or if the source is inhomogeneous.

This relation between the timescale of variability and the upper limit on the size of its source has historical significance because it helped determine the true nature of AGNs. Despite the extreme luminosities that they were observed to possess, they were found to vary significantly over time. This indicates that they are compact and have radii below the pc scale (Matthews and Sandage, 1963).

AGNs are observed to vary at all wavelengths (Ulrich et al., 1997). As X-rays are generated near the innermost part of the accretion disc, close to the ISCO, AGN X-ray light curves vary on the shortest timescales (Nandra et al., 1997; Crenshaw et al., 1996).

In this environment, the extreme gravitational field and the interaction of in-falling matter and radiation produce a variety of unique effects, producing various kinds of variability, many of which still need to be properly understood.

If the size of a variable source is known, Equation 1.6 can be inverted to obtain an estimate of the smallest timescale over which it can show significant variability. For example, a 10 $M_\odot$ BH has a minimum X-ray variability timescale on the order of 10^{-3} s. In contrast, a 10^{10} $M_\odot$ SMBH has an X-ray variability timescale of at least on the order of 10^{6} s (Paolillo et al., 2004).

Previous variability studies of XRBs have considered high-frequency variability (see e.g. Belloni et al., 2012; Belloni and Stella, 2014), at hundreds of Hz, probing the structure very close to the ISCO. Such studies are often limited by instrument sensitivity and the brightness of the source. At the other end of the range, long-term variability studies often involve comparisons of observations made by a variety of different instruments over many years (see e.g. Zheng et al., 2017; Paolillo et al., 2017). Such studies are limited by differences between detectors. Moreover, they can only describe differences in source brightness in narrow observing windows separated by very long gaps, missing all the variability occurring in-between. An analysis of quasar light echoes has even enabled an analysis of AGN variability on a timescale of $> 10^{4}$ years (Lintott et al., 2009). Sartori et al. (2018) have combined multiple methodologies to investigate the optical AGN variability on timescales ranging from several days to tens of thousands of years.

1.2.2 Variability correlations

The normalised excess variance (NEV; Edelson et al., 1990; Nandra et al., 1997) is a parameter that quantifies the degree of variability of a source within a set of measurements in excess of the measurement errors. It will be discussed in more detail in Chapter 3. When comparing the NEV measured from the X-ray light curves of several different AGNs, Nandra et al. (1997) found that it is anticorrelated with the AGN luminosity. Following this discovery, various other correlations with the variability of AGNs were identified. Lu and Yu (2001) found an anticorrelation between the NEV and the BH mass of AGNs. This was further supported by the analysis of a greater number of AGNs, by Papadakis (2004), O'Neill et al. (2005), Ponti et al. (2012b), and Paolillo et al. (2017), among others. Furthermore, Turner et al. (1999) also discovered a positive correlation between the NEV and the photon index of the X-ray spectrum.

O'Neill et al. (2005) argue that the anticorrelation between the NEV and the luminosity, and the correlation between the NEV and the photon index, are merely consequences of the anticorrelation between the NEV and the BH mass. There is particular interest in this relationship due to the possibility of estimating the BH mass from a measurement of the NEV. This could be useful when other methods of measuring BH mass are unavailable, such as in obscured AGNs.

1.2.3 AGN PSDs

The Power Spectral Density (PSD; van der Klis, 1989) describes the distribution of variability power as a function of frequency. A periodogram is an estimate of the PSD of a variable source based on the available data for it. This will be described in more detail in Chapter 3.

The periodograms of variable sources are often dominated by a power law shape: $P(\nu) \propto \nu^{-\alpha}$, where α is the power law index. A power law PSD with $\alpha = 0$ is known as white noise and corresponds to a light curve in which the count rate in every bin is independent of the count rate in any other bin. A red noise PSD is described by a power law with $\alpha = 2$ and is associated with a light curve dominated by long-term trends, in which the count rate in each bin is strongly correlated to that in adjacent bins. A pink noise PSD lies between the two, has $\alpha = 1$, and corresponds to a light curve that is still dominated by long-term trends but in which the correlation between the count rates of adjacent bins is smaller than for red noise.

AGN and XRB X-ray light curves are typically observed to have PSDs, which can be described by a red noise power law of $\alpha \approx 2$ at high frequencies (González-Martín and Vaughan, 2012), and a pink noise power law of $\alpha \approx 1$ at lower frequencies (Papadakis et al., 2002; Papadakis, 2004). The transition from $\alpha \approx 2$ to $\alpha \approx 1$ is usually described by a sharp break but has also been modelled as a gradual bend in the PSD (González-Martín and Vaughan, 2012). This break or bend occurs somewhere between about $10^{-6.4} - 10^{-3.3}$ Hz (González-Martín and Vaughan, 2012). To avoid a divergence in the integral of the PSD, it is also necessary for $\alpha < 1$ at very low frequencies. For XRBs, it is observed that $\alpha = 0$ at very low frequencies (Belloni and Hasinger, 1990). This has not yet been observed for AGNs.

Similar power laws have been identified in the PSDs of XRBs, at much lower frequencies. For those, a low-frequency break from $\alpha \approx 1$ to $\alpha \approx 0$ can be observed (Belloni and Hasinger, 1990) as well. XRB PSDs can also be used to investigate the properties of quasi-periodic oscillations, which are frequently observed to occur.

1.2.4 Quasi-periodic oscillations

Temporal analysis of the X-ray light curves of many BHTs has revealed the existence of strong stochastic noise, along with quasi-periodic oscillations (QPOs). These appear as Lorentzian peaks in periodograms. Greater variability and QPO amplitude at higher X-ray energies suggest that the QPOs originate very close to the BH: in the corona or the inner accretion flow (Belloni et al., 1997). Accordingly, QPOs provide an additional useful window for studying this region and are essential in distinguishing different parts of an outburst. They come in two classes: low (LFQPOs), and high-frequency QPOs (HFQPOs) (Motta, 2016). The LFQPOs have frequencies in the range 0.01–30 Hz, but are typically found below 10 Hz, and have been detected and studied much more than HFQPOs, which commonly have frequencies of several hundreds of Hz. The LFQPOs have themselves been grouped into three different types: A, B, and C (Wijnands et al. 1999, and for a detailed analysis, see Casella et al. 2005, and Motta 2016). The three types are distinguished via their PSD properties. The PSDs of BHTs are often fitted with a sum of several Lorentzian functions (Belloni et al., 2002), which describe both the QPO and the stochastic noise continuum:

$$L(\nu) = \frac{r^2\nu_0}{2\pi Q}\left((\nu-\nu_0)^2 + \left(\frac{\nu_0}{2Q}\right)^2\right)^{-1}.$$

Here, ν_0 is the centroid frequency at which the Lorentzian reaches its maximum; $Q = \nu_0/\mathrm{FWHM}$ is the quality factor, which is used to characterise the width of the Lorentzian; r describes the rms variability ($r = \sqrt{\mathrm{NEV}}$) of a single Lorentzian. The total rms of a PSD is determined by taking the square root of its properly normalised integral between two frequencies: ν_1, and ν_2, and is typically expressed as a percentage: $\mathrm{rms} = \sqrt{\int_{\nu_1}^{\nu_2}\left(\sum_i^n L_i(\nu)\right)d\nu}$, where n is the number of Lorentzian components used in the fit.

Type C QPOs are observed most frequently and are usually detected in the LHS and HIMS, although they can appear in all spectral-timing states (Motta, 2016). They have frequencies spanning the entire LFQPO frequency range: 0.01 Hz $\lesssim \nu_0 \leq$ 30 Hz, are very narrow: $Q \gtrsim 10$, and have a large rms variability: $5\% \lesssim \mathrm{rms} \lesssim 20\%$. The strong broad-band continuum associated with type C QPOs contributes significantly toward their large rms. The continuum appears as a flat top noise extending up to a break frequency, above which it drops steeply. The break frequency has a similar value to, and correlates with, the QPO frequency (Wijnands and van der Klis, 1999). The type C frequency also roughly correlates with the X-ray flux and anticorrelates with the broad-band rms (Motta et al., 2011).

Type B QPOs are the defining characteristic of the SIMS, distinguishing it from the HIMS. (Belloni and Motta, 2016). They are typically located at $\nu_0 \approx 6$ Hz in the early bright phase of the outburst but have also been found at lower frequencies, especially in the low flux horizontal branch traversed in the return to the LHS (Motta et al., 2011). They are also quite narrow: $Q \gtrsim 6$, and have low variability amplitude: $\mathrm{rms} \lesssim 4\%$. The rms is predominantly due to the QPO peak, with a comparatively small contribution from the power law continuum noise (Motta, 2016).

Type A QPOs are observed even less frequently than type Bs, and are the hardest to detect. This is due to their broad and very shallow peak, with $Q \lesssim 3$, at frequencies of 6 Hz $\lesssim \nu_0 \lesssim$ 8 Hz. In addition, type A QPOs have the lowest variability amplitude of $\mathrm{rms} \lesssim 3\%$, a result of the weak QPO and a small contribution from the continuum noise. Type A QPOs do not feature harmonics, unlike types B and C. They are commonly found in the HSS (Homan and Belloni, 2005).

Despite the considerable analysis of LFQPOs, there has yet to be a consensus regarding their physical origin. It is thought that QPOs originate close to the BH, but LFQPO frequencies are substantially smaller than the Keplerian orbital frequencies at these radii, which are several hundred Hz. Motta et al. (2011) and Motta et al. (2012) argue that QPO types B and C have different physical origins, whereas types A and C could be due to the same process. They point to the different correlations between QPO frequency and power law flux of the three QPO types, along with the observation of simultaneous type B and

C QPOs in the PSD of GRO J1655-40.

The relativistic precession model (Stella and Vietri, 1998) identifies the centroid frequency of LFQPOs with the frequency of Lense-Thirring precession at a single radius in the accretion disc. Within this model, pairs of HFQPOs are identified as corresponding to periastron precession and Keplerian frequency at the same radius (Stella et al., 1999).

Ingram et al. (2009), Ingram and van der Klis (2013), and Ingram et al. (2016) describe the precessing inner flow model. In this model, the entire inner flow undergoes Lense-Thirring precession as a rigid body. A variation in the inclination of the inner flow leads to a modulation of the Doppler boost, which subsequently generates the QPO.

The accretion ejection instability model (Tagger and Pellat, 1999; Varnière et al., 2002; Varnière and Tagger, 2002; Varniere and Vincent, 2016) hypothesises an instability with which energy and angular momentum are transported from a magnetised accretion disc to the corona, via a spiral density wave and a Rossby vortex. In this model, QPO types A, B, and C are distinguished by being produced in a relativistic, semi-relativistic, and non-relativistic regime of the instability, respectively.

Another model for the QPOs is the transition layer model (Titarchuk and Fiorito, 2004). In this model, a transition layer separates the outer parts of the accretion disc, which orbit at a Keplerian frequency, from the inner regions, which have a sub-Keplerian orbital frequency. The orbital frequency at this transition layer would equal the QPO frequency.

A fourth hypothesis was developed by Cabanac et al. (2010). It considers a magneto-acoustic wave in the corona, which causes a variation in the efficiency of Comptonisation, thereby generating a QPO.

In summary, QPO properties provide a way to ascertain the physical processes occurring in the vicinity of the BH and are an effective way of classifying the state of the source, providing information that the flux and energy spectra alone cannot supply. However, the physical origin of the QPOs is still debated. While observations suggest a clear distinction between QPO types A, B, and C, it remains unknown whether these are the result of different physical processes or merely different regimes of the same fundamental mechanism.

1.3 X-ray telescopes

A large part of this thesis is dedicated to the analysis of *eROSITA* data. Nevertheless, we also used a variety of other telescopes and instruments to observe variable X-ray sources. We will describe their properties in this, and the following two sections.

As the Earth's atmosphere absorbs X-rays, telescopes sensitive to these energies need to be in space to perform observations of astronomical sources. Unlike at lower energies, X-ray photons cannot easily be reflected or refracted. Compton (1923), however, found that X-rays can be reflected, but only at small glancing angles of $\lesssim 1^\circ$ off of a polished surface. Jentzsch and Nähring (1929) concluded that a pair of mirrors with a different curvature is required to focus X-rays.

These discoveries led to the design of Wolter telescopes (Wolter, 1952). Typically, these feature a combination of hyperboloid and paraboloid mirrors to bring X-rays to a focus on the detector. This pair of conical mirrors enables the telescope to focus X-rays from a wider field of view. However, the limitation on the glancing angle for reflection implies a large focal length. To maximise the photon collection area, X-ray telescopes typically use several nested shells of pairs of mirrors.

1.3.1 *Astrosat*

Astrosat is a telescope sensitive to a wide range of the electromagnetic spectrum, with separate detectors for X-rays, ultraviolet, and optical photons (Singh et al., 2014). It was launched into a low Earth orbit on September 28, 2015. It has now exceeded its planned mission duration of 5 years and is continuing to perform observations.

Astrosat has three Large Area Xenon Proportional Counters (LAXPCs), which are its main X-ray detectors. They are sensitive to photon energies in the range $3 - 80$ keV. The field of view of the LAXPCs is $1^\circ \times 1^\circ$. The LAXPC detectors excel due to their short time resolution of 10 μs (Antia et al., 2017). This resolution is, however, accompanied by a dead time of 42 μs, in the default event analysis mode. In the fast counter mode, the dead time is even reduced to 10 μs, which is useful for observing bright, or flaring sources. The LAXPCs have a combined large effective area of ≈ 6000 cm^2 at 20 keV. (Antia et al., 2017). This effective area, however, drops to ≈ 2500 cm^2 at 20 keV (Basu et al., 2021). These properties enable the LAXPCs to be ideal instruments for timing and variability studies of bright sources in X-rays, especially for high-frequency modulations.

1.3.2 *Chandra*

The *Chandra X-ray observatory* (*Chandra*; Weisskopf et al., 2000) is the longest running, still active X-ray telescope. It was launched into a highly elliptical geocentric orbit on July 23, 1999, and was the heaviest payload ever launched by a space shuttle.

Chandra is sensitive to the X-ray photons in the energy range: $0.2 - 10.0$ keV. It specialises in X-ray imaging, and its optics have an on-axis angular resolution of $0.5''$ (Murray et al., 2000). *Chandra* has a high-resolution camera and three different spectrometers; the advanced CCD imaging spectrometer (ACIS) and the high-, and low-energy transmission grating spectrometers. Below 3 keV, *Chandra* has a resolving power of up to 500 (Weisskopf et al., 2000). ACIS has a time resolution of 16 μs.

1.3.3 *Insight*-HXMT

The *Hard X-ray Modulation Telescope* (*Insight*-HXMT; Zhang et al., 2018) is China's first astronomical satellite. It was constructed with the aim of improving the hard X-ray survey capabilities, especially regarding observations of transient sources, XRBs, and gamma-ray bursts. It was launched into low Earth orbit on June 14, 2017.

Insight-HXMT consists of three X-ray telescopes that are sensitive to three different energy bands; the High Energy X-ray telescope (HE, $20 - 250$ keV), the Medium Energy

X-ray telescope (ME, $5 - 30$ keV), and the Low Energy X-ray telescope (LE, $1 - 15$ keV). Of them, HE has the best time resolution of 25 μs. ME has a time resolution of 280 mus. HE also has a large effective area at this high energy range, of 5000 cm^2 below 80 keV (Zhang et al., 2018).

1.3.4 *NICER*

The *Neutron star Interior Composition ExploreR* (*NICER*; Gendreau et al., 2012) was designed to investigate neutron stars through soft X-ray timing and spectroscopy. It was launched on June 3, 2017, and was attached to the International Space Station on June 13, 2017.

NICER is sensitive to the energy band $0.2 - 12$ keV, with an effective area of 2000 cm^2 at 1.5 keV (Okajima et al., 2016). It has a short time resolution of 300 ns.

1.3.5 *NuSTAR*

The *Nuclear Spectroscopic Telescope ARray* (*NuSTAR*; Harrison et al., 2013) is an X-ray telescope with sensitivity over the wide energy band of $3 - 79$ keV. It is the first X-ray telescope capable of focusing hard X-rays above 10 keV. *NuSTAR* was launched into a low Earth orbit on June 13, 2012, and is still in operation.

NuSTAR consists of two telescopes, which are labelled as Focal Plane Modules A, and B (FPMA, and FPMB, respectively). Data collected from them should be analysed independently, as there are slight differences between them. *NuSTAR* also features a deployable mast, which was used to increase the distance between the Wolter hyperbolic and parabolic mirrors and the detector to 10.15 m, which is required for the focusing of hard X-rays with high efficiency (Harrison et al., 2005).

NuSTAR has a full width at half maximum energy resolution of 400 eV below 50 keV, which increases to 1.0 keV at the highest energies (Harrison et al., 2013). It has a square field of view, with a width of $12'$, and a time resolution of 2 μs.

1.3.6 *Swift*

The *Neil Gehrels Swift observatory* (*Swift*; Gehrels et al., 2004) was constructed for the purpose of studying gamma-ray bursts. It was launched to low Earth orbit on November 20, 2004. *Swift* consists of three telescopes; the Burst Alert Telescope (BAT), the X-Ray Telescope (XRT), and the UltraViolet/Optical Telescope (UVOT).

Gamma rays are also absorbed by the Earth's atmosphere, which is why telescopes sensitive to these energies also need to be located in space. Gamma rays cannot even be deflected by glancing reflection off of polished surfaces, so gamma-ray telescopes have to abandon focusing for image generation entirely. Instead, gamma-ray telescopes use coded apertures to distinguish sources and determine their approximate positions. These can be reconstructed via the shadows of the coded aperture at the detector. Coded aperture telescopes have a lower angular resolution but a large field of view. *Swift*/BAT has a half-coded field of view of 1.4 sr, allowing it to detect gamma-ray bursts quickly. Following the detection, the observatory slews to point to the source position, such that the XRT and UVOT can observe the afterglow (Barthelmy et al., 2005). *Swift*/BAT performs wide area surveys, to be able to quickly detect gamma-ray bursts wherever they occur. It, therefore, frequently observes bright sources throughout the sky, within the hard X-ray to the gamma-ray energy range of $15 - 150$ keV (Barthelmy et al., 2005).

The *Swift/XRT* is sensitive to X-rays within the energy range $0.2 - 10$ keV. It has a square field of view with a width of $23.6'$ and an effective area of ≥ 120 cm^2 at 1.5keV (Burrows et al., 2005). It has four different readout modes for different types of analysis. The photodiode mode, for instance, has a temporal resolution of 0.14 ms but no spatial resolution at all. In contrast, the image mode has a temporal resolution of $0.1 - 2.5$ s but provides the most accurate measurements of the source position and flux (Hill et al., 2004).

1.3.7 *XMM-Newton*

The *X-ray multi-mirror mission* (*XMM-Newton*; Jansen et al., 2001) was launched to an elliptical geocentric orbit on December 10, 1999, and is the largest scientific satellite ever launched by the European Space Agency. It comprises three telescopes, each consisting of 58 concentric mirror shells. Two of these use a metal oxide semiconductor (MOS) charged coupling device (CCD), and the third uses a pn CCD as their imaging detector (Turner et al., 2001; Strüder et al., 2001). *XMM-Newton* can also determine the X-ray spectrum, using two reflection grating spectrometers (RGS).

XMM-Newton has an effective area of up to ≈ 1400 cm^2 at 1.5 keV (Jansen et al., 2001). It is sensitive to X-rays within the energy range of $0.1 - 10.0$ keV and has a field of view of $30'$. It also has an optical monitor, which detects optical and UV photons in the same direction, and is simultaneous to the primary X-ray observations.

Observations by *XMM-Newton* are performed in one of several different modes. The full frame mode enables observations of the entire available field of view. Other modes obtain information for a smaller region. In the large and small window modes, the different CCDs can be set to collect different types of data. The timing mode optimises the time resolution, improving it to 3×10^{-5} s. This is achieved by discarding the spatial information along one of the two CCD pixel axes (Strüder et al., 2001). The pn CCD even has a burst mode, with a time resolution of 7×10^{-6} s, but only detects 3% of all photons.

1.4 The *eROSITA* instrument

The extended ROentgen Survey with an Imaging Telescope Array (*eROSITA*; Predehl et al., 2021) is a new X-ray telescope that aims to map out the entire sky at an unprecedented sensitivity. For this purpose, *eROSITA* has a large field of view of $1.03°$, and a large on-axis effective area of a circle with a diameter of $\approx 1.5 \times 10^3$ cm^2. It is sensitive to the energy range of $0.2 - 8.0$ keV.

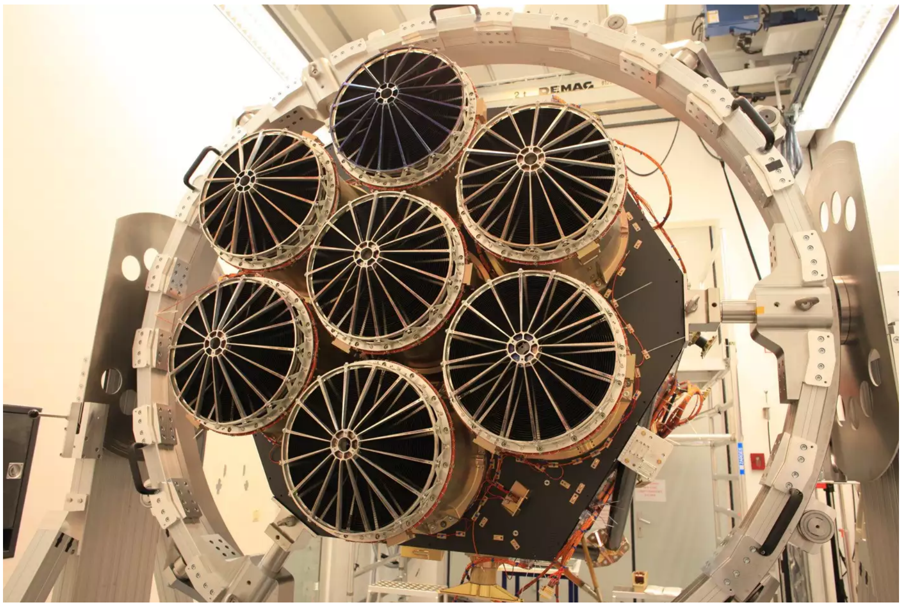

Figure 1.3: An image of the 7 *eROSITA* telescope modules prior to packing and mounting on SRG. Credit: MPE (https://www.mpe.mpg.de/7604815/gallery)

eROSITA follows in the footsteps of the *ROSAT* (an abbreviation of the German word 'ROentgenSATelit') mission (Truemper, 1982), which observed the entire sky in the X-ray energy band $0.1 - 2.4$ keV in the 1990s. *eROSITA* constitutes a considerable improvement over the capabilities of *ROSAT*. The *eROSITA* sensitivity in the soft X-ray band, from $0.2 - 2.4$ keV, exceeds that of ROSAT by 25 times. Additionally, *eROSITA* performs the first ever all-sky survey in the $2.4 - 8.$ keV band (Predehl et al., 2021).

The main aim for *eROSITA* is the detection of $\approx 10^5$ galaxy clusters throughout the sky. This extensive sample will be crucial for testing cosmological models and investigating the large-scale structure in the universe (Predehl et al., 2021). Additionally, it is expected that *eROSITA* will detect $\approx 3 \times 10^6$ AGNs throughout the sky, which will enable a much more detailed investigation of their population properties and evolution throughout cosmic time.

eROSITA is the primary instrument on the Spectrum-Roentgen-Gamma (*SRG*; Sunyaev et al., 2021) spacecraft, which also contains the Astronomical Roentgen Telescope - X-ray Concentrator (*ART-XC*; Pavlinsky et al., 2021) telescope. SRG is a joint German - Russian mission. SRG was successfully launched on July 13, 2019, from the Baikonur cosmodrome in Kazakhstan with a Proton-M rocket. It travelled to the L_2 Lagrange point and entered a 6-month Halo orbit around it on October 21 (Predehl et al., 2021). SRG has performed several orbit corrections since then, but the main properties of the orbit have remained the same.

The *eROSITA* calibration and performance verification phase took place from October 18 to December 8. This included a test of its survey capabilities via the *eROSITA* Final Equatorial Depth Survey (eFEDS; Brunner et al., 2022). This comprised observations of a

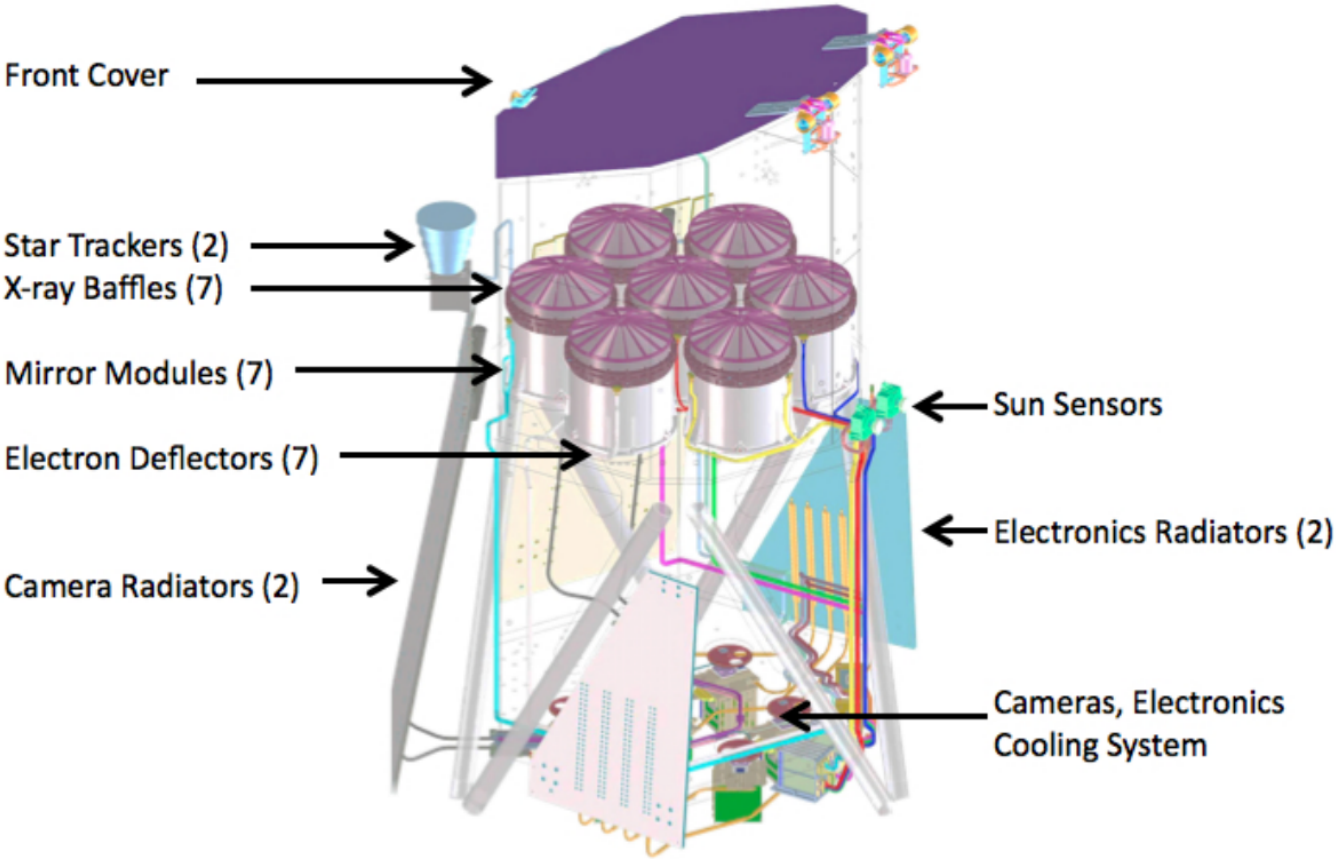

Figure 1.4: Schematic of the *eROSITA* telescope, and its components. Credit: Merloni et al. (2012)

$\approx$ 140 square degree field at an exposure depth comparable to what *eROSITA* is expected to achieve over the course of its all-sky surveys.

eROSITA consists of seven identical Telescope Modules (TMs), which each contain 54 Wolter mirror shells. These can be seen in Fig. Fig. 1.3, which is an image of the *eROSITA* TMs when it was still under construction. Each TM also features a baffle to reduce the stray light detected. The *eROSITA* detectors are seven identical pn-CCDs which consist of 384×384 pixels; each has a size of 75 μm. Fig. 1.4 depicts a schematic of the *eROSITA* telescope with all of its components.

The angular resolution differs slightly between different TMs, but it has an average on-axis value of $\approx 16.0''$. The average angular resolution across the entire field of view is $\approx 26.0''$ (Predehl et al., 2021). *eROSITA* has a temporal resolution of 50 ms, and a spectral resolution of 49 eV at 0.3 keV (Predehl et al., 2021).

On December 12, 2019, the first of eight planned all-sky surveys began. The *eROSITA* All Sky Surveys (eRASSs) each take six months to complete. Each subsequent survey starts with the completion of the previous one. It was planned for all eight surveys to be completed within four years, in December 2023. Following the survey phase, SRG was planned to be used for pointed and survey observations, partly based on proposals from the astronomical community, for at least another 3.5 years. The eRASSs are discussed in greater detail in Chapter 2.

As a consequence of the Russian invasion of Ukraine, *eROSITA* was placed into safe mode on February 26, 2022. Since then, at the time of writing, no more science operations have been performed. By that point, *eROSITA* had completed the first four of its eight planned all-sky surveys and was about halfway through eRASS5 (the fifth eRASS).

The data collected by *eROSITA* is shared equally between a German and a Russian consortium. Each consortium has data rights within one of two hemispheres. The boundary between the two hemispheres is defined by the great circle that intersects both Galactic poles, and the Galactic centre (Predehl et al., 2021). The German *eROSITA* consortium holds data rights to $|l| > 180°$, which includes the south ecliptic pole (SEP), but not the north ecliptic pole (NEP)

1.5 Telescopes in other bands

The primary focus of this work is on exploiting X-ray observations of BHs across the mass range. However, these are complemented by observations in other energy bands, which can provide additional information about the sources, that cannot be obtained in the X-ray band. We, therefore, utilised observations in radio, infrared, optical, ultraviolet, and gamma-ray bands to assist our primary X-rays observations. In this section, we briefly describe the various non-X-ray telescopes used.

1.5.1 *INTEGRAL*

The *INTErnational Gamma-Ray Astrophysics Laboratory* (*INTEGRAL*; Winkler et al., 2003) is a coded aperture telescope focused on the imaging and spectroscopy of gamma-ray sources in the sky. *INTEGRAL* has two primary instruments. The first is a spectrometer (SPI), with an energy resolution of 3 keV at 1.5 MeV, for the energy range of 20 keV−8 MeV, a detector area of 500 cm^2, and a low angular resolution (Vedrenne et al., 2003). The second primary instrument is the imager IBIS, which utilises the entire 15 keV − 10 MeV energy band, has an angular resolution of $12'$, a detector area of up to 2890 cm^2, but a low energy resolution (Ubertini et al., 2003). *INTEGRAL* was launched into geostationary orbit on October 17, 2002. It is still operational despite the loss of its thrusters.

INTEGRAL also contains an X-ray (JEM-X) and an optical (OMC) monitor. The X-ray monitor is sensitive to $3-35$ keV and also uses a coded mask aperture. This enables it to have a large field of view of 4.8°, significantly larger than that of most focusing X-ray telescopes. The field of view of *INTEGRAL* in gamma-rays is even larger, of up to 16°. The OMC is sensitive to the optical Johnson V band and has the greatest angular resolution of $25''$ (Winkler et al., 2003).

1.5.2 *AAT*

The Anglo-Australian Telescope (*AAT*; Hopkins et al., 2012) is an optical telescope with a diameter of 3.9 m and is located at Siding Springs Observatory, Australia. It was opened on October 16, 1974. The *AAT* specialises in optical spectroscopy.

Most prominent among the many spectrographs of the *AAT* is AAOmega (Sharp et al., 2006). It can determine the spectrum of up to 392 different objects simultaneously from equally many input fibres within the wavelength interval of $370 - 950$ nm. A dichroic beam splitter is used to separate the input multi-object spectra into blue and red arms at a wavelength of 570 or 670 nm. AAOmega has a spectral resolution of between 10^3 and 10^4 (Sharp et al., 2006). For optimal background subtracted spectra, it is essential to select several sky fibres that point to empty space. Two fibres are typically also devoted to standard star calibrators for calibrating the other spectra being observed.

The 2degree Field fibre positioner (2dF; Lewis et al., 2002) is commonly used to select the targets for the input to the AAOmega spectrograph. It can collect light from 392 different positions, which are connected to the spectrograph via optical fibres. These can be placed within a large field of view with a diameter of 2°, hence the name. The fibres are placed with an accuracy of $0.3''$ and need to be separated from one another by at least $30''$. The fibres can be reconfigured within 55 minutes, which can be done while a different set of fibres are used for observation. To utilise this large field of view, 2dF uses a corrector lens. It also uses an atmospheric dispersion compensator, which enables the same set of targets to be observed for a duration of up to 2 hours, as long as they are not too low in the sky. For longer time intervals, the zenith angle of a selected region of the sky changes too much, shifting the apparent locations of targets due to atmospheric refraction.

1.5.3 *Gaia*

Gaia (Gaia Collaboration, 2016) is an optical space observatory that is fine-tuned for measuring positions, parallaxes, and proper motions of more than 10^9 sources brighter than magnitude 20.7. It was launched on December 19, 2013, to the Sun-Earth L2 Lagrange point. *Gaia* scans across the entire sky by completing a revolution about itself within 6 h, while the spin axis precesses gradually. By combining measurements of the exact positions of point sources taken over many years, *Gaia* can determine highly accurate positions and parallaxes to within 0.3 marcsec, and proper motions accurate to within 1 marcsec yr^{-1} (Gaia Collaboration, 2016). The parallaxes can subsequently be used to estimate the distance to these sources.

1.5.4 *VISTA*

The *Visible and Infrared Survey Telescope for Astronomy* (*VISTA*; Sutherland et al., 2015) is a ground-based, 4 m wide telescope located at the Paranal Observatory in Chile. It is sensitive to the wavelength range $0.8 - 2.3$ μm. It started operations in December 2009. Here we use the Vista Hemisphere Surveys (VHS; McMahon et al., 2021), which aim to observe the entire southern hemisphere at a 30× greater depth than the previous 2MASS survey (McMahon et al., 2013).

1.5.5 *WISE*

The *Wide-field Infrared Survey Explorer* (*WISE*; Wright et al., 2010) is a 40 cm diameter mid-infrared space telescope, designed to survey the entire sky. It was launched into a sun-synchronous orbit around Earth on December 14, 2009. *WISE* is sensitive to 4 infrared wavelength bands, centered around 3.4, 4.6, 12, and 22 μm. It has detected 1.57×10^5 minor planets, including asteroids, comets, and near-Earth objects.

The primary and secondary solid hydrogen cryogens, required to maintain the infrared detection capabilities, were used up on October 1, 2010, (Mainzer et al., 2011). Afterwards, the bands W1 and W2 were used until February 1, 2011, when the telescope was put into hibernation.

The satellite was reactivated in October 2013 to detect and monitor near-Earth asteroids and comets. It was, therefore, renamed to the *Near Earth Object Wide-field Infrared Survey Explorer* (*NEOWISE*; Mainzer et al., 2014). It has been observing the sky ever since, at a similar sensitivity as the initial post-cryogenic phase. *NEOWISE* detected, among many other objects, the comet NEOWISE, which became the brightest comet visible in the northern hemisphere of Earth since Hale-Bopp in 1997.

CatWISE (Marocco et al., 2021) is an all-sky catalogue combining the observations of *WISE* and *NEOWISE*. The latest public release, CatWISE2020, contains observations performed over an interval 16× longer than the duration of the AllWISE catalogue, which contains the results of the initial *WISE* survey.

1.5.6 *ATCA*

The *Australia telescope compact array* (*ATCA*; Frater et al., 1992) is a radio telescope in New South Wales, that consists of six parabolic dishes, which have a diameter of 22 m each. Out of those, five antennas are movable to create different arrays for different observing patterns. Observations take place in energy bands centered around 1.5, 2.3, 5.0, 5.5, 8.4, 9, 17, and 19 GHz (Wilson et al., 2011). *ATCA* was first opened on September 2, 1988.

1.6 Outline of this thesis

In this thesis, we predominantly focus on the investigation of the X-ray variability properties of X-ray sources, particularly of AGNs and XRBs. We also used observations at other wavelengths to obtain as much information on the source properties as possible.

Chapter 2 discusses the eRASSs, and the *eROSITA* observations of the SEP. In that chapter, we also investigate the distribution of properties of *eROSITA* sources detected within 3° of the SEP.

Special treatment is required for an investigation of the variability of sources observed during the eRASSs. In Chapter 3, we introduce various variability methodologies and discuss how to adapt them for the analysis of *eROSITA* data. We set up procedures for identifying variable sources and for quantifying their variability. We investigate their efficacy on simulated *eROSITA*-like light curves.

This variability analysis is then put to use on real *eROSITA* data in Chapter 4. Significantly variable sources in the SEP field are identified and classified. We investigate their variability properties and highlight a few interesting sources.

In Chapter 5, we discuss the 2018 outburst of the BHT Swift J1658.2-4242. It featured several flip-flops with unique characteristics, which we explored in great detail. We investigated the QPOs, the amplitude of flip-flop flux changes, the timing of transitions, and the changes to the source spectrum.

Finally, Chapter 6 summarises the findings of this thesis and discusses the avenues for future research. We identify what new questions arose as a result of this work and how those might be investigated.

Chapter 2

eROSITA observations of the South Ecliptic Pole

We sought to analyse the variability properties of X-ray sources observed during the eRASSs. This chapter details the properties of the survey, the selection of a field for variability analysis, the source detection, and the distribution of key parameters within the observations. These results will subsequently be used in the definition of the variability methods in Chapter 3, and their application to the *eROSITA* data in Chapter 4.

2.1 The eRASSs

Fig. 2.1 depicts the entire sky, as seen by *eROSITA* after the completion of eRASS1. In the first all-sky survey alone, 1.1×10^6 X-ray sources were identified by *eROSITA*. It is estimated that about 77% of those are likely to be AGNs [1].

During the all-sky surveys, *SRG* rotates about itself at a constant angular velocity of 2.5×10^{-2} deg s^{-1}. It completes one rotation about itself, called an eroDay, every 4 hours. In order to scan over the entire sky, its angular momentum vector moves by an average of $10'$ per eroDay. Furthermore, the angular momentum vector is always retained within the ecliptic plane. This observing pattern ensures that the entire sky is observed within 6 months.

Due to the large field of view of $\approx 1^\circ$, *eROSITA* can observe many sources simultaneously. It also means that sources spend a significant amount of time within the field of view, as *eROSITA* scans over them. Therefore, sources passing right through the centre of the field of view can be observed for up to ≈ 40 s per eroDay. Sources are observed at least six times per eRASS in this way.

As the *SRG* angular momentum vector lies in the ecliptic plane throughout all eRASSs, the SEP and the NEP are observed on every eroDay. Therefore, sources lying close to the ecliptic poles are observed consistently every 4 hours for most of the duration of the eRASSs. These observations enable a detailed investigation into the medium to long-term

---

[1] https://https://www.mpe.mpg.de/7461950/erass1-presskit

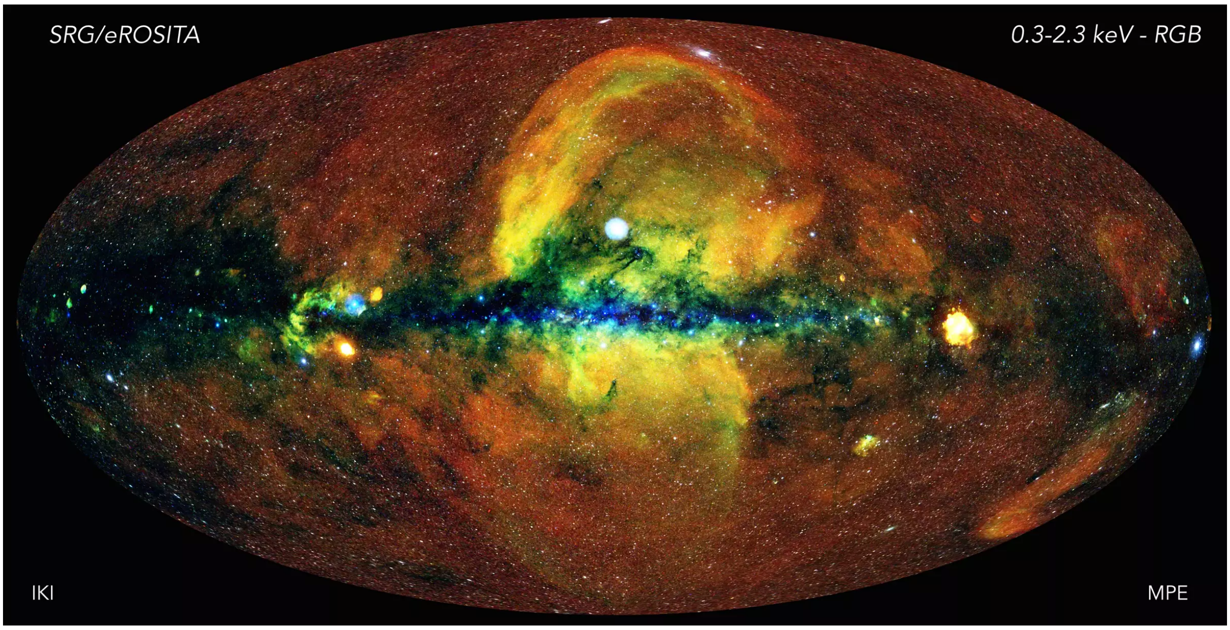

Figure 2.1: An image of the entire X-ray sky, as observed by *eROSITA* during its first all-sky survey. Credit: Jeremy Sanders, Hermann Brunner and the eSASS team (MPE); Eugene Churazov, Marat Gilfanov (on behalf of IKI)

X-ray variability properties of sources in these fields, which we will explore in Chapters 3 and 4.

On average, X-ray sources lying close to the ecliptic plane are observed in 6 consecutive eroDays per eRASS. All these observations occur within 24 hours, after which they will not be observed again by *eROSITA* for the next 6 months. However, some of those observations might be significantly shorter than the maximum of ≈ 40 s per eroDay. They will also be afflicted by significant vignetting at the edge of the field of view.

At greater altitude angles from the ecliptic, the advance of the field of view per revolution is smaller than at the ecliptic equator so that a source can be observed on a greater number of eroDays per eRASS. Of particular interest are the poles of the survey, which are the two points that *eROSITA* scans through during every rotation. The eRASS survey poles approximately correspond to the ecliptic poles, but shift slightly with each eRASS. Due to *eROSITA*'s large field of view, all sources within $\approx 0.5°$ can, in theory, be observed every 4 hours for the duration of the eRASSs.

A quick estimate of the average number of eroDays (N_{ed}) that sources located at an angle θ (expressed in degrees) relative to the ecliptic poles can be observed for per eRASS, can be found from the equation:

$$N_{ed} = \begin{cases} 12\theta \sin^{-1}\left[(2\theta)^{-1}\right] \left(\sin\left[\theta\pi/180\right]\right)^{-1} & \text{if } 0.5° \leq \theta \leq 90° \\ 1080 & \text{if } \theta < 0.5°, \end{cases} \tag{2.1}$$

where sin, and $\sin^{-1}$ are evaluated in radians. This equation assumes that the survey poles always match the ecliptic poles. It is derived by firstly considering that, on average, the field of view advances by $(\sin[\theta\pi/180])/6^\circ$ per eroDay. This advance follows a circular path with a radius of θ. The length of the arc to this circle, with a chord of 1° length, is $2\theta\sin^{-1}\left[(2\theta)^{-1}\right]$. The above equation is found by combining these two results, and inputting the size of the *eROSITA* field of view, of $\approx 1^\circ$. The number of eroDays of observations per eRASS no longer increases for $\theta < 0.5$, as all sources in this region are observed during every eroDay. In practice, a non-uniform advance of the scanning axis with every revolution, computer resets, small effective exposure times at the border of the field of view, a slight movement of the survey pole, and other issues affect this number. Therefore, it should be treated as an estimate, not an accurate measurement.

The solid angle of the region that lies within an angle of θ relative to either of the two survey poles, is defined by the equation:

$$\Omega = 2\pi\left(1 - \cos\theta\right). \tag{2.2}$$

This solid angle decreases rapidly with decreasing θ. At small angles, the solid angle contained between θ and $\theta + d\theta$ is:

$$d\Omega = 2\pi d\theta \sin\theta \approx 2\pi\theta d\theta. \tag{2.3}$$

So there are fewer sources to observe at smaller angles to the poles. As a result, the number of sources that can be observed above a lower flux limit per degree from the poles, decreases approximately linearly with θ.

The consequence of Equations 2.1 and 2.2 is that half of the sky will only be observed on 7 eroDays per eRASS, or less. Three-quarters of the sky will only be observed on 9 eroDays per eRASS, or less. In contrast, the field within 0.5° of the poles is expected to be observed on up to 1080 eroDays per eRASS. However, this region is small and only contains ≈ 300 *eROSITA* detected X-ray sources (see Section 2.3).

The region close to the ecliptic poles is particularly interesting for studying long-term variability. For our variability analysis, we chose to expand the selection of sources to the region within 3° of the SEP. This selection provides a balance between increasing the number of sources investigated, while ensuring that each source is still observed sufficiently often to enable a long-term variability study. This region contains almost all sources observed with more than 100 eroDays per eRASS. From now on, we will refer to this region as the SEP field.

Using equations Eq. 2.1 and 2.2, we find that the sources in the SEP field should, in theory, be observed on 115 to 1080 eroDays per eRASS. This region only makes up 0.137% of the solid angle of a hemisphere, but *eROSITA* spends 3.33% of its time observing it.

In this thesis, we investigated *eROSITA* observations spanning from December 13, 2019, when eRASS1 started, until June 16, 2021, when eRASS3 was completed. In addition, we analysed light curves spanning up to 552 days.

2.2 Source detection and light curve extraction

The data were processed using the *eROSITA* Science Analysis Software System (eSASS; Brunner et al., 2021). Within eSASS, the sky observed by *eROSITA* is subdivided into 4700 rectangular sky tiles, each of which has a size $3.6° \times 3.6°$, and slightly overlaps with the adjacent sky tiles on all sides (Predehl et al., 2021). The SEP field is completely covered by the seven *eROSITA* sky tiles: 082159, 085153, 087156, 090159, 092153, 093156, and 098159, with some overlap.

X-ray sources were detected in each sky tile using the three-band detection procedure described in Brunner et al. (2021). This includes both point-like and extended sources found with a detection likelihood (`DET_LIKE`) of greater than 6. Sources were labeled as extended, if they had an extent likelihood (`EXT_LIKE`) larger than 14 (Liu et al., 2022b). All other sources were considered to be point-like. The eSASS version and the detected catalogue used in this work are preliminary. They will differ slightly from the final release of eRASS catalogues that are still in preparation (Merloni et al. in prep.). However, these minor differences are expected to mainly affect the faintest sources near the detection limit, which are of little interest to this work. For the faintest sources to be detected as variable, they must feature a large degree of variability.

In this preliminary source detection, two sky tiles have small corner regions at the boundary of the SEP field masked out in eRASS2 and 3 because of the large exposure gradients. Sources in these regions were omitted in this work.

We merged the catalogues detected in the seven sky tiles into one. The sources located in the overlapping regions of the sky tiles are contained in both catalogues. To avoid analysing the same sources multiple times, we identified the matching sources in the overlapping sky tiles within $15''$ and removed the duplicates. The choice of this separation is based on the angular resolution of *eROSITA* (Predehl et al., 2021). We used three independent catalogues for the three eRASSs we investigated, and initially performed independent variability analyses on the three eRASS data sets. We describe the merging of the variable source catalogues in Chapter 4. A combined eRASS:3 (the colon indicates that this includes all eRASSs up to that number) data set was not yet available at the time of writing.

In each eRASS scan, we extracted source light curves at the uncorrected source positions using the eSASS task `srctool`. We used a bin size of 40 s, and the four energy bands selected for source detection (Liu et al., 2021). These four bands consist of three narrow bands ($0.2 - 0.6$ keV, $0.6 - 2.3$ keV, and $2.3 - 5.0$ keV) and a full band ($0.2 - 5.0$ keV). We predominantly used the full $0.2 - 5.0$ keV band for the variability detection and analysis. Energies above 5 keV are less useful for variability study because of the low effective area and strong background (Predehl et al., 2021). In the light curve extraction, we adopted the `srctool` option `lctype="REGULAR-"`, which ensures regular spaced bins, and removes all empty bins. This option also reduces the output file sizes.

The light curves generated in this way sometimes contain two bins for the same eroDay, as the exposure is split between two 40.0 s bins. To ensure that there is precisely one bin per eroDay of observations, we merged such bins into one.

A key parameter for investigating *eROSITA* light curve is the fractional exposure, ϵ, which denotes what fraction of the chosen bin duration is fully exposed as if it had been observed on-axis. Therefore, a fractional exposure of 0.5 for a 40.0 s bin denotes that the exposure of the source in this bin is equivalent to it having been observed for 20.0 s on-axis in a pointing mode observation. The fractional exposure is calculated to be the product of the fractional time, ϵ_t, and the fractional area, ϵ_a. The fractional time denotes the fraction of the bin duration during which the source was contained inside the *eROSITA* field of view. With a circular field of view of 1.03°, the fractional time for a 40.0 s bin, is found from the length of the chord passing within ϕ degrees of the center, with: $\epsilon_t = \sqrt{1.03 - \phi}$. The fractional area corrects for vignetting effects away from the centre of the field of view. The effective area drops rapidly with increasing off-axis angle. The fractional area utilises the known drop-off in the effective exposure as a function of the off-axis angle, coupled with the path of the source through the field of view. Both the amount of time a source spends within the field of view, and the fraction of that time spent in strongly and weakly vignetted regions varies from eroDay to eroDay. Therefore, the fractional exposure is different in every eroDay, and can have a wide range of values.

A source can spend up to 41.2 s inside the field of view. Therefore, the fractional time associated with a 40.0 s bin that contains all data collected on a source within a single eroDay can exceed a value of 1.0. However, even in such instances, the fractional area always causes the fractional exposure to be significantly less than 1.0. The effective exposure time, which is the product of the bin duration and its fractional exposure, can be a more useful quantity, than the length of time that the source spent inside the field of view.

We discarded all bins with a fractional exposure of less than 0.1. This corresponds to an effective on-axis exposure time of merely 4.0 s, and is a factor of 7 smaller than the maximum observed fractional exposure of almost 0.7 (corresponding to an effective on-axis exposure time of 28.0 s).

As the vast majority of sources we observed have a low count rate in the range of $0.001 - 1.0$ cts/s, single eroDay observations were not split further to look into < 40 s variability, as there is not enough information available for that. However, the variability techniques described in Chapter 3 could also be adapted to explore sub-eroDay variability for sufficiently bright variable sources.

In the following analysis, unless otherwise stated, we denote the number eroDays of observation with $\epsilon > 0.1$ as the number of bins of the light curve, N_b. For rebinned light curves, the number of bins correspond to a fraction of the number of eroDays of observation.

2.3 Properties of the SEP field

The *eROSITA* source detection pipeline (Brunner et al., 2021) detected 8728 X-ray sources in eRASS1, 7984 in eRASS2, and 7770 in eRASS3 in the SEP field. Despite being located within a small region of the sky, the properties of these observations are inhomogeneous, and vary a lot as a function of angle from the SEP.

2.3.1 Number of eroDays of observation

In the idealised and simplified model of the survey, the number of eroDays of observations of sources in the SEP field per eRASS ranges from 115 to 1080. In practice, orbit corrections, computer resets, calibration observations, and other events can prevent continuous survey mode operation. A non-uniform rotation rate of the scanning plane around the ecliptic equator also implies that sources at the same ecliptic latitude, but different ecliptic longitudes, are observed a different number of times. Removing bins with a fractional exposure of less than 0.1 further reduces the number of eroDays of observation with sufficient exposure times used in the variability analysis. Another effect to consider, is the slight shift in the position of the scanning pole. This reduces the size of the region within which every source is observed on every eroDay of survey mode.

The combination of these effects is illustrated in Fig. 2.2. In this, and several subsequent figures, we separate out the sources identified as variable using the methods developed in Chapter 3. These sources are discussed further in Chapter 4. Fig. 2.2 displays the distribution of the number of eroDays during which each X-ray source detected in the SEP field was observed with a fractional exposure of at least 0.1, in eRASS1, 2, and 3. The greatest number of sources per logarithmic interval on the number of eroDays occurred at 145, 149, and 113 eroDays, in eRASS1, 2, and 3, respectively. This is because the solid angle of sources that lie within θ and $\theta + \delta\theta$ grows rapidly with increasing θ (see Equation 2.3). So the greatest density of sources per logarithmic interval on the number of eroDays of observation is detected close to the lower limit on this parameter. eRASS3 has a maximum at a slightly lower value than in the previous two surveys, as it suffered from more days of telescope downtime, which caused the entire distribution to be shifted to lower values.

Fig. 2.2, and the other figures in this chapter also show the distribution of the relevant parameter for variable sources. These are identified using the methods discussed in Chapter 3. Chapter 4 discusses these results in more detail. The distribution of various parameters is relevant for defining the variability methods in Chapter 3, which is why we are discussing them here.

The distribution of the number of eroDays of observation per source per eRASS with $\epsilon > 0.1$ extends significantly below the value of 115 eroDays which we calculated for the idealised model of the survey. The minimum number of eroDays during which a source was observed in the SEP field was 28, 41, and 32, in eRASS1, 2, and 3, respectively. However, only a few sources were observed for fewer than 50 eroDays. This tail of the distribution is due to sources at the edge of the field of view. These sources lost a significant fraction of their observing interval due to a combination of detrimental effects.

Using Equation 2.2, the number of sources observed within logarithmically spaced intervals on the number of eroDays ($\rho(N_{ed})$), can be estimated to decrease with the inverse square of the number of eroDays (N_{ed}): $\rho(N_{ed}) \propto N_{ed}^{-2}$. This assumes a constant detected source density per unit solid angle. It only applies to small angles relative to the ecliptic poles, for which Eq. 2.3 is accurate. It also does not apply to angles smaller than 0.5°, as sources lying closer to the ecliptic poles are observed approximately the same number of

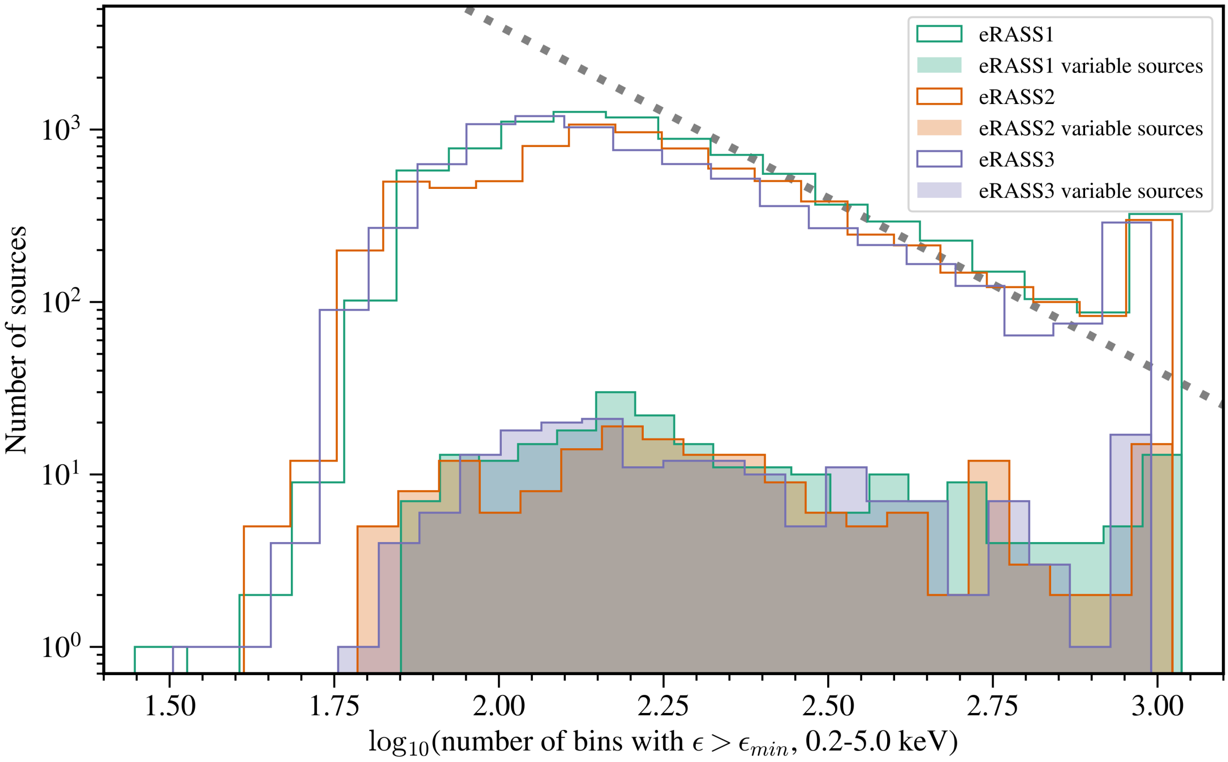

Figure 2.2: The distribution of the number of eroDays, that all sources in the SEP field were observed for, in eRASS1, 2, and 3. The coloured in histogram depicts the distribution of the number of eroDays of observation of variable sources. The grey dotted line depicts an inverse square relationship, which approximately matches the decrease in the number of sources observed at a particular logarithmic interval on the number of bins. We only counted bins with a fractional exposure larger than $\epsilon_{min} = 0.1$.

times. Additionally, a greater number of eroDays of observation implies a greater exposure depth and, therefore, an enhanced ability to detect fainter sources. As a result, the source density is not identical throughout this region, but increases notably towards the SEP. The interplay between these two effects explains most of the shape of the distribution of Fig. 2.2. However, the effect of a smaller solid angle closer to the SEP dominates over the effect of enhanced source detection ability. The resulting distribution is slightly less steep than an inverse square relationship. This is shown in Fig. 2.2, in which an inverse square relationship is displayed as a grey dotted line.

The spike at the highest number of bins is caused by sources lying within 0.5° of the SEP. We observed a maximum number of eroDays of observation per source, of 1088, 1054, and 978, in eRASS1, 2, and 3, respectively. The sources at this maximum are observed on almost every eroDay of operational survey mode during the eRASS in question. At angles of less than 0.5° to the survey pole, there no longer is any increase in the number of eroDays of observation, so all sources within 0.5° are observed approximately the same number of times. eRASS2 suffered 34 more eroDays of telescope downtime than eRASS1. eRASS3 was even more strongly affected, with 110 eroDays more downtime than in eRASS1. This can be seen by the shift of the entire eRASS3 distribution on the number of eroDays in Fig. 2.2. This is also illustrated in Figure 2.3, where we find the central 0.5° region to have an approximately constant colour, corresponding to an approximately constant exposure time and a constant number of eroDays of exposure.

Another point of interest is the shift of the survey pole away from the SEP. The SEP has coordinates of RA = 06h 00m 00.00s, Dec = $-66^\circ\ 33'\ 38.55''$. In contrast, the average direction of the pole in eRASS1, 2, and 3 did not perfectly match the coordinates of the SEP. We identified the average position of the survey pole to be located at: RA = 05h 58m 27s, Dec = $-66^\circ\ 33'\ 41''$ (0.1535° away from the SEP) in eRASS1, RA = 06h 01m 02s, Dec = $-66^\circ\ 32'\ 33''$ (0.1050° away from the SEP) in eRASS2, and RA = 05h 58m 29s, Dec = $-66^\circ\ 33'\ 54''$ (0.1512° away from the SEP) in eRASS3. In all three eRASSs, the separation of the survey pole from the SEP is comparable. The average eRASS1 and eRASS3 survey poles only have a small separation of $18''$. In contrast, the pole of eRASS2 is located approximately on the other side of the SEP, compared to the poles of eRASS1 and 3. These differences in the direction of the survey pole are small, but they affect the frequency of observation of sources as a function of the angle to the SEP. This contributes to the variation in the number of eroDays of observation of the same source in different eRASSs. As the shift in the average scanning pole in eRASS1-3 is smaller than the size of the *eROSITA* field of view, there still is a region within which sources are observed during every scan of *eROSITA*. However, it is smaller than the circle of radius 0.5°, which it would be if the scanning pole were constant throughout the eRASSs.

Additionally, the requirement that $\epsilon > \epsilon_{min}$ results in some eroDays being discarded. Only sources located very close to survey poles are unaffected by this.

Fig. 2.3 depicts the location of all sources detected by *eROSITA* in the SEP field in eRASS1, 2, and 3. Colours denote the logarithm of the total exposure time, which is found using $\sum_{i=1}^{N_{ed}} 40\epsilon_t$. This remains reasonably constant at small angles to the SEP, but starts to drop off rapidly for angles $> 0.5^\circ$ from the survey pole.

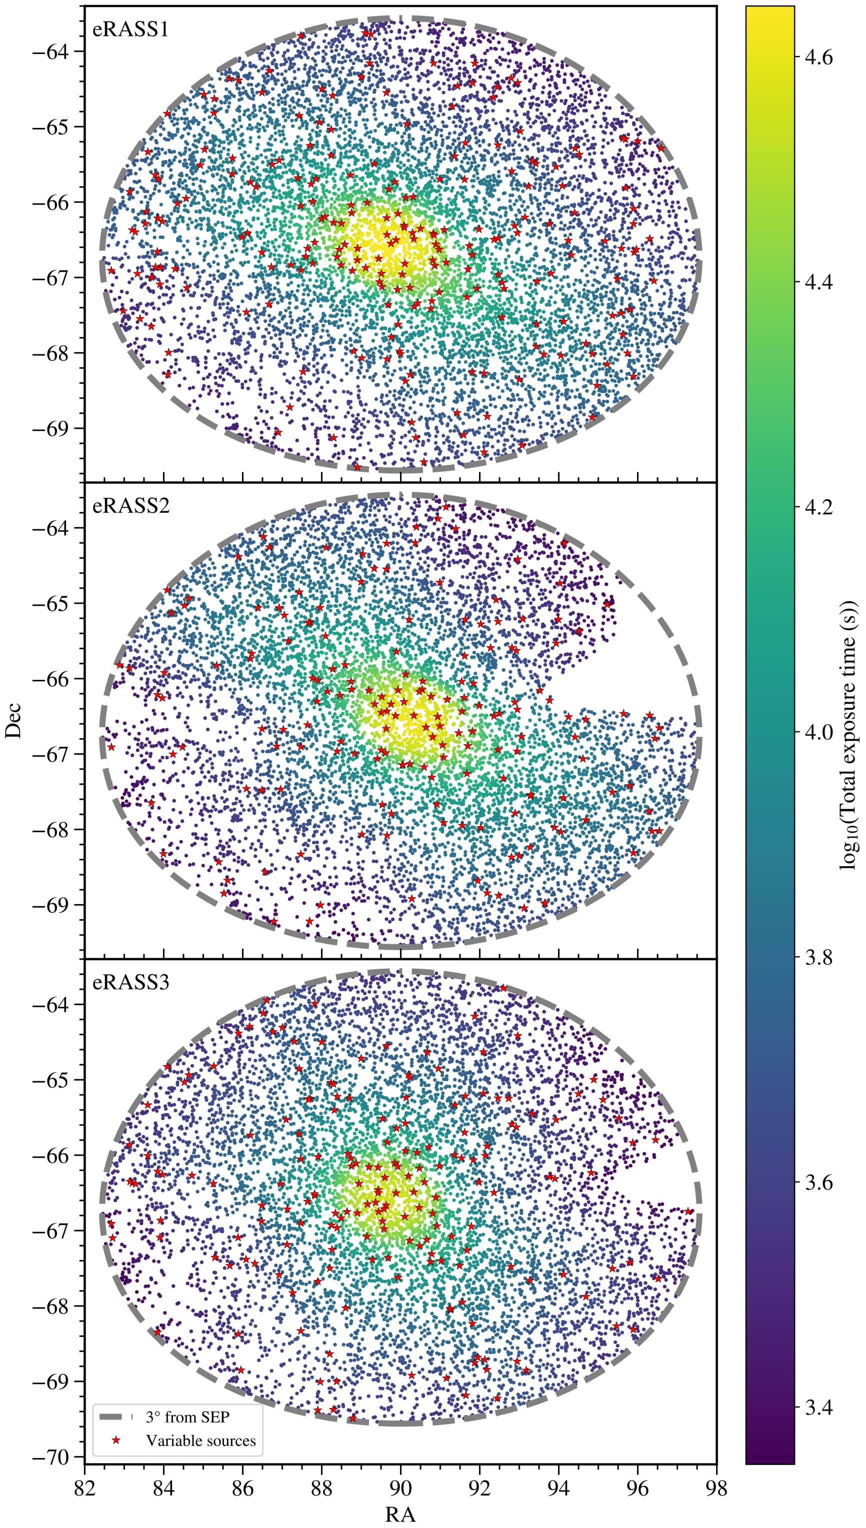

Figure 2.3: Location of all X-ray sources detected by *eROSITA* in the 3° SEP field, in eRASS1, 2, and 3. Colours indicate the total exposure time of each source in that specific eRASS. The location of variable sources is indicated via red stars. The slight shift of the central bright yellow region indicates the change in the location of the pole centre. The blank triangular field in eRASS2 and 3 is caused by the source detection pipeline failing for large exposure gradients.

There is a noticeable increase in the source density towards the survey poles. For example, the annulus within $2.75 - 3.00°$ of the SEP was observed to have a source density per unit solid angle (in square degrees) in eRASS1, 2, and 3, of: 244, 204, and 209, respectively. In contrast, the circle of radius 0.25°, centred on the SEP itself, has a source density per square degree, which is more than three times as large, of 708, 667, and 662, in eRASS1, 2, and 3, respectively.

The source detection pipeline failed in some regions of the positive RA end of the SEP field in eRASS2 and 3, due to the large exposure gradients in this region. This caused the white triangular holes, in which no sources were detected by the pipeline. These source detection issues in eRASS2 and 3 will be resolved in a future version of the eSASS pipeline.

Another notable feature of Figure 2.3 is the non-isotropic distribution of exposure time. The advance of the field of view with each subsequent rotation of SRG about itself is not constant, but varies over the course of each eRASS. Therefore, we see an enhanced source density and greater exposure depth in the direction of negative RA, positive Dec, and positive RA, negative Dec, in contrast to the positive RA, positive Dec, and negative RA, negative Dec directions. Additionally, the presence of the Large Magellanic Cloud (LMC) at the low RA end of the SEP field causes an enhanced background count rate, and therefore corresponds to a reduction of the source density further than the equivalent exposure time region on the opposite side of the SEP, at more positive RA.

The colour of the central region indicates that the maximum exposure time differs between the three eRASSs. Fig. 2.3 also depicts the shift of the survey pole, via the location of its bright central region.

2.3.2 Fractional exposure

Fig. 2.4 displays the distribution of the fractional exposure of all sources in the SEP field, for all eroDays of observation, in eRASS1. The figure also depicts the different distributions of the 4 energy bands investigated. The $0.2 - 0.6$ keV, $0.6 - 2.3$ keV, and $0.2 - 5.0$ keV energy bands all have a similar distribution. However, the $2.3 - 5.0$ keV energy band has a distribution shifted to lower fractional exposure values due to a stronger vignetting effect at higher energies. In all four energy bands, the peak of the distribution in the fractional exposure, grouped into linear bins, occurs towards a value of 0. Nevertheless, when grouped in logarithmic intervals, the peak of the distribution occurs at $\epsilon = 0.30$. About 26% of all bins lie below the fractional exposure threshold we selected, of $\epsilon_{min} = 0.1$. We note that the distribution of the fractional exposure within this field is reasonably uniform between 0.1 and 0.45, but drops rapidly at larger fractional exposures, with a maximum of about 0.7. There is a small spike in the fractional exposure at a value of 0.32, probably due to the geometry of the scans close to the SEP. The distribution seen in eRASS2 and 3 is almost identical to that in eRASS1, but shifted to a lower total number of eroDays, as those surveys contain fewer sources, and fewer eroDays of observation. The average fractional exposure per source, per eRASS, follows an approximately normal distribution, with a mean of 0.30, and a standard deviation of 0.041.

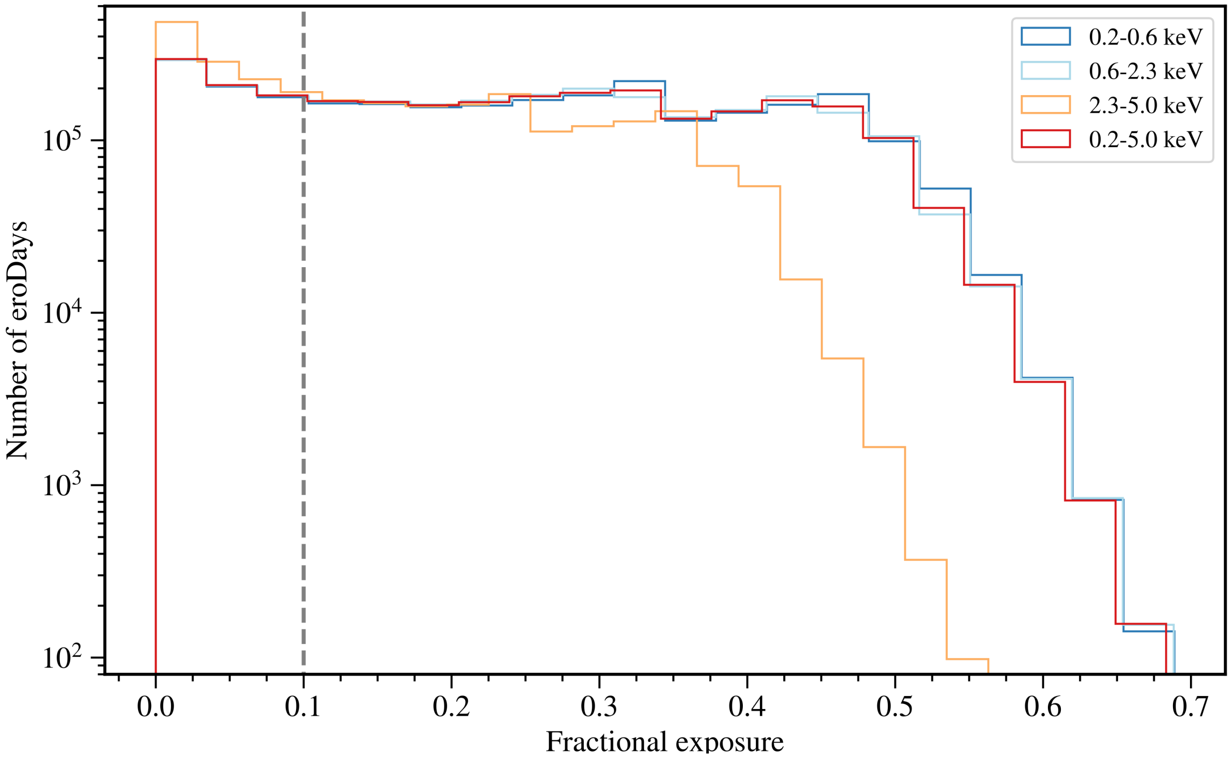

Figure 2.4: The distribution of the fractional exposure per eroDay, for the 4 different energy bands, for all sources in the SEP field in eRASS1. The vertical dashed line at $\epsilon = 0.1$ indicates our chosen lower limit.

2.3.3 Exposure time

Fig. 2.5 depicts the distribution of the total exposure time for sources in the SEP field, in eRASS1, 2, and 3. This parameter describes the total time that each source was inside the field of view during each of the surveys, and is calculated as the product of the fractional time and the maximal duration of an eroDay observation, summed over all eroDays: $\sum_{i=1}^{N_{ed}} 40\epsilon_t$. It does not consider any vignetting effects and is, therefore, closely related to the number of bins. So the distribution of the number of bins per source (Fig. 2.2.), and the distribution of the total exposure time, are similar. A noticeable difference is that the peak at the highest exposure times is broadened. Even though sources within 0.5° of the survey pole centre are observed during almost every revolution of *eROSITA* about itself, sources further away from the survey pole spend less time in the field of view, on average. A source located at the survey pole passes through the centre of the field of view during every rotation, such that the total exposure time is approximately $40N_{ed}$. In contrast, a source that lies 0.5° away from the survey pole will be observed during every rotation, but sometimes only briefly at the edge of the field of view.

We observe the total exposure time of sources in the SEP field to range between $2.6 - 44.1$ ks in eRASS1, $2.2 - 41.6$ ks in eRASS2, and $2.3 - 36.3$ ks in eRASS3. The maximum of the distribution of total exposure time per eRASS, in logarithmic intervals, occurs at 5.77 ks in eRASS1, 6.10 ks in eRASS2, and 3.83 ks in eRASS3.

The total effective exposure time denotes how long an observation of the selected source would have to be performed, if it had been observed at the centre of the field of view throughout, to yield the same exposure depth as the one observed. It describes the product of the fractional exposure with the duration of each observation, summed over all eroDays of observation: $\sum_{i=1}^{N_{ed}} 40\epsilon$. Fig. 2.6 depicts the distribution of the total effective exposure time for the full energy band, $0.2 - 5.0$ keV. The overall shape is still similar to the distribution of the number of eroDays (Fig. 2.2) of sources in the SEP field, but with two notable differences. Firstly, the low exposure tail is more pronounced, as vignetting effects exacerbate the issues causing this tail. Secondly, the high exposure peak corresponding to sources observed almost every eroDay is gone in Fig. 2.6. Even though the number of eroDays of observation hardly changes at angles $< 0.5°$ from the survey pole, the average fractional exposure per eroDay of these sources keeps increasing with decreasing angle from the SEP. Sources lying 0.5° away from the survey pole are observed with a low fractional exposure on some eroDays. In contrast, a source located directly at the survey pole will be observed with maximal fractional exposure on every eroDay.

Therefore, the effective exposure time of sources in the SEP field varies between $0.40 - 22.88$ ks in eRASS1, $0.35 - 21.75$ ks in eRASS2, and $0.36 - 19.11$ ks in eRASS3. The peak of the distribution of the effective exposure time, per logarithmic interval, was observed at 1.866 ks in eRASS1, 1.983 ks in eRASS2, and 1.367 ks in eRASS3. The vignetting effect at large off-axis angles significantly reduces the effective exposure time compared to the total exposure time.

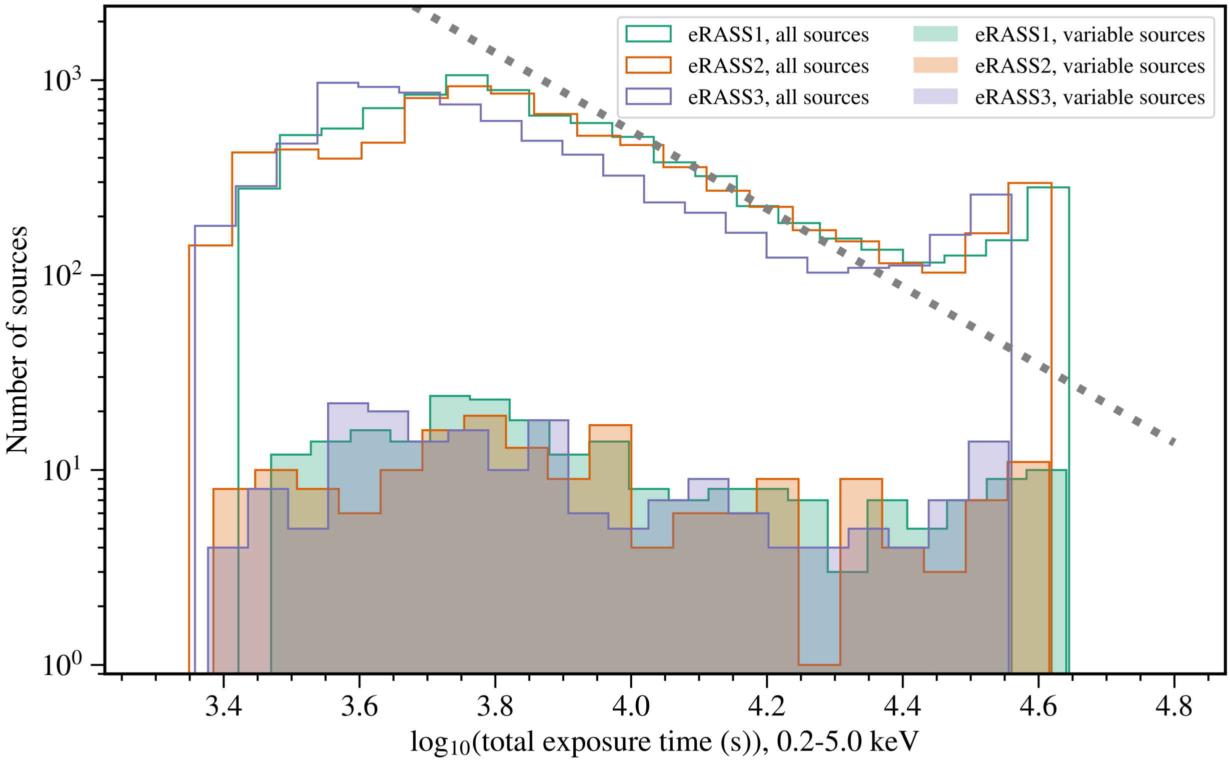

Figure 2.5: The distribution of the exposure time of sources in the SEP field, for eRASS1, 2, and 3. This quantity describes the total time that a source was inside the field of view, and does not consider vignetting effects. The distribution of the total exposure time for the sources we identify as variables is also shown. The grey dotted line depicts an inverse square relationship, showing that the drop in the number of sources within a logarithmic interval on the total exposure time is dominated by the decrease in the solid angle at a decreasing angle from the SEP.

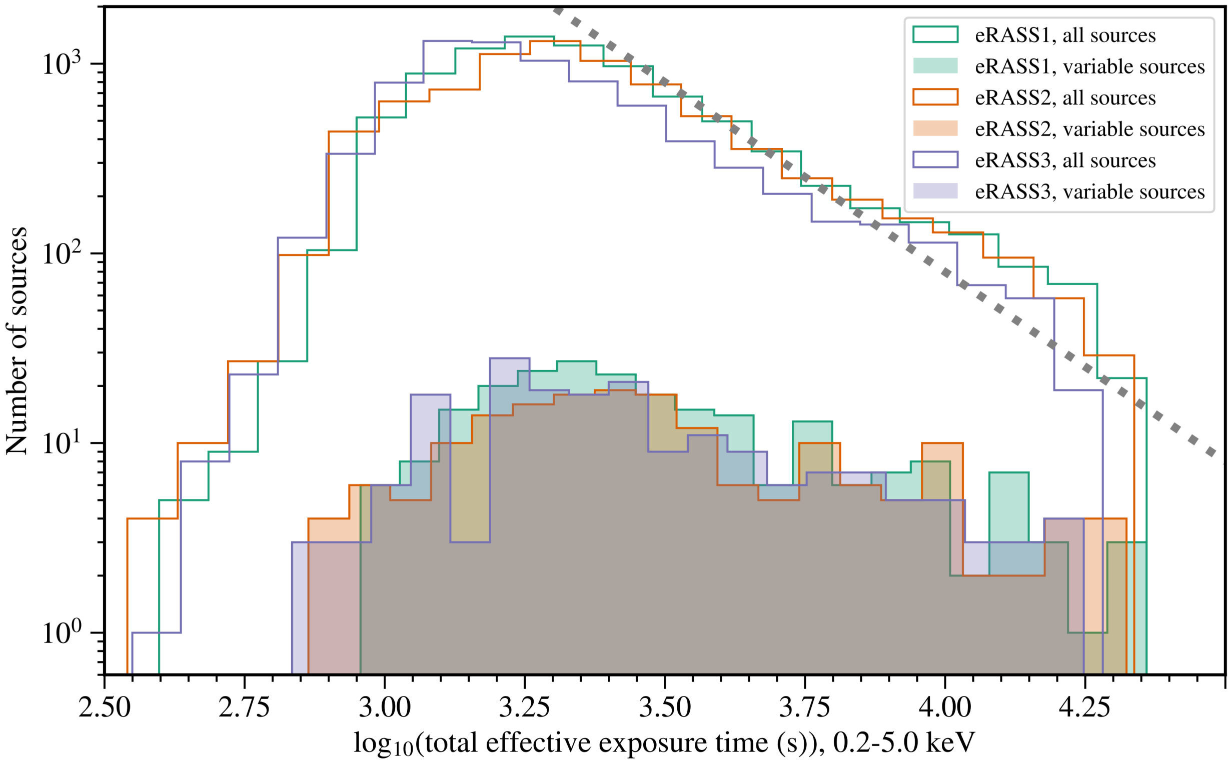

Figure 2.6: The distribution of the total effective exposure time. This is found by summing the product of the maximum duration of an observation in survey mode (40 s), with the fractional exposure. The distribution of the effective exposure time of the variable sources identified in each eRASS is also indicated. The grey dotted line depicts an inverse square relationship, showing that the drop in the number of sources within a logarithmic interval on the total effective exposure time is dominated by the decrease in solid angle at a decreasing angle from the SEP.

2.3.4 Source count rate

The standard calculation for the count rate in an individual eroDay bin is performed as follows. Within a $\Delta t = 40$ s bin at time t_i (corresponding to the i$^{\text{th}}$ eroDay), $C(t_i)$ counts were measured in the source extraction region, $B(t_i)$ counts were measured in the background extraction region, which covers an area $1/A(t_i)$ times as large as the source extraction region. The effective exposure time is obtained by multiplying the fractional exposure ($\epsilon(t_i)$) by Δt. All of these parameters can then be used to estimate the source count rate in this bin, $R_S(t_i)$, with:

$$R_S(t_i) \approx \frac{C(t_i) - A(t_i)B(t_i)}{\epsilon(t_i)\Delta t}. \tag{2.4}$$

Due to varying fractional exposures, a simple measure of the difference in the source and background counts divided by the bin time is an inaccurate estimator of the source flux. Instead, this equation calculates the exposure corrected count rate, which is a reasonable estimator of the source flux. However, it only works well if the measured values for $C - AB$, are large, and if the background count rate is uniform and equivalent within both the source and background extraction regions. However, most sources have a relatively low count rate. Even bright sources are observed with small ϵ on some eroDays. The varying fractional exposure, therefore, creates challenges for accurately estimating the actual flux of sources in each eroDay. This effect needs to be taken into account when identifying variable sources.

A different method for calculating the count rate in each bin is to use the posterior probability distribution on the count rate, calculated using the inverse incomplete Gamma function, as described by Buchner et al. (2022), and Knoetig (2014). The application of this method to *eROSITA* data is discussed by Buchner et al. (2022). In this method, the count rate for each bin, $R_S(t_i)$, is estimated to be located at the peak of the posterior distribution. Following Kraft et al. (1991), the credible confidence interval in the count rate of each bin is set by the following two conditions on the posterior probability distribution of the source count rate:

$$\int_{R_{S,min}(t_i)}^{R_{S,max}(t_i)} f\left(R_S(t_i)\right) dR = CL \tag{2.5}$$

$$f\left(R_{S,min}(t_i)\right) = f\left(R_{S,max}(t_i)\right), \tag{2.6}$$

where $R_{S,min}(t_i)$, and $R_{S,max}(t_i)$ are the upper and lower bounds of the confidence interval that we selected for the count rate of the bin at time t_i. We used a confidence level of $CL = 1\sigma \approx 0.68269$. At low count rates, the two error bar sizes, $\sigma_{+err} = R_{S,max}(t_i) - R_S(t_i)$, and $\sigma_{-err} = R_S(t_i) - R_{S,min}(t_i)$ are strongly asymmetric.

For the analysis of *eROSITA* data, we used the posterior distribution to define the estimate of the count rate of the source in all bins, rather than Eq. 2.4. This method of calculating the exposure-corrected count rate is less strongly affected by the varying fractional exposure in the eRASS light curve, as compared with Eq. 2.4. All *eROSITA*, and simulated *eROSITA*-like light curves were generated using this method.

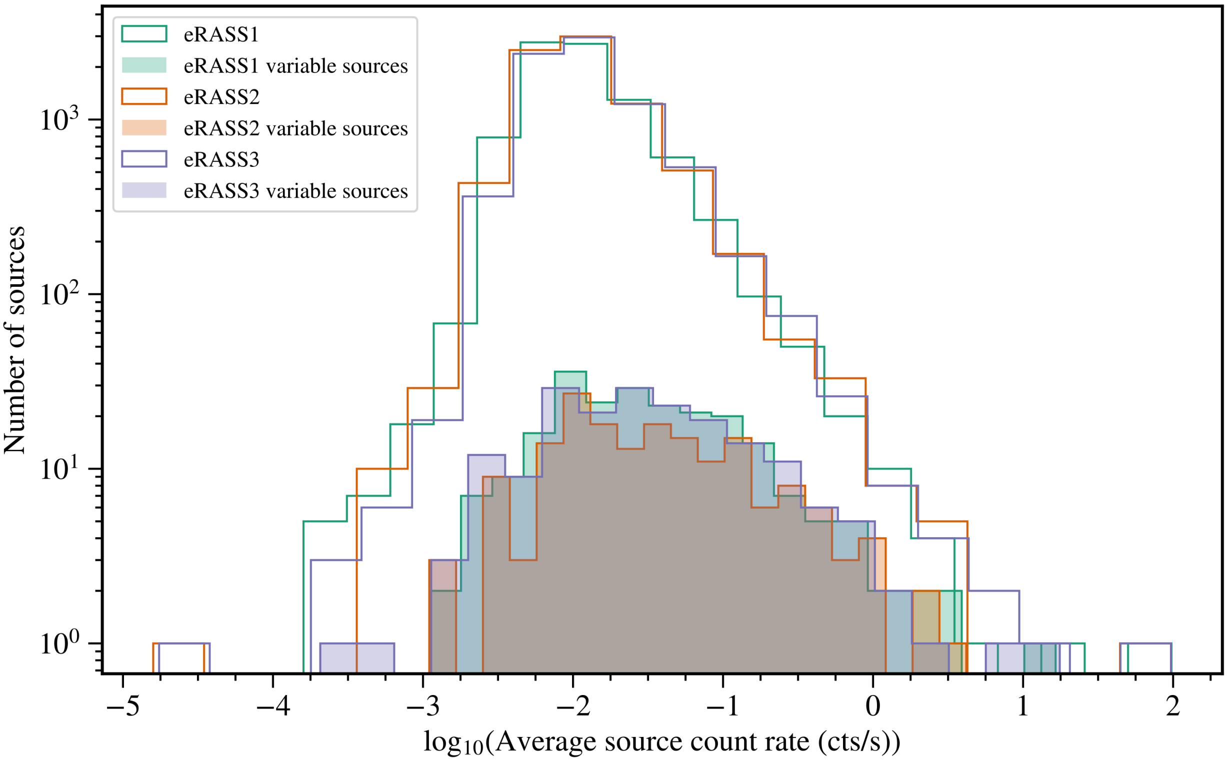

Figure 2.7: The distribution of the average count rates of sources in the SEP field, in eRASS1, 2, and 3. The average count rates are determined by Equation 2.7. Due to the high exposures in the SEP field, *eROSITA* can detect sources at much lower count rates than in the rest of the sky. This figure does not show the six sources with a slightly negative average count rate.

Fig. 2.7, shows the distribution of the average source count rate measured by *eROSITA* for all sources in the SEP field, within $0.2 - 5.0$ keV. The average source count rate is calculated by determining the peak of the posterior distribution of the count rate for a single bin containing the data of all observations. This involves summing over all source counts, summing the product of the background counts and background ratio in every bin, and summing the fractional exposures.

Except at low count rates, or a small number of bins in the light curve, the average source count rate is approximately equal to:

$$\overline{R_S} \approx \frac{\sum_{i=1}^{N_b} C(t_i) - A(t_i)B(t_i)}{\sum_{i=1}^{N_b} \epsilon(t_i)\Delta t}. \tag{2.7}$$

This is also approximately equal to the weighted average of the count rate of each eroDay, using the fractional exposure as a weight.

The brightest source in the SEP field is the supernova remnant SNR B0535-66.0, for which we measured an average source count rate of ≈ 97.5 cts/s. Unfortunately, this estimate is also no longer accurate at such high count rates, due to pileup effects.

Using Eq. 2.7, we found six sources with slightly negative average source count rates. These sources are extremely faint. Despite being observed for several ks, only very few source counts were detected for them, slightly below the average background count rate. These lowest count rate estimates are not only unreliable, but they are also from potentially spurious detections.

Almost all detected sources in the SEP field have average source count rates in the range $1.8 \times 10^{-4} - 4.2$ cts/s. The most common average source count rate, per logarithmic interval, is about 8.9×10^{-3} cts/s. The distribution increases rapidly from low count rates towards the maximum. It decreases less rapidly, with an almost constant power law with an index of -1.2, to higher count rates, up to about 10 cts/s. Unlike previous plots, the distribution of the average source count rate in eRASS1, 2, and 3 is very similar, which is an independent check of the consistency of the observations across this interval, despite differences in the number of eroDays of observations for all sources.

In Fig. 2.8, the distribution of the total measured counts in the source extraction region for the energy band $0.2 - 5.0$ keV is displayed. The distribution features a steep rise from the lowest number of counts, of merely 4, to a peak at 25 counts. That is the value at which there are the greatest number of sources per logarithmic interval in the total number of source counts. This indicates that many observed sources are very faint. Above the peak, the distribution is dominated by a decline that approximately follows a power law, with an index of approximately -1.2. This power law continues until about 5600 counts per source, per eRASS. The greatest number of source counts measured for a single source in one eRASS was 3.19×10^5 for the supernova remnant SNR B0535-66.0 in eRASS1. The strongly positively skewed distribution, with a peak at just 25 counts, results from the positive skew of both the average count rate distribution, and the effective exposure time distribution, both of which feature most sources lying above the peak.

2.3.5 Background count rate

The background ratio describes the ratio of the source extraction area to the background extraction area. Fig. 2.9 shows that the background extraction region is always significantly larger than the source extraction region. The most common value of the background ratio, per logarithmic interval, is 0.0089. Nevertheless, the distribution of background ratios is broad, spanning the interval from 0.0017 to 0.074, and features a second prominent peak at a value of 0.0040.

Fig. 2.10 depicts the observed distribution of the background count rate for the $0.2 - 5.0$ keV interval, measured in the background extraction region for each detected source. Similar to the source count rate, the average background count rate, $\overline{R_B}$, can be estimated to be:

$$\overline{R_B} \approx \frac{\sum_{i=1}^{N_b} B(t_i)}{\sum_{i=1}^{N_b} \epsilon(t_i)\Delta t}. \tag{2.8}$$

Similar to the distribution of the observed average source count rate within the source extraction region, this follows a sharp rise from the lowest background count rate of

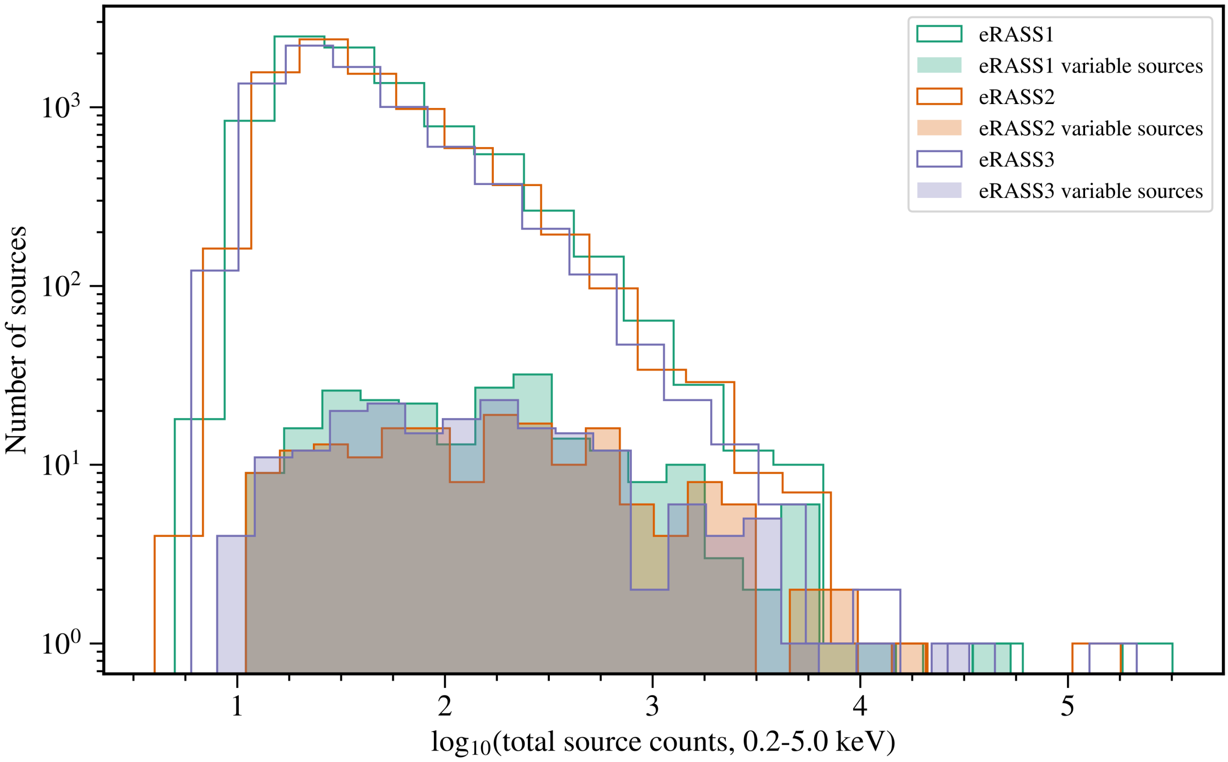

Figure 2.8: The distribution of the total source counts for all sources in the SEP field, in eRASS1, 2, and 3, for the energy band $0.2 - 5.0$ keV.

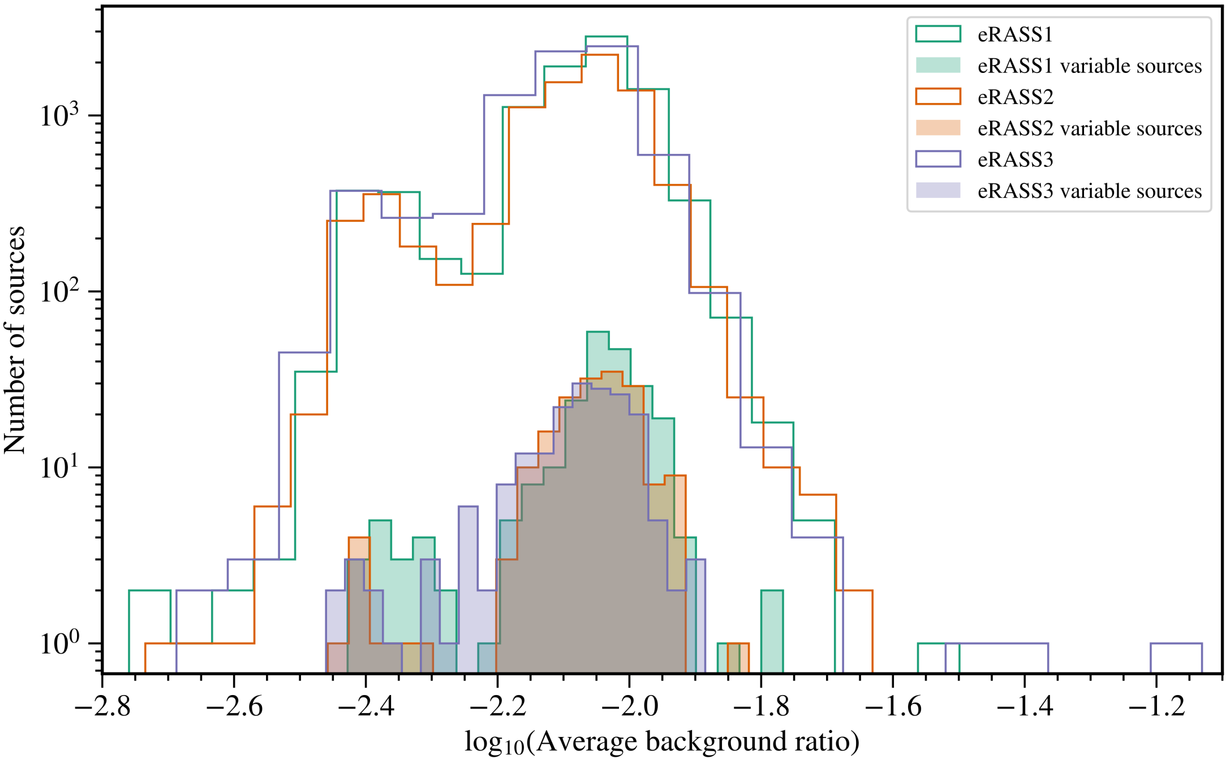

Figure 2.9: The distribution of the average ratio of the area of the source to the background extraction region, for each source detected in the SEP field, in eRASS1, 2, and 3. The background ratio exhibits slight variations throughout each eRASS, and across eRASSs, but usually remains close to the average value.

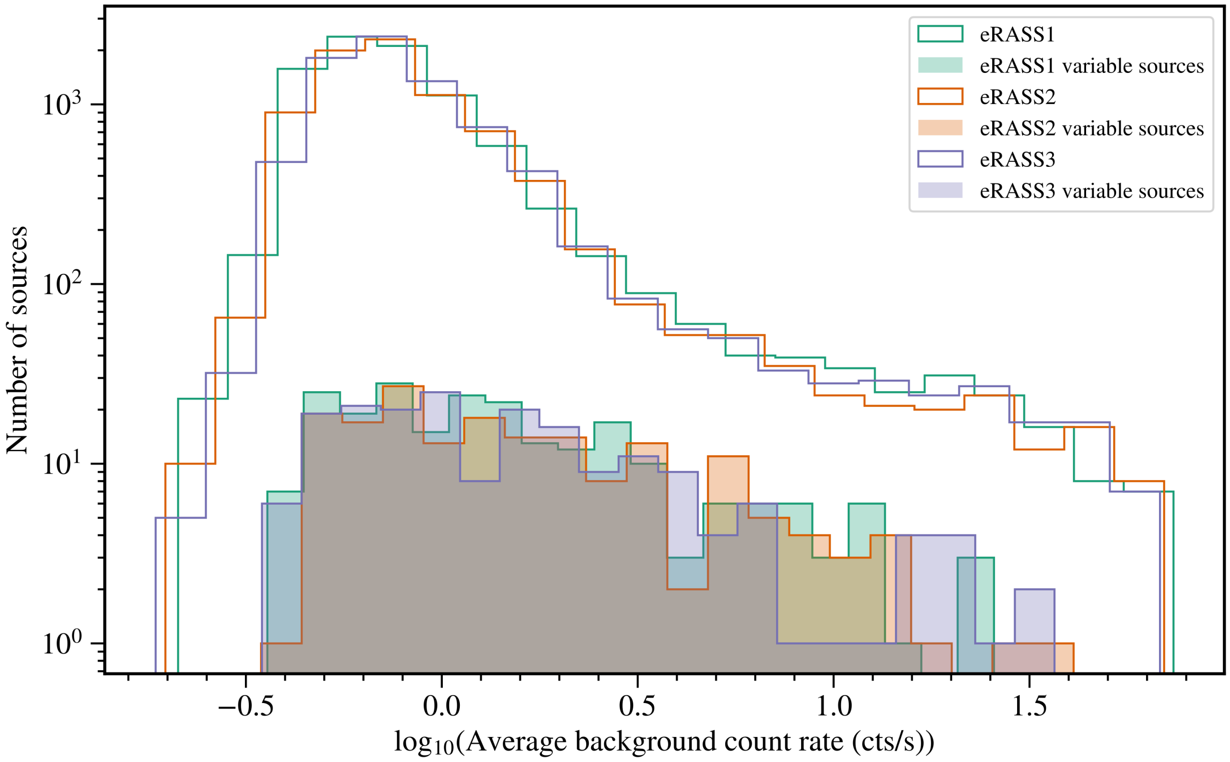

Figure 2.10: The distribution of the average background count rate in the background extraction regions used to determine the background subtracted source count rates. The background extraction region is typically more than ≈ 100 times larger than the source extraction region.

0.186 cts/s up to a peak at 0.71 cts/s. However, this distribution is much more positively skewed, almost reaching a plateau at the highest background count rates of up to 73.7 cts/s. These background count rates are often much larger than the source counts rates, as they are obtained over a much larger area than the source count rate (See Fig. 2.9). This is done to accurately estimate the background count rate in the much smaller source extraction region. The part of the SEP field that intersects within the LMC has a higher background count rate than the rest of the field, which accounts for parts of the high background count rate end. The width of this distribution is also due to the spread in the background ratio, which causes a large variety of differently sized areas to be used to estimate the background count rate. This graph does not represent the distribution of the background count rate per square degree within this field.

One additional benefit of observing sources close to the SEP is that there is no need to modify the time of observations with a barycentric correction. The light travel time of the sources observed is not affected by the orbit around the Sun. As viewed from the solar system barycenter, *eROSITA* observations of the SEP field are taken regularly every 4 hours, and this interval does not change throughout each eRASS.

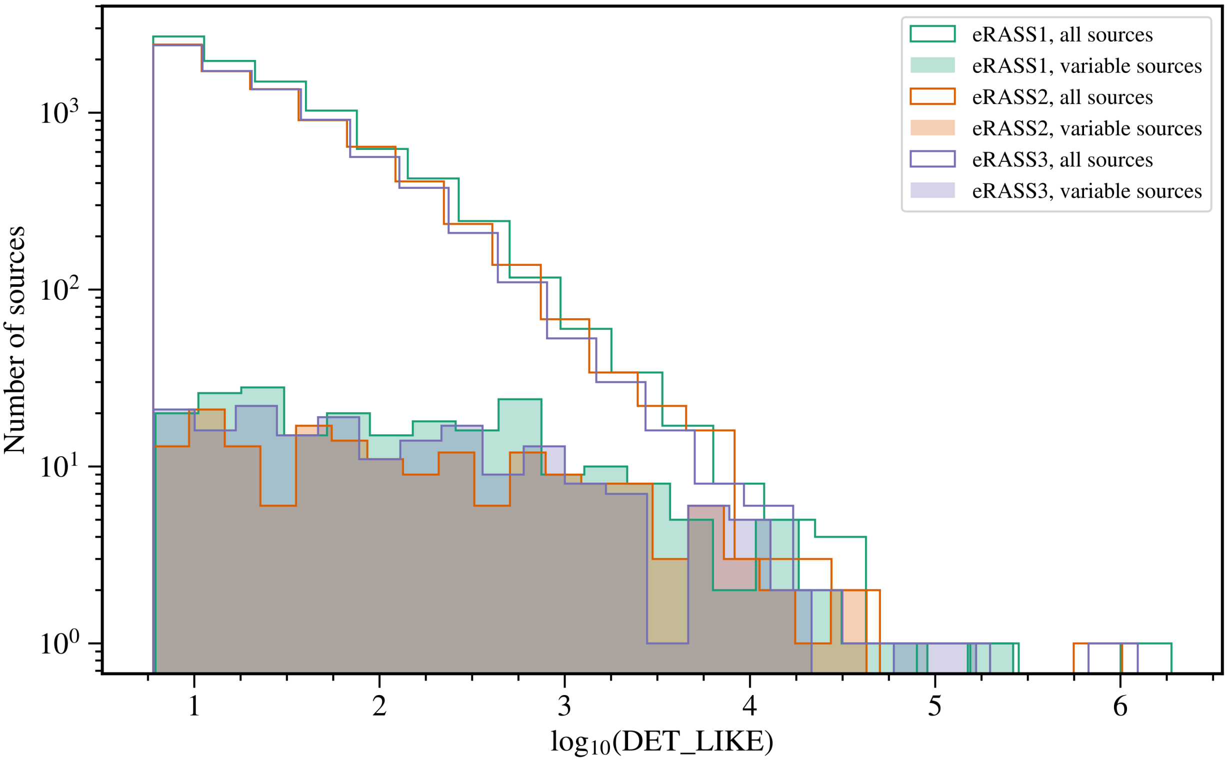

Figure 2.11: The distribution of the detection likelihood parameter of all sources lying in the SEP field in eRASS1, 2, and 3, for the energy band $0.2 - 5.0$ keV.

2.3.6 Detection likelihood

The `DET_LIKE` parameter quantifies the reliability of the source detection. The greater the `DET_LIKE` value is, the greater the probability that a source is genuine. Fig. 2.11 shows the distribution of the detection likelihood for all detected sources in the SEP field in eRASS1, 2, and 3, for the $0.2 - 5.0$ keV energy band. We observe the maximum of the distribution at the lowest likelihoods that are still accepted as a genuine source by the source detection pipeline at `DET_LIKE` $= 6$. This indicates that many sources are barely above the detection threshold and should still be considered as potentially spurious sources. For a detailed discussion of this, see Liu et al. (2022b). The distribution drops rapidly but extends to very high `DET_LIKE`, of up to 1.89×10^6 in eRASS1. Low `DET_LIKE` sources are observed with just a few counts above the expected background level. The three eRASSs feature a similar distribution of the detection likelihood parameter in the SEP field.

Chapter 3

Characterising *eROSITA* light curve variability

During its all-sky survey phase, the *eROSITA* X-ray telescope onboard SRG scans through the ecliptic poles every 4 hours. This extensive data set of long duration, frequent, and consistent observations of thousands of X-ray sources is ideal for a detailed X-ray variability analysis on frequency scales only rarely probed.

Sources exhibiting significant X-ray variability could be AGNs, XRBs, or stars. The X-ray variability in AGNs and XRBs is produced in the innermost part of the accretion disc and the corona (Czerny, 2004). Many stars also exhibit X-ray variability, typically in the form of flares. These are generated by particle acceleration in the stellar corona due to chromospheric evaporation (Antonucci et al., 1984). X-ray variability studies of stars can provide constraints on these processes.

However, a variability analysis of *eROSITA* data faces various challenges, including varying exposure times, low count rates, and long gaps between short observations. In this chapter, we describe the variability methods we used. We also modified some of these methods to minimise the detrimental effects of survey mode observations, and to improve upon previous methods.

3.1 Variability methods

This section introduces various methods for analysing variable sources, particularly BHTs and AGNs. These methods can be used to distinguish strongly from weakly variable sources and quantify their degree of variability.

3.1.1 Normalised excess variance and the normalised intrinsic variance

A frequently used parameter for quantifying the degree of variability of X-ray light curves of AGNs and XRBs is the normalised excess variance (NEV; Edelson et al., 1990; Nandra et al., 1997):

$$\mathrm{NEV} = \frac{\sigma^2_{obs} - \overline{\sigma^2_{err}}}{\overline{R_S}^2}, \tag{3.1}$$

where the excess variance, $\sigma^2_{obs} - \sigma^2_{err}$, describes how much the observed variance in the source count rates, σ^2_{obs}, exceeds the expected variance due to the Poisson errors, σ^2_{err}. The excess variance is then normalised by the square of the average of the measured source count rate, $\overline{R_S}^2$, to form the NEV. The normalisation is necessary to ensure that the variability measurement does not depend on the brightness of the source and that the measured NEV of a variable source can be compared between instruments of different sensitivities.

The NEV is an estimator of the normalised variance of the brightness of a variable source for a particular set of bins and a particular sampling pattern. The quantity that the NEV estimates is intrinsic to the source but also depends on the properties of the observation. It describes the variability of the source at the particular times when it was observed. Therefore, we refer to it as the normalised intrinsic variance (NIV). The NIV is equal to the NEV that would be calculated if the observations were obtained with infinite accuracy and negligible Poisson noise, but with the same observing pattern over the same period, with the same bin size, observed at the same time. It is found from Eq. 3.1, when $\overline{\sigma^2_{err}} = 0$. The NIV is unaffected by the source brightness. However, the NIV of variable sources is not constant but changes over time, even if the mechanisms causing the variability remain unchanged. The geometric mean of the NIV of an infinite number of light curves, all covering the same frequency interval and caused by the same variability, is defined as NIV_∞. The NIV of any finite light curve segment varies around the NIV_∞.

For N_b bins of a variable source at times t_i, where $i \in \{1, 2, .., N_b\}$, with a measured, background subtracted, exposure-corrected source count rate of $R_S(t_i)$, and associated measurement uncertainties of $\sigma_{err}(t_i)$, the quantities in Eq. 3.1 are defined by:

$$\begin{aligned} \overline{R_S} &= \frac{1}{N_b} \sum_{i=1}^{N_b} R_S(t_i) \\ \sigma^2_{obs} &= \frac{1}{N_b - 1} \sum_{i=1}^{N_b} \left(R_S(t_i) - \overline{R_S} \right)^2 \\ \overline{\sigma^2_{err}} &= \frac{1}{N_b} \sum_{i=1}^{N_b} \sigma^2_{err}(t_i)\,. \end{aligned} \tag{3.2}$$

If the true source flux at all the bins at times t_i were known, the same equations could be used to calculate the NIV, but with $\sigma^2_{err}(t_i) = 0$. The lowest NIV that a source can have is 0. This occurs when its flux does not change at all during the observing window, so $\sigma^2_{obs,NIV} = 0$. The highest value that the NIV can have occurs when the flux of a variable source is 0 in all bins, except one. For such a light curve, $\overline{R_{S,NIV}} = R_0/N_b$, where R_0 is the only non-zero true count rate of the source. Therefore, the variance of this true light curve is $\sigma^2_{obs,NIV} = R_0^2/N_b$, so NIV $= N_b$. Therefore, the NIV can only have values within the $[0, N_b]$ range.

In contrast, the NEV is sometimes measured to lie outside this range. The definition of the NEV, as described in Eqs. 3.1, and 3.2 works well for estimating the NIV of sources with light curves consisting of many bins, each containing many source counts above the background level. It is accurate as long as the probability distribution of the measured source counts in each bin approximately follows a normal distribution. However, for light curves consisting of a few bins or a few source counts per bin, it is not unlikely to measure a negative excess variance when the observed count rates vary less than their uncertainties indicate that they should. The NEV is most useful for investigating the variability of light curves consisting of $\gtrsim 20$ bins (Turner et al., 1999).

3.1.2 Root mean square variability

The root mean square variability (rms) is a similar variability quantifier to the NEV. Whereas the NEV is commonly used to investigate the variability of AGNs, XRB variability analysis frequently uses the rms instead. It is defined as:

$$\text{rms} = \sqrt{\text{NEV}}, \tag{3.3}$$

and is often expressed as a percentage. It estimates the true, normalised standard deviation of variations in a light curve within the frequency space observed. Even though the NEV and the rms can be larger than 1, the type of variability typically exhibited by XRBs in X-rays is constrained to be below 1. Using the rms as an estimator relies on the NEV being positive.

3.1.3 Maximum amplitude variation

Another method for quantifying the degree of variability of a variable source is the maximum amplitude variation significance (AMPL_SIG; Boller et al., 2016). The standard definition of the maximum amplitude variation (AMPL_MAX) uses the bins in which the highest and lowest count rates were measured, which we denote to have occurred at times t_{max}, and t_{min}, respectively. Then AMPL_MAX, and its significance, AMPL_SIG, are defined as:

$$\begin{aligned}\text{AMPL\_MAX} &= [R_S(t_{max}) - \sigma_{err}(t_{max})] - [R_S(t_{min}) + \sigma_{err}(t_{min})] \\ \text{AMPL\_SIG} &= \frac{\text{AMPL\_MAX}}{\sqrt{\sigma^2_{err}(t_{max}) + \sigma^2_{err}(t_{min})}},\end{aligned} \tag{3.4}$$

where $\sigma_{err}(t_i)$ is the uncertainty in each measured count rate at time t_i. It is a more useful variability quantifier for short light curves and sources exhibiting flares.

This may be a more useful variability quantifier for analysing short light curves or ones that feature flares. However, AMPL_MAX only considers the two most extreme points of a light curve, regardless of how long it is. The definition of AMPL_SIG in Equation 3.4 is invalid in the Poisson regime.

3.1.4 Bayesian excess variance

A third method we will use to quantify the variability is the Bayesian excess variance[1] (bexvar; Buchner et al., 2022). Bexvar uses a hierarchical Bayesian model to determine a posterior probability distribution for the mean and standard deviation of an assumed log-normal source count rate distribution. Background, instrument, and Poisson variability are modelled out. Bexvar models the variability intrinsic to the source, in addition to Poisson variability, background and instrument sensitivity variations with a hierarchical Bayesian model. The source variability is assumed to follow a log-normal count rate distribution, with the mean indicating the logarithm of the mean count rate) and the standard deviation, σ_l, quantifying the strength of the intrinsic variability. Bexvar infers the posterior probability distribution of these parameters. We refer to samples from the posterior probability distribution of σ_l as σ_b. Buchner et al. (2022) also introduced the quantity SCATT_LO, the 10% quantile of the distribution of the σ_b samples, which is useful for distinguishing between variable and non-variable sources. The geometric mean of the samples ($\sigma_{b,i}$) found by bexvar;

$$\overline{\sigma_b} = 10^{\sum_i \log\left(\sigma_{b,i}\right)} \tag{3.5}$$

provides an estimate of σ_I. The 15.87% and 84.13% quantiles of the σ_b distribution were selected as the upper and lower bounds of the estimate of σ_I. These define the 1σ uncertainties in the measurement.

The NIV and the σ_I are independent quantifiers of the intrinsic variability of a source. Both are unaffected by Poisson noise or measurement uncertainties but depend on the timing of the observations: the start time, the bin size, and the separation of the bins. Although σ_I is not normalised, as the NIV is, it is also invariable to a multiplicative scaling of the flux, by being defined on a logarithmic scale. The NIV describes a variance, and σ_I describes a standard deviation. Nevertheless, as the NIV is defined for the linear flux, and σ_I is defined for a logarithmic flux, the two quantities are not related by a square; $\sigma_I \neq \sqrt{\mathrm{NIV}}$.

The strength of bexvar lies in the self-consistent Bayesian handling, modelling the observed counts with a Poisson distribution and propagating the probability distributions. Unlike the NEV, which is often found to have negative values, which are useless for further analysis, σ_b and $\overline{\sigma_b}$ can never be negative. Bexvar uses the Poisson probability distributions of the measured count rate in each bin.

---

[1] https://github.com/JohannesBuchner/bexvar

3.1.5 Power spectral density and periodograms

None of the variability quantifiers discussed so far (NEV, rms, AMPL_SIG, and bexvar) consider the time ordering of the bins in the light curve. They do not require the light curve to consist of regularly spaced bins and are unaffected if the bins and their associated count rates are rearranged. They investigate to what extent the measurements of the count rate deviate from one another.

In contrast, the PSD is a variability quantifier which depends on the time ordering of bins in a light curve. It is a potentially more informative method for investigating variable sources, which details how the total observed variability is caused by modulation at individual frequencies. The greater the value of the PSD at a particular frequency, the more significantly the overall variability is driven by variation at that frequency. The PSD describes the correlation between bins of the light curve as a function of their temporal separation.

For a light curve consisting of N_b consistently sampled bins, with a constant separation between the start of one bin and the start of the next one of τ, the periodogram estimates the variability power within the frequency range from $(N_b\tau)^{-1}$ to $(2\tau)^{-1}$, as a function of the frequency. This frequency range stretches from the inverse of the duration of the light curve to the Nyquist frequency, which is equal to half of the sampling frequency. The periodogram is calculated as the modulus square of the Fourier transform of the count rates observed (van der Klis, 1989).

If each of the bins of a light curve features data collected over an exposure time of Δt (which may or may not be equal to τ), then the periodogram is found via (van der Klis, 1989):

$$a_j = \sum_{k=0}^{N_b} \left(R_S(t_k) \Delta t \; e^{2\pi k j / N_b} \right) \tag{3.6}$$

$$P'_j \propto |a_j|^2 , \tag{3.7}$$

where P'_j is the periodogram power at a frequency of $j/(N_b\tau)$. The periodogram frequencies are defined for values of $j = 1, 2, ..., N_b/2$, where $N_b/2$ is rounded down if N_b is odd. A is a normalising constant, which can depend on $\overline{R_S}$ (as defined by Equation 3.2), N_b, Δt, and τ.

There are different ways in which a PSD and periodogram can be normalised. We predominantly used the fractional rms normalisation (Belloni and Hasinger, 1990). We chose this normalisation due to the useful feature that the NEV is equal to the integral of the Poisson noise subtracted periodogram between the frequencies of $(N_b\tau)^{-1}$ and $(2\tau)^{-1}$:

$$\text{NEV} \approx \int_{1/\tau N_b}^{1/2\tau} P(\nu) \; d\nu, \tag{3.8}$$

where ν is the frequency associated with the Poisson noise subtracted periodogram power, $P(\nu)$. This follows from Parseval's theorem (van der Klis, 1989).

There is an inherent variability to all measured count rates due to Poisson statistics. This increases all measured powers by a constant value independent of frequency. It acts like a background that needs to be removed to estimate the nature of the source variability. The Poisson noise depends on the selected periodogram normalisation. For the fractional rms normalisation, the Poisson noise (N_P) in the periodogram has a value of (Belloni and Hasinger, 1990):

$$N_P = \frac{2}{\overline{R_S}}. \tag{3.9}$$

To accurately estimate the PSD by computing a periodogram requires more information on the source flux at different times than is required to estimate the NIV, the σ_I, or for measuring AMPL_MAX. Therefore, for *eROSITA* observations, this detailed analysis will only be possible for bright variable sources close to either of the two ecliptic poles.

The Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982) is a slight modification to the standard periodogram methodology. Its primary function is to detect periodicities in the data. It highlights frequencies that dominate in the overall variability observed. Unlike the standard periodogram, the Lomb-Scargle periodogram does not require consistently sampled data. It is also usually calculated over a larger frequency interval than the periodogram.

Periodograms of light curves were computed with the Stingray[2] timing package (Huppenkothen et al., 2016). We particularly used the Stingray `Powerspectrum` function, which computes periodograms using a fast Fourier transform algorithm, and can be set to different normalisations.

3.1.6 Band-limited power

The integral of the true PSD of a variable source within a selected frequency range is a useful quantity to seek to estimate. The true PSD describes the variability power of an infinitely long, true light curve of the variable source, unaffected by sampling, or a limited frequency range. We refer to the integral of the true PSD between two frequencies as the band-limited power. It describes how variable a source is within a particular frequency interval. It differs from other variability measures described so far, as it is a stationary intrinsic quantity of the source that does not depend on windowing, sampling, or the strength of the source variability at higher or lower frequencies.

The band-limited power is a constant quantity as long as the true PSD does not vary with time. The NIV of a light curve varies around the NIV$_\infty$, which is also a constant quantity but is offset from the band-limited power due to the power leakage. The properties of the observations can induce variability power at both lower (red noise leak) and higher frequencies (aliasing) than those being investigated, to leak into the $[(N_b\tau)^{-1}, (2\tau)^{-1}]$ frequency space, and increase the power measured within it. This will be discussed in more detail in Section 3.7.1.

---

[2] https://docs.stingray.science/

When comparing the variability of different sources, when considering the scaling of the variability with other properties, such as the BH mass or the AGN luminosity, or when investigating whether the variability of an individual source changed over time, it is preferable to compare estimates of the band-limited power. The NEV, AMPL_SIG, $\overline{\sigma_b}$, and the periodogram are all non-intrinsic, measurement-dependent quantities prone to significant sampling errors. Comparing these quantities between different sets of observations of the same source or between different sources, can lead to changes caused by windowing or sampling effects being falsely identified as changes in the source variability. Nonetheless, these variability measures can be used to estimate the band-limited power. But to do so, it is imperative to consider the biases and sampling errors of such estimates.

The first step to calculate the band-limited power is to estimate the NIV. This could be done by calculating the NEV, integrating the periodogram, or using $\overline{\sigma_b}$, along with the conversion from σ_I to the NIV, that is described in Section 3.5. Then, by considering the distribution of the NIV of observations performed at different times (which will be discussed in Sections 3.7.2 and 3.7.3), the NIV estimate can act as an estimator of the NIV_∞, with a larger uncertainty, to account for the sampling error. Finally, this can be converted into an estimate of the band-limited power, by quantifying the strength of the power leakage from higher and lower frequencies into the $[(N_b\tau)^{-1}, (2\tau)^{-1}]$ frequency space, and subtracting that from the estimate. However, this final step requires several assumptions about the shape of the true source PSD in the frequency space not investigated. That can introduce uncertainty in the estimate of the band-limited power. Therefore, unless the strength of the power leakage can be reliably estimated without bias, it can be preferable to compare estimates of NIV_∞ instead.

3.2 Challenges for eRASS light curve variability analysis

3.2.1 Varying fractional exposure

A key challenge faced when analysing the variability properties of *eROSITA* light curves, and the key feature that differentiates an *eROSITA* variability analysis from many other variability analyses, is the varying fractional exposure in the light curves. One way in which this affects the light curve is that it makes it appear to bend upwards towards a higher count rate when the fractional exposure decreases. Combined with the Poisson nature of the light curves, this makes it challenging to identify variable light curves by visual inspection.

Unless the source is very faint or strongly variable, the highest count rates are almost always measured at the smallest fractional exposures. For sources located further than 0.5° away from either of the survey poles, the fractional exposure is lowest towards the start and end of the interval during which it is observed. This effect makes light curves appear to have a 'U'-like shape, as shown in Fig. 3.1. Intrinsically non-variable sources can appear visually similar to significantly variable sources.

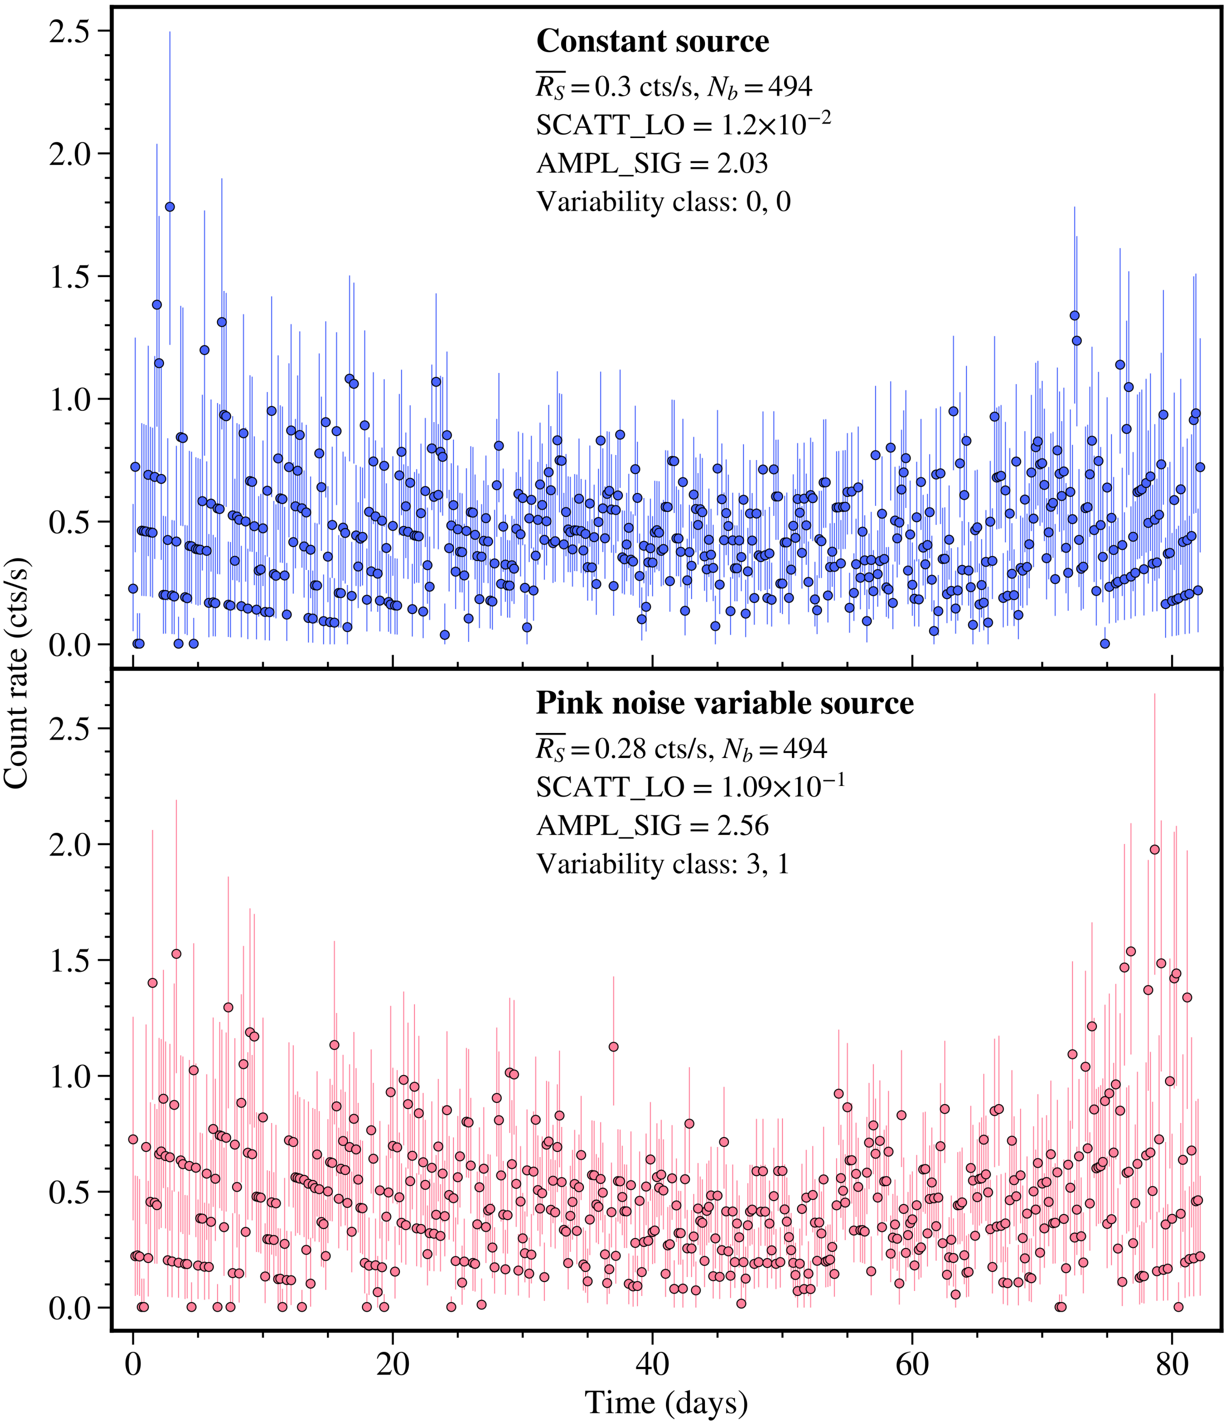

Figure 3.1: Simulated *eROSITA*-like light curves of sources close to either of the two ecliptic poles. These light curves are computed using a typical *eROSITA* ϵ evolution, as observed in actual *eROSITA* light curves of sources lying close to the ecliptic poles. The top panel shows an intrinsically constant source, and the lower panel depicts an intrinsically variable source exhibiting pink noise variability. The variability class describes how the value of SCATT_LO and AMPL_SIG compares to the variability thresholds. The pink noise variable source is identified to have a SCATT_LO value above the 3σ threshold but an AMPL_SIG value only above the corresponding 1σ threshold.

The ‘U’-like shape of the light curves can be reduced by rebinning it over several eroDays. This allows for easier visual identification of the long timescale intrinsic variability of the sources themselves, but is only possible for sources located close to the ecliptic poles.

However, rebinning implies a loss of information on eroDay to eroDay variability. Therefore, for the variability analysis, it is preferable to avoid rebinning where possible. Instead, we develop methods to best deal with the effects described here to minimise their impact on the variability estimate while retaining all information on the eroDay to eroDay variability. Furthermore, any effective variability analysis of eRASS sources needs to cope with the significantly varying exposure times per bin experienced by every source and the effect this has on the observed light curves.

Fig. 3.1 shows two simulated *eROSITA*-like light curves of sources lying close to one of the two ecliptic poles for 494 bins with $\epsilon > 0.1$ in one eRASS, with an average true source count rate of ≈ 0.3 cts/s. For both the constant and the variable source, the ‘U’-like shape dominates the light curve, making it difficult to distinguish between them.

Another challenge for analysing *eROSITA* light curves is that they consist of short observation times (of 41.2 s or less), separated by long but consistent gaps (4 h). For *eROSITA*, the gaps are at least 350 times longer than the duration of each bin. This means that variability which occurs during the long gaps will go undetected.

The *eROSITA* observations, particularly in the SEP field, also constitute an amazing opportunity for variability analysis. Observations of tens of thousands of X-ray sources close to the ecliptic poles occur frequently and consistently over several years. *eROSITA* observations enable a long-term X-ray variability analysis of a large, unbiased sample of AGNs, XRBs, and stars. For this purpose, we first need to decide on the variability quantifiers that are most useful for an *eROSITA* variability analysis.

The previously established variability detection and quantification methods introduced in Section 3.1 need to be considered in a different light when analysing *eROSITA* data. Here we discuss what particular features need to be taken into account.

3.2.2 NEV

Special care must be taken when using the NEV for *eROSITA* analysis. For low count rate sources, the probability distribution of the measured number of counts per bin is strongly Poissonian, differing significantly from the Gaussian distribution that is assumed by Eqs. 3.1, and 3.2. In such instances, the uncertainties in the measured count rates in each bin are asymmetric and do not have constant intervals between $n\sigma$ and $(n+1)\sigma$.

The NEV relies on accurate estimates of the uncertainties in the measured count rates. We usually used the asymmetric uncertainties in the count rate found by Eq. 2.5. However, to calculate the NEV, it is necessary to define a single error (Eq. 3.1). We chose this single error, σ_{err}, to be equal to the size of the uncertainty in the direction of the mean. If $R_S(t_i) > \overline{R_S}$, we used $\sigma^2_{-err}(t_i)$ as the contribution of bin i to $\overline{\sigma^2_{err}}$. In contrast, if $R_S(t_i) < \overline{R_S}$, we used $\sigma^2_{+err}(t_i)$ instead.

The NEV calculated by Eqs. 3.1 and 3.2 remains an inaccurate estimate of the NIV whenever the probability distribution on the measured counts per bin significantly deviates from a normal distribution. At low count rates, small differences between the variance in the light curve and the expected variation from the measurement uncertainties are pushed to very positive or negative values by the small denominator of the NEV. It is common for faint sources to have negative NEVs. Light curves with negative NEVs are frequently discarded from further variability analysis in literature (see e.g. O'Neill et al., 2005; Lu and Yu, 2001).

The NEV, as defined by Eqs. 3.1, and 3.2, also faces several additional challenges when applied to *eROSITA* data. Due to the varying fractional exposure, the amount of information obtained about the intrinsic source brightness varies from bin to bin. However, the parameters of the NEV, as described by Eq. 3.2 depend equally on all bins. In light curves featuring varying fractional exposures, the parameters in Eq. 3.2 are biased by, and depend too strongly on, the bins with the lowest fractional exposures. This results in an inaccurate estimate of these quantities, including the NEV.

Almost all sources observed by *eROSITA* lie in the Poisson regime at least during some eroDays at low fractional exposures. Placing a higher limit on the minimal fractional exposure risks losing valuable data on the source brightness. Furthermore, the incorrect assumption of a normal probability distribution for the measured number of source counts per bin causes the excess variance, $\sigma_{obs}^2 - \overline{\sigma_{err}^2}$, to be negative even more frequently than for light curves with a constant fractional exposure. Therefore, as it is usually defined, the NEV is only of limited use for analysing eRASS data.

The NEV is also not useful for analysing the variability of light curves consisting of very few bins. The ability of the NEV to reflect the true variability of a source increases with the number of bins at a fixed average number of source counts per bin. We only calculate the NEV of light curves consisting of 20 bins or more. When combining four or more eRASSs of observation, every detected X-ray source in the sky will have been observed at least 24 times. This is, however, not relevant to this work.

3.2.3 AMPL_SIG

AMPL_SIG can quickly determine large differences in the count rate observed within a light curve, by only using the two most extreme data points. However, the highest and lowest count rates measured by *eROSITA* often occur in bins with the lowest fractional exposures and the largest uncertainties. This is not ideal for variability analysis, which should not significantly depend on the bins that contain the least information about the brightness of the source. The standard methodology could significantly underestimate or overestimate the variability of an *eROSITA* light curve. It can easily lead to a genuinely variable source to be quantified to have a negative AMPL_MAX if the errors of the highest and lowest count rate measurements overlap. Additionally, this definition does not consider the different uncertainties in the positive and negative directions.

We, therefore, modified the equation to a more useful form for investigating light curves featuring varying exposure times. Rather than using the standard method of comparing the two bins with the highest and lowest measured count rates, we instead use the two bins with the highest lower bound and the lowest upper bound confidence interval. Therefore, in order to optimally analyse *eROSITA* data near the SEP, we redefine t_{max}, t_{min}, AMPL_MAX, and its significance, AMPL_SIG, as follows:

$$\begin{aligned}
&R_S(t_{max}) - \sigma_{-err}(t_{max}) = \max\left[R_S(t_i) - \sigma_{-err}(t_i)\right] \\
&R_S(t_{min}) + \sigma_{+err}(t_{min}) = \min\left[R_S(t_i) + \sigma_{+err}(t_i)\right] \\
&\mathrm{AMPL\_MAX} = \left[R_S(t_{max}) - \sigma_{-err}(t_{max})\right] - \left[R_S(t_{min}) + \sigma_{+err}(t_{min})\right] \\
&\mathrm{AMPL\_SIG} = \frac{\mathrm{AMPL\_MAX}}{\sqrt{\sigma^2_{-err}(t_{max}) + \sigma^2_{+err}(t_{min})}}.
\end{aligned} \tag{3.10}$$

Where σ_{+err}, and σ_{-err} denote the 1σ errors of the measured count rate in the positive and negative directions, respectively. Henceforth, we will use this modified definition of AMPL_SIG.

The AMPL_SIG can be calculated for all *eROSITA* detected sources, regardless of how often they were observed. However, the more bins there are in a light curve, the more efficient AMPL_SIG is at detecting variability. The likelihood of false positive detection also increases with more bins in the light curve. Therefore, we did not rebin any light curves for the AMPL_SIG variability detection and analysis.

3.2.4 bexvar

Bexvar estimates the excess variability power on the timescale of the binning, assuming that each bin has an independently drawn count rate. We used a uniform prior within the $[-2, 2]$ interval for $\log(\sigma_b)$, as this is the range of values we expect to be able to measure for it. Greater degrees of variability are not expected, except for very long light curves featuring a single non-zero flux measurement. Smaller degrees of variability are possible but are unlikely to be distinguished from non-variability in eRASS light curves. As σ_b is defined on a logarithmic interval, $\overline{\sigma_b}$ can never be negative.

Unlike other methods, like the NEV or the AMPL_SIG, which assume a normal probability distribution on the measured count rate, bexvar can yield helpful information about the variability properties of low count rate sources with Poisson statistics.

Nevertheless, a bexvar variability analysis can also benefit from rebinning the light curves of very faint sources, which consist mostly of bins with 0 source counts. Bexvar is more computationally expensive than other variability estimators, and its computation time increases linearly with the number of bins in the light curve. Unless a very faint source exhibits brief, large flares, it is hard to determine a precise or accurate estimate of its variability at the timescale of the separation of the bins. Therefore, rebinning the light curves of very faint sources usually does not reduce the ability to investigate their variability, but reduces the computation time.

Therefore, we chose to rebin light curves of faint sources until an average of at least one source count was contained in every two bins, for detecting variability with bexvar. We also required that the rebinned light curve still consisted of at least 20 bins. Rebinning should not be problematic as long as most of the variability power contained within the frequency interval of the original light curve is maintained below the Nyquist frequency of the rebinned light curve. We expect this to be true for most sources, except ones that exhibit brief flares. For flares, it is preferable to use AMPL_SIG to detect and characterise their variability (Buchner et al., 2022).

3.3 Simulations

Throughout this chapter, we used simulations to assess the effects of survey mode observations on the reliability of various variability methods when applied to such data. The simulations were also used to probe the available parameter space, and optimise the methods for use on actual data.

We performed a variety of different types of simulations, each selected to investigate a particular topic optimally. All simulations start with the source flux of the simulated variable or non-variable source. These light curves do not include any background or Poisson noise. The flux of simulated variable sources was determined by selecting the PSD of the source, and then using the Timmer and Koenig (1995) method to generate a light curve from it. We typically selected pink noise PSDs for this purpose. The break in AGN PSDs, from an $\alpha \approx 1$ power law, to an $\alpha \approx 2$ power law has been observed to often occur within, or above, the frequency range probed by *eROSITA* (González-Martín and Vaughan, 2012). The aliasing effect often counteracts the steeper power law at high frequencies, and flattens the PSD (see Section 3.7.1). Therefore, the PSDs of typical *eROSITA* observed AGNs could, to first order, be assumed to approximately follow an $\alpha = 1$ power law. The accuracy of this first order assumption is also supported by the periodograms of actual *eROSITA* observations of AGNs, as will be discussed in Section 4.5. We also simulated power law PSDs with $\alpha = 0$ (white noise) and $\alpha = 2$ (red noise) for comparison, to determine how strongly the variability methods depend on the PSD shape. Non-variable sources that have a constant flux level were also simulated.

We selected intervals of simulated light curves of at most 1050 bins, as that approximately matches the upper limit on the number of observations of a source that can be made per eRASS. To simulate the red noise leak, we selected input PSDs that extended with the same power law to frequencies of one order of magnitude below the inverse of the total light curve duration. We selected random starting points for the light curves generated by the Timmer and Koenig (1995) method.

The variable light curves thus simulated were scaled such that the mean average source flux matched the desired source count rate at the detector. The NIV in the light curves was varied by modifying the range of fluxes of the simulated light curve, at a constant mean. This scaling procedure also ensured that the simulated source flux stayed above 0. The light curves generated in this way do not contain any background or Poisson noise. Therefore, we refer to them as the true light curves of the simulated sources. The NIV values were computed directly from the mean and variance of the true light curves, using Eq. 3.1, with $\overline{\sigma_{err}^2} = 0$.

To investigate the ability to identify variable sources, and the reliability of estimating the NIV from a light curve, we used the true light curves as a basis for generating observed light curves, with properties as similar as possible to those detected by *eROSITA*. For this purpose, we selected a background count rate of 0.71 cts/s and a background area of 0.0089, which are equal to the mean value of both parameters found for sources detected by *eROSITA* close to the SEP (see Section2.3.5). We also selected each simulated bin to have a 40 s duration. Next, we randomly selected a fractional exposure for each bin from the distribution observed for sources in the SEP field (Section 2.3.2). We cropped this distribution to avoid fractional exposures below 0.1, as we did for the actual data. By modifying the distribution, we ensured that the number of bins of these simulated light curves would equal the number of bins that could be used for variability analysis. In so doing, the dependence of various parameters on the number of *eROSITA* bins with $\epsilon > 0.1$ could be established. The light curves simulated in this way look similar to those in Fig. 3.1, except that they feature a random assortment of fractional exposures. These simulated light curves were only used in conjunction with the SCATT_LO, and AMPL_SIG methods, which do not depend on the time ordering of bins. These types of simulated light curves are referred to as being *eROSITA*-like.

For each bin of the simulated *eROSITA*-like light curve, we randomly selected a measured number of source counts from the Poisson distribution with a mean of $(R_{S,t}(t_i) + R_{B,t}(t_i)A(t_i)B(t_i))\epsilon\Delta t$. In this equation, $R_{S,t}(t_i)$, and $R_{B,t}(t_i)$ are the true source and background count rates at time t_i, converted from the true source and background fluxes to the properties of an on-axis *eROSITA* observation. We, similarly, selected a measured number of background counts from the Poisson distribution of $R_{B,t}(t_i)\epsilon\Delta t$. Together with the fractional exposure, the time bin size, and the background area, we determined what count rate would be measured in each bin, given the number of source and background counts selected in this simulation.

We also simulated light curves with fractional exposure distributions that were not based on any observed distribution. This was done to investigate the dependence of the noise level in a periodogram on the mean and variance of the fractional exposure distribution in the light curve. To span a large parameter space in both parameters, we first selected a minimum and maximum fractional exposure, drawn from a grid of 40 equally spaced points between 0.1 and 1.0, covering all possible combinations.

We generated scenarios for a maximum, minimum, and intermediate fractional exposure variance for each range of fractional exposures, at a fixed mean. For the maximum variance scenario, we assigned half of all bins to the maximum fractional exposure, and half to the minimum fractional exposure. The ordering of the fractional exposures does not influence the periodogram noise level. For the minimum variance scenario, we assigned one randomly chosen bin with the maximum fractional exposure, one bin with the minimum fractional exposure, and kept the rest at the value halfway between the two. For the intermediate variance scenario, we selected fractional exposures to cover the interval with a constant incremental increase from start to end. We refer to the simulated light curves created in this way as the patterned fractional exposure light curves.

In this and the next chapter, we investigated various analytical models. These were fitted with the nested sampling Monte Carlo algorithm MLFriends (Buchner, 2014, 2019) through the UltraNest[3] package (Buchner, 2021).

3.4 Methods for identifying varying sources

Out of the millions of X-ray sources detected by *eROSITA* (Predehl et al., 2021), we intend to select a much smaller set of sources whose count rate changes significantly throughout the observing interval. These can subsequently be investigated individually, which allows their variability properties to be studied in more detail. As described in Section 2.3, eRASS observations of X-ray sources feature a large variety of different properties. We aim to be able to detect significantly variable sources throughout the observed parameter space, and not restrict ourselves to only analysing the brightest sources observed for the greatest number of eroDays with the largest average fractional exposures.

We did not seek to optimally divide the sample into likely variable and likely non-variable sources. Instead, we aim to select variable sources at a low false positive rate. In addition, we intend the variability thresholds to not be biased towards particular types of variability, and to be able to detect unusual types of variability.

Buchner et al. (2022) investigated the ability of the AMPL_SIG, NEV, SCATT_LO, and Bayesian block methods to detect flaring, white noise ($P \propto \nu^0$), and red noise ($P \propto \nu^{-2}$) variability, for the *eROSITA* Final Equatorial-Depth Survey (eFEDS). Of the four methods, Buchner et al. (2022) found that SCATT_LO is almost always the most sensitive to detecting variability, regardless of the type of variability. They, however, also found that AMPL_SIG is slightly better at detecting flaring sources at high count rates than SCATT_LO. Both AMPL_MAX and SCATT_LO were designed to quantify variability, rather than to distinguish variable from non-variable sources. Nevertheless, they can both be used for that purpose as well. Following the conclusions of Buchner et al. (2022), we decided to use both SCATT_LO, and AMPL_SIG to distinguish likely variable from likely non-variable sources in eRASS data sets as well.

To use SCATT_LO and AMPL_SIG for variability detection in the eRASSs, we need to properly define thresholds on both of these quantities, to distinguish likely variable from likely non-variable sources. The thresholds identified by Buchner et al. (2022) were derived for typical eFEDS observations, which differ from eRASS observations in several ways. The exposure depth varies more across the sky in the eRASSs than it does within the eFEDS observations, spanning $6 - 1080$ eroDays of observation per source, per eRASS. Close to the poles, there is also an enhanced sensitivity to detect lower count rate sources than in eFEDS. The thresholds for distinguishing likely variable from likely non-variable sources using SCATT_LO, and AMPL_SIG, depend on both the count rate and the number of bins. Since the range of both of these parameters in the eRASSs differs from what was observed in eFEDS, we set out to define new thresholds for distinguishing variable and non-variable sources, using SCATT_LO, and AMPL_SIG, specifically tuned for use in the eRASSs.

---

[3] https://johannesbuchner.github.io/UltraNest/

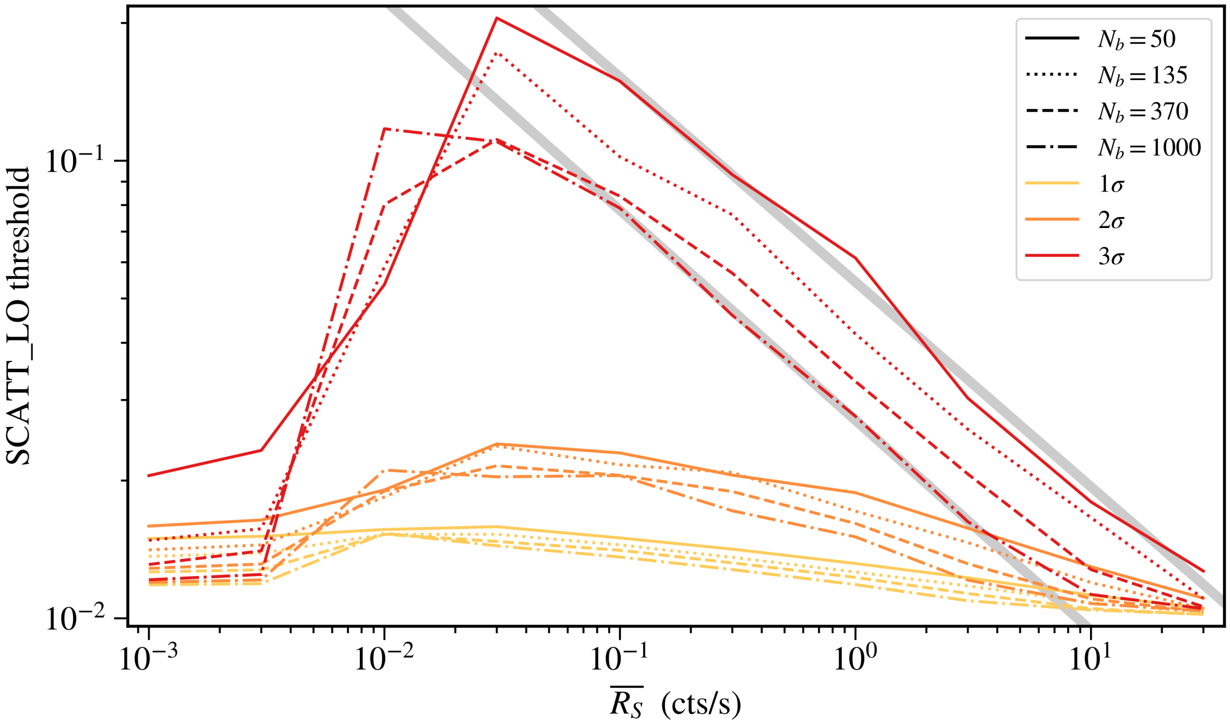

Figure 3.2: The 1, 2, and 3σ thresholds on SCATT_LO, for identifying variable sources. The thresholds are displayed as a function of the count rate. The dependence on the number of bins is illustrated by using different line styles. The grey lines indicate the best fit power law relationship to the decreasing thresholds with increasing count rate.

To do this, we simulated 4×10^5 *eROSITA*-like light curves of intrinsically non-variable sources, as discussed in Section 3.3. We simulated 10^4 iterations of light curves for 40 sets of combinations of intrinsic source count rates of $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1.0, 3.0, 10, 30\}$ cts/s, and number of eroDays of observation with $\epsilon > 0.1$, of $\{50, 135, 370, 1000\}$. This range of count rates and number of bins was selected to be representative of eRASS sources. When considering a single eRASS, this broad parameter space applies to most X-ray sources within ≈ 7° of either of the two ecliptic poles. When combining eight eRASSs of observations, all sources in the sky will have been observed on at least 48 different eroDays. So these thresholds can be used throughout almost the entire sky for an eRASS:8 data set. However, sources close to the poles will have been observed far more often than the upper limit of 1080 bins and will require separate thresholds not considered in this work.

For each simulated light curve, we computed the SCATT_LO and AMPL_SIG parameters. From the resulting distribution of values, we determined one-tailed 1σ (84.13%), 2σ (97.72%), and 3σ (99.865%) equivalent quantiles for each input count rate and number of bins. These are displayed in Figures 3.2, and 3.3, for SCATT_LO, and AMPL_SIG, respectively.

One of these three thresholds needs to be selected, depending on the desired sensitivity to detect variability, and the purity of the selection. For instance, variable sources could be selected as those whose light curves have SCATT_LO values above the 3σ threshold. For that selection, only 0.135% of all intrinsically non-variable sources would be wrongly identified as variable. However, it also excludes moderately variable sources.

The thresholds on both SCATT_LO, and AMPL_SIG, are significantly affected by both the count rate and the number of bins. The thresholds appear to have a stronger dependence on the count rate. This may partly be because the range of simulated count rates spans 4.5 orders of magnitude, whereas the range of number of bins only spans 1.3 orders of magnitude. However, the number of bins of the light curve still affects the value of the threshold, which should be accounted for.

The SCATT_LO thresholds initially rise sharply with increasing count rate. They reach a peak at $0.01-0.03$ cts/s, before gradually declining over three orders of magnitude in the count rate. The peak is most prominent for the 3σ threshold. The 1σ threshold only has a very weak peak and does not change much as a function of the count rate or the number of bins. At count rates above the peak, the three thresholds converge to one another, and towards SCATT_LO $= 10^{-2}$. This is a consequence of the choice of prior for $\log(\sigma_b)$, which has a minimum value of -2. We found that there were no benefits to reducing the minimum value of the prior further, as sources with such a low standard deviation are effectively non-variable sources within the duration of the observations. The convergence of the thresholds to this limit is an indication of the sensitivity of the method.

The decline of the 3σ thresholds above $\overline{R_S} = 0.1$ cts/s approximately follows a power law of SCATT_LO $\propto \overline{R_S}^{-0.45}$, which is shown via the grey lines for light curves of 50 and 1000 bins in Fig. 3.2. With an increasing number of bins, the SCATT_LO thresholds shift to slightly lower values. The amplitude of the power law decline as a function of the count rate approximately depends on the number of bins as $\propto N_b^{-0.22}$. This power law index is about half of that describing the thresholds as a function of the count rate. Therefore, they can approximately be described as SCATT_LO $\propto \left(\overline{R_S}/\sqrt{N_b}\right)^{-0.45}$. Both of these trends occur because more bins and a higher average number of counts per bin improve the ability to accurately determine the degree of variability of the source, resulting in a narrower posterior distribution and a lower value of SCATT_LO. We also observe that the peak in the SCATT_LO thresholds appears to shift to lower count rates at a greater number of bins.

The AMPL_SIG thresholds have a very different dependence on the count rate and the number of bins. The three thresholds are much closer together than the SCATT_LO thresholds are. This makes this method more prone to measurement uncertainties of AMPL_SIG. The thresholds are dominated by a gradual rise with an increasing count rate until a shallow peak is reached at a count rate of 1.0 cts/s. The thresholds drop slightly at higher count rates, before plateauing towards the highest count rates we investigated. This general shape of the thresholds can be understood to be a consequence of the accuracy of the assumptions of normal probability distributions on the measured count rate per bin. The peak of the 1σ thresholds appears to be shifted to slightly lower count rates for light curves with a greater number of bins.

The dependence of the AMPL_SIG thresholds on the number of bins is the opposite of

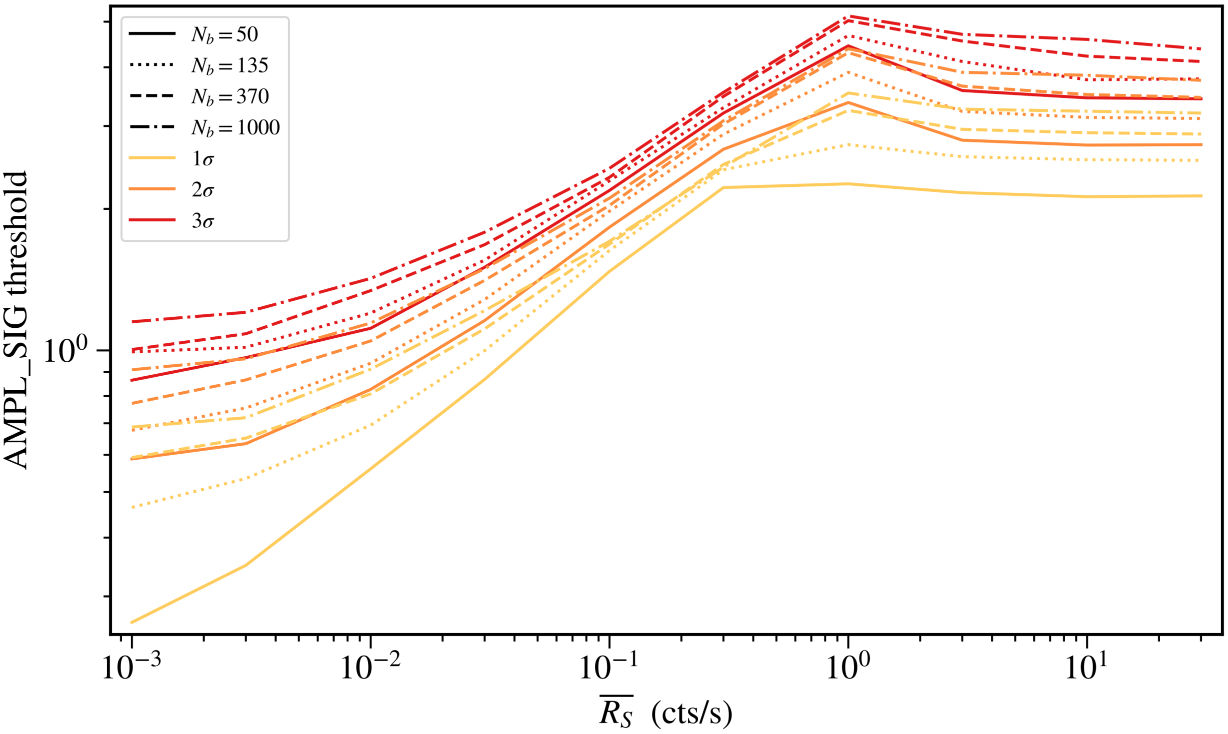

Figure 3.3: The 1, 2, and 3σ threshold conditions on AMPL_SIG, for identifying variable sources. These thresholds are shown as a function of the count rate, with the dependence on the number of bins being illustrated through lines with different line styles, as in Fig. 3.2.

what was observed for the SCATT_LO thresholds. The more bins there are, the greater the probability that at least one bin from an intrinsically constant source is measured to have a significantly higher or lower count rate than expected. Since AMPL_SIG only considers the two most extreme data points, a higher value of AMPL_SIG is found for intrinsically non-variable sources with longer light curves, on average.

These results are qualitatively similar to what was found by Buchner et al. (2022). Their thresholds are different from the ones presented here, as they were calculated for a different number of bins in the simulated light curves. They also used a different false positive rate for defining the thresholds. The dependence of the SCATT_LO thresholds on the count rate in Fig. 3.2 is very similar to what Buchner et al. (2022) found. But there is a noticeable difference between the dependence of the AMPL_SIG thresholds on the count rate in these two instances. Buchner et al. (2022) also found that the AMPL_SIG thresholds increase with increasing count rate. However, while the thresholds shown in Fig. 3.3 reach a plateau at the highest count rates, Buchner et al. (2022) instead found that they keep increasing with increasing count rate. The likely reason for these differences is that the two thresholds use different definitions of AMPL_SIG. We modified the AMPL_SIG function to be better suited for analysing *eROSITA*-like light curves featuring variable fractional exposures (Eq. 3.10), whereas Buchner et al. (2022) used the standard definition of AMPL_SIG (Eq. 3.4).

We applied the two variability identification methodologies to the intrinsically non-variable and variable simulated *eROSITA*-like light curves of Fig. 3.1. Both methods located the non-variable source, shown in the top panel, in the variability class $< 1\sigma$, indicating a low variability significance that might not be distinguishable from non-variability. In contrast, the intrinsically variable source exhibiting pink noise variability, shown in the lower panel, was identified as variable above the 3σ threshold by SCATT_LO. AMPL_SIG instead identified it between the 1 and 2σ thresholds. This is a consequence of SCATT_LO being more sensitive to detecting pink noise variability.

3.5 Intrinsic variance estimation

The thresholds on SCATT_LO and AMPL_SIG were set up principally to identify variable sources at a given false positive rate. They do not necessarily indicate the strength of the variability of a given source within a set of observations. In this section, we quantify the degree of variability, rather than determine the variability significance.

A successful method for estimating the NIV should be relatively unaffected by the features of eRASS light curves and remain accurate at both low and high count rates. It should adequately determine the error in the estimate of the NIV, and apply to all types of variability with all kinds of PSDs.

Here, we investigate the possibility of converting the bexvar $\overline{\sigma_b}$ estimate of σ_I into an estimate of the NIV of a light curve. Unlike the standard NEV methodology, bexvar does not treat all bins identically. Instead, it takes the varying fractional exposures into account. Furthermore, Bexvar does not assume that the measured count rate in each bin follows a normal distribution, unlike the standard NEV method, and instead uses the Poisson probability distributions for the count rate in each bin to determine the degree of variability. These two features enable bexvar to estimate σ_I and the error in the measurement accurately, for a log-normal white noise ($P \propto \nu^0$) light curve, over a wide range of count rates, as was shown by Buchner et al. (2022), their Figure A.1.

We simulated *eROSITA*-like light curves of intrinsically variable sources exhibiting pink noise variability, as described in Section 3.3, in order to investigate the ability of $\overline{\sigma_b}$ to accurately estimate the σ_I for these types of sources. We investigated light curves consisting of $\{20, 75, 150, 400, 1050\}$ bins, with mean input count rates of $\{0.015, 0.15, 1.5, 15\}$ cts/s. This range of count rates and number of bins was chosen to cover the parameter space eRASS light curves.

The results of these simulations are shown in Fig. 3.4. The σ_I was determined from the true light curve, before adding Poisson noise and a background count rate. The $\overline{\sigma_b}$ was estimated from the set of simulated source and background extraction region counts, the background area, the time bin size, and the fractional exposure. As this figure shows, $\overline{\sigma_b}$ is also an accurate estimator of σ_I for sources exhibiting pink noise variability. This method is accurate across a wide range of count rates, number of bins, and values of σ_I. Even for very low count rate light curves, for which σ_I cannot be accurately determined, bexvar still estimate an accurate interval within which σ_I is most likely to be found. This is very different to the ability of the standard NEV methodology to estimate the NIV. However, we also found that $\overline{\sigma_b}$ might systematically underestimate large σ_I for pink noise light curves, even at high count rates.

For a given count rate and number of bins, there is a minimum value of σ_I that bexvar can measure. The bexvar methodology is unable to detect lower degrees of variability, and instead computes an upper bound on the σ_I estimate.

The bexvar methodology allows for a more accurate estimate of σ_I than the NEV is at estimating the NIV, at low count rates (see Section 3.6). The σ_I is a new variability quantifier and can be a useful quantity to seek to estimate. However, it is unclear how it relates to other ways of describing the intrinsic variability of a source for a particular set of measurements, such as the NIV or the PSD.

In contrast, the NIV is more easily interpretable, as it can be associated with the integral of the PSD, and is a measure of the variance of the linear flux distribution. There are methods for converting estimates of the NIV into estimates of the constant band-limited power, which is intrinsic to the source, and not dependent on any properties of the observation. However, at the moment, there is no established method for determining a similar stationary quantity, independent of the observations, that relates to σ_I. Furthermore, the impact of the red noise leak and the aliasing effect on σ_I are still unknown.

Some of the problems of estimating the NIV at low count rates, and the lack of interpretability of σ_I, can be resolved by determining the relationship between these two parameters. If a conversion can be found, the accuracy of the bexvar estimate of σ_I could be used to enable more accurate estimates of the NIV than with the NEV methodology. The relationship could also allow them to be interpreted in relation to one another. It could also be used to explore the influence of power leakage on σ_I, and determine a stationary parameter equivalent to the band-limited power.

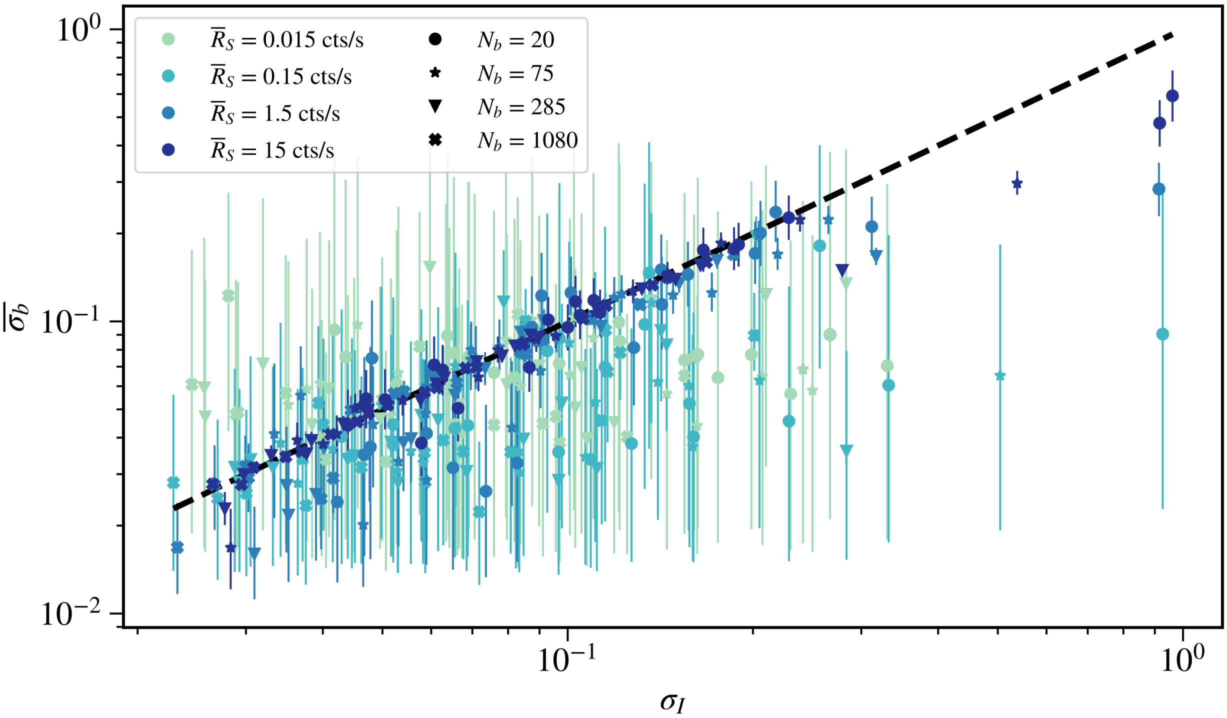

Figure 3.4: The ability of $\overline{\sigma_b}$ to estimate σ_I of a pink noise light curve. Each data point represents a single simulated *eROSITA*-like light curve. The black dashed line indicates the 1 : 1 relationship between the two parameters.

We investigated the possibility of estimating the NIV directly from σ_b, instead of determining a conversion from σ_I to the NIV. As Fig. 3.4 shows, even though the σ_I estimate is mostly accurate, it might face issues at both low and high degrees of variability for pink noise light curves.

While a light curve with a larger NIV also tends to have a larger σ_I, the exact relationship between those two parameters is significantly more complex than the first-order estimate of NIV $\propto \overline{\sigma_b}^2$. To investigate the nature of the relationship, we simulated 3×10^4 *eROSITA*-like pink noise light curves. These consist of 600 simulated light curves for all combinations of the number of bins within the set $\{20, 50, 135, 370, 1000\}$, and count rates of $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1.0, 3.0, 10.0, 30.0\}$cts/s (see Section 3.3). We note that sources with significantly different PSDs might exhibit a different relationship between these two parameters. We did not investigate the correlation between $\overline{\sigma_b}$ and the NIV for other types of variability, such as red noise.

At the lowest count rates, and the smallest number of bins, the simulations produce so few source counts, that it is not meaningful even to use the measured $\overline{\sigma_b}$ as an estimator of the degree of variability of the source. Nevertheless, we still included these instances, to investigate the ability to use the measured σ_b distribution to determine uncertainties and upper limits on an estimate of the NIV.

Fitting the relationship between the measured $\overline{\sigma_b}$ of an *eROSITA*-like light curve and the NIV of the true light curve, enables a conversion from one to the other. We denote the NIV estimated in this way as NEV_b. It might be expected that NEV_b is a more accurate method for estimating the NIV in a Poisson regime than the NEV is, as it uses a Bayesian approach that deals with Poisson probability distributions, unlike the NEV. However, as this is a new method, we first evaluated its strengths and weaknesses, to determine the limits to its applicability.

Fig. 3.5 shows the relationship between $\overline{\sigma_b}$ and the NIV in several simulated light curves for two particular average source count rates and number of bins. The two parameters are very tightly correlated for sources with high count rates, observed for many bins, as shown in the top panel. However, many more variable eRASS sources in the SEP field have count rates and number of consistently, adjacently spaced bins similar to the values used for the simulations whose results are shown in the lower panel ($R_S = 0.3$ cts/s, $N_b = 135$). The panel indicates that there will be significant uncertainty in the NEV_b estimate of the NIV, for most sources.

At low variabilities, low count rates, and a small number of bins, $\overline{\sigma_b}$ reaches a plateau, indicating that bexvar could not determine the NIV, based on the available information. Therefore, whenever this level of $\overline{\sigma_b}$ is measured, it should be treated as an upper limit measurement of the true variability. This plateau depends on both the count rate and the number of bins of the light curve.

At greater variabilities than the upper limit plateau, there is an approximately linear relationship between $\log(\overline{\sigma_b})$, and log(NEV), which has a gradient of approximately 0.5. This is expected, as $\overline{\sigma_b}$ estimates a standard deviation (albeit of the logarithmic count rate), whereas the NIV describes a variance. The best fit linear relationship has a similar gradient and constant term across the range of simulated count rates and number of bins. In an initial fit of the relationship between $\overline{\sigma_b}$ and the NIV, we assumed that only the low variability plateau depends on the average count rate and the number of bins of the light curve. We defined the relationship as follows:

$$\begin{aligned} y_0 &= m\log(N_b) + M\log(\overline{R_S}) + k \\ y_1 &= b\log(\mathrm{NIV}) + c \\ \log(\overline{\sigma_b}) &= \begin{cases} y_1 & \text{if } y_1 \geq y_0 \\ y_0 & \text{if } y_1 < y_0, \end{cases} \end{aligned} \tag{3.11}$$

where y_0 denotes the value of $\log(\overline{\sigma_b})$ at the plateau at the lowest NIVs, y_1 is the main function relating $\log(\overline{\sigma_b})$ to log(NIV) above the plateau, and m, M, k, b, and c are the parameters of the fit. We defined the upper limit plateau as a level in $\log(\overline{\sigma_b})$, rather than as a function of log(NIV), below which it cannot be determined accurately. We decided on this because we want to determine the value of $\overline{\sigma_b}$ at which NEV_b becomes merely an upper limit. By defining the equation in this way, we also reduce the degeneracy between the fit parameters.

We fitted $\overline{\sigma_b}$ from all of the simulations as a function of the NIV using Eq. 3.11. We used the uncertainties in each of the measurements of $\overline{\sigma_b}$ from the simulated light curves

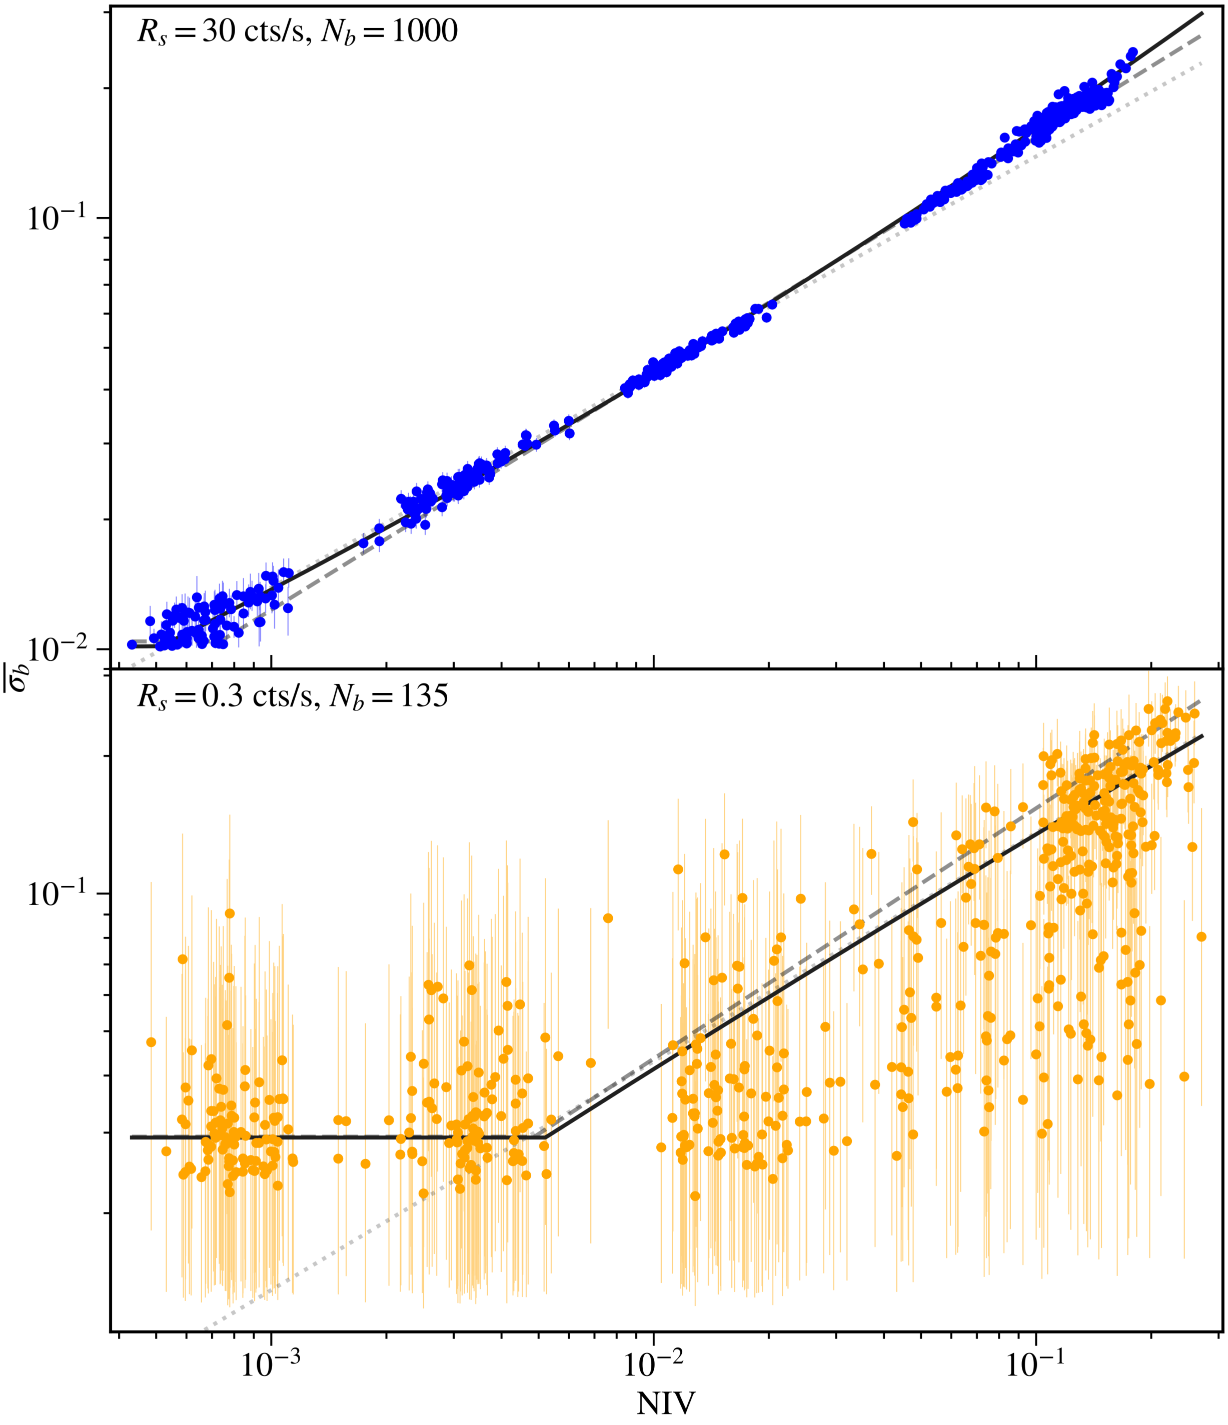

Figure 3.5: The relationship between $\overline{\sigma_b}$ and the NIV in simulated pink noise *eROSITA*-like light curves. The relationship depends on both the number of bins and the count rate, so this figure showcases two examples, one for 1000 bin light curves with a mean count rate of 30 cts/s (top panel), and another for light curves of 135 bins, with a mean count rate of 0.3 cts/s (bottom panel). The solid black line depicts the best fit of $\overline{\sigma_b}\left(\mathrm{NEV}, \overline{R_S}, N_b\right)$, using Eq. 3.13 and the parameters listed in Table 3.1. The dashed line represents using Eq. 3.11 with its best fitting parameters. The dotted line shows the fit of $\mathrm{NEV}_b \propto \sqrt{\overline{\sigma_b}}$ for simulated values of NIV above the detection limit.

in the fit. The NIV does not have an error, as it is an intrinsic quantity of the source and the timing properties of the observations. However, the measured value of $\overline{\sigma_b}$ has an uncertainty, so we fitted $\overline{\sigma_b}$(NIV), rather than NIV ($\overline{\sigma_b}$). The relationship between the two parameters was fitted with $b = 0.54769 \pm 0.00060$, which is close to the first-order estimate of $\overline{\sigma_b} \propto \sqrt{\text{NIV}}$, but not consistent with it. The dashed lines in Fig. 3.5 show this best fit for the two selections of the number of bins and average count rates in the simulations. This best fit accurately describes the relationship between $\overline{\sigma_b}$ and the NIV.

However, the top panel of Fig. 3.5 shows that a linear relationship underestimates $\overline{\sigma_b}$ at both small and large values of the NIV. It is insufficient to describe the exact nature of this relationship. The gradient between $\log(\overline{\sigma_b})$ and log(NIV) increases at large values of the NIV for high count rate sources. This effect can also be seen in the underestimation of large values of σ_I by $\overline{\sigma_b}$, in Fig. 3.4. To account for this effect while keeping the model as simple as possible, we added a quadratic term to Eq. 3.11:

$$\begin{aligned} y_0 &= m\log(N_b) + M\log(\overline{R_S}) + k \\ y_1 &= a\,(\log(\text{NIV}))^2 + b\log(\text{NIV}) + c \\ \log(\overline{\sigma_b}) &= \begin{cases} y_1 & \text{if } y_1 \geq y_0 \\ y_0 & \text{if } y_1 < y_0. \end{cases} \end{aligned} \tag{3.12}$$

This equation describes the relation between $\log(\overline{\sigma_b})$ and log(NIV) better. However, a more careful inspection of the actual relationship between the two variability parameters at different count rates and number of bins, reveals differences in the relation at $y_1 \geq y_0$, particularly at the highest degrees of variability. Therefore, we considered a third possibility, in which the parameters a, b, and c are functions of the logarithm on the number of bins and the logarithm of the count rate. For this third fit, we updated Eq. 3.12 to:

$$\begin{aligned} y_0 &= m_y\log(N_b) + M_y\log(\overline{R_S}) + k_y \\ y_1 &= a\,(\log(\text{NIV}))^2 + b\log(\text{NIV}) + c \\ a &= m_a\log(N_b) + M_a\log(\overline{R_S}) + k_a \\ b &= m_b\log(N_b) + M_b\log(\overline{R_S}) + k_b \\ c &= m_c\log(N_b) + M_c\log(\overline{R_S}) + k_c \\ \log(\overline{\sigma_b}) &= \begin{cases} y_1 & \text{if } y_1 \geq y_0 \\ y_0 & \text{if } y_1 < y_0. \end{cases} \end{aligned} \tag{3.13}$$

We fitted all the simulated data regarding the relationship between $\log(\overline{\sigma_b})$ and log(NIV) a third time. Table 3.1 lists the best fit parameters. To reduce degeneracy and improve the fit, we rescaled the NIV as: $\log(\text{NIV}') = \log(\text{NIV}) - \overline{\log(\text{NIV})}$, where $\overline{\log(\text{NIV})}$ is the average NIV over all simulations. The parameter values presented in the table have, however, been rescaled back, to describe the dependence of $\log(\overline{\sigma_b})$ on log(NIV).

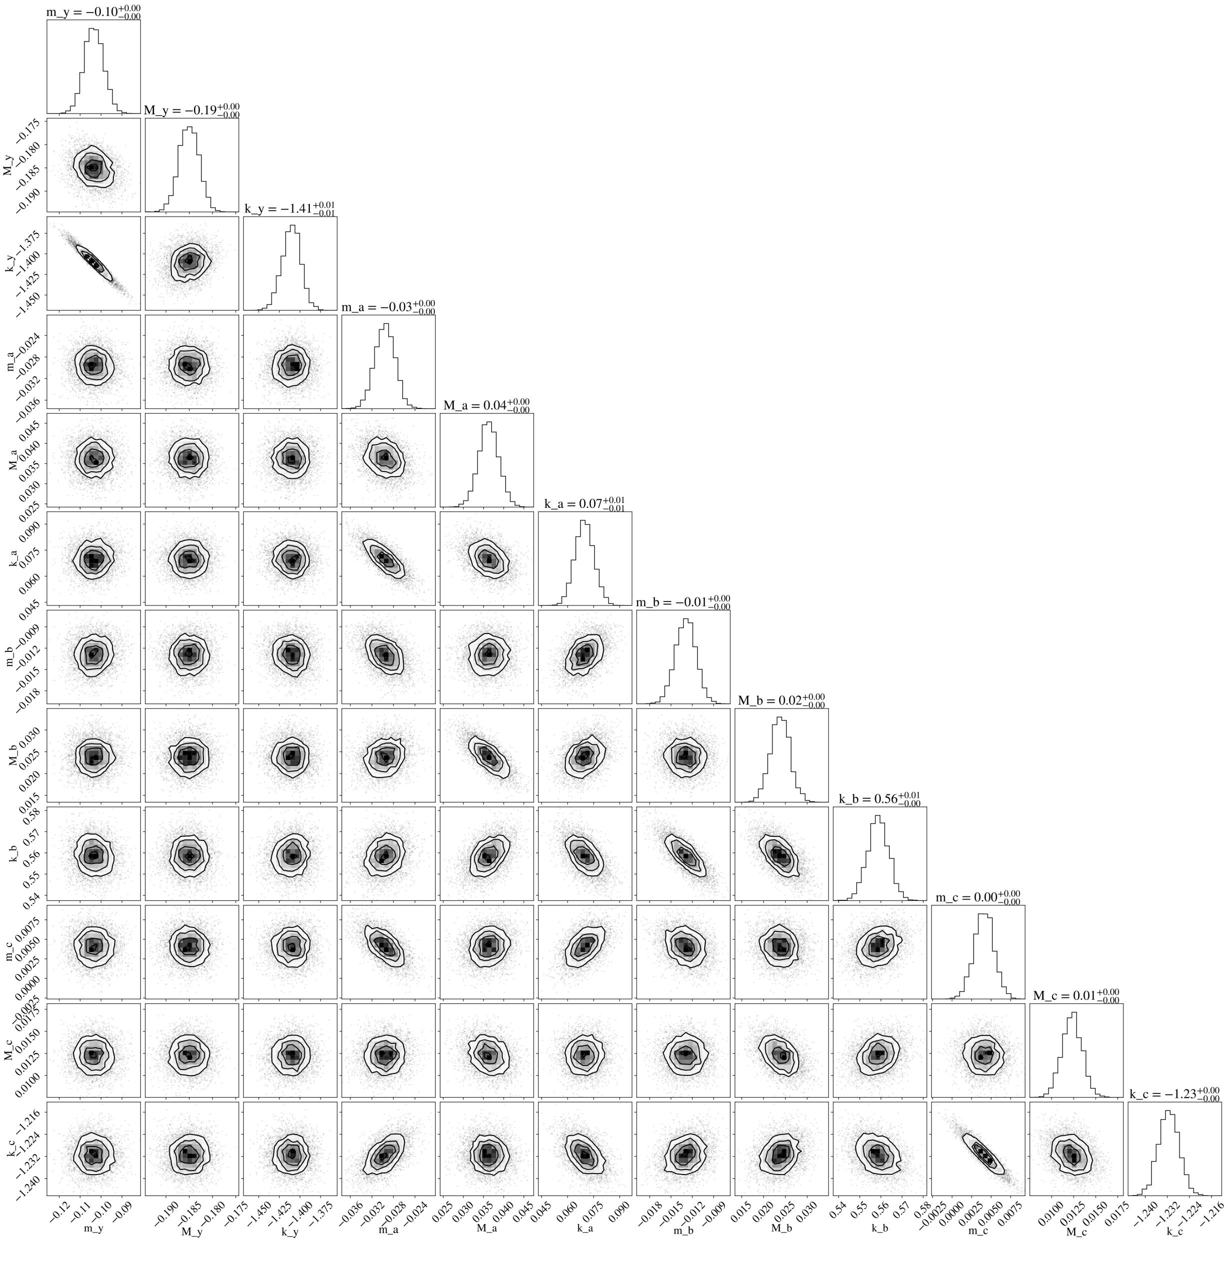

Figure 3.6: The corner plot of the best fitting parameters of Eq. 3.13, for fitting $\log(\overline{\sigma_b})$ as a function of the rescaled true input $\log(\mathrm{NIV}')$, across all simulations spanning count rates from 0.001 cts/s to 30 cts/s, and number of bins from 50 to 1000. The values of the best fit parameters shown here differ from those displayed in Table 3.1, as the fit was performed for $\log(\overline{\sigma_b})$ as a function of the rescaled $\log(\mathrm{NIV}') = \log(\mathrm{NIV}) - \overline{\log(\mathrm{NIV})}$. Table 3.1 instead shows the parameter values for the fit of $\log(\overline{\sigma_b})$ as a function of $\log(\mathrm{NIV})$, and those are the ones that should be used for estimating NEV_b from $\overline{\sigma_b}$.

Parameter	**Value**
m_y	-0.1034 ± 0.0047
M_y	-0.1850 ± 0.0021
k_y	-1.410 ± 0.0011
m_a	-0.0296 ± 0.0019
M_a	0.0363 ± 0.0025
k_a	0.0698 ± 0.0055
m_b	-0.1137 ± 0.0066
M_b	0.1474 ± 0.0088
k_b	0.796 ± 0.019
m_c	-0.1036 ± 0.0061
M_c	0.1580 ± 0.0083
k_c	-0.078 ± 0.018

Table 3.1: Table of the best fitting parameters of $\log(\overline{\sigma_b})$ as a function of the true input log(NIV). These were determined for the wide range of count rates from 0.001 cts/s to 30 cts/s, number of bins from 50 to 1000, and normalised excess variances from 5×10^{-4} to 3×10^{-1}, using Eq. 3.13.

In this fit, we found that the three parameters a, b, and c, which we had previously assumed to be constant in Eq. 3.12, all significantly depend on both the count rate and the number of bins. Therefore, an accurate treatment of the dependence of $\overline{\sigma_b}$ on the NIV cannot consider a, b, and c to be constant. The quadratic term, whose amplitude is defined by a in Eq. 3.13, only weakly depends on the count rate and the number of bins, but it is nonetheless non-negligible. It even changes sign, by being positive for light curves with a large average count rate, and few bins, and negative for light curves with a low average count rate, and a large number of bins. The linear and the constant terms, whose strength is defined by parameters b and c, respectively, depend strongly on both the count rate and the number of bins and will have different values for different eRASS sources. The three parameters that describe the relationship above the plateau (a, b, and c) decrease with an increasing number of bins and increase with an increasing count rate. For these three parameters, the dependence on the count rate is similar to the negative of their dependence on the number of bins. This suggests that all three parameters predominantly depend on $\overline{R_S}/N_b$. The dependence of c on the count rate is similar to that of parameter b, within their respective uncertainties. The same applies to the dependence of the two parameters on the number of bins.

Fig. 3.6 is the corner plot of the best fit of Eq. 3.13 to all the simulated data relating $\overline{\sigma_b}$ to the rescaled NIV in the light curve, log(NIV′). There are some degeneracies, most notably between m_y and k_y and between other m and k parameters. This is probably because the number of bins does not change as much as the count rate within the sample of simulated light curves, so the m parameters can often act similarly to the constant k terms.

There is also a slight negative degeneracy between M_a and M_b, as well as between some parameters of a and c. Nevertheless, the fitting parameters are otherwise well constrained. We tried to keep the degeneracies between parameters as minimal as possible, and only maintained parameters necessary to the fit. Even though we fitted for a total of 12 different parameters, only very few combinations of them exhibit some degree of degeneracy. These parameters are necessary to accurately represent the relationship between $\overline{\sigma_b}$ and the NIV. There might, however, be the possibility of simplifying Eq. 3.13 by setting $m_b = m_c$ and $k_b = k_c$.

This model can accurately fit the observed relationship across the range of count rates, number of bins, and degrees of variability, for which we simulated these light curves. That is the parameter space most relevant for an eRASS variability analysis. The ability of this model to fit the relationship between $\overline{\sigma_b}$ and the NIV is shown with the solid black lines in Fig. 3.5. The Bayes factor comparing the best fit using Eq. 3.11 to the best fit with Eq. 3.13 strongly favours the latter model. Therefore, we decided to use Eq. 3.13 with best-fit parameters detailed in Table 3.1 as a basis for the conversion of an estimate of $\overline{\sigma_b}$ into an estimate of the NIV, under the assumption that light curves exhibit a pink noise variability.

The Eq. 3.12 can be rearranged to determine the bexvar estimate of the NIV, NEV_b, from the measurement of $\overline{\sigma_b}$ as follows:

$$\log(\mathrm{NEV}_b) = \frac{-b + \sqrt{b^2 - 4a(c - \log(\overline{\sigma_b}))}}{2a} \tag{3.14}$$

$$\log(\overline{\sigma_{b,l}}) = m_y \log(N_b) + M_y \log(\overline{R_S}) + k_y, \tag{3.15}$$

where $\overline{\sigma_{b,l}}$ is the approximate value of the lower limit of $\overline{\sigma_b}$ that is measurable at that particular count rate and number of bins. The equation for NEV_b should only be considered as an upper limit estimate of the NIV if $\overline{\sigma_b} \approx \sigma_{b,l}$. The expressions, and values to use for a, b, c, m_y, M_y, and k_y, are listed in Eq. 3.13 and Table 3.1. This conversion is only accurate for a set of equally spaced bins in the light curve.

The above function can fail in two particular instances. Firstly, it is undefined whenever the denominator, $2a = 0$. When $-0.0296\log(N_b) + 0.0363\log(\overline{R_S}) = -0.0698$, Eq. 3.13 reverts to the linear Eq. 3.11, with parameters y_0, b and c as expressed in Eq. 3.13, and Table 3.1. Secondly, this conversion fails when the contents of the square root are negative, so when $\log(\overline{\sigma_b}) < c - (b^2/4a)$ if $a > 0$, or $\log(\overline{\sigma_b}) > c - (b^2/4a)$ if $a < 0$. However, in all of the simulations we used here, we never found $\overline{\sigma_b}$ to satisfy either inequality. When $a > 0$, $\overline{\sigma_{b,l}}$ is always larger than $c - (b^2/4a)$. When $a < 0$, $\log(\overline{\sigma_b}) > c - (b^2/4a)$ only for extremely variable sources, for which the assumption of a pink noise PSD is no longer accurate. In those cases, this conversion from $\overline{\sigma_b}$ to NEV_b should not be used. These statements apply to the parameter space we investigated; for light curves of between 20 and 1080 bins, and average count rates in the range $[0.001, 30.0]$ cts/s.

The function y_1, with parameter values as listed in Table 3.1 can also be used to convert between σ_I, and the NIV. It remains an accurate description of their relationship within the parameter space investigated, and as long as $\log(\sigma_I) \lesssim y_0\left(N_b, \overline{R_S}\right)$.

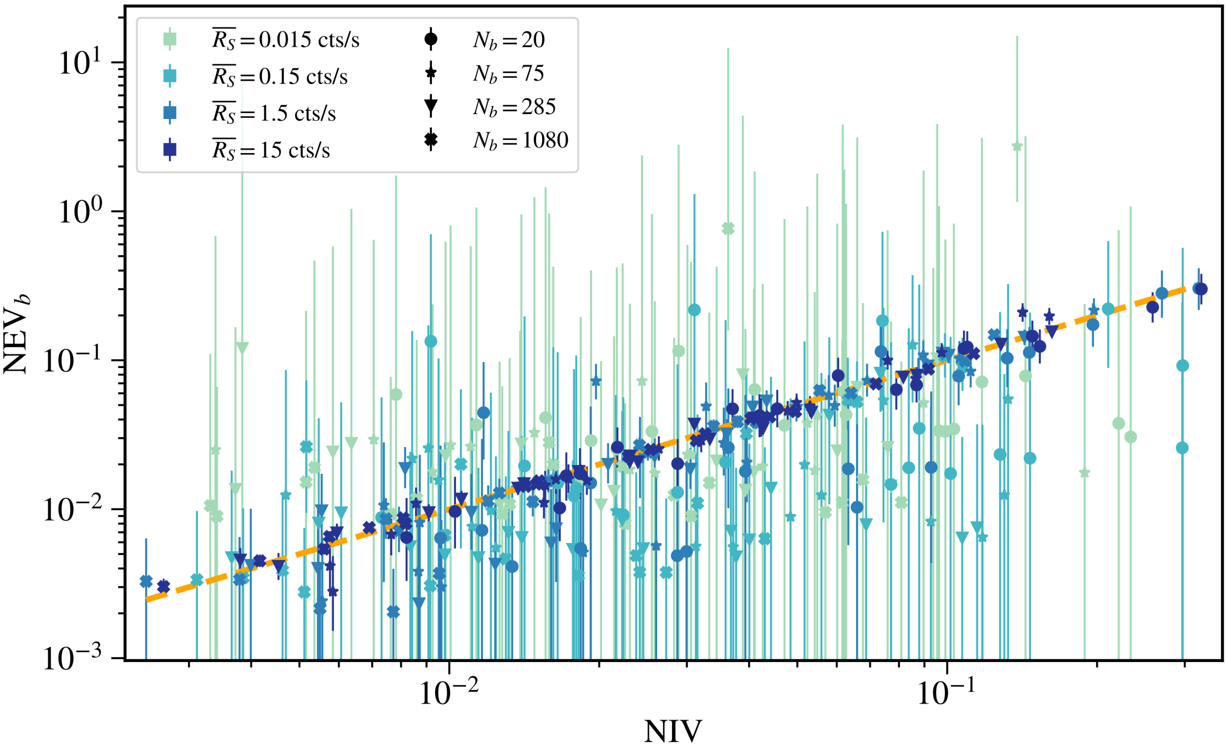

Figure 3.7: Comparing the NEV$_b$ measurements with the actual NIV of the light curve. The NEV$_b$ estimate is obtained by computing $\overline{\sigma_b}$, and using the conversion function from $\overline{\sigma_b}$ to NIV. This conversion uses Eq. 3.14, with parameters described by Eq. 3.13, having values as listed in Table 3.1. Light curves were simulated based on the assumption of pink noise variability. The orange dashed line indicates the 1 : 1 relationship between NEV$_b$, and the NIV. The closer the data points are to this line, the better the estimate of the NIV in the light curve is.

Using these results to convert $\overline{\sigma_b}$ to NEV$_b$, we then investigated the ability of NEV$_b$ to estimate the NIV of a light curve. Just as for Fig. 3.4, we simulated *eROSITA*-like light curves of pink noise variable sources, consisting of $\{20, 75, 150, 400, 1050\}$ bins, and average count rates of $\{0.015, 0.15, 1.5, 15\}$ cts/s. We specifically selected different values for the number of bins and the average count rate than in the simulations used for determining the relationship between $\overline{\sigma_b}$, and the NIV. In this way, we could independently verify the usefulness of this method, for other parameters not previously investigated. However, we used light curves of 20 bins both for defining the method and testing it, as that is the lower limit we chose for this type of variability analysis.

Fig. 3.7 depicts how accurately NEV$_b$ can estimate the NIV. It shows that NEV$_b$ is mostly accurate in its estimate. Contrary to other methods, NEV$_b$ always has a positive value.

The uncertainties of NEV$_b$ depicted in this figure are determined by converting the upper and lower bounds of the 1σ confidence interval of $\overline{\sigma_b}$ into NEV$_b$ values, using Eq. 3.13, and the parameters of Table 3.1. Whenever the lower bound error on $\overline{\sigma_b}$ extends below 10^{y_0}, the lower bound error on NEV$_b$ is extended to a value of 0. This is due to the inability to determine lower degrees of variability for those particular light curves.

At very low count rates, when there is insufficient information available to properly constrain the NIV of a light curve, NEV$_b$ is still able to provide an accurate confidence interval within which the NIV of the light curve is likely to be. This can be seen in Fig. 3.7 which shows that the errors of NEV$_b$ contain the NIV in most instances, regardless of the count rate.

3.6 Comparing methods for NIV estimation

In this section, we consider two other methods that could be used to estimate the NIV of the light curve of a variable source, and compare them against the NEV$_b$ method. The first of these estimates the NIV by integrating the periodogram of the light curve, which we label as NEV$_i$. The second methodology is based on using Eqs. 3.1 and 3.2, but with an adjustment for reducing the effect of the lowest exposure bins on the estimate. We denote this estimate of the NIV as NEV$_{eq}$. Finally, we evaluate the strengths and weaknesses of these three methods.

3.6.1 NEV$_i$

Parseval's theorem states that in the high count rate limit, when the assumption of a normal probability distribution of the measured count rate in every bin is true, the definition of the NEV as shown in Eq. 3.1, corresponds to the integral of the Poisson noise subtracted, rms normalised periodogram (van der Klis, 1989):

$$\mathrm{NEV}_i = \int_{(N_b\tau)^{-1}}^{(2\tau)^{-1}} (P(\nu) - N_{P,\epsilon}) d\nu. \tag{3.16}$$

Therefore, it is also possible to estimate the NIV of the light curve of a variable source in this way. The Poisson noise of a light curve with variable fractional exposures is discussed in more detail in Section 3.8.1.

This method is more computationally expensive than calculating NEV$_{eq}$. Nevertheless, it can be useful for estimating the NIV for a different frequency range. We also use it to compare the other two methods against.

Fig. 3.8 shows the ability of NEV$_i$ to estimate the NIV in simulated pink noise light curves, when subtracting the standard Poisson noise (Eq. 3.25) from the periodogram. It can be compared against Fig. 3.7, to see the different accuracies of the two methods. This shows that NEV$_i$ is only a good estimator of the NIV at the highest count rates and the greatest NIVs. At lower count rates and NIVs, NEV$_i$ levels off, diverging from the 1:1 line that the data points should ideally be distributed around. The minimum level of NEV$_i$ decreases with an increasing count rate. However, it seems largely independent of the number of bins in the light curve. At the lowest count rates considered for this figure, this plateau level even lies above the maximum NIV of the simulated light curves. The plateaus are caused by varying fractional exposures. The excess noise this creates has not been subtracted from the periodograms before integrating them to estimate the NIV. This indicates the significance of the fractional exposure noise, showing that it cannot be ignored.

Therefore, there is a need for an accurate modelling of the excess noise in the periodogram of light curves with varying fractional exposures, as will be discussed in Section 3.8.1. Fig. 3.8 shows that the excess fractional exposure noise increases inversely proportionally to the count rate. This helped motivate the expression of the total Poisson and fractional exposure noise of Eq. 3.28.

When subtracting the combined Poisson and fractional exposure noise from the periodograms of the simulated light curves, using Eqs. 3.28, and 3.29, as well as the parameters listed in Table 3.4 in Section 3.8.1, the plateaus disappear, as the noise in the periodogram is correctly estimated, and removed. This can be seen in Fig. 3.9, in which the NEV$_i$ is a much more accurate estimator of the NIV than in Fig. 3.8. Henceforth, we will use NEV$_i$ to refer to the Poisson and fractional exposure noise subtracted integral of the periodogram.

At high count rates ($\overline{R_S} \gtrsim$ 1.5 cts/s), NEV$_i$ is an accurate estimator of the NIV, for the entire range of values of the NIV, and number of bins in the light curve, that we investigated. However, its accuracy decreases rapidly with a decreasing count rate. At $\overline{R_S} \lesssim$ 0.15 cts/s, the NEV$_i$ often differs significantly from the NIV, and is frequently found at negative values. 87% of the light curves we simulated at a count rate of 0.015 cts/s had a negative NEV$_i$. At 0.15 cts/s, 34% of the light curves are still found at a negative NEV$_i$, and even at 1.5 cts/s, 7% of light curves were found to have a negative NEV$_i$.

At the lowest count rates, the NIV is frequently underestimated by NEV$_i$. This could partially be due to the uncertainty in the fractional exposure noise calculated in Section 3.8.1, which might not be completely accurate at low count rates. It can also result from insufficient data to determine an accurate periodogram. The NEV$_i$ methodology also treats each bin identically, regardless of the fractional exposure.

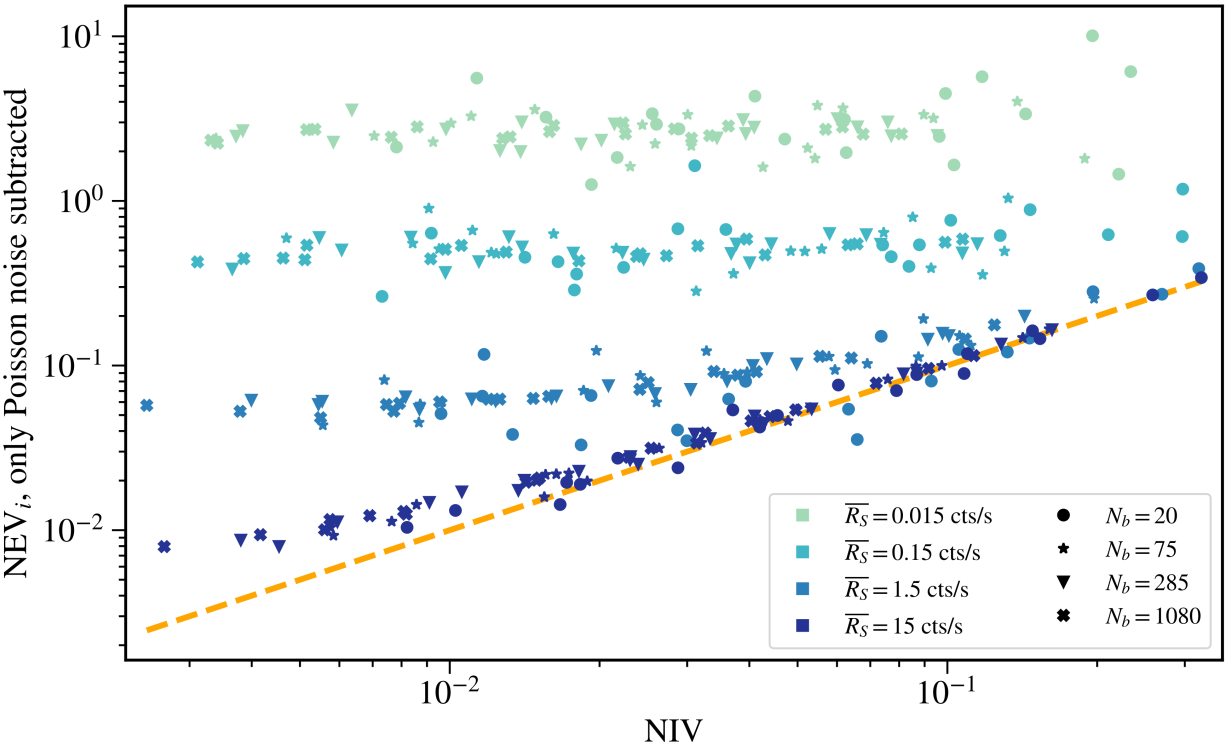

Figure 3.8: The ability to estimate the NIV by integrating the periodograms of simulated *eROSITA*-like pink noise light curves. The dependence is determined as a function of the NIV. Each data point indicates the result of a single simulation. The closer a data point lies to the orange dashed 1:1 line, the better the estimate of the NIV is. For this plot, we merely subtracted the standard Poisson noise (Eq. 3.25) from the periodograms. This is insufficient to describe the noise level in periodograms of exposure-corrected light curves.

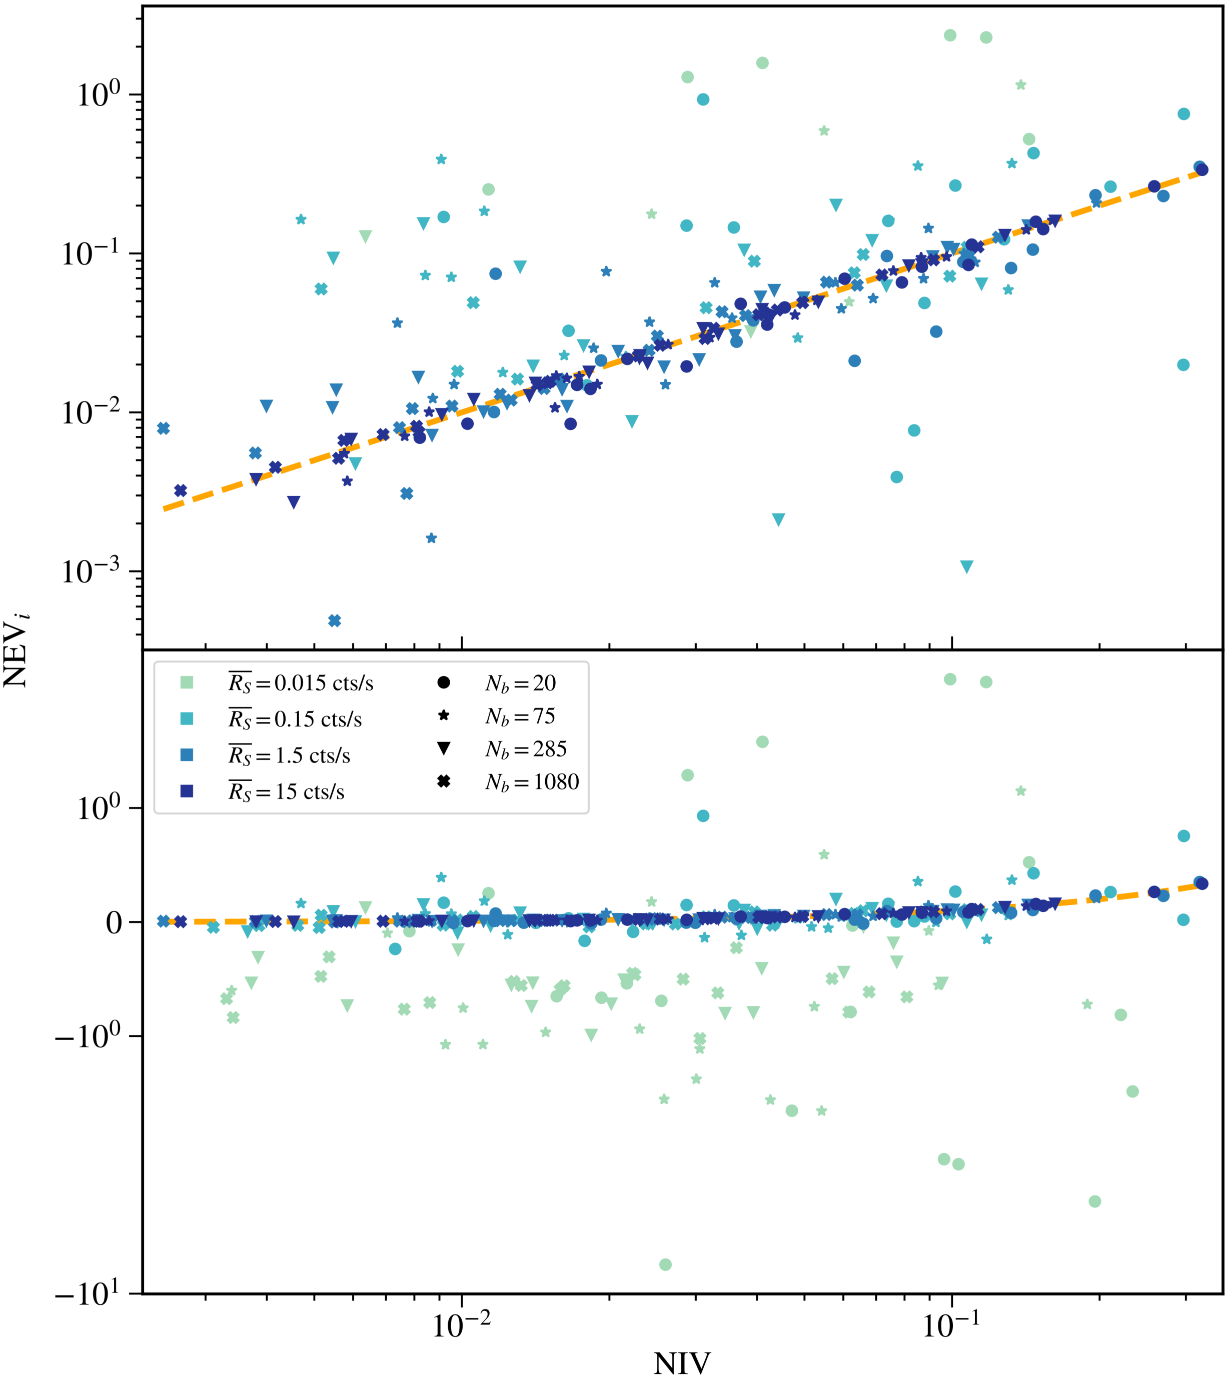

Figure 3.9: The accuracy of estimating the NIV by integrating periodograms of simulated *eROSITA*-like pink noise light curves. We subtracted both the Poisson noise and the fractional exposures noise for this plot, using Eqs. 3.28, and 3.29, with parameters listed in Table 3.4 from the measured periodograms. This figure shows the relationship on both a logarithmic (top panel) and a symmetric logarithmic (lower panel) scale, to showcase how close the individual NEV$_i$ measurements are to the NIV, as well as the complete range of values measured for NEV$_i$. The orange dashed line indicates the 1 : 1 relationship. The closer the data points are to this line, the better the estimate of the NIV in the light curve is. By subtracting the excess noise, the accuracy of the NIV estimate is significantly improved. There is no longer a systematic offset between the NEV$_i$ and the NIV.

3.6.2 NEV$_{eq}$

The standard NEV, as defined in Eqs. 3.1 and 3.2 is only valid in case of a normal probability distribution for the count rate in each bin. When the components used to calculate the NEV (Eq. 3.1) are defined as in Eq. 3.2, the NEV treats every bin identically, regardless of their fractional exposure. This can cause a significant portion of the excess variance detected in an *eROSITA*-like light curve to originate from the bins with the smallest exposure times.

One way of reducing the effect that the lowest exposure bins have on the measured NEV, is to weight each of the terms of the NEV by the fractional exposure itself. Rather than using the definitions described in Eqs. 3.2, we would use:

$$\begin{aligned}
\overline{R_S} &= \frac{\sum_{i=1}^{N_b} \epsilon\left(t_i\right) R_S\left(t_i\right)}{\sum_{i=1}^{N_b} \epsilon\left(t_i\right)} \\
\sigma_{obs}^2 &= \frac{N_b}{N_b-1} \frac{\sum_{i=1}^{N_b} \epsilon^2\left(t_i\right)\left(R_S\left(t_i\right)-\overline{R_S}\right)^2}{\sum_{i=1}^{N_b} \epsilon^2\left(t_i\right)} \\
\overline{\sigma_{err}^2} &= \frac{\sum_{i=1}^{N_b} \epsilon^2\left(t_i\right) \sigma_{err}^2\left(t_i\right)}{\sum_{i=1}^{N_b} \epsilon^2\left(t_i\right)} .
\end{aligned} \tag{3.17}$$

The first of these equations is a more accurate estimate of the average count rate in a light curve with varying fractional exposures than the equivalent part of Eq. 3.2 is. If the background area and count rate are constant, then the first part of Eq. 3.17 is approximately equal to the average count rate found by 2.7. The other equations are similarly more accurate measures of the variance and the mean error in the data. By weighting the terms, the equations effectively deal with the individually measured source counts, rather than count rates.

Since the uncertainties in the measured count rates are asymmetric, it is worth considering which error to use for $\sigma_{err}^2\left(t_i\right)$. The excess variance compares the difference between the contribution of each measured count rate to the total variance, with the error of that measurement. Therefore, we decided to use the error in the direction of the mean count rate, similar to the definition of the modified AMPL_SIG in Section 3.2.3. So if $R_S\left(t_i\right) \leq \overline{R_S}$, we set $\sigma_{err}\left(t_i\right)=\sigma_{+err}$, and if $R_S\left(t_i\right)>\overline{R_S}$, we set $\sigma_{err}\left(t_i\right)=\sigma_{-err}$. When using these modifications, we label the resulting NIV estimate as NEV$_{eq}$, to distinguish it from the standard NEV method. The NEV$_{eq}$ is still calculated by using Eq. 3.1, but with $\overline{R_S}$, σ_{obs}^2, and $\overline{\sigma_{err}^2}$ as determined using Eq. 3.17.

These modifications to the standard NEV method enable a more accurate estimate of the NIV in a light curve afflicted by varying fractional exposures. However, they do not adequately deal with the inherent problems of assuming the uncertainty in the measured count rates in one direction to be equivalent to the standard deviation of an assumed normal distribution.

Fig. 3.10 shows the performance of using NEV$_{eq}$ as an estimator of the NIV, using Eq. 3.1, and the weighted parameters of Eq. 3.17. This method works well at high count rates when there are enough source counts per bin to make the probability distribution for the count rate approximately Gaussian. At intermediate and low count rates, weighting the bins by the fractional exposure improves the estimate of the NIV in most cases, when compared to the standard NEV. However, even though it reduces the scale of the problem, it does not remove it entirely. 95% of the simulated light curves at a count rate of 0.015 cts/s, and 29% of the light curves with a count rate of 0.15 cts/s were found to have a negative NEV_{eq}, some at extremely negative values. At a count rate of 1.5 cts/s, 3% of simulated light curves still have a negative NEV_{eq}. The NIV is still systematically underestimated at the lowest count rates by NEV_{eq}. The NEV_{eq} estimates are, however, less likely to significantly exceed the NIV than the NEV_i estimates do.

We used the results of Vaughan et al. (2003), their Equation 11, to determine the uncertainties in the individual measurements of the NEV_{eq}, which we plot in Fig. 3.10. This equation might not be accurate for light curves with varying fractional exposures or for the modifications to the standard method discussed above. It also does not reliably estimate the uncertainty in the measured NEV_{eq} when it is less than 0, as can be seen by the lower panel in Fig. 3.10. Accurate error bars should extend to above 0 in most instances. Using these uncertainties, negative NEV_{eq} values are useless for estimating the NIV, or the NIV_∞, and cannot even be used to determine a linear average, computed over multiple segments.

Figs. 3.4, 3.8, 3.9, and 3.10, which depict the ability of various methods to estimate the NIV, are all based on the same set of simulated light curves. These were generated to have $\{20, 75, 150, 400, 1050\}$ bins, and mean input count rates of $\{0.015, 0.15, 1.5, 15\}$ cts/s. Comparing Fig. 3.4 to Figs. 3.9 and 3.10, shows that NEV_b is significantly more accurate at estimating the NIV in the light curve at low count rates, than NEV_i or NEV_{eq} are.

3.6.3 Comparison

We compared the ability of the three methods previously discussed (NEV_b, NEV_i, and NEV_{eq}), to provide an accurate estimate of the NIV in the light curve. The accuracy of the three methods was investigated as a function of the count rate, the number of bins, the degree of variability, and the type of variability observed. Fig. 3.11 depicts the absolute magnitude of the difference between the measured NIV estimate (which we label as NEV_m to refer to the three methods), and the actual NIV, normalised by the NIV of the simulated light curves: $|\mathrm{NEV}_m - \mathrm{NIV}| / \mathrm{NIV} = |D_m|$. The figure shows the dependence of $|D_m|$ on the NIV in each plot, as the ability to estimate the NIV strongly depends on its value. The three methods were applied to the same simulated light curves for a fair comparison. Three different types of variability were considered; white noise ($P \propto \nu^0$), pink noise ($P \propto \nu^{-1}$), and red noise ($P \propto \nu^{-2}$). We again selected mean count rates of $\{0.015, 0.15, 1.5, 15.0\}$ cts/s, and a number of bins of $\{20, 75, 150, 400, 1050\}$ for the construction of the simulated light curves used in this analysis.

Fig. 3.11 shows that NEV_b represents a significant improvement over NEV_i and NEV_{eq} at low count rates, and low degrees of variability. At the lowest count rates we investigated, of 0.015 cts/s, NEV_b is always the most accurate NIV estimate, for the three types of PSDs,

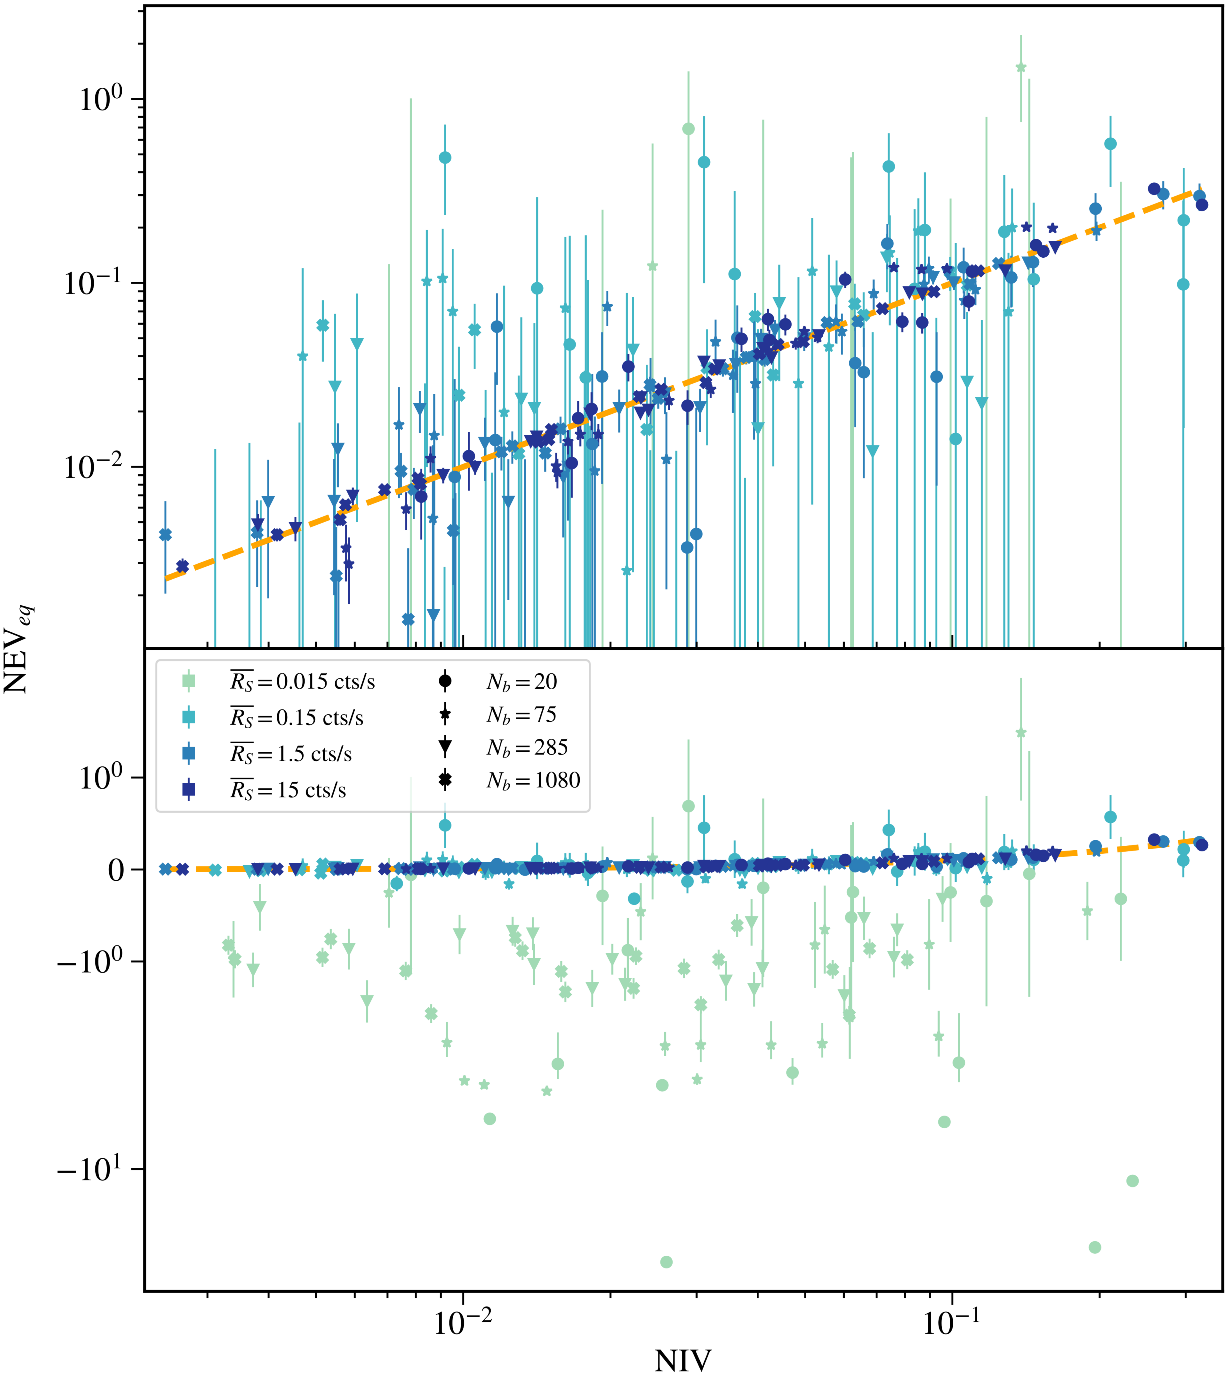

Figure 3.10: Comparison of the NEV$_{eq}$ obtained by using the weighted parameters of Eq. 3.17 in Eq. 3.1, with the NIV. This relationship is shown for simulated pink noise *eROSITA*-like light curves. This figure depicts the dependence of the estimate on the average count rate, the number of bins, and the degree of intrinsic variability. The relationship between the NIV estimate and the true NIV is depicted on both a logarithmic (top panel) and a symmetric logarithmic (lower panel) scale. The orange dashed line indicates the 1 : 1 relationship. The closer the data points are to this line, the better the estimate of the NIV in the light curve is.

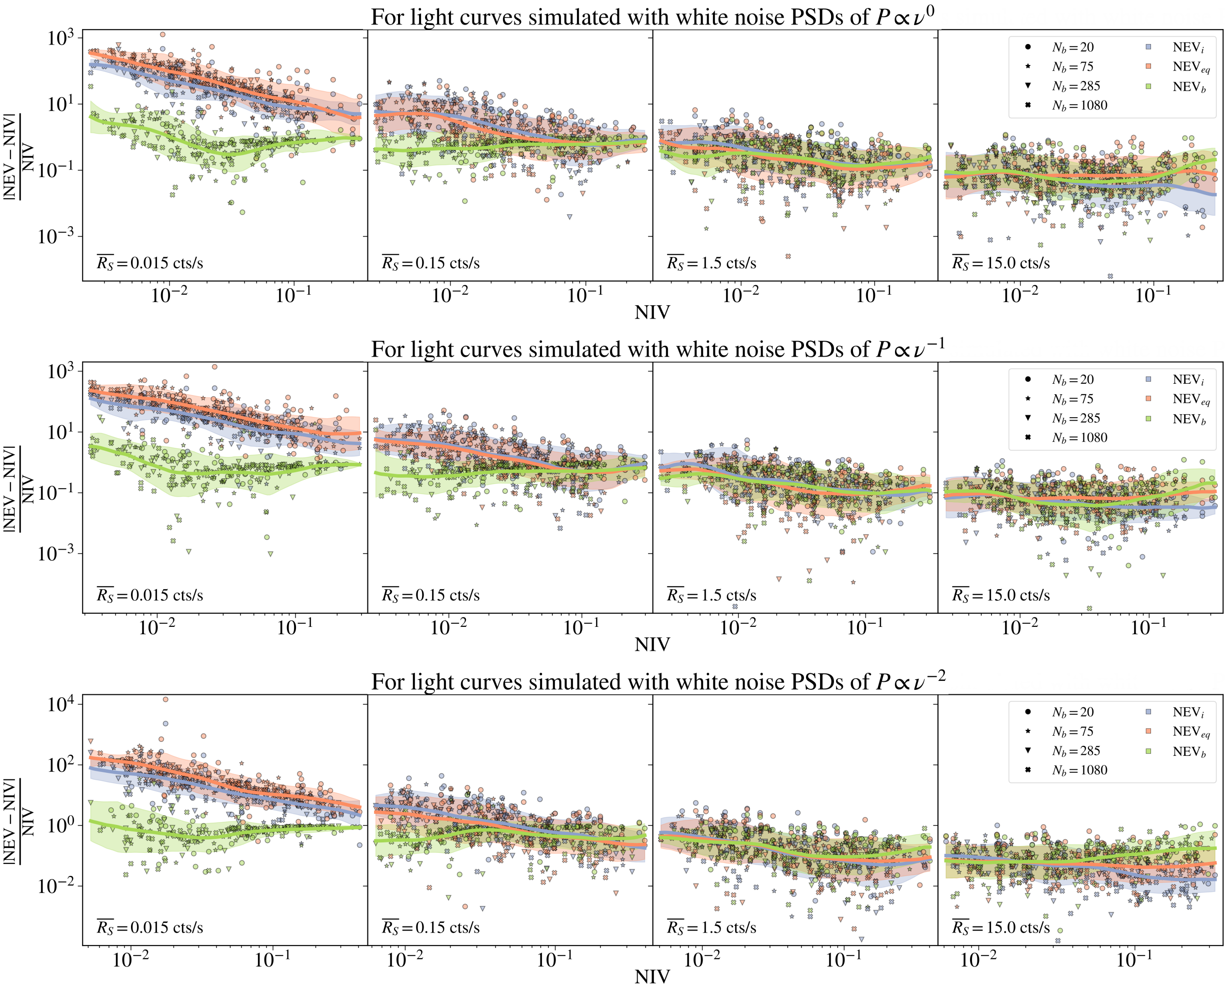

Figure 3.11: Comparison of the ability of the integral of the periodogram (NEV_i), the modified NEV equation (NEV_{eq}), and the bexvar NIV estimate (NEV_b), to accurately estimate the NIV of *eROSITA*-like light curves. This figure plots the absolute value of the difference between the measured NEV_m (which refers to either NEV_i, NEV_{eq}, or NEV_b), and the NIV (NIV), normalised by the NIV, as a function of the NIV of the simulated light curves. The lower the value of this parameter, the more accurate the estimate of the NIV is. Each data point represents one instance of a simulated light curve. The same light curves were used for computing NEV_i, NEV_{eq}, and NEV_b. Each panel represents a different count rate. The dependence on the number of bins is showcased via different symbols. Each row of panels represents a different type of variability; white, pink, and red noise. Solid lines are plotted on top of the distribution of points, corresponding to the geometric mean of the distribution of $\log\left(|\mathrm{NEV}_m - \mathrm{NIV}|/\mathrm{NIV}\right)$ as a function of log(NIV). Transparent areas indicate the part of the distribution corresponding to one standard deviation from the mean, in either direction.

for all values of the NIV, and the entire range of the number of bins that we investigated. At this count rate, NEV_b is often 2 orders of magnitude more accurate than the other two methods at estimating the NIV of an observed light curve.

At 0.15 cts/s, NEV_b is more accurate at low variabilities than either NEV_i or NEV_{eq}, but all three methods are approximately equally accurate at high degrees of variability. In general, sources featuring a larger range of fluxes within the duration of the observations, have a larger NIV. Therefore, at a fixed average source flux, strongly variable sources will be observed to have significantly more source counts than expected, in a few bins. That helps improve the NIV estimates using the NEV_i or NEV_{eq} methodology.

For count rates of 1.5 cts/s and 15 cts/s, the measured number of source counts per bin can reasonably accurately be assumed to have a normal probability distribution for most, or all eroDays. As a result, all three methods have similar accuracies at these count rates, across the entire range of simulated NIVs. However, NEV_b provides a less accurate estimate of the NIV of high count rate sources (15 cts/s) exhibiting red noise variability of NIV $\gtrsim 5 \times 10^{-2}$), as compared to NEV_i or NEV_{eq}.

The NEV_i and NEV_{eq} methods estimate the NIV with similar accuracy. The $|D_m|$ parameter follows a similar anti-correlation as a function of the NIV, for both NEV_i and NEV_{eq}. The anti-correlation is predominantly caused by the normalisation of $|D_m|$. This indicates that the average difference between the estimate, and the true value of the NIV is mainly independent of the NIV, for these two methods, at low count rates. In contrast, the mean difference between the NEV_b estimate, and the NIV, are maintained at an almost constant ratio of the NIV.

The ability to estimate the NIV substantially improves with an increasing count rate for NEV_i and NEV_{eq}, as $|D_m|$ drops by about 3 orders of magnitude from 0.015 cts/s to 15 cts/s. NEV_i is slightly better than NEV_{eq} at accurately estimating the NIV at 0.015 cts/s and 15 cts/s, but NEV_{eq} is slightly better at 0.15 cts/s and 1.5 cts/s. As expected, the ability to estimate the NIV improves with an increasing number of bins, for all three methods.

Even though the conversion from $\overline{\sigma_b}$ to NEV_b was determined specifically for pink noise PSDs, the conversion is still very accurate for light curves generated from white and red noise PSDs. Even for those types of variability, NEV_b is significantly more accurate than NEV_i, or NEV_{eq}. For high count rate white noise light curves, NEV_b remained equally accurate as NEV_i, or NEV_{eq}. Only at high count rates and high degrees of variability, do NEV_i, or NEV_{eq} perform slightly better than NEV_b.

Another relevant feature to consider when comparing different methods of estimating the NIV in a light curve is the ability to estimate the uncertainty in the estimate accurately. To determine how accurate the NEV_{eq}, and NEV_b errors are, we investigated the ratio of the absolute value of the difference between the estimate, and the real value of the NIV, to the estimated error: $|\mathrm{NEV}_m - \mathrm{NIV}| / \sigma_{\mathrm{NEV},m} = |Z_m|$. Here, $\sigma_{\mathrm{NEV},m}$ represents the error in NEV_{eq}, or NEV_b. Fig. 3.12 shows how this parameter varies as a function of the NIV, the mean count rate, the number of bins, and the type of variability (white noise, pink noise, and red noise). The figure display $|Z_m|$ for the two NIV estimators NEV_{eq}, and NEV_b. The NEV_i is not considered here, as we found the computation of its error to be unreliable.

Assuming that the individual NIV estimates are distributed normally, with a mean equal to the NIV, and a standard deviation equal to the measurement uncertainty, $|Z_m| = \sqrt{2/\pi}$ (Geary, 1935). This is the value we compare the two methods against, indicated in Fig. 3.12 as a dashed black line. If the mean value of $|Z_m|$ for either method is found to lie below $\sqrt{2/\pi}$, then the uncertainties in the measurement are overestimated. If, instead, the value of this parameter lies above $\sqrt{2/\pi}$, then either the mean NIV estimate differs significantly from its actual value, or the uncertainties are underestimated. The closer the average value of $|Z_m|$ lies to $\sqrt{2/\pi}$, the better the method is at correctly estimating the measurement uncertainties.

At low count rates, the errors in NEV$_b$ are slightly overestimated. For example, at an average source count rate of $\overline{R_S} = 0.015$ cts/s, the errors in the NEV$_b$ estimates of the NIV are accurate at low variabilities. However, they are overestimated for light curves with NIV $\gtrsim 1 \times 10^{-2}$, which can be seen in the panels on the far left of Fig. 3.11.

In contrast, the NEV$_{eq}$ uncertainties are significantly underestimated at $\overline{R_S} = 0.015$ cts/s. The measured NEV$_{eq}$ values significantly differ from the NIV, and are found to be negative for 95% of all simulated light curves. There are even some instances, when $\sigma^2_{obs} < \overline{\sigma_{err}}^2/2$, such that NEV$_{eq} = -\overline{\sigma_{err}}^2/2$. When this happens, the error in NEV$_{eq}$ cannot be calculated. The figure does not include light curves with these properties as data points. The NEV$_{eq}$ errors are inaccurate, whenever NEV$_{eq} < 0$. This leads to $|Z_m| > \sqrt{2/\pi}$ for almost all simulated red, pink, and white noise light curves at a mean source count rate of $\overline{R_S} = 0.015$ cts/s. This can be seen in the panels on the far left of Fig. 3.11.

The accuracy of the uncertainties improves with an increasing average source count rate, such that the errors are reasonably reliable for count rates of $\overline{R_S} \gtrsim 0.015$ cts/s, for both methods. This can be seen in the second, third, and fourth panels from the left in Fig. 3.11. At $\overline{R_S} = 0.15$ cts/s, the NEV$_{eq}$ uncertainties more accurately represent the error in the measurement than the NEV$_b$ ones. At $\overline{R_S} = 1.5$ cts/s, and $\overline{R_S} = 15$ cts/s, both sets of errors perform about equally well, for all three different types of variability investigated here. The NEV$_b$ errors are, however still more likely to be slightly overestimated, and the NEV$_{eq}$ are still more likely to be slightly underestimated. The NEV$_b$ errors are also reasonably reliable for red and white noise variability, even though they are based on the assumption of pink noise variability.

These errors only depict the measurement errors of estimating the NIV for the selected set of observations. They do not contain the sampling errors of using NEV$_b$ or NEV$_{eq}$ to estimate NIV$_\infty$, as will be discussed in Sections 3.7.2, and 3.7.3.

We note that NEV$_b$ is the best method for estimating the NIV of *eROSITA*-like light curves. Its estimates are reliable across a wide range of degrees of variability, count rates, number of bins, and types of variability. It never estimates the NIV to be negative, as NEV$_i$, and NEV$_{eq}$ do. It can identify when an estimate of the NIV is merely an upper limit. The uncertainties of NEV$_b$ are well-defined and close to the size they should be. However, they are overestimated at low count rates. The downside to NEV$_b$ is that it is more computationally expensive than NEV$_i$, and NEV$_{eq}$. We recommend using NEV$_b$ to estimate the NIV in all light curves, except when a source has $\gtrsim 20$ source counts in every

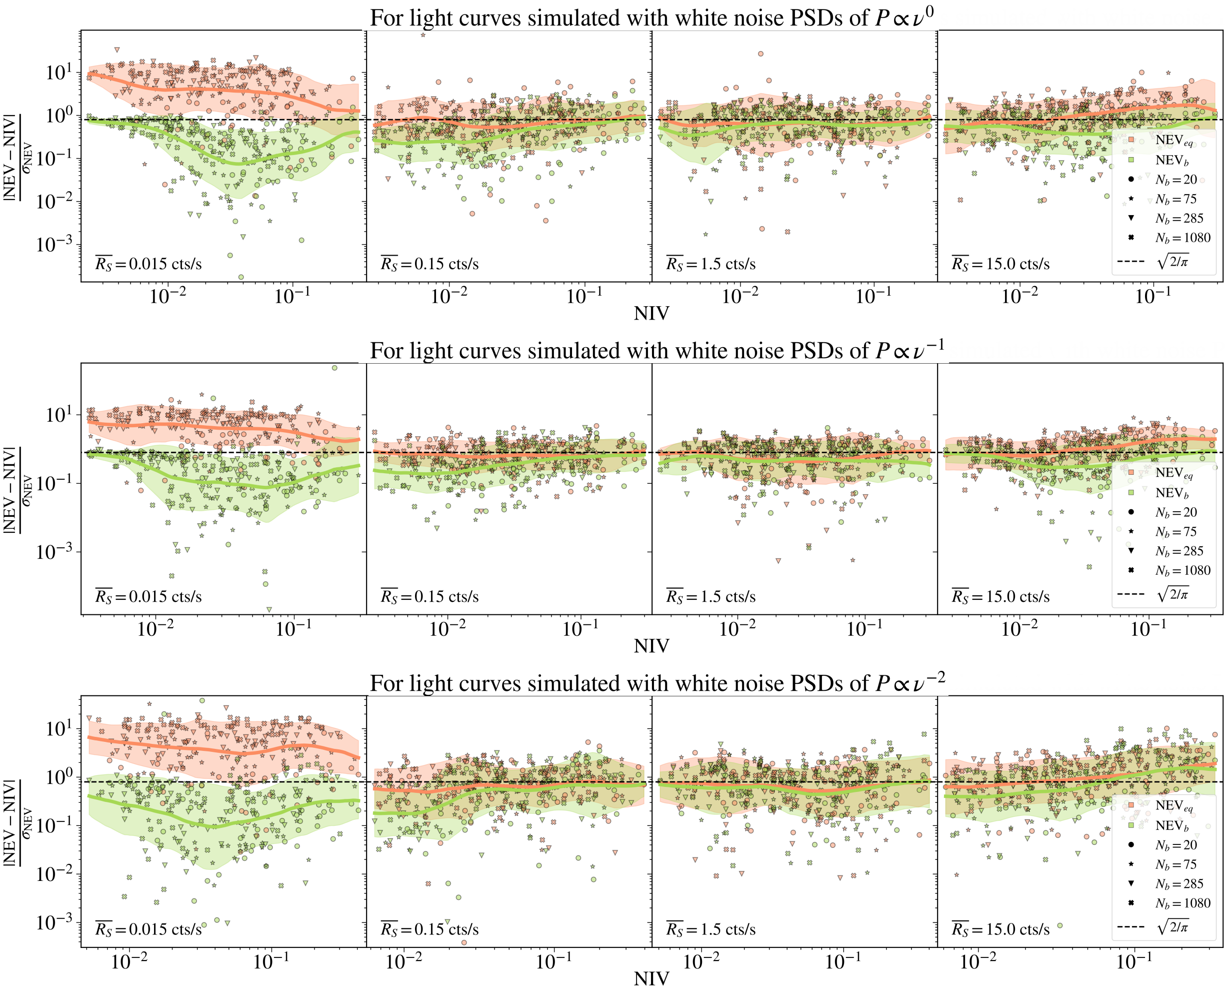

Figure 3.12: The accuracy of the uncertainties of NEV_{eq} and NEV_b, for a range of different count rates, number of bins, and types of PSDs. This figure used the results of the same simulated light curves displayed in Fig. 3.11. The parameter NEV_m denotes a NIV estimate, which is either NEV_b, or NEV_{eq}. The parameter σ_m denotes the uncertainty in the measurement. The black dashed line at a value of $\sqrt{2/\pi}$, indicates the level that $|\mathrm{NEV}_m - \mathrm{NIV}| / \sigma_{NEV,m} = |Z_m|$ should have. This figure excludes points from simulated light curves with such negative values of NEV_{eq}, that the error in those measurements could not be determined. Solid lines are plotted on top of the distribution of points, corresponding to the geometric mean of $|Z_m|$ as a function of log(NIV). Transparent areas indicate the interval of the distribution corresponding to one standard deviation from the mean, in either direction.

bin, or when it shows variability that strongly differs from pink noise.

The methods for computing NEV_i and NEV_{eq}, as they are described in Sections 3.6.1 and 3.6.2 can be used for any other instrument, as they were not explicitly developed for *eROSITA*. The modifications to the standard version of these estimates of the NIV are necessary whenever a light curve consists of varying fractional exposures.

3.7 Systematic and sampling errors of estimating the band-limited power of variable sources

We have developed and evaluated methods for accurately estimating the NIV of variable sources. However, the NIV is not a constant and can change rapidly. It is also offset from the band-limited power by aliasing and the red noise leak. In this section, we will discuss these effects, and how to estimate the degree to which they afflict the variability measurement.

3.7.1 Aliasing and the red noise leak

All the variability power contained at frequencies above the inverse of the bin duration is integrated out within each bin. If the bins of a light curve are adjacent, having no gaps in between, the average count rate in each bin is not affected by power above the sampling frequency, so there is no aliasing effect. However, for light curves consisting of gaps of a constant duration between bins, power at frequencies between the sampling frequency and the inverse of the duration of each bin increase the flux difference from one bin to the next. This affects both the NIV and the periodogram. It is known as the aliasing effect (see van der Klis, 1989; Kirchner, 2005), and it is particularly strong for *eROSITA* light curves, as the gaps between observations are at least 360 times as long as the observations themselves.

Even without Poisson noise, aliasing causes each *eROSITA* observation of a variable source to be less correlated with the previous or subsequent observation than if we had observed each source continuously and had averaged all observations into 4-hour bins. This is true whenever there is non-negligible power above the Nyquist frequency. If the individual bins are substantially shorter than the gaps between them, and if there is negligible power at frequencies larger than the inverse of the duration of each bin, the timing of the observations can be mathematically described as delta functions. In that case, we can simplify the mathematical description of the impact that aliasing has on a PSD, to:

$$P_a(\nu) = P(\nu) + \sum_{k=1}^{\infty} \left(P(\nu - k\nu_s) + P(\nu + k\nu_s) \right), \tag{3.18}$$

where ν is the frequency, $P_a(\nu)$ is the observed periodogram affected by aliasing, $P(\nu)$ is the true PSD that we want to determine accurately, and ν_s is the sampling frequency. For a derivation of this equation, see Kirchner (2005).

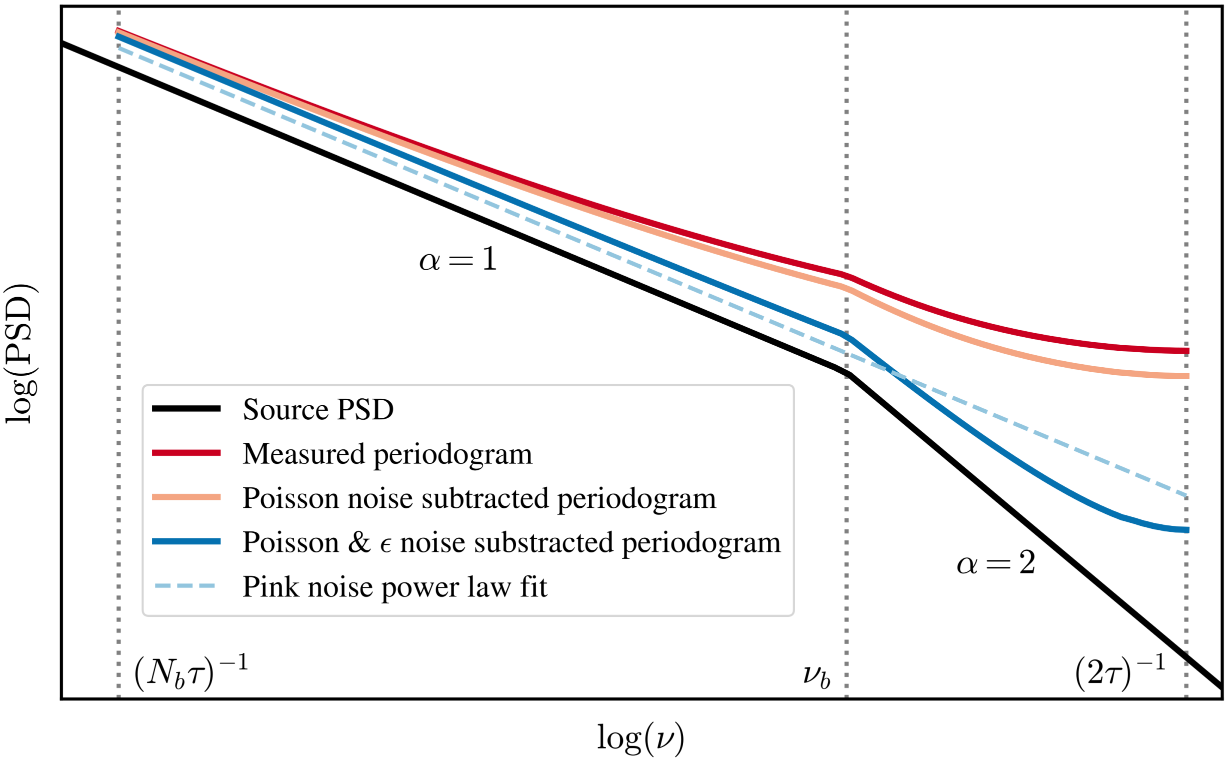

Figure 3.13: A comparison of a typical AGN PSD, and how a measured periodogram differs from it, within the frequency space that can be explored by *eROSITA* for sources lying close to the ecliptic poles. All measured and corrected periodograms are measured within the frequency range from the inverse of the duration of the entire light curve, $(N_b\tau)^{-1}$, to the Nyquist frequency, $(2\tau)^{-1}$. The intrinsic PSD is depicted as following a broken power law of $P \propto \nu^{-1}$ at $\nu < \nu_b$, and $P \propto \nu^{-2}$ at $\nu \geq \nu_b$, and continues beyond the frequency range probed by the observations. The measured periodogram is offset from the intrinsic PSD by the Poisson noise, the ϵ noise, aliasing, and the red noise leak. The ϵ noise (Eq. 3.28) is often significantly larger than the standard Poisson noise (Eq. 3.25). The flattening of the Poisson and ϵ subtracted periodogram at the highest frequencies is due to aliasing. The red noise leak increases the amplitude of the power laws. $P \propto \nu^{-1}$ is a reasonably accurate first-order approximation for the overall shape of the aliased periodogram. The relative strength of the Poisson, fractional exposure noise, aliasing, and the red noise leak, depends on the shape of the PSD, the average count rate, and the mean, and variance of the fractional exposure.

The first requirement for using the above equation, that the gaps between bins are significantly longer than the bins themselves, is true for all sources observed during the eRASSs. The second requirement of negligible variability above a frequency of 2.5×10^{-2} Hz does not apply to all variable sources observed by *eROSITA*, especially not for XRBs, which vary on much shorter timescales than *eROSITA* can detect. However, this requirement is satisfied by AGNs, as they do not show variability at shorter timescales (Tennant and Mushotzky, 1983). A variability analysis in the frequency space surveyed by *eROSITA*, of $2.11 \times 10^{-8} - 3.47 \times 10^{-5}$ Hz is also most useful for studying AGNs. This is the interval within which they vary strongly and within which we can expect the break from a pink to a red noise power law to occur in many sources. Therefore, Eq. 3.18 accurately describes the effect that aliasing has on the periodograms of AGNs observed by *eROSITA*.

For AGN PSDs described by power laws, in which variability power decreases with increasing frequency, aliasing predominantly affects periodogram frequencies close to the Nyquist frequency. It causes the power law slope to flatten at the highest frequencies observed. Fig. 3.13 depicts various effects that afflict a measured periodogram and offset it from the source PSD. The blue line in Fig. 3.13 depicts the flattening of the periodogram towards the Nyquist frequency, caused by the aliasing effect. Aliasing manifests itself in a light curve generated from a power law PSD, by reducing the correlation in the count rate of adjacent bins. For eRASS AGN light curves, this means that frequently only the long-term trends, spanning multiple eroDays, can be identified.

Aliased periodograms can be fitted using Eq. 3.18 to determine an estimate of the intrinsic shape of the PSD of the variable source, if there were no aliasing. However, this requires some assumptions about the shape of the PSD above the Nyquist frequency, for which no data are available. Fortunately, we can rely on previous analyses of AGN periodograms at higher frequencies to inform this assumption (Papadakis et al., 2002; González-Martín and Vaughan, 2012). We suggest fitting *eROSITA* periodograms with either one or two power laws, featuring a sudden break in between. A maximum frequency needs to be defined, above which we expect not to find any power. This is essential to avoid summing over infinitely many terms, and is also grounded in realistic physical constraints of the light-crossing timescale. We recommend extending the higher frequency power law up to a maximum frequency of $\nu_{max} \approx 1 \times 10^{-2}$ Hz, as AGNs have not been observed to show significant variability at higher frequencies (Tennant and Mushotzky, 1983; McHardy, 1988). The corresponding PSD model that can be used as an input to the aliasing effect for fitting *eROSITA* AGN periodograms, is thus defined as either:

$$P_1 = \begin{cases} A\nu^{-\alpha} & \text{if } \nu \leq \nu_{max} \\ 0 & \text{if } \nu > \nu_{max}. \end{cases} \tag{3.19}$$

or as:

$$P_2 = \begin{cases} A\nu^{-\alpha_1} & \text{if } \nu \leq \nu_b \\ A\nu^{-\alpha_2}\nu_b^{\alpha_2-\alpha_1} & \text{if } \nu_b < \nu \leq \nu_{max} \\ 0 & \text{if } \nu > \nu_{max}. \end{cases} \tag{3.20}$$

The selected PSD model can then be used in conjunction with Eq. 3.18 to determine the shape of the aliased version of it. *eROSITA* periodograms can then be fitted with this model, to investigate the properties of the source PSD, or test to what extent these models accurately represent it.

Another effect causes variability power at lower frequencies to leak into the observed frequency space, enhancing the measured degree of variability. This is known as the red noise leak and results in the NIV being larger than the band-limited power, on average. However, whereas the aliasing of $\alpha > 0$ power laws predominantly affects the variability measured at the highest frequencies, the red noise leak increases the power at all frequencies. For power law PSDs, the red noise leak mainly increases their normalisation, but does not affect the power law slopes (Zhu and Xue, 2016). This is shown in Fig. 3.13, via the offset between the intrinsic source PSD, and the low-frequency part of the Poisson and ϵ noise subtracted periodogram.

Aliasing can be modelled, and its effect on the periodogram, and the NIV estimate can be removed. As the NIV can be equated to the integral of the periodogram of a set of observations obtained with infinite accuracy, the aliasing effect increases the NIV above the level it would otherwise have had. This effect can be quantified, and its influence on the measured NIV can be subtracted from it. However, the impact of the red noise leak cannot be determined as easily. For very long light curves, it is possible to estimate the effect that low frequency power has on a smaller subset of the observed frequency range. This would involve comparing the periodogram of the entire light curve with the averaged periodogram of smaller sections of it. This can, however, only approximately remove the red noise leak up to the frequency of the inverse of the total observing duration. It would also be biased by using the same data set to determine the red noise leak, and the periodogram or NIV estimate. To properly remove the red noise leak requires assumptions to be made about the shape of the low-frequency tail of the PSD.

Even without removing the impact that the red noise leak and aliasing have on the periodograms of light curves of a set of variable sources, it is still possible to compare their estimates of NIV_∞, if it can be assumed that the combination of the two effects has a similar impact on the sources being compared. However, to do so, it is also necessary to determine the sampling errors involved when trying to estimate the NIV_∞.

3.7.2 The sampling error of estimating NIV_∞

The NIV is unaffected by uncertainties due to Poisson statistics and measurement errors. Nevertheless, it can still change over time, even if the PSD and the mechanism generating the variability remain unchanged. This has been shown by Vaughan et al. (2003) and Allevato et al. (2013). Fig. 3.14 demonstrates how strongly the NIV can vary between light curves observed at different times, even if they are generated by the same PSD with the same band-limited power. In this particular example, a 5×10^3 bin pink noise light curve was simulated, of which a 500 bin segment is shown. The NIV was computed for a sliding window interval of 50 bins each. Even without any measurement errors, the NIV was spread out over more than one order of magnitude between its maximal (2.1×10^{-1})

and minimal (1.5×10^{-2}) value. As this figure shows, the NIV can change significantly in less time than the duration of the observing window. For steeper power laws, this sampling effect is even more pronounced (Vaughan et al., 2003).

In Fig. 3.14, the orange dashed line depicts the band-limited power of the source, for the frequency range corresponding to a 50 bin light curve, which has a value of 3.5×10^{-2} for this example. The NIV_{∞} is offset from the band-limited power by aliasing and the red noise leak. In this example, the aliasing effect was not included, so the offset is only caused by the red noise leak, which amplifies NIV_{∞} to a value of 4.3×10^{-2}. Due to aliasing, light curves consisting of gaps between bins would have an even larger value for NIV_{∞}. Fig. 3.14 also shows that log(NIV) follows a normal distribution with a mean of NIV_{∞}.

In this section, we will discuss calculating the sampling error of using a single measurement of the NEV_b as an estimator of NIV_{∞}. To determine the band-limited power from NIV_{∞} requires the subtraction of the aliasing and red noise leak, which is accompanied by additional uncertainties, and biases. We label the logarithm of the sampling error when using an individual measurement of NEV_b (or NEV_{eq}, or NEV_i) as an estimate of NIV_{∞} as Δ_s. It is equivalent to the standard deviation of the distribution of log(NIV), which can be seen in Fig. 3.14.

We simulated 3.2×10^5 light curves like the one shown in Fig. 3.14 to determine the standard deviation of the distribution of NIVs for light curves consisting of $\{5, 10, 20, 30, 40, 50, 60, 75, 90, 100, 120, 140, 170, 200, 500, 1000\}$ bins. We mainly focused on the range of bins most useful for an eRASS variability analysis. We extended the range somewhat, to better determine the dependence of the sampling error on the number of bins of the light curve. The light curves were simulated using an input pink noise PSD that extended between the frequencies $(2\tau)^{-1}$ and $(10^5\tau)^{-1}$. A longer interval was chosen to ensure that the NIV_{∞} remained constant for all intervals selected from the single long light curve. This was necessary to determine the standard deviation in the distribution of the NIV of individual segments of the same light curve with the same degree of intrinsic variability.

Unlike previous simulations, we did not apply Poisson noise, a background count rate, or a fractional exposure. We only considered the true source flux variability. Therefore, the NEV calculated for a particular interval is equivalent to its NIV. The observed distribution of the NIV was only caused by the variation of the NIV around the NIV_{∞}.

We randomly selected starting positions within the simulated light curves and calculated the NIV for each selected interval. We investigated different degrees of intrinsic variability by scaling the range of variations in the simulated light curves, while keeping the mean flux constant. A lower range at a constant mean, results in a lower NIV_{∞}. For each set of 5000 such simulated light curve intervals, which all have the same number of bins, and were generated from the same PSD, with the same NIV_{∞}, we determined the standard deviation of the log(NIV) distribution. We assumed that that value is similar to the logarithm of the sampling error, Δ_s, and equated the two. We subsequently investigated the dependence of Δ_s on various parameters.

The sampling error was found to depend on both the NIV_{∞}, and the number of bins of the light curve. These two parameters are correlated with one another, as NIV_{∞} increases when the frequency interval it is determined for increases, if the variability is caused by a

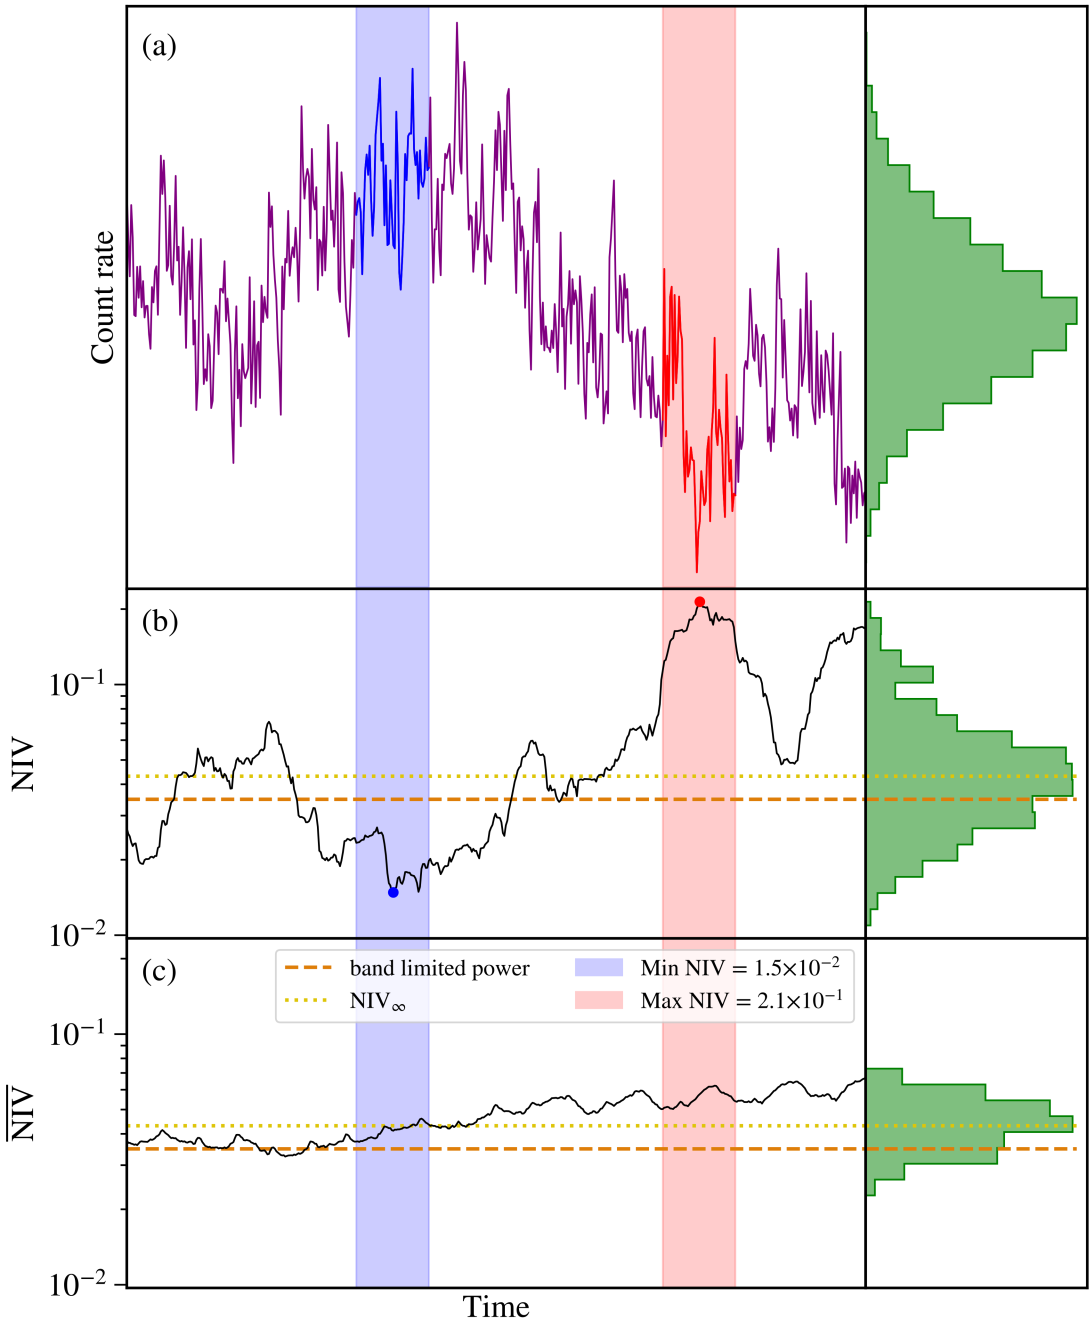

Figure 3.14: The NIV of a variable source can change even if the mechanism generating the variability, the PSD, and the band-limited power remain unchanged. Panel (a) shows a 500-bin segment of a simulated pink noise light curve without measurement errors or Poisson noise. Panel (b) shows the NIV of a sliding window interval of 50 bins from the light curve in panel (a). Each NIV data point is placed at the time of centre of the selected interval. The maximum and minimum NIV of this interval of the observed light curve of the variable source is indicated, and the parts of the light curve for which they were obtained are highlighted. Panel (c) shows the average NIV obtained by computing the geometric mean NIV of 10 adjoined segments, each consisting of 50 bins. The panels on the right show the distribution of these parameters, throughout the entire 5×10^3 bin simulated light curve.

power law PSD with $\alpha > 0$. Nevertheless, we still treat them as independent parameters, as the NIV_∞ can also vary across a wide range at a fixed number of bins of the light curve, depending on how intrinsically variable the source is.

The dependence of Δ_s on the NIV_∞ of a section of a light curve of a variable source exhibiting pink noise variability at a fixed number of bins, was found to approximately follow a quadratic relationship. Fig. 3.16 depicts the dependence of Δ_s on the number of bins and the NIV_∞. The sampling error decreases with decreasing $\log(\mathrm{NIV}_\infty)$. But the rate of decrease of Δ_s slows down at lower values of $\log(\mathrm{NIV}_\infty)$. Eventually, at very low values of $\log(\mathrm{NIV}_\infty)$, the sampling error increases again with decreasing $\log(\mathrm{NIV}_\infty)$. We therefore fitted the relationship between these two parameters as $\Delta_s = a\left(\log(\mathrm{NIV}_\infty)\right)^2 + b\log(\mathrm{NIV}_\infty) + \gamma$. We investigated the dependence of the parameters a, b, and γ on the number of bins of the light curve. Within the parameter range we had investigated, a and b seem to be independent of the number of bins. However, γ does depend on the number of bins, and this dependence can approximately be described as $\gamma(N_b) = (c/N_b)^d + f$. We found that we could accurately describe the dependence of Δ_s on the NIV_∞ and the N_b using:

$$\Delta_s = a\left(\log(\mathrm{NIV}_\infty)\right)^2 + b\log(\mathrm{NIV}_\infty) + cN_b^{-d} + f. \tag{3.21}$$

As Fig. 3.16 shows, the range of Δ_s values at a fixed number of bins does not change significantly with an increasing number of bins in the light curve. That is why we modelled the dependence of Δ_s on NIV_∞ and N_b as being additive, rather than multiplicative.

The largest value of NIV_∞ is limited by the number of bins of the light curve. Very short light curves cannot have the same maximum value of NIV_∞ as significantly longer light curves do, if they exhibit pink noise variability. This prevents the low N_b, high $\log(\mathrm{NIV}_\infty)$ corner of this figure from being populated with data points. The sampling error depends more strongly on the number of bins of the light curve, but the dependence on NIV_∞ cannot be ignored. For variabilities of $\mathrm{NIV}_\infty \approx 10^{-3}$, the sampling error only has a very wek dependence on NIV_∞. But for large variabilities of $\mathrm{NIV}_\infty > 10^{-2}$ The sampling error increases significantly with increasing NIV_∞ values.

To reduce the degeneracy between the two parameters a and b, the NIV was rescaled, so that Eq. 3.21 was instead fitted as a function of $\log(\mathrm{NIV}'_\infty) = \log(\mathrm{NIV}_\infty) - \overline{\log(\mathrm{NIV}_\infty)}$, where $\overline{\log(\mathrm{NIV}_\infty)}$ denotes the average $\log(\mathrm{NIV}_\infty)$ over all simulations. When fitting Eq. 3.21 to all the values of $\Delta_s(\mathrm{NIV}'_\infty, N_b)$ that had been obtained from the simulations, we obtained the corner plot for the best fit parameters a, b', c, d, and f', that is shown in Fig. 3.15. Table 3.2 lists the non-rescaled best fit parameters of $\Delta_s(\mathrm{NIV}_\infty, N_b)$, that were determined using: $b = b' - 2a\,\overline{\log(\mathrm{NIV}_\infty)}$, and $f = f' - b'\,\overline{\log(\mathrm{NIV}_\infty)} + a\left(\overline{\log(\mathrm{NIV}_\infty)}\right)^2$.

We had expected Δ_s to depend on $1/\sqrt{N_b}$, and indeed we found d to be close to 0.5. There is some degeneracy between c and d, and between d and f. Nevertheless, all five parameters were essential to fit the dependence of Δ_s on N_b and $\log(\mathrm{NIV}_\infty)$. We did not find sufficient indications for a more complicated fit to determine Δ_s more accurately.

A single measurement of NEV_b has a significant sampling error when used as an estimator of NIV_∞. As Fig. 3.16 shows, the sampling error increases towards a smaller number

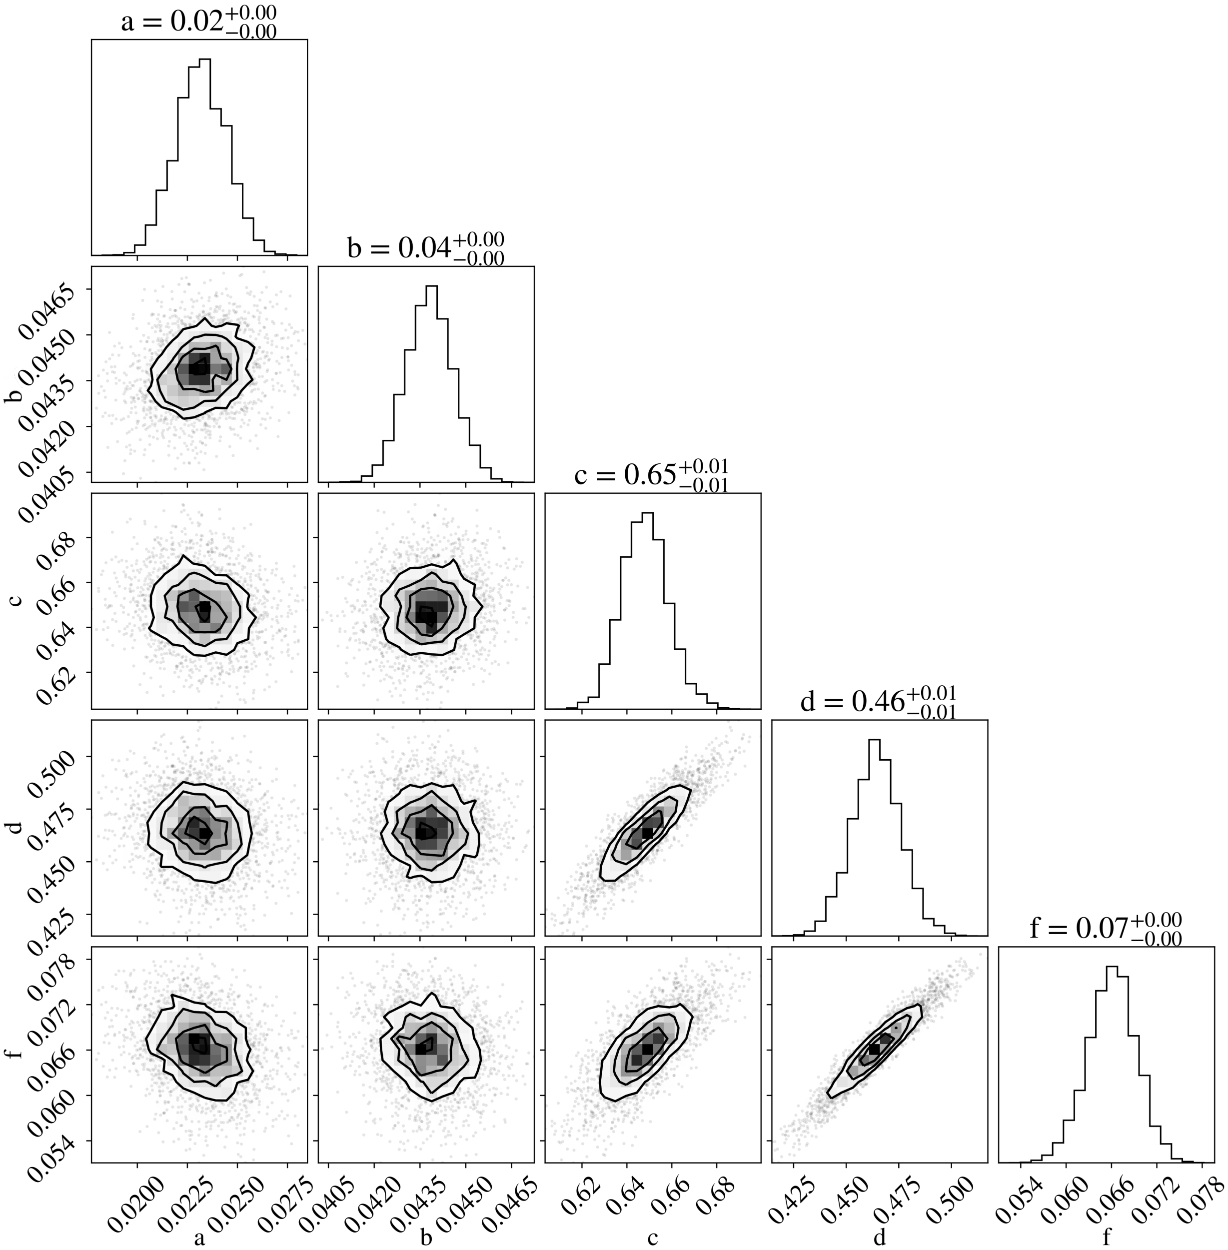

Figure 3.15: Corner plot for the best fitting parameters of a, b', c, d, and f', for fitting Δ_s, as a function of the N_b, and $\log(\mathrm{NIV}'_\infty)$. This figure is the result of 3.2×10^5 simulated pink noise light curves.

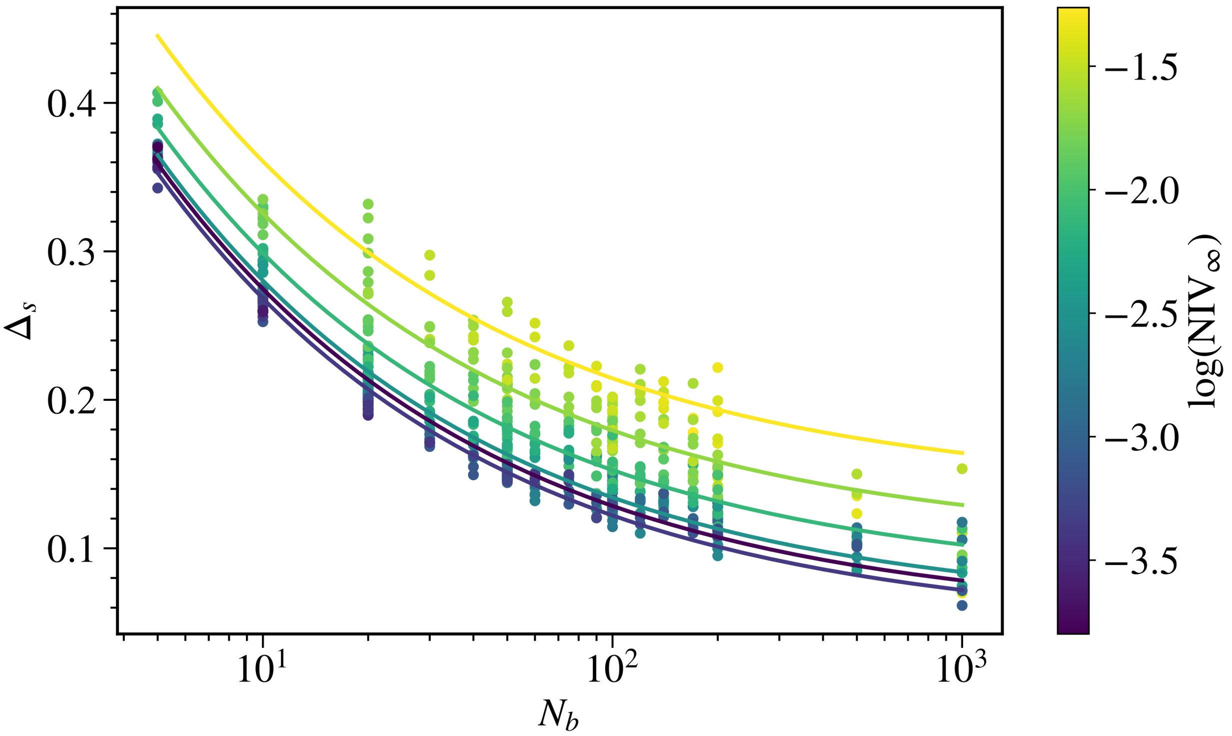

Figure 3.16: The dependence of Δ_s on the number of bins, and the NIV_∞. Each data point represents the sampling error determined from 5000 individual light curves, randomly selected from a 10^5 bin simulated pink noise light curve produced from a constant input PSD. The curves show the best fit relation of $\Delta_s(N_b)$, at fixed levels of NIV_∞. These curves are described by Eq. 3.21 with parameter values as contained in Table 3.2. There is a significant scatter among the data points due to the probabilistic nature of this estimate.

Parameter	**Value**
a	0.0232 ± 0.0013
b	0.1513 ± 0.0061
c	0.648 ± 0.010
d	0.464 ± 0.012
f	0.2920 ± 0.0079

Table 3.2: Table of the best fitting parameters of Δ_s, as a function of N_b, and $\log(\mathrm{NIV}_\infty)$, described by Eq. 3.21.

of bins in the light curve. However, the sampling error is large, even for long light curves. The difference between the errors in the positive and negative directions results from Δ_s being defined on a logarithmic scale. These errors are shown on a linear scale.

In practice, when calculating the sampling error, NIV_∞ is unknown. This means that it has to be approximated to be equal to the NIV estimate. The sampling error is not dependent on the method used to estimate the NIV. However, as NEV_b is the most accurate estimator of the NIV, we also used it as an estimator of the value of NIV_∞ in Eq. 3.21. As the NIV_∞ is always positive, negative values for NEV_i or NEV_{eq} cannot be used as estimators of NIV_∞. This means that the size of the sampling errors found by Eq. 3.21 are likely to frequently be under-, or overestimated. This is especially a problem for strongly variable sources, for which Δ_s depends strongly on the accuracy of the NIV_∞ estimate. This issue can only be reduced by decreasing the measurement and sampling errors.

3.7.3 Reducing the sampling error by averaging over multiple segments

As the sampling error of using a NEV_b measurement to estimate the NIV_∞ is large, it is of great interest to attempt to reduce it as much as possible. The best way to do this is to continue observing the variable source for a longer time, to produce a longer light curve.

Without additional observations, the sampling error can instead be reduced by splitting the total observed light curve into several segments. Instead of computing one NIV estimate for the entire frequency range from $(N_b\tau)^{-1}$ to $(2\tau)^{-1}$, a more precise estimate over a smaller frequency range can be found by computing the geometric mean over the NIV estimates found in N_{seg} segments of N_b/N_{seg} bins each:

$$\log\left(\overline{\mathrm{NEV}_b}\right) = \frac{1}{N_{seg}}\sum_{j=1}^{N_{seg}} \log(\mathrm{NEV}_{b,j}), \tag{3.22}$$

where $\mathrm{NEV}_{b,j}$ is the value of NEV_b for segment number j of the total light curve.

This method of averaging the NIV estimates found in multiple segments requires that they each cover the same frequency range. The $\overline{\mathrm{NEV}_b}$ estimates NIV_∞ with a smaller sampling error, but applies to a reduced frequency interval in which variability can be investigated and quantified; from $N_{seg}/N_b\tau$ to $(2\tau)^{-1}$.

Fig. 3.14 panel (c), shows the geometric mean of the NIV, calculated for ten adjoined segments of 50 bins each. The adjacent histogram shows that the average NIV does not vary much between different parts of the light curve. The greater the number of segments used to estimate NIV_∞ is, the more precise it becomes. However, for a finite set of observations of a variable source, a greater number of segments implies a smaller number of bins in each segment. The shorter the segments, the less valuable the NIV estimate is, as it corresponds to a smaller frequency interval. Shorter segments also suffer from having a larger variance of the NIV estimates in each. So a choice of the number of segments into which a light curve is split needs to find a balance between maximising the precision and maximising the frequency space investigated.

Vaughan et al. (2003), and Allevato et al. (2013) describe the use of a linear ensemble average of the NEV. However, if the distribution of the NIV is log-normal, a geometric mean provides a more accurate estimate of the NIV_∞. If the NIV follows a log-normal distribution, an arithmetic mean NEV will tend to overestimate the NIV_∞. This means that the linear average NEV will likely feature an additional positive variability offset from the band-limited power, in addition to the red noise leak and the aliasing effect. However, calculating a geometric mean NEV is not possible if the NEV of any segment is negative. As the NEV_b is never negative, the $\overline{\mathrm{NEV}_b}$ can always be computed, regardless of the mean count rate, or degree of variability in different segments.

We distinguish between the logarithmic sampling error of a single measurement of NEV_b, which we label as Δ_s, and the logarithmic sampling error of a geometric mean $\overline{\mathrm{NEV}_b}$, which we label as $\Delta_{s,n}$. We investigated the dependence of the sampling error of using $\overline{\mathrm{NEV}_b}$ as an estimator of NIV_∞ on various parameters, through additional simulations, identical to those described in section 3.7.2. We found that it depends on the NIV_∞, the number of bins per segment, and the number of segments.

The sampling error is also affected by the placement of the segments in relation to each other. If all the segments used to compute $\overline{\mathrm{NEV}_b}$ are spaced randomly, such that there are long, and inconsistent gaps between each of them, then the NIV in each segment is independent of the NIV of any of the other segments. In that case,

$$\Delta_{s,n} = \frac{\Delta_s}{\sqrt{N_{seg}}}. \tag{3.23}$$

In this equation, Δ_s is described by Eq. 3.21, with parameter values as listed in Table 3.2. The parameter N_b in Eq. 3.21 now refers to the number of bins in each of the segments of the light curve. The NIV_∞ is estimated to be $10^{\overline{\log(\mathrm{NEV}_b)}}$, using Eq. 3.22.

In most cases, the NIV in one segment of the light curve is dependent on the NIV in other segments. This is particularly true if the segments are adjoined, with one starting immediately after the previous one ended. When splitting a continuous set of observations into segments of equal duration, with each segment starting right when the previous one ended, the NIV in each segment depends, to some extent, on the NIV of the previous and subsequent segments. This implies that the ability to reduce the sampling error of estimating the NIV_∞, through determining the $\overline{\mathrm{NEV}_b}$ of adjoined segments, is diminished compared to using segments with long gaps in between. This can, for instance, be seen in Fig. 3.14(c), in which the geometric mean NIV is determined for ten adjoined segments consisting of 50 bins each. The range of values of the average NIV, computed over ten adjoined segments of 50 bins, is smaller than the range of the NIV in 50 bins, but not by a factor of $1/\sqrt{10}$.

For *eROSITA* observations of sources spanning multiple eRASSs, there is a natural way to separate the light curve into segments during which the source was observed every eroDay, separated by intervals with no observation. The gaps between sets of observations are not random, but are, in most cases, considerably longer than the duration of the observations, so the sampling error of estimating the NIV_∞ is approximately described

Parameter	Value
m	0.1782 ± 0.0027
k	0.0824 ± 0.0063

Table 3.3: Table of the best fitting parameters describing the sampling error of $\overline{\mathrm{NEV}_b}$, as a function of N_{seg} and NIV_∞, if the segments are adjoined. They relate to Eq. 3.24. These parameter values can be used alongside those listed in Table 3.2 to determine the sampling error, when combing Eqs. 3.21, and 3.24.

by Eq. 3.23. However, this does not apply to sources within 0.5° of the ecliptic poles. This procedure would also result in NIV_∞ estimates for a significantly different number of bins, from source to source. A different possibility is to select a smaller number of bins per segment, such as 20, 50, or 100. Subsequently, the light curves of all sources could be split into segments of that length. In that case, the NIV_∞ estimates are more comparable between different sources, but the sampling errors of the estimate could be larger, as many of the segments will be adjoined.

To determine how the sampling error, $\Delta_{s,n}$ depends on the number of adjoined segments, we simulated longer light curves in the same way as previously described, and selected adjoined segments for computing the $\overline{\log(\mathrm{NEV_b})}$. There is a significant spread of estimates from individual simulated light curves, complicating the estimate of the dependence on N_{seg}, as can be seen in Fig 3.17. These simulations investigated segments of 20, 50 and 100 bins, and a number of segments ranging from 1 to 200. The product of N_b and N_{seg} was constrained to be less than or equal to 1000. This was chosen to approximately match the maximum number of bins that the light curve of a source could have in a single eRASS. We also investigated whether the NIV_∞ impacted the ability of $\Delta_{s,n}$ to be reduced by averaging the NEV_b over multiple adjoined segments of a light curve. At smaller values of $\log(\mathrm{NIV}_\infty)$, the sampling error of the NIV estimate decreases more rapidly with increasing N_{seg}. It was expected that $\Delta_{s,n}$ would also decrease with increasing N_{seg}, following a power law relationship, as: $\Delta_{s,n}(N_{seg}) = \Delta_s N_{seg}^{-\gamma}$. The scaling factor is set by the requirement that $\Delta_{s,n} = \Delta_s$ when $N_{seg} = 1$.

Visual exploration of the simulated data suggested that γ depends approximately linearly on $\log(\mathrm{NIV}_\infty)$, but does not depend on the number of bins of each segment of the light curve. Therefore, we attempted to describe $\Delta_{s,n}$ as:

$$\Delta_{s,n} = \Delta_s N_{seg}^{m\log(\mathrm{NIV}_\infty)+k}, \tag{3.24}$$

where Δ_s is described by Eq. 3.21, and the parameters listed in Table 3.2. The NIV_∞ is estimated by $\overline{\mathrm{NEV}_b}$, using Eq. 3.22. We fitted the results of the simulations of the drop in the sampling error as a function of the number of adjoined segments with this equation and obtained best-fit values for m and k as listed in Table 3.3.

This relationship is not necessarily reliable for light curves caused by PSDs that differ from pink noise. For these best fit parameters, Eq. 3.24 can no longer be used if $\overline{\mathrm{NEV}_b} < 5.75 \times 10^{-4}$, or if $\overline{\mathrm{NEV}_b} > 0.263$, as in those instances, $m \log\left(\overline{\mathrm{NEV}_b}\right) + k < -0.5$, and

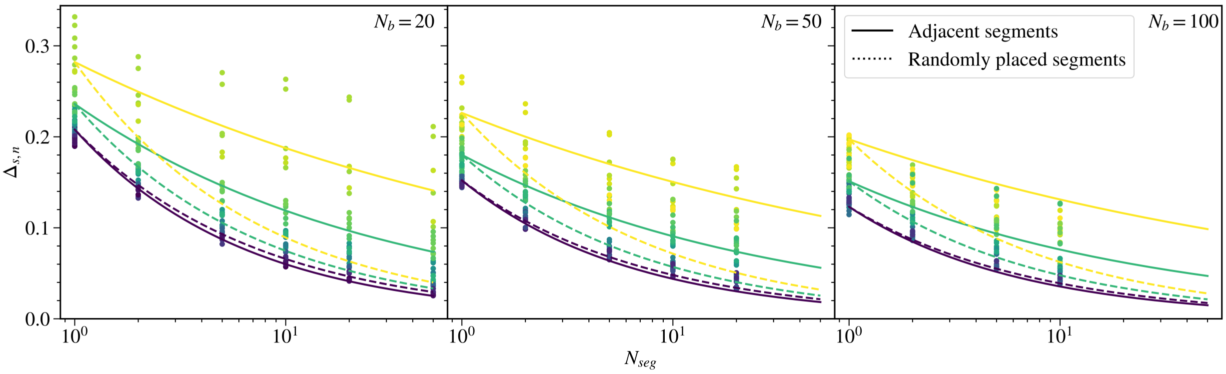

Figure 3.17: The sampling error of estimating the NIV_∞, by determining the geometric mean NEV_b over N_{seg} segments, each consisting of N_b bins. This figure shows that the sampling error of the measurement can be significantly reduced by averaging it over several segments. The data points indicate the results of individual simulations. The lines depict the best fits to the simulations. We only show data points for the simulations of the sampling error when averaging the NEV_b of adjoined segments. The solid lines show the dependence for adjoined segments, and are described by Eqs. 3.24 and 3.21, and the parameter values in Tables 3.2, and 3.3. The dotted lines depict $\Delta_{s,n}$ for NIV_∞ estimates, that are determined by calculating the geometric mean NEV_b of distant segments. These lines are described by Eqs. 3.23 and 3.21, with parameters in Table 3.2.

$m \log\left(\overline{\mathrm{NEV}_b}\right) + k > 0$, respectively, which should not happen. However, pink noise light curves of significantly variable sources should have $\log\left(\overline{\mathrm{NEV}_b}\right)$ values well within these limits. A source with $\overline{\mathrm{NEV}_b} < 5.75 \times 10^{-4}$ has a hardly measurable degree of variability, and will be indistinguishable from a non-variable source in eRASS light curves. In order to reach a high average variability of $\overline{\mathrm{NEV}_b} > 0.263$ within sets of 20, 50, or 100 bin segments, the PSD for that source would have to significantly differ from pink noise, in which case this estimate of $\Delta_s(N_b, N_{seg}, \mathrm{NIV}_\infty)$ is likely to be unreliable, and should not be used. Just as in previous instances, the sampling error can be estimated in practise, by estimating NIV_∞ to be equal to $\overline{\mathrm{NEV}_b}$.

Fig. 3.17, displays how $\Delta_{s,n}$ decreases with increasing N_{seg}, for the two different arrangements of segments used to determine the geometric mean NEV_b considered here. At the highest NIV_∞ values, the sampling error is reduced much more slowly if the NEV_b values are determined for adjoined segments, especially when N_b is small. At the lowest NIV_∞ values we simulated, Δ_s drops similarly with increasing N_{seg} for adjoined and randomly positioned segments. There is a significant scatter in the data points. Increasing the precision would have required a significantly greater number of simulations than were used.

Even though the sampling error does not drop as rapidly with an increasing number of segments if they are all adjoined, splitting a light curve apart can significantly reduce the sampling error of estimating NIV_∞. Splitting a light curve apart into segments of equal length can also enable a better comparison of the variability of different sources, even if they were observed with a different number of bins. If each of these light curves is individually separated into segments consisting of the same number of bins, the $\overline{\mathrm{NEV}_b}$ determined for them corresponds to the same frequency interval, thereby removing the effect that different observing durations have on the measurement.

3.8 Periodogram analysis

A periodogram is a more detailed variability quantifier. However, it requires significantly more information to be an accurate estimator of the intrinsic source PSD, than the NEV_b requires to estimate the NIV accurately. Variable sources need to be bright, and exhibit a high variability significance (see Section 3.4) for their periodograms to be accurately determined.

Using an observed periodogram as an estimator of the intrinsic source PSD requires understanding the impact of the observation characteristics on it. For eRASS observations, the long but consistent gaps between observations lead to the aliasing effect, which was described in Section 3.7.1. In addition, the varying fractional exposures within a light curve also generate an additional noise component in the data, similar to the Poisson noise, which increases all PSD powers, and needs to be subtracted to determine an accurate estimate of the PSD of the source.

3.8.1 Fractional exposure noise

Poisson noise in the count rate measurements affects the observed periodogram by increasing all powers by a constant value. In the fractional rms normalisation, the Poisson noise level for continuous pointing mode observations with a constant exposure depth in each bin is defined as $N_P = 2/\overline{R_S}$ (Belloni and Hasinger, 1990). In this equation, N_P is the Poisson noise for a light curve with a mean source count rate of $\overline{R_S}$. However, since the count rates are determined within bins of duration Δt, but the PSD frequencies are based on the constant temporal separation of the bins of τ, we need to multiply this equation by $\tau/\Delta t$, to give:

$$N_P = \frac{2\tau}{\overline{R_S}\Delta t}. \tag{3.25}$$

This calculates the Poisson noise level for a light curve that features gaps of a constant length between fully exposed bins. Here, $\overline{R_S}\Delta t$ is the average number of source counts observed per bin, $\overline{C_S}$. For light curves that do not have $\epsilon = 1$ all the time, the average number of source counts per bin is instead found via:

$$\begin{aligned}\overline{C_S} &= \frac{1}{N_b}\sum_{i=0}^{N_b}\epsilon(t_i)R_S(t_i)\Delta t\\ \overline{C_S} &= \overline{\epsilon R_S}\Delta t,\end{aligned} \tag{3.26}$$

based on Eq. 2.7. For light curves with fractional exposures that are not always 1, we replaced $\overline{R_S}\Delta t$ with $\overline{\epsilon R_S}\Delta t$ in Eq. 3.25, resulting in:

$$N_P = \frac{2\tau}{\overline{\epsilon R_S}\Delta t}. \tag{3.27}$$

This would be the correct description of the Poisson noise in the periodogram if the fractional exposure were constant during all bins. However, the fractional exposure varies throughout a set of survey mode observations. This manifests itself in the periodogram as well, as the varying fractional exposure acts as an additional noise component caused by the nature of the observations, rather than any excess variability of the source.

The varying fractional exposure affects the measured power at all frequencies in the same way; via an additive constant. This additional noise is, therefore, similar to the Poisson noise. Therefore, it can be described as effectively increasing the Poisson noise level.

This effect can, alternatively, be regarded as the consequence of computing a periodogram of a light curve of exposure corrected count rates, rather than more direct and consistent measures of the source flux. The Poisson noise equation applies to sets of consistently determined count rates, rather than exposure corrected count rates. As a result, the actual Poisson noise of an *eROSITA*-like light curve will be larger than in the standard description.

The excess noise is caused by varying fractional exposures that differ from a constant value of 1. It, therefore, has to depend on the mean ($\overline{\epsilon}$), and the variance (σ_ϵ^2) of the fractional exposure. We did not detect a dependence of the excess noise on the time ordering of the fractional exposures. The total Poisson noise ($N_{P,\epsilon}$) in a fractional rms normalised periodogram of a light curve with varying fractional exposures can, therefore, be described as:

$$N_{P,\epsilon} = \frac{2\tau}{\overline{\epsilon R_S}\Delta t}\left(1 + f(\overline{\epsilon}, \sigma_\epsilon^2)\right). \tag{3.28}$$

When defined in this way, f only depends on the mean and variance of the fractional exposure, and amplifies all noise levels by the same fraction, regardless of the average count rate in the light curve. The fractional exposure noise is $2\tau f/\overline{\epsilon R_S}\Delta t$. In this definition, the combination of Poisson and fractional exposure noise can easily be adjusted to other PSD normalisations, by modifying the term $2\tau/\overline{\epsilon R_S}\Delta t$ in Eq. 3.28, to the corresponding Poisson noise term. For Eq. 3.28 to be accurate both for light curves with, and without varying fractional exposures, we require that $f(\overline{\epsilon}, \sigma_\epsilon^2 = 0) = 0$.

We could not find an analytical expression for $f(\overline{\epsilon}, \sigma_\epsilon^2)$. So instead, we sought to determine a numerical expression for it from simulations. For this purpose, we simulated

Parameter	**Value**
c_a	188.48 ± 0.96
d_a	9.836 ± 0.019
g_a	1.9451 ± 0.0057
c_b	1159.9 ± 9.7
d_b	6.174 ± 0.022
g_b	-12.74 ± 0.18

Table 3.4: Table of the best fitting parameters of the fit from Eq. 3.29 to the excess periodogram noise level induced by the variable fractional exposure.

2×10^6 instances of 1000 bin, high count rate (30 cts/s) patterned fractional exposure light curves (see Section 3.3), of intrinsically constant sources with a negligible background. We also repeated this process for different count rates and number of bins. As expected, these two parameters did not impact $f(\overline{\epsilon}, \sigma_\epsilon^2)$, except when the light curves lacked the necessary information to be able to determine the noise level accurately. We varied the mean and variance of the fractional exposure in the simulations, to determine the influence of these two parameters on f, using the method described in Section 3.3. For each selected range in the fractional exposure, and for each variance scenario, 1000 light curves were generated. We determined the fractional exposure noise level and averaged it over all instances. In so doing, the simulated range of fractional exposure means and variances contained the observed distribution of both parameters within *eROSITA* observations.

At a constant $\overline{\epsilon}$, the extra noise due to varying fractional exposure could be approximated by a quadratic equation of the form $f(\sigma_\epsilon^2) \approx a\sigma_\epsilon^2 + b\sigma_\epsilon^4$. The lack of a zeroth order term in this expression is due to the requirement that $f(\overline{\epsilon}, \sigma_\epsilon^2 = 0) = 0$. In contrast, at a constant σ_ϵ^2, we observe that the extra noise can be well fitted with an exponential and a constant: $f(\overline{\epsilon}) \approx ce^{-d\overline{\epsilon}} + g$. Therefore, combining these two effects, we decided to fit all simulated excess noise with the following model:

$$f(\overline{\epsilon}, \sigma_\epsilon^2) = \left(c_a e^{-d_a\overline{\epsilon}} + g_a\right)\sigma_\epsilon^2 + \left(c_b e^{-d_b\overline{\epsilon}} + g_b\right)\sigma_\epsilon^4. \tag{3.29}$$

This description provided a good estimate of the dependence of f on $\overline{\epsilon}$, and σ_ϵ^2).

To reduce the degeneracy between the fitting parameters, we fitted f as a function of σ_ϵ^2, and $\overline{\epsilon}' = \overline{\epsilon} - \sum_i(\overline{\epsilon}_i)$, where $\overline{\epsilon}_i$ is the mean fractional exposure of datapoint i used in the fitting. The best fitting parameter values are described in Table 3.4, for the non-rescaled parameters $\overline{\epsilon}$, and σ_ϵ^2. Fig. 3.18 depicts the correlation between the parameters in the fit. There is a degeneracy between c_a and d_a, and between c_b and d_b, as these pairs of parameters increase the effect of the exponential functions involved.

Within the parameter space of $\overline{\epsilon}$, and σ_ϵ^2 relevant to eRASS observations, Eq. 3.29, combined with the best fitting parameter values contained in Table 3.4 provides a good description of the excess noise due to varying fractional exposures. However, this description slightly underestimates the excess noise at the highest fractional exposure variances, and

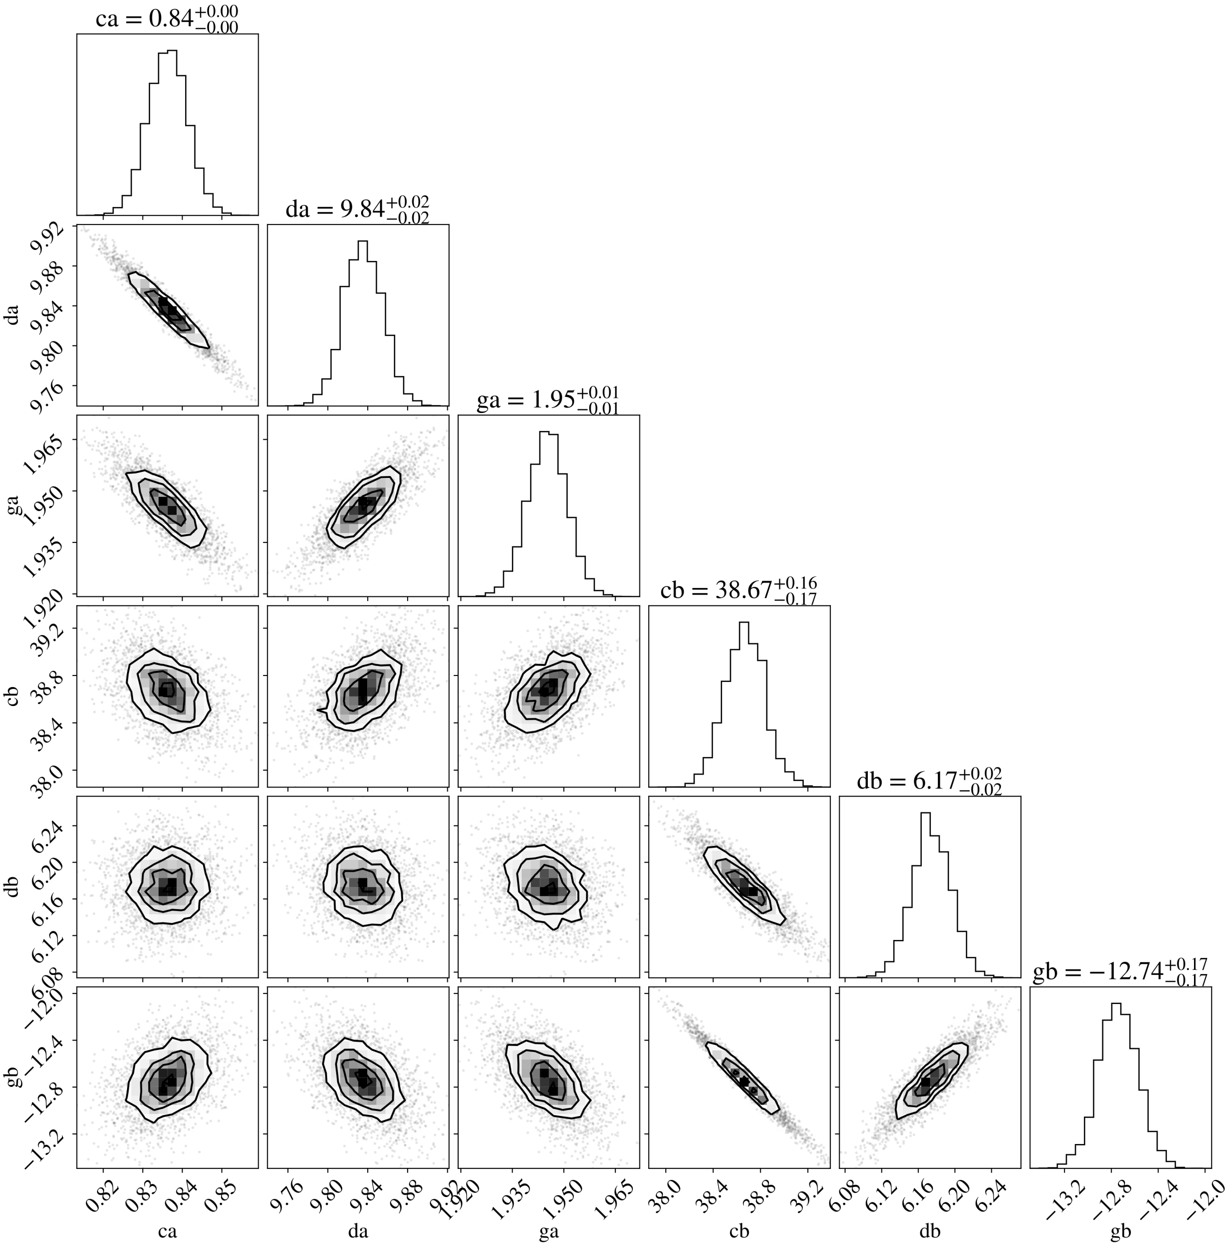

Figure 3.18: Corner plot for the best fitting parameters of c_a, d_a, g_a, c_b, d_b, and g_b for fitting the excess periodogram noise due to variable fractional exposures with Eq. 3.29.

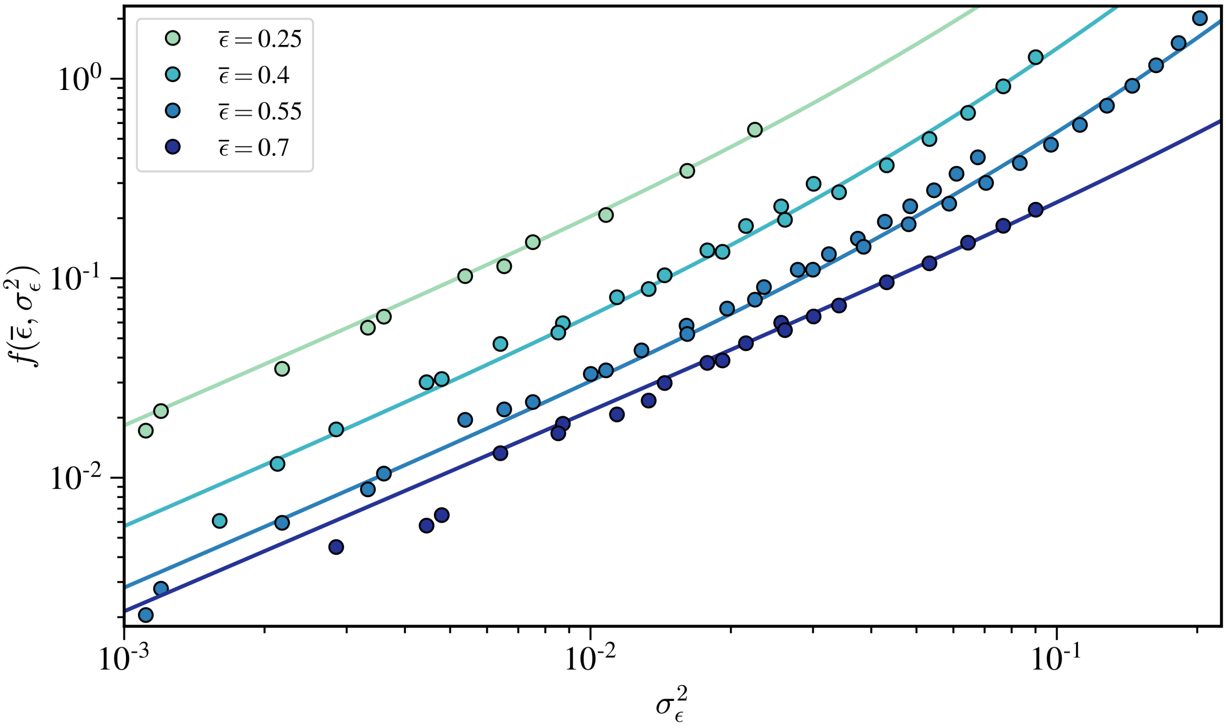

Figure 3.19: The dependence of the excess noise in a periodogram caused by varying fractional exposures, as a function of $\bar{\epsilon}$, and σ_ϵ^2. The data points indicate a subset of the results of the simulations used to determine $f(\bar{\epsilon}, \sigma_\epsilon^2)$ (Eqs. 3.28 and 3.29). The lines indicate the best fit to all data, for four particular values of $\bar{\epsilon}$, using Eq. 3.29, and the parameter values in Table 3.4.

slightly overestimates it at the lowest values. This could potentially be solved by including a term dependent on σ_ϵ^6 in Eq. 3.29. However, the range of the variance of the fractional exposure within eRASS light curves does not extend into the region where this slight offset of the excess noise from this function is relevant, so we do not consider this further.

The necessity of determining the excess noise and the accuracy of this estimate is demonstrated in Section 3.6.1 when estimating the NIV of an observed *eROSITA*-like light curve, by integrating the periodogram, after having removed the Poisson and varying fractional exposure noise. Fig. 3.13 shows that the ϵ noise can be significantly larger than the standard Poisson noise, for the typical values of $\bar{\epsilon}$, and σ_ϵ^2 corresponding to *eROSITA*.

3.9 Discussion

In this chapter, the thresholds for identifying variable sources using the two variability quantifiers SCATT_LO and AMPL_SIG were determined specifically for applicability to *eROSITA* observations of sources close to the SEP. They apply to light curves of between 50 and 1000 bins, for average count rates of between 0.001 and 30 cts/s. As neither of the two variability quantifiers uses the timing of the observations, it is possible to combine data collected over multiple eRASSs. The greater the number of bins of the light curve, at a constant bin duration, the easier it is to identify intrinsically variable sources in the frequency interval probed. Therefore, it is recommended to combine as many measures of the source flux as possible. For *eROSITA*, this means combining all observations of the same source obtained over multiple eRASSs, rather than performing multiple independent variability detection tests on individual eRASSs. After eRASS:8, all X-ray sources in the sky will have been observed on $\gtrsim 50$ eroDays. Therefore, the thresholds described in Section 3.4 can be used to identify likely variable sources throughout the entire sky, as observed by *eROSITA*. A caveat is that the thresholds might need to be adjusted for regions of the sky with significantly stronger background count rates, especially at low source count rates.

Selecting variable sources in observations by other instruments may require these thresholds to be modified. They will be impacted by the properties of the observation, the fractional exposure of the light curve bins, and the strength of the background count rate. The methods and simulations used to determine the thresholds for identifying variable sources in this chapter can be adapted to other surveys, by changing the relevant parameters.

The NEV_b method is the most accurate way to estimate the NIV of a variable source. We encourage its use in future variability analyses, to avoid the issues of the NEV method. The conversion from $\overline{\sigma_b}$ to NEV_b (Section 3.5), was made with the assumption of a pink noise PSD, and for typical properties of eRASS observations of SEP sources. The conversion was determined for the range of source count rates and number of bins typical of eRASS light curves in the SEP field. However, the conversion between σ_I and the NIV, as well as the NEV_b estimate of the NIV, can still be used for light curves obtained by other instruments. The exact value of the parameters used in the expression might require slight tweaking at low count rates. A similar set of simulations can also be performed to improve the accuracy of a similar NEV_b estimate for different types of PSDs generating the variability. Nevertheless, even without such changes, it is likely that NEV_b will still provide the most accurate estimate of the NIV of a variable source in most instances. Exceptions include sources with non-power law PSDs, or sources with consistently high count rates, with PSDs that significantly differ from pink noise.

The NEV_b method provides an accurate estimate of the NIV even for short light curves consisting of only 20 bins. We did not consider shorter light curves, as the NIV is less useful as a quantity the shorter the light curve is. We only recommend estimating the NIV for eRASS light curves consisting of at least 20 consecutive bins. That corresponds to sources within $\approx 17^\circ$ of the ecliptic poles. We recommend using AMPL_SIG as a variability measure for shorter light curves instead.

In order to compare the NIV estimate of different sources, the relevant light curves should cover the same frequency interval. This can be achieved by cropping longer light curves to be the same length as shorter ones. Alternatively, all light curves can be split into smaller segments of the same length. Then the geometric mean NEV_b can be computed for all segments of the same source, and these quantities can then be compared. This method also helps reduce the sampling error of the NIV_∞ estimate.

The equations that were determined for the sampling error of estimating NIV_{∞}, in Sections 3.7.2 and 3.7.3, apply to light curves of sources with approximately pink noise PSDs. They are not specific to *eROSITA* and can be used for light curves of any instrument. If the source PSD significantly differs from pink noise, the dependence of the sampling error on the number of segments, the number of bins per segment, and the NIV_{∞} might need to be modified. We did not investigate how strongly the PSD shape affects the sampling errors.

The description of the excess noise in the periodogram in Section 3.8.1, applies to all light curves from any single instrument or collection of instruments. It is not specific to *eROSITA*, or the fields closest to the ecliptic poles.

Periodograms require a lot of information to represent the PSD accurately. Therefore, they are only reliable for light curves consisting of many consistently spaced bins, with a large average number of source counts per bin. Therefore, it will not be possible to compute accurate PSDs for the vast majority of *eROSITA* sources.

For applicability to other surveys, the count rates quoted in this chapter can be converted to an average number of source counts per bin. The conversion is: $\overline{C_S} \approx 12\overline{R_S}$, where $\overline{C_S}$ is the average number of source counts per bin. This is a consequence of the choice of bin size of 40 s, and the mean of the fractional exposure distribution above 0.1, of ≈ 0.30 (see Chapter 2).

For optimal variability analysis of eRASS observations, we recommend first reducing the extensive data set to a more manageable size, with the methods detailed in Section 3.4. After that, we recommend using the NEV_b methodology to estimate the NIV of all variable sources. To reduce the sampling errors in the NIV_{∞} estimates and to compare the variability strength of different sources, the light curves can be split into multiple segments of the same length. Different segment lengths can be considered for different purposes; to minimise the sampling error or to maximise the frequency interval investigated. Finally, to compare the NIV_{∞} estimates of different sources, the sampling errors (Sections 3.7.2 and 3.7.3) need to be calculated. The total error is then found by adding the measurement and sampling errors in quadrature.

Periodograms can be computed for the brightest variable sources observed for the longest continuous intervals by *eROSITA*. In order to analyse the periodograms of *eROSITA* sources, it is first necessary to subtract the Poisson and fractional exposure noise from them (Section 3.8.1), and consider the effects of aliasing (Section 3.7.1).

Chapter 4

Variability properties of the *eROSITA* SEP field

After having defined variability methods for the investigation of *eROSITA* light curves in Chapter 3, we next applied them to the eRASS1, 2, and 3 data sets in the SEP field. The properties of these observations were described in Chapter 2. Here, we aim to identify variable sources in the SEP field, characterise their variability properties, and investigate individual sources of particular interest.

4.1 Variable sample in the SEP

As a first step, we sought to identify likely variable sources. Utilising the variability thresholds determined in Section 3.4, we computed the SCATT_LO and AMPL_SIG values for all sources detected in the SEP field. Next, we classified the significance of the variability of all observed *eROSITA* sources in this data set into the four variability classes: $< 1\sigma$, $\geq 1\sigma$, $\geq 2\sigma$, and $\geq 3\sigma$. These classifications are based on the thresholds defined in Section 3.4. These thresholds depend on the count rate and the number of bins of the light curve, so they are different for every source. We estimated the value of the thresholds for each source by interpolating between the values found from the simulations in Section 3.4. This was done by assuming that $\left(\frac{\partial \log(T)}{\partial \log(N_b)}\right)_{\overline{R_S}}$, and $\left(\frac{\partial \log(T)}{\partial \log(\overline{R_S})}\right)_{N_b}$ are constant within each region bounded by four points on this grid. In these expressions, T represents the threshold parameter, which is either SCATT_LO or AMPL_SIG.

In our analysis, the *eROSITA* observations were split into different eRASSs for the source detection and light curve generation. Due to slight boresight inaccuracies in the preliminary data of each eRASS, the three catalogues have positional offsets of a few arcseconds. This could lead to sources being wrongly matched across the three eRASSs. Wrongly matched sources are much more likely to be identified as variable. To avoid having a significant fraction of the variable sample consist of wrongly matched sources, we analysed each eRASS data set individually. However, this also means that this analysis is less sensitive to long-term variability.

A small number of sources had count rates or number of bins slightly outside of the interval for which the thresholds were determined from simulations. For those sources, we assumed that the dependence of the thresholds on the count rate and number of bins was identical to the nearest bound rectangle on the two parameters. For instance, SNR B0535-66.0 was observed with an average count rate of ≈ 97.5 cts/s. This significantly exceeds the interval for which we determined variability thresholds. However, at this high count rate, there are additional issues arising from *eROSITA* photon pileup, requiring careful treatment. When extrapolating the thresholds observed at lower count rates, it was identified as not significantly variable.

The combination of SCATT_LO and AMPL_SIG enabled us to look for different types of variability from different types of sources. For example, Buchner et al. (2022) found that AMPL_SIG performs significantly better for stellar-type flaring variability than for the typical continuous power law variability observed in AGNs.

Figs. 4.1 and 4.2, depict the distribution of the SCATT_LO and AMPL_SIG parameter for sources in eRASS1 relative to the 1, 2, and 3σ variability thresholds. On this logarithmic scale, sources that only vary by a negligible amount have very similar values of SCATT_LO, but a wide range of values of AMPL_SIG. Some sources were also observed to have negative AMPL_SIG values, which are not shown in this figure. This means that most of the range of values observed for SCATT_LO distinguishes between the different variability classes. In contrast, most of the logarithmic range of AMPL_SIG values covers the distribution below 1σ. The three AMPL_SIG thresholds are comparatively close together.

Many sources have a different AMPL_SIG and SCATT_LO variability significance. This is to be expected, as the two methods are sensitive to different types of variability, and have different sensitivities as a function of count rate. In the following analysis, we assigned sources to the higher of the two variability classes. For instance, if AMPL_SIG classifies a source in the $\geq 3\sigma$ variability class, but SCATT_LO places it between the 2 and 3σ thresholds, the source is considered to have a $\geq 3\sigma$ variability significance.

Using both methods increases the total number of sources classified as variable. However, it also affects the false positive rate. For one method alone, the thresholds were defined to keep the likelihood of a false positive detection at 15.9%, 2.28%, and 0.135%, for the 1, 2, and 3σ levels, respectively. If SCATT_LO and AMPL_SIG were two independent variables, then the false positive rate would almost double, to 29.2%, 4.50%, and 0.270%, for the 1, 2, and 3σ thresholds, respectively. If SCATT_LO and AMPL_SIG were to always yield the same variability significance for all sources, the false positive rate would be unaffected. In reality, the rate of false positives lies somewhere between those two extremes, but is likely much closer to being unaffected. Most sources that SCATT_LO identified as lying above 3σ were also detected at this significance in AMPL_SIG, and vice-versa.

Table 4.1 lists the number of sources in each of the three eRASSs that were found above the 1, 2, and 3σ thresholds for either the SCATT_LO or the AMPL_SIG parameter. Of the 8728, 7984, and 7770 sources detected in eRASS1, 2, and 3, we identified 4900 (56.1%), 4588 (57.5%), and 4512 (58.1%) sources to lie below the 1σ line for both variability measures.

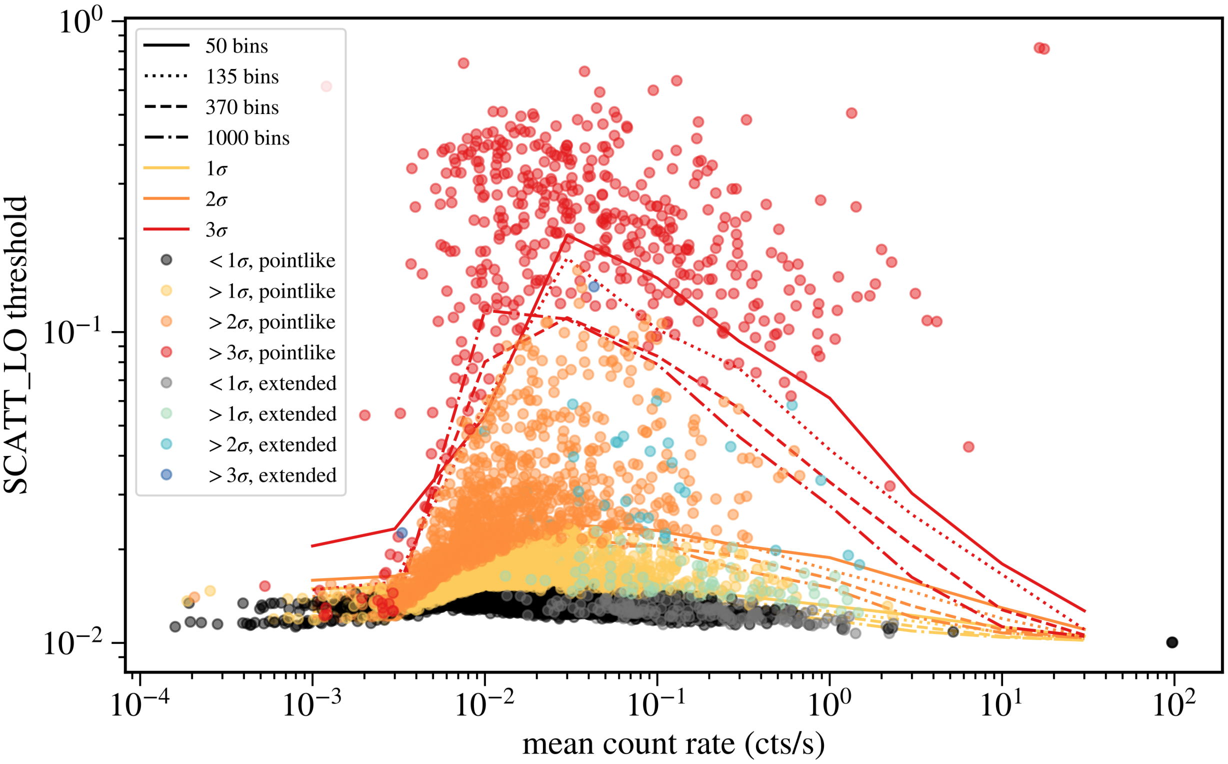

Figure 4.1: The measured SCATT_LO parameter of the light curves of all sources detected in eRASS1, 2 and 3, relative to the SCATT_LO variability thresholds. The colours denote what variability class each source was assigned. This figure also distinguishes between point-like and extended sources.

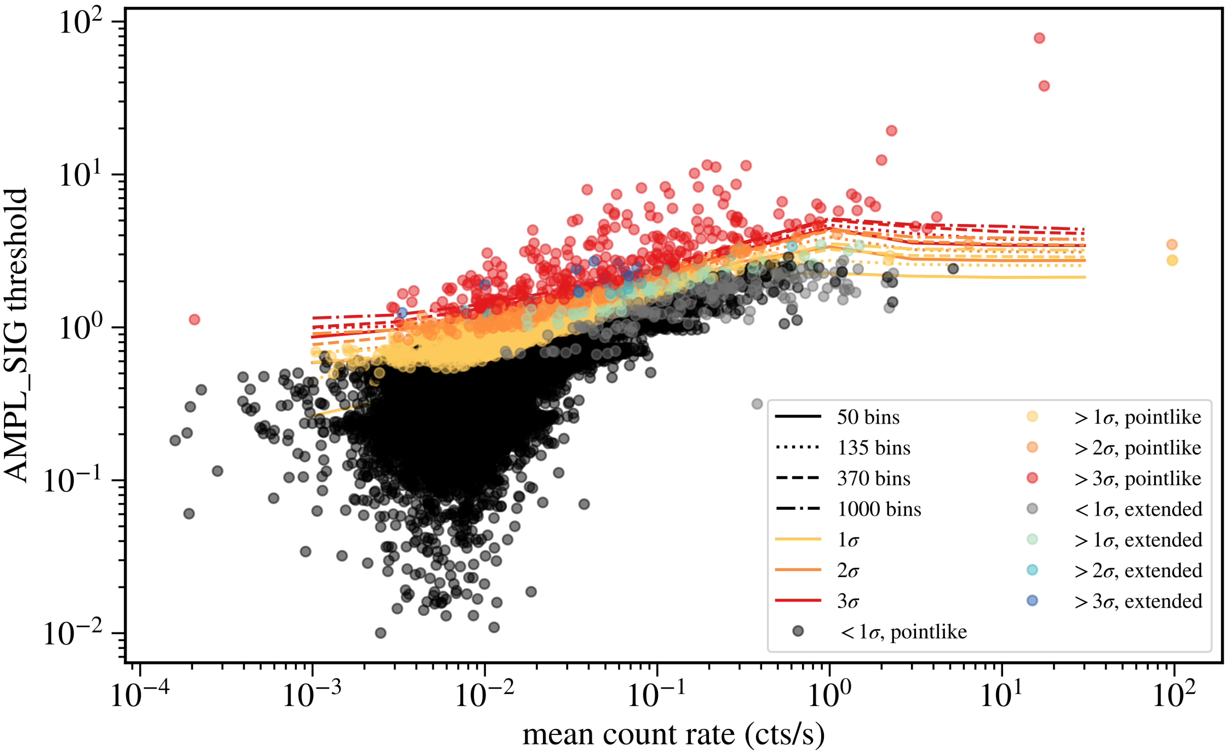

Figure 4.2: The measured AMPL_SIG parameter of the light curves of all sources observed in eRASS1, 2, and 3. The colours denote the variability class we assigned each source, based on the thresholds of AMPL_SIG. This figure also distinguishes between point-like and extended sources.

Number of variable sources	**eRASS1**	**eRASS2**	**eRASS3**
Above 1σ	3828	3396	3258
Of which: *only in SCATT_LO*	*2375*	*2138*	*2056*
only in AMPL_SIG	*711*	*596*	*554*
variable in both	*742*	*662*	*648*
expected FP	*1385*	*1267*	*1233*
extended	*96*	*77*	*67*
Above 2σ	914	822	743
Of which: *only in SCATT_LO*	*555*	*518*	*475*
only in AMPL_SIG	*183*	*153*	*106*
variable in both	*176*	*151*	*162*
expected FP	*199*	*182*	*177*
extended	*33*	*27*	*14*
Above 3σ	226	176	188
Of which: *only in SCATT_LO*	*83*	*67*	*82*
only in AMPL_SIG	*46*	*37*	*35*
variable in both	*97*	*72*	*71*
expected FP	*12*	*11*	*10*
extended	*6*	*6*	*3*

Table 4.1: Table summarising the results of the variability detection analysis applied to the *eROSITA* eRASS1, 2, and 3 observations of the SEP field. FP is short for false positives. It indicates the number of sources that can be expected to be found above the particular threshold, if all sources were intrinsically non-variable. All sources identified at each level, that are not characterised as being extended, are point-like.

These fractions are all noticeably less than the 70.8 – 84.1% we expected to find, if all sources were intrinsically non-variable. This indicates that at least 19% of all sources observed are inconsistent with being non-variable in the frequency space we investigated.

The number of sources classified as variable above a certain variability threshold drops rapidly with each additional σ, from an average of 3494 per eRASS above 1σ, to 826 above 2σ, and to 197 above 3σ. In contrast, the purity of the variable sample increases significantly with increasing variability significance. At least 37.1% of all sources detected above 1σ are expected to be false positives. However, the fraction of false positives drops to an estimated 22.5% above 2σ, and down to 5.59% above the 3σ thresholds. If we had investigated higher variability significances, we would have obtained even lower fractions of likely false positives. However, the vast majority of sources that would be excluded above the 3σ threshold are likely to be intrinsically variable sources. Therefore, we decided to use the 3σ thresholds to distinguish between variable and non-variable sources. Henceforth, we will label all sources detected above either the SCATT_LO or the AMPL_SIG 3σ

thresholds as variable.

Out of the 8728, 7982, and 7770 sources detected in eRASS1, 2, and 3, we found 184, 160, and 166 extended sources, respectively. These were not excluded from the data set, because they can act as an independent check of the variability detection technique we used. Extended sources should only feature significant variability, if they were wrongly classified as extended, if they contain a very bright point source. For instance, a bright AGN in a cluster of galaxies could be identified as a variable extended source. The fraction of extended sources detected above the variability thresholds should approximately correspond to the false positive rate.

The fraction of extended sources detected above either the AMPL_SIG or SCATT_LO 1, 2, and 3σ thresholds is 47.1%, 14.5%, and 2.94%, respectively. This noticeably exceeds even the false positive rate calculated under the assumption of complete independence between SCATT_LO and AMPL_SIG. This excess of extended sources identified as variable is predominantly due to AMPL_SIG. All 15 instances of an extended source in one eRASS being classified as variable were found above the 3σ threshold on the AMPL_SIG parameter. In contrast, merely 2 of those were also detected as variable by SCATT_LO. This occurred even though SCATT_LO detected 32% more sources above the 3σ threshold, than AMPL_SIG. This indicates that the false positive rate is presumably larger than 0.1350%, especially for AMPL_SIG. This could be a consequence of the assumptions on the value of the thresholds between the simulated points. Therefore, there are probably more than ≈ 11 sources wrongly identified as variable in each eRASS.

Fig. 4.3 compares the number of sources that were placed into each of the four variability categories for the two variability detection methods. Even though both SCATT_LO and AMPL_SIG have thresholds defined by excluding the same fraction of intrinsically non-variable sources, SCATT_LO always identified more sources above a particular variability threshold. Of the 590 instances of a source being detected as variable in one of the three eRASSs using either the SCATT_LO or the AMPL_SIG methodology, 80.0% were identified using SCATT_LO only, compared to 60.7% using only AMPL_SIG. 40.7% of all variable sources were identified above 3σ by both methods. In contrast, 39.3% of the variable sources were identified as such by SCATT_LO, but not by AMPL_SIG. Only 20.0% were identified as variable by AMPL_SIG, but not SCATT_LO. It is worth noting that the second most common of the seven categories of sources identified as variable, as shown in Fig. 4.3, amounting to 21.2% of all variable sources, were detected above the 3σ threshold in SCATT_LO, but below the 1σ threshold in AMPL_SIG.

Of these two methods, SCATT_LO is better at identifying variable sources within *eROSITA* observations. It finds 32% more sources above 3σ than AMPL_SIG does. Nevertheless, it simultaneously appears to be more reliable, as it finds 87% fewer extended sources above 3σ than AMPL_SIG. Nevertheless, AMPL_SIG identified significantly variable sources that SCATT_LO did not find. Therefore, we recommend using both methods for variability detection.

Figs. 2.2, and 2.5 - 2.11 show the distribution of various observation parameters for the variable sources identified using either AMPL_SIG or SCATT_LO, alongside the total distribution of all sources in the SEP field. The distribution of the number of eroDays of

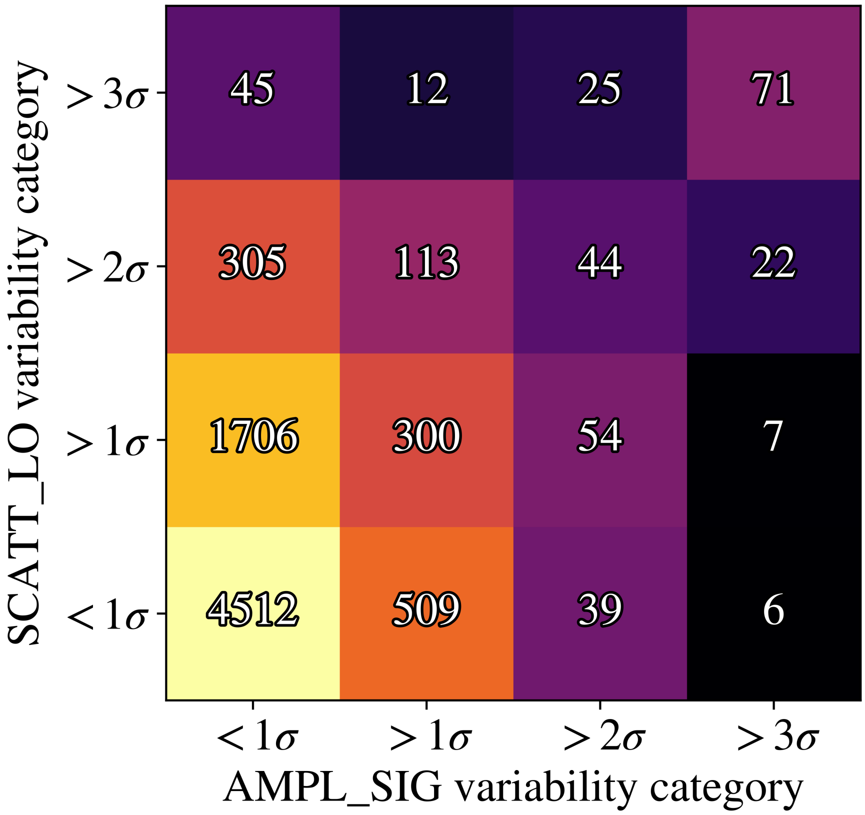

Figure 4.3: This figure depicts how many sources were identified in each of the four variability categories, for the two variability identifiers SCATT_LO and AMPL_SIG. We added the numbers of sources in each of the 16 categories from the three separate analyses of the three individual eRASSs. This means that many sources will appear up to three times in this figure, possibly in different regions. All four categories exclude sources at a higher variability significance. For example, sources that are shown to lie in the bin $> 1\sigma$ do not include the sources observed with a variability significance of $> 2\sigma$.

observation (Fig. 2.2), the total (Fig. 2.5) and the effective exposure time (Fig. 2.6) per eRASS, for the variable sources identified here, is similar to that observed for all sources in the SEP field. However, the likelihood of a source being identified as variable increases with increasing exposure time. This caused a slightly greater fraction of high exposure sources to be identified as variable. Furthermore, there is a lack of variable sources at the lowest exposure times, and the shortest light curves.

Similarly, the distribution of the background area (2.9), and the background count rate (2.10) of variable sources is very similar to that of all sources in this field, as expected. However, there are noticeable differences in the distribution of the average source count rates (Fig. 2.7) and the total source counts (Fig. 2.8), between variable and non-variable sources. The higher the average count rate and the greater the total number of source counts are, the more likely it is for a source to be identified as variable.

Variable sources have an almost uniform DET_LIKE distribution, as shown in Fig. 2.11. This contrasts the DET_LIKE distribution of all detected sources, which decreases rapidly towards larger DET_LIKE values. Therefore, the higher the DET_LIKE of a source, the more likely it is to be classified as variable.

Additionally, Fig. 2.3 shows the sky locations of the sources identified to be variable in each eRASS. The variable sources were also found to have a higher source density at smaller angles to the survey pole.

Despite some differences in these distributions, variable sources still span almost the same range as all SEP sources in these parameters. Therefore, the methods of detecting variability used here are sensitive to variability throughout the parameter space spanned by the sources in our sample.

4.2 Matching variable sources between eRASSs

After having identified variable sources in each of the three eRASSs individually, we next sought to investigate how many of them were also identified as variable in the other two eRASSs. This could allow us to create a list of unique sources that featured significant variability in at least one of the three eRASSs.

We matched the variable sources identified in each of the three eRASSs to the entire catalogue of sources in the other two eRASSs. It was not yet possible for the individual event files and their corresponding source catalogues to feature a boresight correction. Instead, we used a boresight-corrected eRASS:2 SEP catalogue (T. Liu, priv. comm.), to estimate the average boresight correction needed to determine more accurate X-ray positions for the sources detected in each of the c946 catalogues of individual eRASSs. We did this by determining the distribution of the shift in RA and Dec between the closest matching sources in the catalogues eRASSi (with i being 1, 2, or 3), and eRASS:2. We subtracted the median shift in RA and Dec between the two catalogues from the eRASSi catalogue positions. Next, this procedure was repeated using the updated positions of sources in the eRASSi catalogues. The previous steps were repeated at least two more times, until the change in the median shift relative to the previous iteration was 0. The total shift in the source positions was then used to estimate the boresight correction that needed to be applied to the individual eRASSi data sets. We performed this estimate individually for each sky tile and each eRASS. In this way, 93% of all sources detected in the SEP field in eRASS1 or 2, could be matched to a source in the boresight-corrected eRASS:2 catalogue within $30''$. Even 87% of eRASS3 sources were matched to the eRASS:2 catalogue within $30''$.

After subtracting this estimate of the average boresight offset in the source positions, we matched the sources in each of the three eRASSs to the boresight-corrected eRASS:2 SEP catalogue. The requirements for this matching were that each source in the boresight-corrected eRASS:2 catalogue could only be matched to one source in each of the individual eRASS catalogues, and vice versa. Additionally, source A_i in catalogue A, and source B_j in catalogue B were only matched, if A_i was the closest source to source B_j in catalogue A, and if source B_j was the closest source to source A_i in catalogue B. We also set an upper limit on their separation of at most $30''$. This separation was chosen to be about twice as large as the angular resolution of *eROSITA* (see Predehl et al., 2021). Sources in different eRASSs, that were matched to the same eRASS:2 source using this procedure were subsequently considered to be the same source, and their light curves and the information about them were combined.

Two sources separated by merely $38.6''$ were found to show significant variability in all three eRASSs. These sources had very similar light curves, caused by the overlapping of their point spread functions. The observed variability was most likely only caused by one of the two sources. We identified the source exhibiting a larger degree of variability, and discarded the other source from the list of variable sources. This reduced the total number of instances of a source being identified as variable within a single eRASS, from 590 (as shown in Table 4.1), to 587.

By matching all sources with the eRASS:2 catalogue, the 587 instances of a source exhibiting significant variability within one of the three eRASS, were identified with 453 unique sources. Of those, 39 sources were independently identified as variables in all three eRASSs. A further 56 sources were found to be variable in two of the three eRASSs. The remaining 358 sources, the vast majority of sources we identified as being variable, were only found to be variable in one of the three eRASSs. Of these, 148 sources were found to only be variable in eRASS1, 98 in eRASS2, and 112 in eRASS3.

There are several reasons why sources might only have been detected to have significant variability in one eRASS. In many instances, the variable sources showed a smaller degree of variability in other eRASSs, and were subsequently not found above the 3σ thresholds. Flaring sources may only have exhibited flares in one of the three eRASSs, and maintained a near-constant luminosity during the others. There were also instances of sources being so dim that they were not even detected in other eRASSs. Finally, there might also be some issues with the matching of sources, such that some variable sources could have been wrongly matched to a different source in a different eRASS.

By visually inspecting the light curves of matched sources, we identified four instances in which a variable source was almost certainly matched to a wrong source in a different eRASS. This check is only possible for sources that were observed during the transition from one eRASS to the next. In these four instances, we detected a sudden sharp rise or fall in the count rate during the transition from one eRASS to the next, even though these observations were only separated by 4 hours. The previously stated number of matched variable sources already incorporated these results. These kinds of changes in the light curve were not observed to happen at any other time for those sources. Therefore, these effects were almost certainly caused by an incorrect matching of sources. It is likely that other variable sources were also wrongly matched between eRASSs. Many sources feature significantly different average count rates in different eRASSs. However, it is to be expected that variable sources could have significantly different average count rates on either side of a gap spanning several months. Therefore, unless the transition from one eRASS to the next is observed, a change in the average count rate is possibly more likely to have been caused by the variable nature of the source than by a wrong matching of sources.

4.3 Multiwavelength source classification of the variable sample

4.3.1 Methods

After having identified 453 unique variable sources exhibiting significant variability, we next sought to investigate what these sources are, to identify the mechanism causing the observed variability. For this purpose, we determined the most likely counterparts of the variable sources in catalogues at other energy bands. We used the spectral, photometric, and parallax information available for the matched sources, to determine their nature.

We first used NWAY[1](Salvato et al., 2018) to find the most likely counterparts to the X-ray source positions of the eRASS:2 boresight-corrected catalogue of the SEP field (described in Section 4.2), in the CatWISE2020 (Marocco et al., 2021) catalogue. An eRASS:3 boresight-corrected catalogue was not available at the time of writing. We used NWAY to select only the most likely counterpart to each eRASS:2 X-ray source. The parameter `match_flag` sorts matched sources by the likelihood that they are the true counterpart. Therefore, for the following analysis, we only selected sources with `match_flag` $= 1$. NWAY is a robust Bayesian tool for counterpart identification across different wavelength bands. The CatWISE2020 positions of matched *eROSITA* variable sources in the SEP field were subsequently also cross-matched with the NSC data release 2 (Nidever et al., 2021), the VHS data release 5 (McMahon et al., 2021), and the GAIA early data release 3 (Gaia Collaboration, 2021) catalogues, all to within $2''$.

Salvato et al. (2021) introduced two methods for classifying *eROSITA* sources. In particular, they distinguished between Galactic and extragalactic sources. The first of these methods uses the X-ray flux in the $0.5-2$ keV interval, and the IR band W1, at a wavelength of 3.4 μm. Salvato et al. (2021) found that most sources detected with $W1 + 1.625\log(F_{0.5-2\text{ keV}}) + 6.101 > 0$ are extragalactic, whereas most sources at negative

---

[1] https://github.com/JohannesBuchner/nway

values are Galactic. The second relation they found utilises the W1 magnitude, and the optical g, r, and z band magnitudes. Most sources with $z - W1 - 0.8(g - r) + 1.2 > 0$ are extragalactic. To use these relationships, we converted the W1 magnitudes obtained from the CatWISE2020 counterparts into AB magnitudes. We used both of these linear relations whenever all the required magnitudes in the various energy band were known, and labelled sources which satisfied either $W1 + 1.625 \log(F_{0.5-2\ \mathrm{keV}}) + 6.101 > 0$, or $z - W1 - 0.8(g - r) + 1.2 > 0$ as likely extragalactic sources. Sources that do not meet either requirement were instead labelled as likely Galactic sources.

We also distinguished between Galactic and extragalactic sources, utilising the method described by Kovács and Szapudi (2015). They found that almost all sources whose W1 and J band magnitudes satisfy the inequality $W1 - J < -1.7$ are extragalactic. This boundary is likely too strict, as many extragalactic sources are excluded in this selection. To use this method, we converted the J magnitudes obtained from VHS to match those used by 2MASS.

A third way to distinguish between Galactic and extragalactic sources is to use the measured parallax significance from GAIA. For $p/\sigma_p > 5$ (where p is the parallax, and σ_p is its uncertainty), we have a high degree of confidence that the source is a Galactic source. Extragalactic sources should not be detected to have a parallax significantly larger than the error in the measurement. Sources with a less significant parallax can still be either Galactic or extragalactic, and need to be investigated further using the other methods we described. We also investigated the possibility of using the CatWISE2020 proper motions for distinguishing between likely Galactic and extragalactic sources. However, we found that those were not ideal for this task, as we found many AGNs with measured proper motions several times larger than the associated errors. There was also an insufficiently large difference in the distribution of the proper motion significance between Galactic and extragalactic sources.

We followed up the observations of a subset of *eROSITA*-detected sources in the SEP by obtaining optical spectroscopy of them with the 2 degree Field fibre positioner (2dF; Lewis et al., 2002) and the AAOmega (Saunders et al., 2004; Sharp et al., 2006) dual beam spectrograph, on the Anglo-Australian Telescope (AAT; Hopkins et al., 2012). The observations with good visibility took place on February 8 - 11, September 5, November 6, 2021, and January 5, 2022. Using the three methods described above, we selected likely extragalactic sources, and the positions of their respective optical counterparts, for observation. This selection was carried out over the entire SEP field, for all sources, not just for the variable sample. However, we did assign a higher priority to observe variable sources. To maximise the quality of the optical spectra, we limited the r-band magnitudes of the optical counterparts we selected, to between 17.0 and 22.5. These limits maximised the number of likely extragalactic sources we could observe, and ensured that we had enough time to get sufficiently good spectra of them. We targeted a total of 2644 likely extragalactic X-ray sources in these observations.

These observations are part of an ongoing survey of the optical spectroscopic properties of *eROSITA* detected, likely extragalactic sources, near the SEP. A detailed description of the selection of fields and targets for the observations, the data analysis, the results, and a catalogue of the X-ray variability, and the optical spectroscopic properties of the observed sources in the SEP field will be described in future work. However, we will mention some results of the spectroscopic analysis relating to the variable sources in the following sections. The observed optical spectra can clearly distinguish between Galactic and extragalactic sources, if they are sufficiently well determined. By matching the observed absorption and emission lines to common stellar, galactic, and quasar spectral types, we could determine the source type and redshift. Sources with typical stellar absorption lines at a negligible redshift are Galactic sources. Sources with galactic emission lines at a positive redshift are extragalactic sources.

We applied these four methods to separate the selected variable sources into the three groups of: likely Galactic, likely extragalactic, and unknown sources. Fig. 4.4 depicts the flowchart that we used to categorise all sources into one of these three groups. For this methodology, we had to prioritise between these four methods, so that sources could be categorised even if the methods disagreed about the source type. We decided that the reliability of the Galactic-extragalactic distinction based on AAT spectra with unambiguously fitted spectral lines was the greatest. Extragalactic sources have a non-detectable parallax, so sources that GAIA measured to have $p/\sigma_p > 5$ are almost certain to be Galactic. We also decided that the $W1 - J < -1.7$ methodology is slightly more reliable than the $W1 + 1.625 \log(F_{0.5-2\ \mathrm{keV}}) + 6.101 > 0$, and $z - W1 - 0.8(g - r) + 1.2 > 0$ distinction, as the latter uses the X-ray flux, which can vary a lot for the variable sources we investigated.

We associate a Galactic-extragalactic distinction parameter, D, to all sources, to indicate which of the three source types it was identified as, in each of the four methods we describe here. To unambiguously associate each value of D with exactly one combination of classifications by all methods, we define it as:

$$D = \sum_{j=0}^{3} 3^j d_j$$
$$d_j = \begin{cases} 1 & \text{if likely extragalactic} \\ 0 & \text{if unknown type} \\ -1 & \text{if likely Galactic,} \end{cases} \tag{4.1}$$

In this equation, $j = 0$ refers to the $W1 + 1.625 \log(F_{0.5-2\ \mathrm{keV}}) + 6.101 > 0$, and $z - W1 - 0.8(g - r) + 1.2 > 0$ methodology. $j = 1$ corresponds to the $W1 - J < -1.7$ distinction. $j = 2$ describes the $p/\sigma_p > 5$ method, and $j = 3$ corresponds to the results of the AAT spectral analysis. In evaluating this parameter, we label all sources with $p/\sigma_p \leq 5$ as 'likely extragalactic sources', even though many of them are likely to be Galactic sources. We set $d_j = 0$ whenever the data required to determine whether the source is likely Galactic or likely extragalactic in methodology j is unknown. For instance, if the J magnitude of the source is not known, we set $d_2 = 0$.

In general, sources with positive values of the distinction parameter were identified as likely extragalactic sources, sources with negative values as likely Galactic sources, and

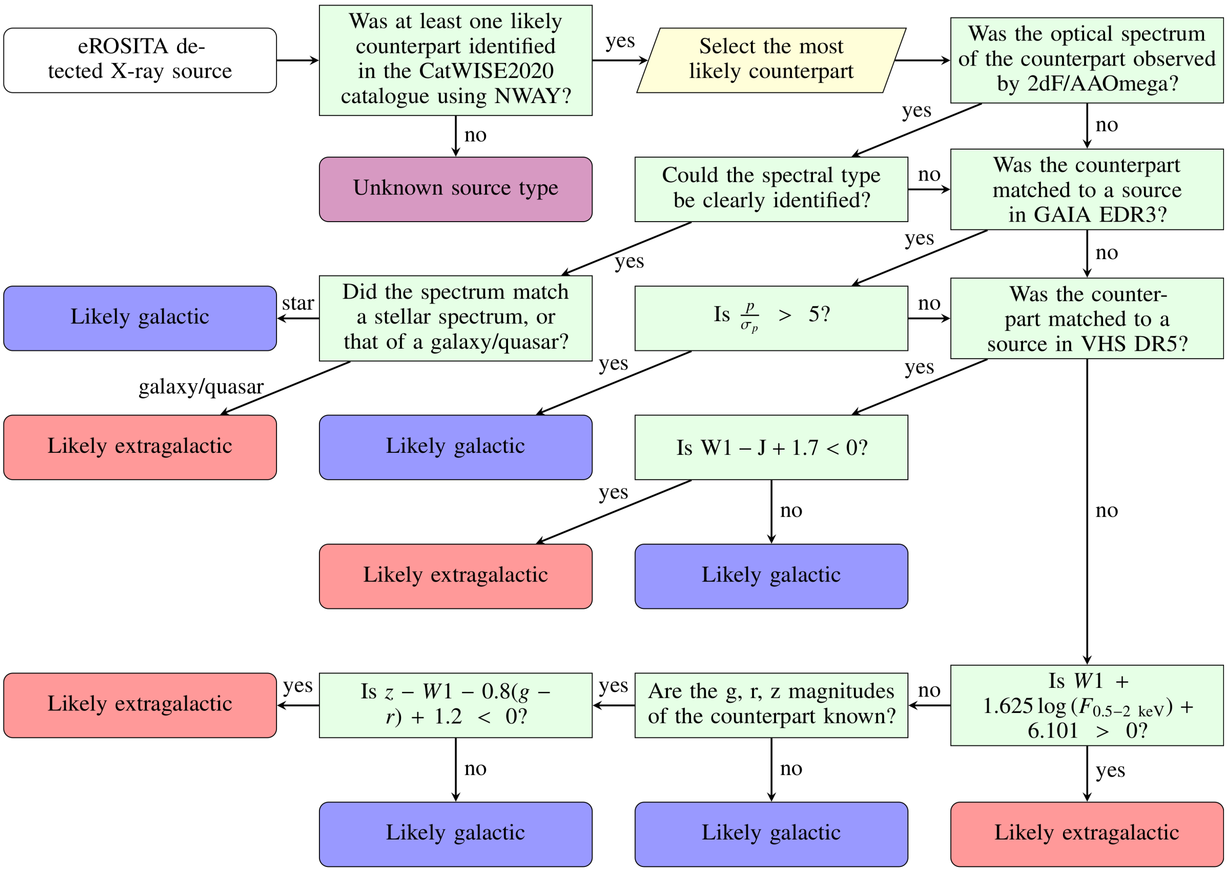

Figure 4.4: This flowchart showcases the methodology used to separate the selected variable X-ray sources into the three groups of likely Galactic, likely extragalactic sources, and sources of unknown type. We used the X-ray flux, and information that was collected on the most likely counterparts to the X-ray detections by CatWISE2020, GAIA, VHS, and the optical spectra observed by the AAT.

sources with 0 as sources of unknown type. However, because sources with $p/\sigma_p \leq 5$ could be both Galactic and extragalactic sources, all sources with $5 \leq D \leq 8$ were still identified as likely Galactic sources.

4.3.2 Identification and classification of variable sources

There are several issues to be aware of regarding the matching of eRASS3 sources. Variable sources that lie below the detection limit in eRASS1 and 2, but became significantly brighter in eRASS3 will not be contained in the eRASS:2 catalogue. Those sources will either be matched to wrong counterparts or be labelled as unknown sources. The estimate of the boresight correction that we used to match eRASS3 sources to eRASS:2 is insufficiently accurate for matching unknown sources to catalogues in other parts of the electromagnetic spectrum.

Out of the 453 individual variable sources we identified in Sections 4.1 and 4.2, 403 were matched to a source in the CatWISE2020 catalogue. Therefore, following Fig. 4.4, we classified the 50 remaining variable sources as sources of unknown type. Out of those, 23 were only detected in eRASS3. The other 27 unknown sources were detected in other eRASSs as well, but were identified as such due to an ambiguous matching between the source positions in various catalogues.

During the AAT observing campaign, we obtained the optical spectra of 120 variable *eROSITA* sources with a match in the CatWISE2020 catalogue. Of those, 81 spectra were identified to be sufficiently well observed that an unambiguous determination of their redshift, and source classification could be performed. The parallax significance was determined for 292 sources. A total of 79 sources were matched to the VHS DR5 catalogue. These are the sources for which we could determine a Galactic-extragalactic distinction using $W1 - J < -1.7$. The $z - W1 - 0.8(g - r) + 1.2 > 0$ method was used to classify 224 variable sources, for which all the magnitudes were known. Finally, all 403 of the matched variable sources could be classified using the $W1 + 1.625\log(F_{0.5-2\ \mathrm{keV}}) + 6.101 > 0$ method.

Using all data available on them, and the flowchart of Fig. 4.4, we first classified the sources with a counterpart in the CatWISE2020 catalogue into 181 likely extragalactic, and 222 likely Galactic sources. In addition to the classification described in Section 4.3.1, we analysed each source individually, searching through various source catalogues at the detected position of the X-ray source and its matched infrared counterpart. We found one AGN that was wrongly identified as a likely Galactic source, because it had $p/\sigma_p = 5.14$. We also incorrectly classified the bright star η^2 Dor as a source of unknown type. It is so bright that CatWISE2020 masked it out, so no counterpart to the X-ray detection could be identified. Additionally, six of the likely extragalactic sources, and one unknown source, were identified to be XRBs in the LMC.

For most variable sources, there was insufficient literature information available to perform an independent check of their source classification. Therefore, there could be other sources that were wrongly classified. Nevertheless, out of the 81 variable sources for which we had obtained unambiguous optical spectra with the AAT, 14 sources (17.3%) were identified to have a different source class than what was found when using the other methods. After applying these corrections to the source classification whenever possible, the 453 variable sources were finally classified into 179 likely extragalactic sources, 224 likely Galactic sources, 7 XRBs in the LMC, and 43 sources of unknown type.

Fig. 4.5 displays the classification of all variable sources with counterparts and known $W1$, g, r, or z magnitudes. It shows that all four methods are mostly reliable and consistent with each other, but that there are still some disagreements between them. All of the sources identified to satisfy both $W1 + 1.625\log(F_{0.5-2\ \mathrm{keV}}) + 6.101 < 0$ and $z - W1 - 0.8(g - r) + 1.2 < 0$ were found to be Galactic by all other methods as well. However, some of the sources identified as likely extragalactic using these two distinctions were found to be likely Galactic by other methods.

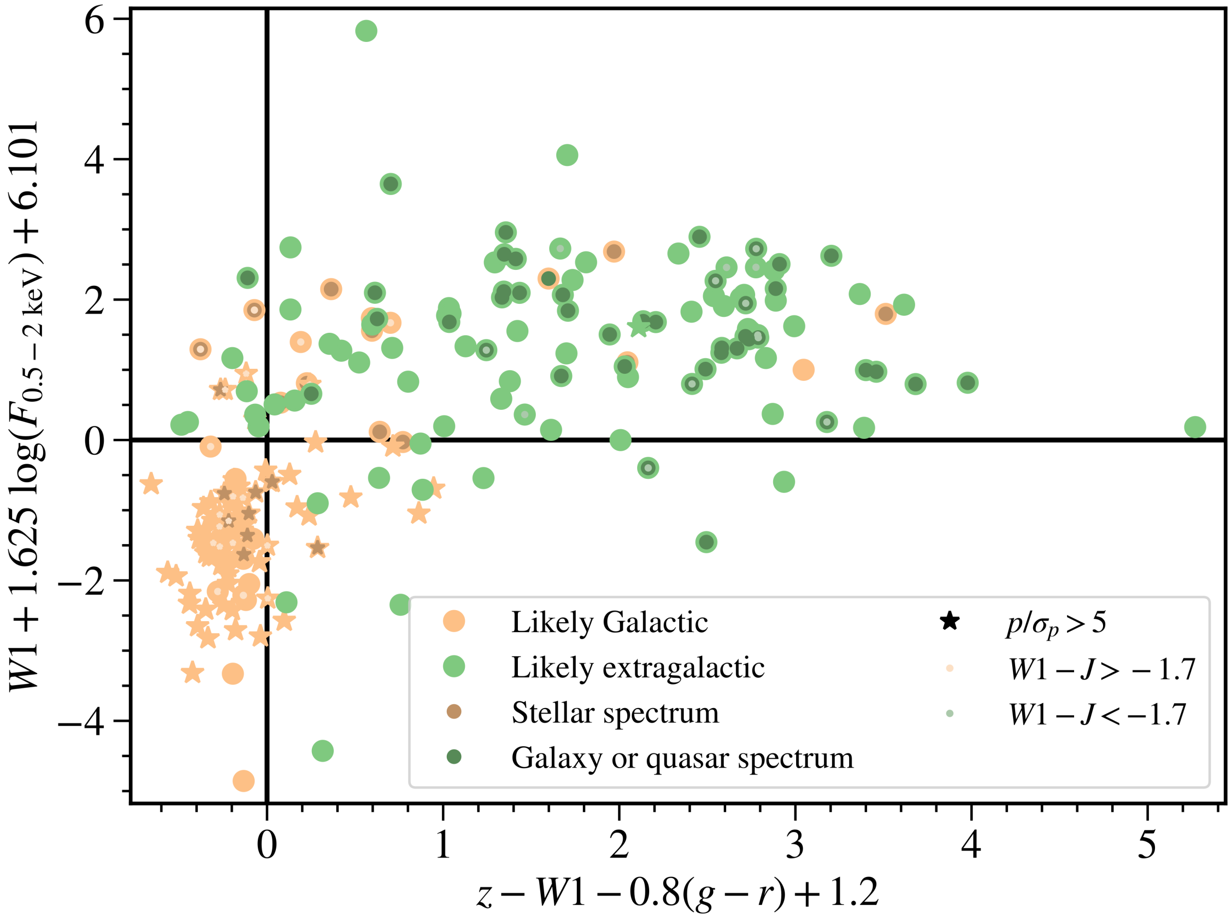

Figure 4.5: The classification of variable sources, according to the four methods outlined in Section 4.3.1. Sources that were not matched to an optical counterpart, or whose counterparts lacked information on the $W1$, g, r, or z magnitudes, are not shown.

4.4 Properties of variable sample

For many sources located between 0.5 – 3° of the SEP, the set of observations performed within each eRASS consists of two parts; one at the start, and one at the end of the eRASS, with a long break in between. Even the light curves of sources located very close to the SEP feature gaps, when the survey mode operation was interrupted for a few days for calibration observations, orbit corrections, or telescope downtime. For those sources, it is more sensible to split the light curves into the segments within which observations occurred every eroDay, rather than to split them by eRASS. This is particularly relevant for computing periodograms, which require a consistent sampling throughout. It is also relevant for computing NEV_b values, if these are to accurately describe the degree of variability within a particular frequency interval, especially if their values are compared between different sources. To remain consistent throughout, we performed the following variability analysis on individual segments of the eRASS:3 light curves, regardless of how these are located relative to the start and end times of each eRASS.

In contrast, neither SCATT_LO nor AMPL_SIG depends on the timing of the observations, whether the sources were regularly or irregularly sampled, or whether they feature a long gap in which no data was collected. Long gaps between observations can be preferable for detecting variability using these two methods, as it enables a search for possible changes over a longer timescale. Therefore, the detection of variable sources was instead performed on light curves separated by eRASS, rather than by segment.

To compare the variability strength of different sources over the same frequency range, we split the segments into smaller parts of 20, 50, and 100 consecutive bins. The NEV_b was determined for each of these smaller segments individually. These were subsequently used to calculate a single geometric average NEV_b for the entire observed eRASS:3 light curve of each variable source. These values can be compared between different sources, and feature a smaller sampling error than the NEV_b measurements of the entire segments (see Section 3.7.3).

We associated sampling errors to the $\overline{\mathrm{NEV}_b}$ measurements in accordance with the findings of Section 3.7. As many of the smaller segments are adjacent, we used the sampling errors of adjoined segments rather than those for randomly spaced segments. However, long gaps exist between some of the segments used in determining $\overline{\mathrm{NEV}_b}$, so these sampling errors might be overestimated.

Fig. 4.6 depicts the distribution of the geometric averaged NEV_b over all segments of 20 eroDays in the combined eRASS1, 2, and 3 light curves of all variable sources. This figure distinguishes between the distributions of this parameter for likely Galactic, likely extragalactic, and unknown sources. The NEV_b distribution for both the likely Galactic and the likely extragalactic sources is approximately Gaussian. However, some likely extragalactic sources have significantly lower $\overline{\mathrm{NEV}_b}$ values than most variable sources. In contrast, there are some likely Galactic sources with a significantly larger degree of variability than is commonly observed. The variable likely Galactic sources were detected within the $\overline{\mathrm{NEV}_b}$ range of $1.18 \times 10^{-2} - 2.43 \times 10^{-1}$, with a geometric mean of 4.43×10^{-2}. The likely extragalactic sources were instead found in the range $2.66 \times 10^{-3} - 1.13 \times 10^{-1}$,

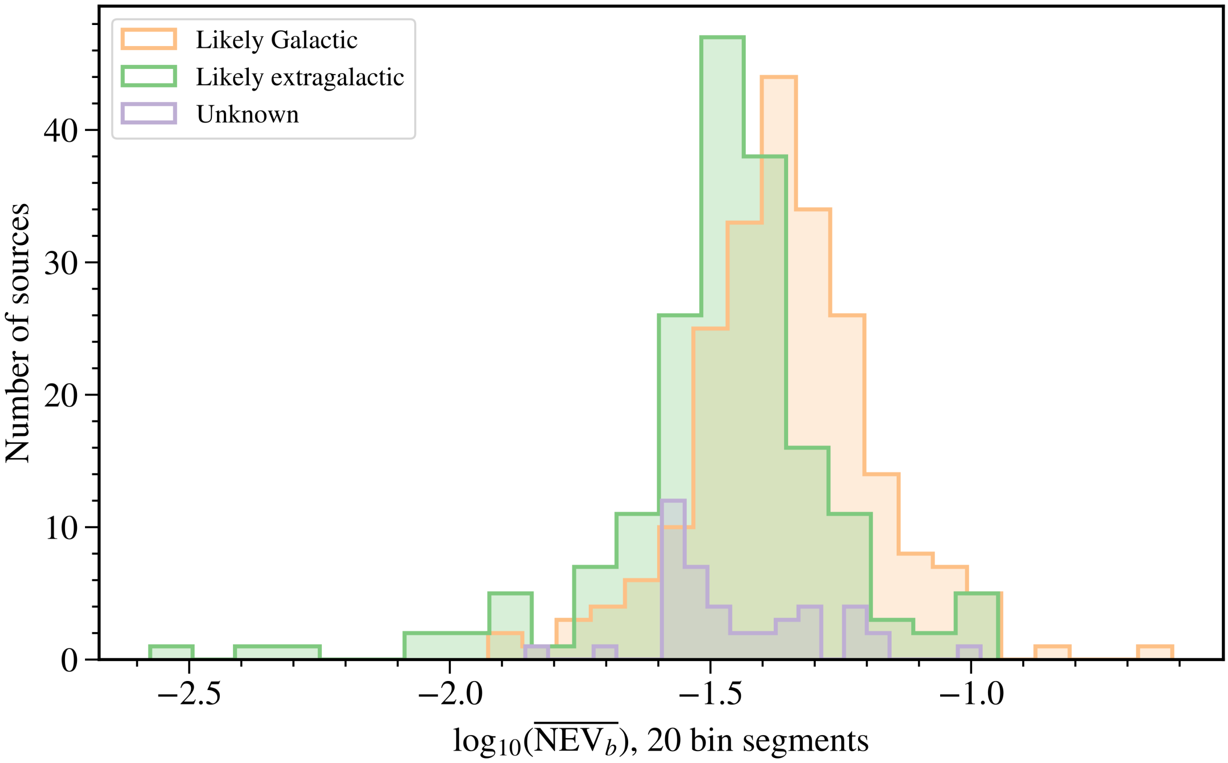

Figure 4.6: Distribution of the log averaged NEV_b of 20 bin segments of the variable sources identified in the SEP field, distinguished into likely Galactic, likely extragalactic and unknown sources.

with a geometric mean of 3.30×10^{-2}. We used Welch's t-test to determine whether there is a difference in the mean of the $\overline{\mathrm{NEV}_b}$ distribution for likely Galactic and likely extragalactic sources. We found $t = 6.28$ and an associated p-value of $p = 1.06 \times 10^{-9}$, so these are indeed distinct distributions.

We did not analyse the properties of the XRBs we detected to show variability. These will be investigated in Kaltenbrunner et al. in prep. These XRBs are some of the brightest, most variable sources that we identified, and are often found towards the upper right corner in Figs. 4.1 and 4.2.

The differences in the Galactic and extragalactic distributions of $\overline{\mathrm{NEV}_b}$ are mainly due to the different types of variability exhibited by stars and AGNs. Stellar variability is dominated by infrequent flares. The NIV of such light curves is typically larger than for continuously variable sources with power law PSDs, which usually applies to AGNs. The most variable sources have light curves which feature infrequent large flares over a low continuum.

This distinction in the variability properties of stars and AGNs is also the cause of the difference in the distribution of the SCATT_LO and AMPL_SIG parameters for them, which is shown in Fig. 4.7. Variable sources classified as likely Galactic tend to be distributed with larger AMPL_SIG values than likely extragalactic sources. By comparing

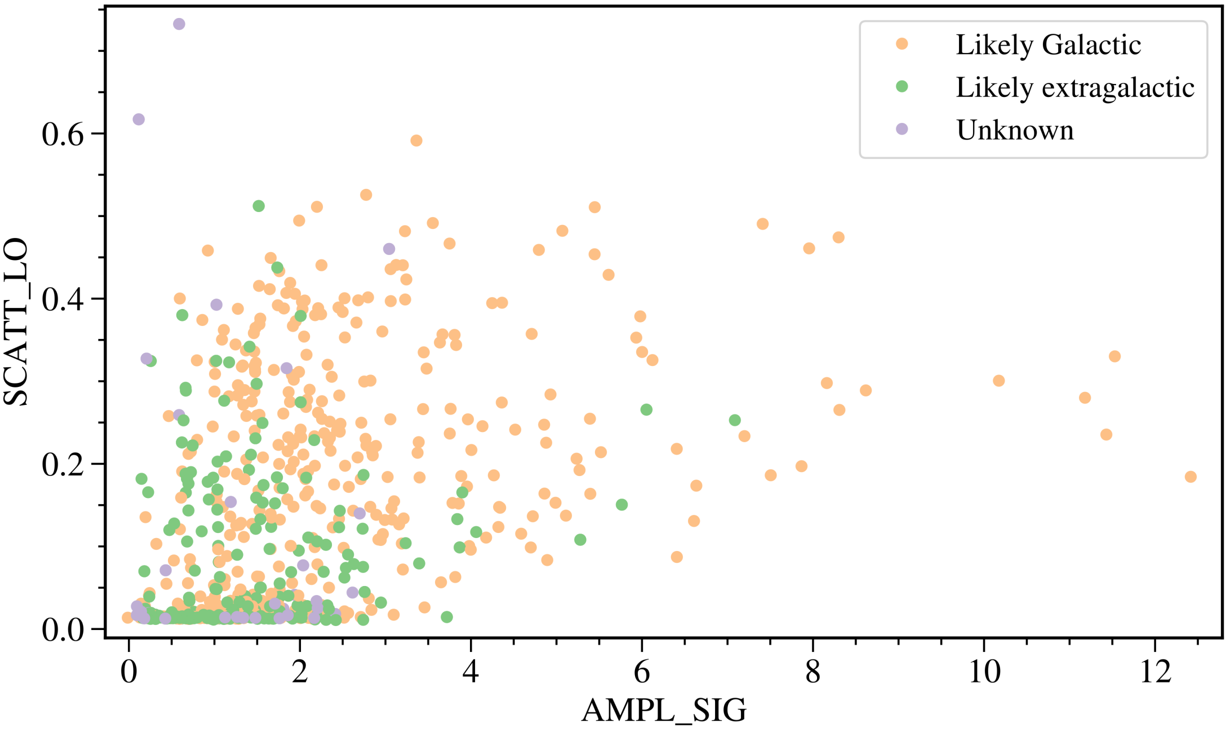

Figure 4.7: Comparing the SCATT_LO and AMPL_SIG values found for different types of variable sources in eRASS1, 2, and 3. Each dot represents a single SCATT_LO and AMPL_SIG value of a variable source in one of the three eRASSs. It also shows the values of the variability parameters for eRASSs in which the variable sources were not found above the 3σ thresholds. A likely galactic source at an AMPL_SIG value of 78, and a SCATT_LO value of 0.82, is omitted from this figure, for display clarity.

the SCATT_LO and AMPL_SIG values of variable sources, it could be possible to classify sources of unknown type. However, as Fig. 4.7 shows, the distinction is imperfect and can often lead to erroneous results. The position of a source on the SCATT_LO and AMPL_SIG plane can vary a lot from eRASS to eRASS. It can also depend significantly on the number of bins in the light curve, the average flux of the source, and whether there is a long gap between observations within an eRASS. Nevertheless, 88% of the sources identified as variable in an individual eRASS by exceeding both the SCATT_LO and AMPL_SIG 3σ thresholds, whose likely source type could be determined, were found to be likely Galactic. In contrast, 62% of sources located above the 3σ SCATT_LO threshold, but below the 3σ AMPL_SIG threshold, that were matched to an optical counterpart, were found to likely be extragalactic.

We collated a list of the variable X-ray sources identified in the SEP field in eRASS1, 2, and 3, at the webpage: https://projects.mpe.mpg.de/heg/erosita/SEP_var/. It lists some of the most relevant properties that we determined for these sources. It includes the boresight-corrected locations of the *eROSITA* detected X-ray sources, as well as the positions of their most likely optical counterparts. In addition, we included the most likely redshifts, spectral type, and measure of the spectral quality, for the sources whose optical spectra were observed by the AAT. The website also lists the Galactic-extragalactic distinction parameter and the most likely source classification. It additionally includes the number of eroDays of observation, the average source count rate in each segment, the SCATT_LO and AMPL_SIG variability classifications of the eRASS1, 2, and 3 data sets, the NEV$_b$ values of all the segments of the light curve, and the geometric average NEV$_b$ of all segments of 20 consecutive bins within the eRASS:3 light curve. Images of the rebinned, and non-rebinned light curves of all variable sources are also included.

4.5 Periodograms of variable likely extragalactic sources

We investigated the properties of the periodograms of variable sources identified to likely be extragalactic. We only selected the 20 sources with at least 2×10^3 background-subtracted source counts in the first three eRASSs for this analysis, as there was insufficient information to constrain the variability power at different frequencies for most sources fainter than that. We generated separate periodograms of each segment of the light curves of these sources. We only selected continuous segments of at least 20 consecutive eroDays, without any breaks in between. The Poisson and fractional exposure noise were subtracted from the periodograms following the results of Section 3.8.1.

We jointly fitted all periodograms of all segments of a variable source using the methodology described by Vaughan (2010). This methodology uses a Bayesian approach that defines the best fit as the function that minimises the deviance (d):

$$d = -2 \log L = 2 \sum_j \left\{ \frac{P_j}{S_j} + \log S_j \right\}, \tag{4.2}$$

where L is the likelihood function, S_j is the power of the model at a frequency of $j/(N_b \tau)$, and P_j is the periodogram power. This function, however, only works for positive values of P_j. Due to the statistical nature of *eROSITA* data, some periodogram powers are negative, especially at higher frequencies. Therefore, we first deleted these points to fit the periodograms using this function.

We jointly fitted all periodograms using the aliased version of the single power law function, Eq. 3.19 (Kirchner, 2005). Table 4.2 lists the best fitting power law index, α, for each of the 20 brightest variable likely extragalactic sources. It also lists several other properties of them.

Most periodograms of the variable, bright, likely extragalactic sources fitted with a single aliased power law, were found with a best-fitting index of $\alpha \approx 1$. This agrees with previous findings by other instruments, which described AGN periodograms to mostly follow a pink noise power law within the frequency range probed by *eROSITA* (Papadakis et al., 2002; Papadakis, 2004; González-Martín and Vaughan, 2012).

For instance, eRASSt J055033.4-663653, whose light curve is shown in Fig. 4.8, has periodograms (Fig. 4.9) that are best fitted with $\alpha = 1.093 \pm 0.050$. This fit is a good approximation of the low-frequency interval of the periodograms of this source. However,

Variable source	z	**Total source counts**	**NEV**$_b$	α
eRASSt J061148.4-662435		45328	$2.66^{+0.96}_{-0.56} \times 10^{-3}$	0.984 ± 0.096
eRASSt J054641.8-641524		19343	$4.3^{+2.2}_{-1.1} \times 10^{-3}$	0.87 ± 0.1
eRASSt J055943.0-660909	0.547	13419	$1.31^{+0.62}_{-0.23} \times 10^{-2}$	0.991 ± 0.030
eRASSt J061504.1-661717	0.178	10490	$8.0^{+6.1}_{-3.3} \times 10^{-2}$	1.369 ± 0.079
eRASSt J055357.6-665003	0.311	10420	$1.06^{+0.53}_{-0.20} \times 10^{-2}$	0.951 ± 0.049
eRASSt J055033.4-663653		9898	$2.0^{+1.1}_{-0.5} \times 10^{-2}$	1.093 ± 0.050
eRASSt J061543.9-653153		7261	$4.1^{+2.8}_{-1.5} \times 10^{-2}$	1.10 ± 0.10
eRASSt J060420.2-670234	0.678	5401	$1.42^{+0.95}_{-0.28} \times 10^{-2}$	0.844 ± 0.064
eRASSt J055934.5-653833	0.297	5191	$1.21^{+0.81}_{-0.28} \times 10^{-2}$	0.785 ± 0.066
eRASSt J060534.6-640045		4836	$5.3^{+4.3}_{-1.8} \times 10^{-3}$	1.258 ± 0.095
eRASSt J054913.6-642931	0.317	4514	$9.5^{+7.5}_{-2.7} \times 10^{-3}$	0.69 ± 0.11
eRASSt J054334.4-642300		4484	$1.29^{+0.99}_{-0.40} \times 10^{-2}$	1.027 ± 0.083
eRASSt J060139.2-655756	0.603	3612	$1.6^{+1.3}_{-0.4} \times 10^{-2}$	0.821 ± 0.045
eRASSt J055801.4-665655	0.230	3305	$2.2^{+2.0}_{-0.5} \times 10^{-2}$	0.638 ± 0.061
eRASSt J055745.4-664453	0.224	3230	$2.9^{+2.5}_{-0.7} \times 10^{-2}$	0.565 ± 0.049
eRASSt J055333.8-665751	0.121	3221	$2.0^{+3.9}_{-0.5} \times 10^{-2}$	0.928 ± 0.058
eRASSt J054750.3-672803		2895	$9.5^{+8.5}_{-2.8} \times 10^{-3}$	0.67 ± 0.10
eRASSt J055419.9-663340	0.251	2829	$2.6^{+7.3}_{-0.6} \times 10^{-2}$	0.560 ± 0.048
eRASSt J061932.9-645955	0.255	2396	$1.07^{+1.2}_{-0.4} \times 10^{-2}$	0.90 ± 0.17
eRASSt J053423.0-652004		2047	$2.7^{+2.8}_{-1.0} \times 10^{-2}$	1.21 ± 0.14

Table 4.2: Properties of the brightest variable likely extragalactic sources with at least 2×10^3 background-subtracted source counts observed by *eROSITA* in eRASS1, 2, and 3 in the SEP field. This table lists the redshift, z, determined from AAT spectroscopy, and the total background-subtracted source counts in eRASS:3. The redshift is only listed if it could be determined unambiguously. The sources are sorted by the number of total background-subtracted source counts observed. The table also lists the geometric average NEV$_b$, for 20 bin segments in the light curve. The associated error includes both measurement and sampling errors. Finally, the table also lists the best fitting power law index, α, when fitting the periodograms of the *eROSITA* light curves with a single aliased power law model, Eq. 3.19.

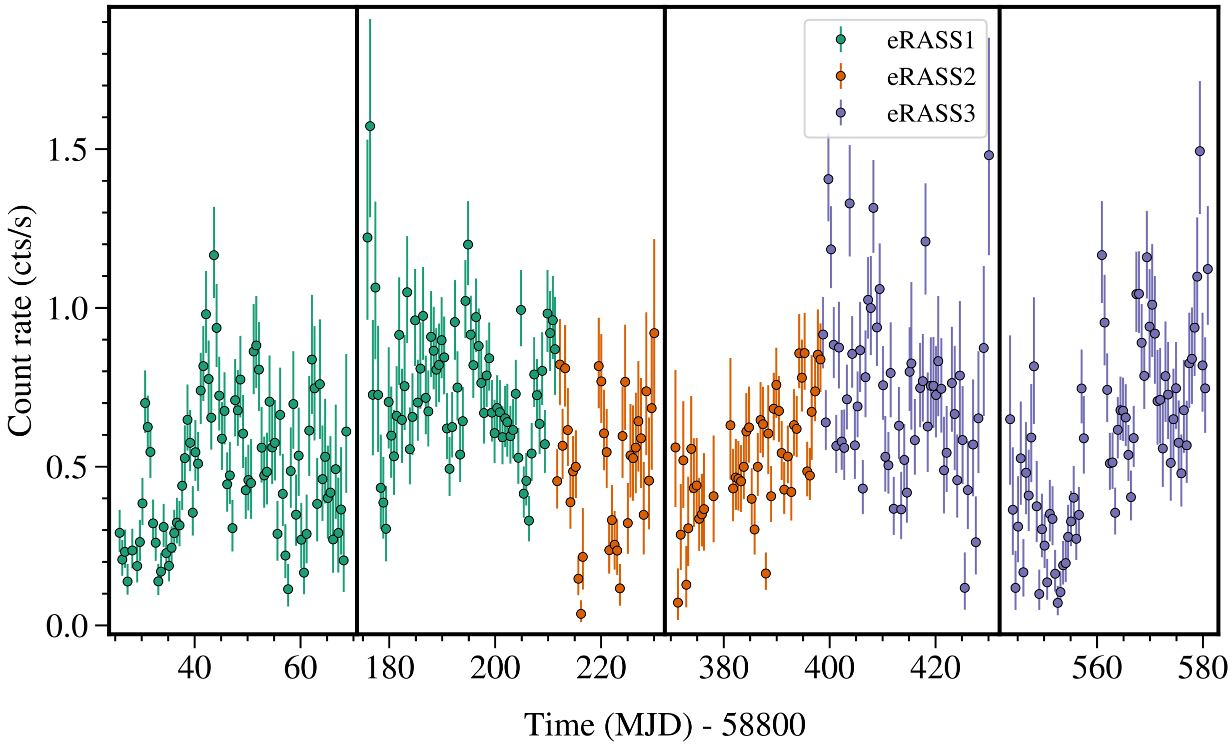

Figure 4.8: Light curve of the likely extragalactic variable source eRASSt J055033.4-663653, as observed by *eROSITA* in eRASS1, 2, and 3. The light curve has been rebinned by a factor of 3, for display clarity.

at high frequencies, the observed periodograms feature less power than expected from this aliased power law model. This could indicate a break in the PSD to $\alpha \approx 2$, which would match the high-frequency periodogram shapes observed in many AGNs (González-Martín and Vaughan, 2012). Some periodograms, like these ones, might be better fitted by an aliased broken power law.

However, not all periodograms conform to this standard shape. For example, the light curve of the likely extragalactic source eRASSt J061504.1-661717, shown in Fig. 4.10, features variability best described by a single power law with $\alpha = 1.369 \pm 0.079$ (Fig. 4.10). This is almost halfway between the two power law indices commonly observed in AGN periodograms, and is inconsistent with $\alpha = 1$. The source eRASSt J060534.6-640045 also has similar periodograms, which are best fitted by $\alpha = 1.258 \pm 0.095$. These best fits by steeper power laws could also be caused by typical AGN PSDs that feature a break within the observed frequency range. The average power law index for such PSDs lies between -1 and -2.

The periodograms of eRASSt J055419.9-663340 are best fitted by a power law with $\alpha = 0.560 \pm 0.048$. Comparably low power law indices were found for the periodograms of several other sources. In these cases, the periodograms are dominated by stochastic aliased high-frequency power, whose slope cannot be well determined. Therefore, the best

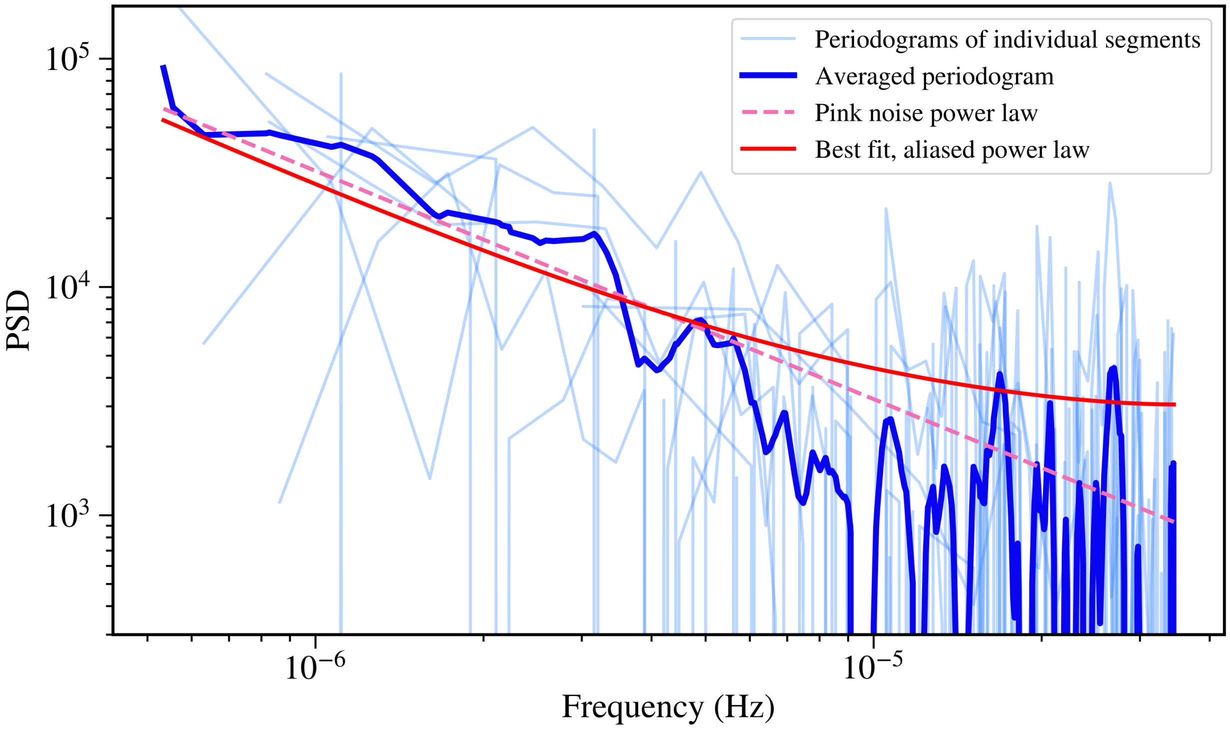

Figure 4.9: Periodograms of the likely extragalactic source eRASSt J055033.4-663653, whose light curve is shown in Fig. 4.8. The periodograms of individual segments of the light curve are shown in light blue. The averaged periodogram is displayed in dark blue. The red line denotes the best fit to all the periodograms of segments of the light curve, using the aliased single power law (Eq. 3.19), with $\alpha = 1.093 \pm 0.050$. The figure also features a dashed pink line, which depicts a pink noise trend.

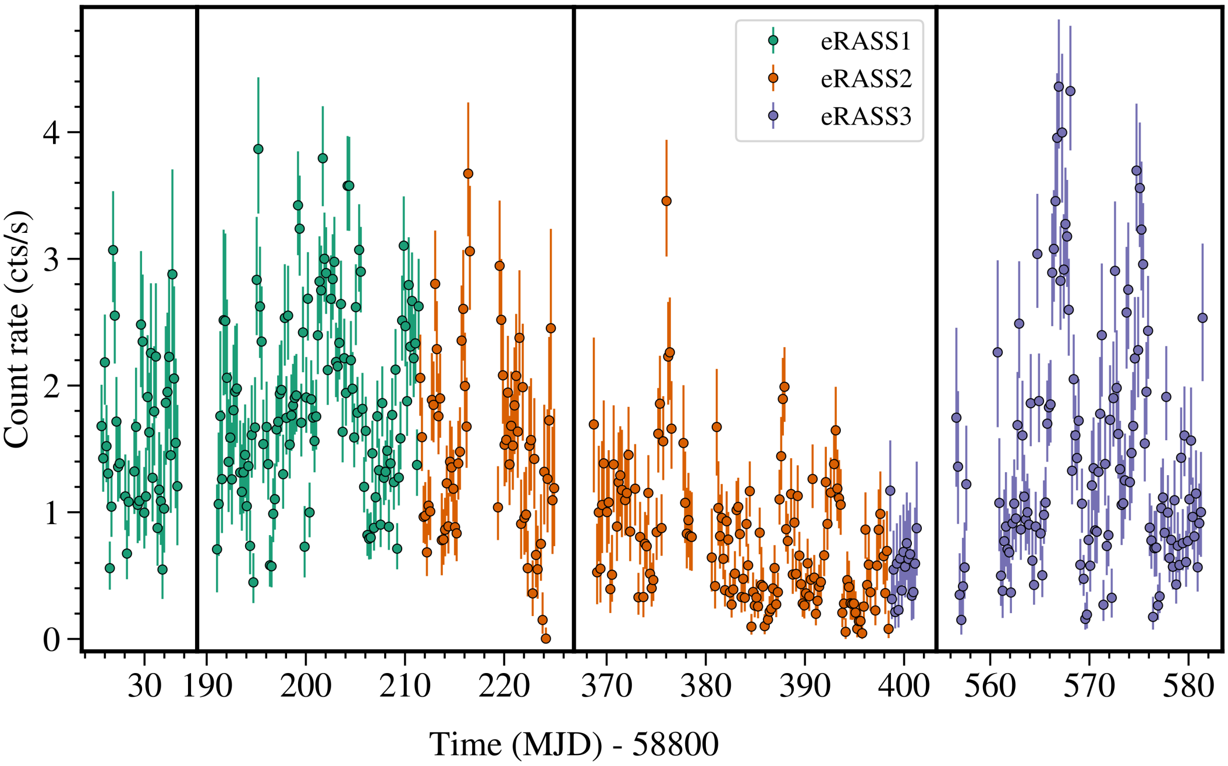

Figure 4.10: Light curve of the likely extragalactic source eRASSt J061504.1-661717, as observed by *eROSITA* in eRASS1, 2, and 3. The light curve has not been rebinned.

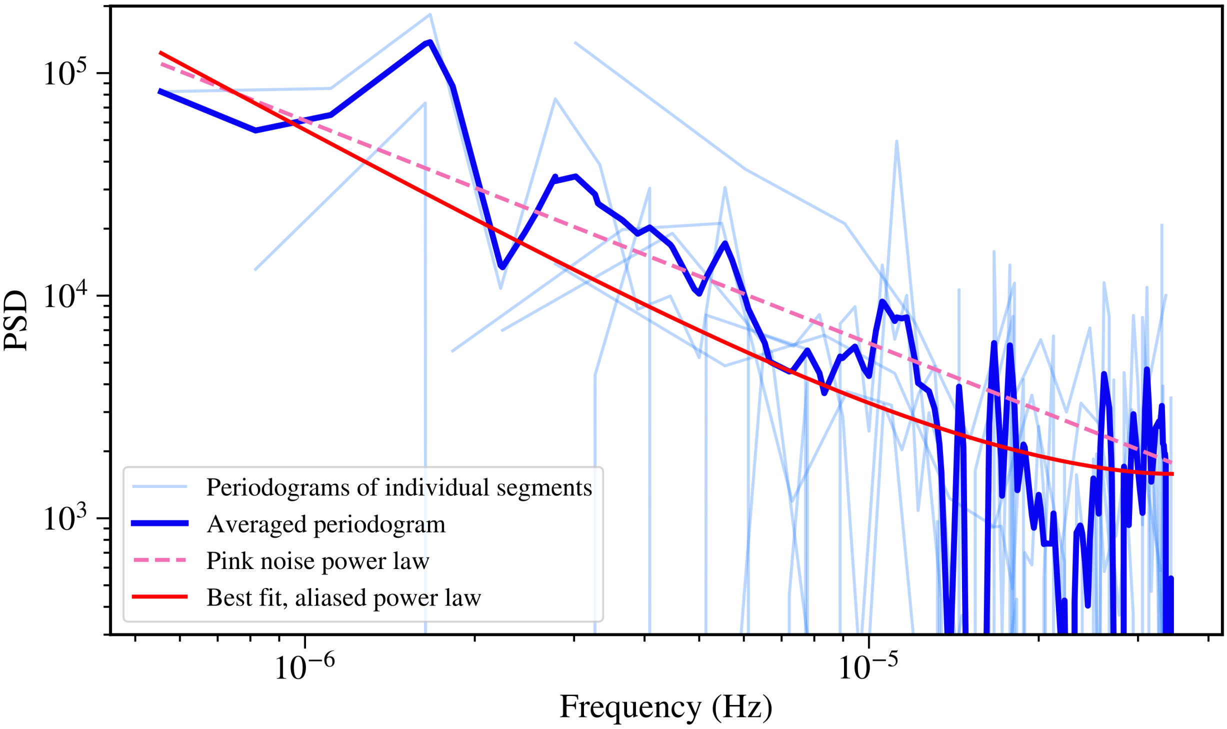

Figure 4.11: Periodograms of the likely extragalactic source eRASSt J061504.1-661717. The *eROSITA* light curve this is computed for is depicted in Fig. 4.10. The light blue lines indicate the periodograms of individual segments. The dark blue line depicts the averaged periodogram. The periodograms are best fitted by a single aliased power law model, with $\alpha = 1.369 \pm 0.079$, which is shown by the red line. The dashed pink line indicates a pink noise power law, to compare against.

fitting power law indices do not accurately represent the low frequency interval of the periodograms. The fraction of sources whose periodograms are not well fitted increases as the total number of source counts decreases. That is why decided to not investigate the periodograms of sources with fewer than 2×10^3 total source counts.

Figs. 4.9, and 4.11 indicate that a pink noise relationship is a good first-order approximation of the frequency dependence of the variability power of AGNs observed by *eROSITA*. In most cases, when fitting periodograms of bright, likely extragalactic sources with a single power law function, and no aliasing effect, the best fitting index is $\alpha \approx 1$. Therefore, the assumptions that we used to define NEV_b, and the sampling errors of estimating NIV_∞ are accurate for AGNs observed by *eROSITA*.

4.6 Variable X-ray sources of particular interest

This section discusses the variability properties of some of the most interesting variable X-ray sources in the SEP field. We primarily investigated variable likely extragalactic sources but also inspected likely Galactic sources for unusual variability.

4.6.1 Flip-flopping source eRASSt J061148.4-662435

The variable X-ray source eRASSt J061148.4-662435 is the brightest likely extragalactic variable source in the SEP field. It also features several intriguing variability properties. Its flux was observed to vary around one of two distinct levels. This is shown in Fig. 4.12, which depicts the *eROSITA* X-ray light curve from eRASS1, 2, and 3, split into four segments and rebinned by a factor of 3. Additionally, there is some indication of long-term periodicity in the light curve.

The most likely optical counterpart of this source was observed to have a GAIA parallax significance of $p/\sigma_p = 0.886$. We did not observe its optical spectrum, but found its W1 and J magnitudes to satisfy: $W1 - J + 1.7 = -1.19$. Furthermore, we found that it has: $W1 + 1.625 \log(F_{0.5-2\ \mathrm{keV}}) + 6.101 = 2.46$, and $z - W1 - 0.8(g - r) + 1.2 = 2.99$. These methods establish eRASSt J061148.4-662435 as a likely extragalactic source, with $D = 13$. This source is also matched to the AGN 2MASS J06114834-6624337, at a separation of $1.63''$. It has a redshift of 0.23 (Masetti et al., 2008).

In the first, third, and fourth segments of the light curve, the mean count rates are very similar, having values of 5.208 ± 0.069 cts/s, 5.126 ± 0.039 cts/s, and 5.332 ± 0.056 cts/s, respectively. Even though the count rate within these three segments varies throughout the observations, their mean count rates are remarkably consistent, differing by merely $4.0 \pm 1.3\%$ between the dimmest and brightest of them. However, the mean count rate in the second segment was only 3.084 ± 0.027 cts/s. This means that the mean count rate in segments 1, 3, and 4 is larger than that in segment 2, by $69.3 \pm 1.8\%$. This significant change in the count rate is not due to an accidental wrong matching of sources between eRASSs. This is because it is observed at both the high and low flux levels in both eRASS1 and eRASS2.

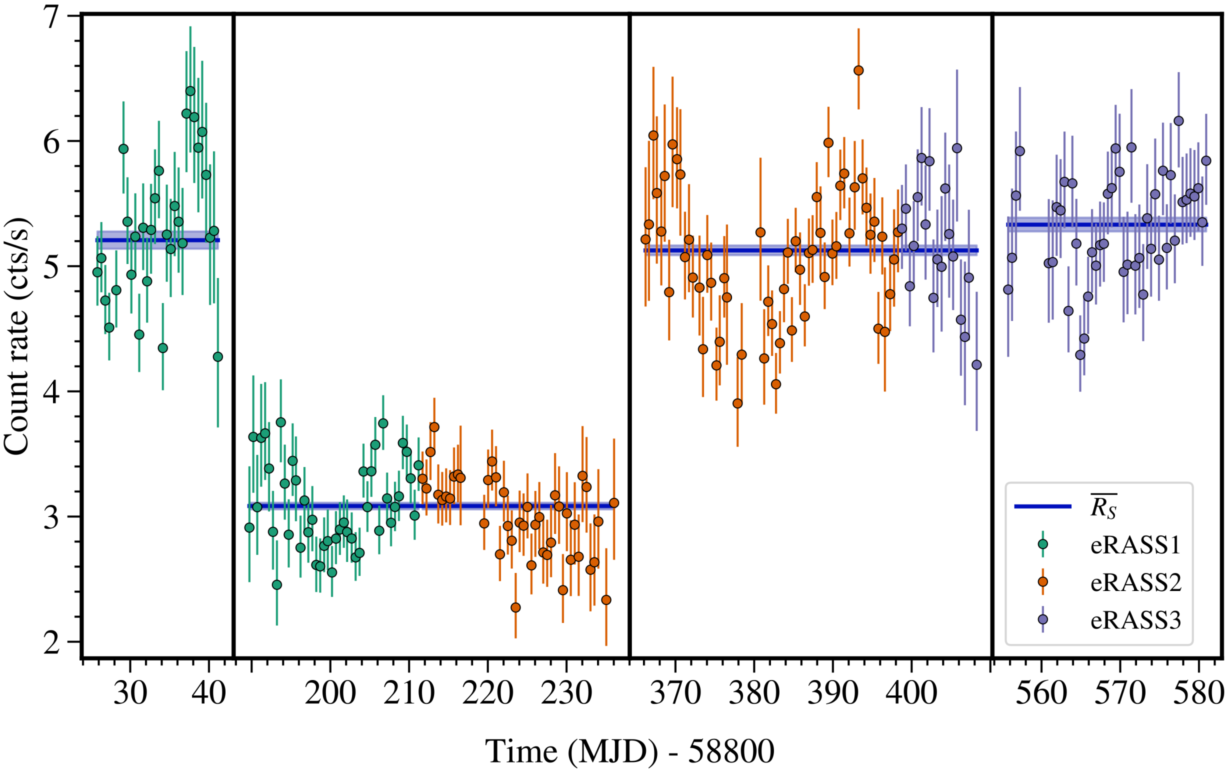

Figure 4.12: Light curve of eRASSt J061148.4-662435, observed by *eROSITA* in eRASS1, 2, and 3. The dark blue lines indicate the mean count rate in each of the four segments. The light curve was rebinned by a factor of 3 for display clarity. This means that each datapoint in this figure corresponds to data collected within 3 eroDays.

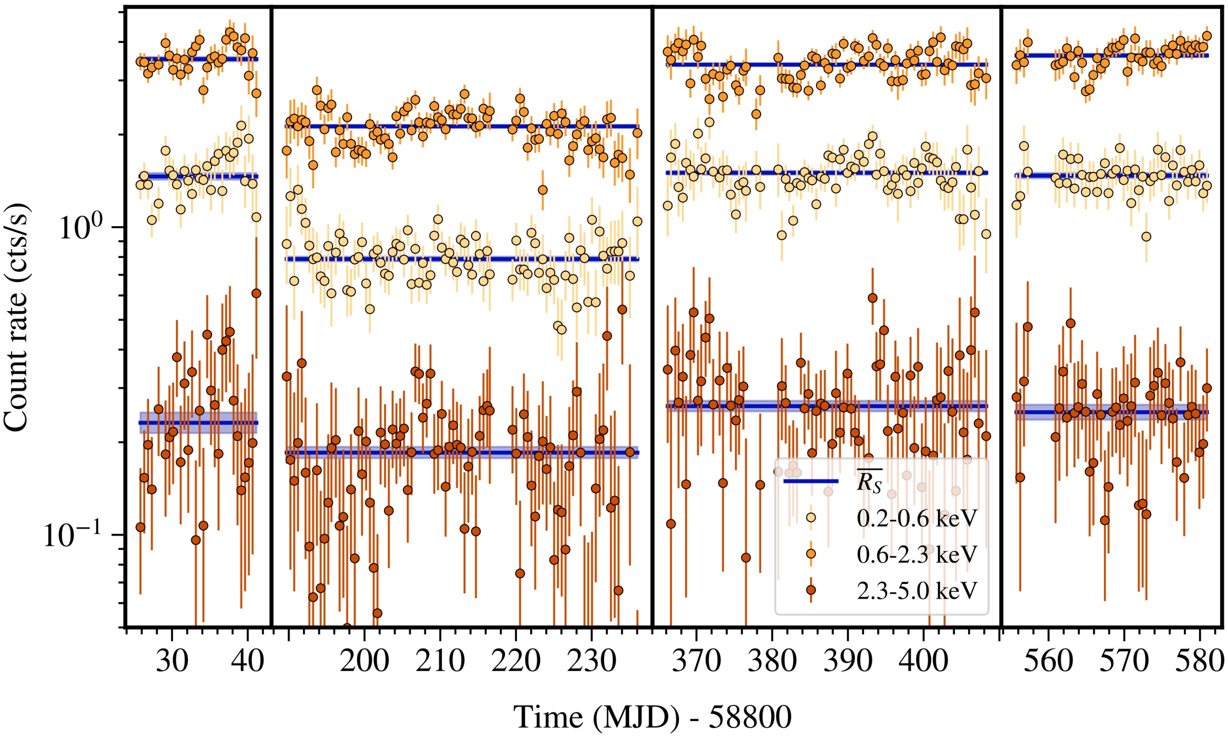

Figure 4.13: Light curve of eRASSt J061148.4-662435, in the three energy bands 0.2 – 0.6 keV, 0.6 – 2.3 keV, and 2.3 – 5.0 keV. This light curve was rebinned by a factor of 3 for display clarity.

However, the difference in the count rate between the dim and bright segments was not uniform across the X-ray spectrum. Fig. 4.13 depicts the light curve of eRASSt J061148.4-662435 in the three energy bands $0.2-0.6$ keV, $0.6-2.3$ keV, and $2.3-5.0$ keV. The $0.2-0.6$ keV energy band featured an $87.7 \pm 3.9\%$ increase in the count rate between segments 2, and the mean of segments 1, 3, and 4. In contrast, the $2.3-5.0$ keV energy band experienced an increase of only $34.0 \pm 7.3\%$. The intermediate energy band, $0.6-2.3$ keV is the brightest, and features an increase in the count rate between the dim and bright intervals of $64.7 \pm 2.1\%$. The X-ray spectrum of the dim state, in segment 2, is harder than in the bright states. This implies that at least part of the difference in the count rates could be due to a temporary occultation, in which a greater fraction of low, rather than high-energy X-ray photons were absorbed and scattered.

However, despite the existence of significant variability within each segment, it appears to occur around a constant flux level. There is no apparent long-term trend to higher or lower fluxes in any segment. We know that there must have been intervals of greater variability between segments 1 and 2, and between segments 2 and 3. In these intervals, the mean count rate dropped by $40.77 \pm 0.93\%$ and increased by $66.2 \pm 2.2\%$. For the four segments that we did observe, we calculated NEV$_b$ values of: $1.7^{+2.3}_{-1.7} \times 10^{-3}$, $6.0^{+3.3}_{-2.3} \times 10^{-3}$, $5.4^{+2.7}_{-2.0} \times 10^{-3}$, and $1.4^{+1.5}_{-1.4} \times 10^{-3}$, respectively (using the methodology described in Section

3.5). The errors quoted here combine both the measurement and the sampling errors. The NEV_b values of the individual segments are not directly comparable with each other, as they are computed for a different number of bins, so a different frequency interval. Even though there are significant differences in the mean count rates between various segments, the variability strength measured within the individual segments is comparatively small. Each segment features some variations in the count rate, but it is mainly constrained to small variations around the mean.

We sought to compare the degree of variability of the source in the bright and the dim states. To do so requires that the NIV be estimated for the same frequency interval. Therefore, we computed a geometric mean NEV_b of sets of 20 bins (covering the frequency range of $3.47 - 34.72 \times 10^{-6}$ Hz) within the bright and dim parts of the light curve. We calculated a geometric mean NEV_b of $2.31^{+0.99}_{-0.58} \times 10^{-3}$ and $3.5^{+2.3}_{-1.2} \times 10^{-3}$ for the bright and dim states, respectively. The errors shown here include the sampling errors for estimating NIV_∞ in adjoined segments, as discussed in Section 3.7.3. The NEV_b values are consistent within their errors. However, the sampling error of the $\overline{\mathrm{NEV}_b}$ measurement for the bright state is likely to have been overestimated, as not all segments used to calculate its value are adjacent.

We investigated the likelihood of *eROSITA* observing an AGN at this distance from the SEP with a light curve that exhibits similar properties to this one. In doing so, we can determine whether this behaviour is a unique phenomenon or whether it was likely to have been observed. We simulated light curves based on a single $P \propto \nu^{-1}$ power law PSD, to match the periodogram measured for this source (see Section 4.5). We do not know what the shape of the PSD is outside of the interval we observed. However, a reasonable first order estimate is that the power law observed in the periodogram continues towards lower frequencies. AGN PSDs do not typically get steeper at lower frequencies, but are expected to become less steep (González-Martín and Vaughan, 2012). However, the steeper the power law, the greater the fraction of the total variability contained at lower frequencies. Therefore, we chose this model, as it is realistic, consistent with previous studies of low-frequency AGN PSDs, and can achieve significant low-frequency variability. We used this model to simulate light curves ten times as long as the temporal difference between the first and last eroDay of observations of eRASSt J053942.0-653038.

The light curves we simulated assumed infinite accuracy, and no Poisson noise. These assumptions are reasonable for eRASSt J061148.4-662435, as it is a bright source, whose mean count rates and NEV_b values are only minimally affected by Poisson noise. However, these assumptions are not valid for most *eROSITA* sources. The conclusions drawn from these simulations only apply to bright sources, such as this one.

We randomly chose a starting time, then selected bins from the simulated light curves, corresponding to the times of observation of eRASSt J061148.4-662435. We also split the simulated light curves into four segments, and analysed each individually, just as we had done for the actual data. Finally, we determined the fluxes and NIV values for each of the four segments, and compared them with each other.

A light curve was deemed to be similar to eRASSt J061148.4-662435 if the mean flux in three of the four segments was similar, but the fourth one significantly differed from them. The second condition we set is that the variability within each segment is smaller than expected from the difference in the mean fluxes. This condition is necessary to avoid selecting light curves that vary more within, than between segments. Such sources do not feature variability around a constant flux level in each segment. Therefore, they are not comparable to the light curve of eRASSt J061148.4-662435, even if they have a similar relationship between the mean fluxes in the individual segments.

In this search, we did not care which of the four segments featured an anomalous mean flux, or whether that anomaly was towards higher or lower fluxes. That is because any source with a single anomalous segment would have been worth exploring further. We set the threshold for satisfying this condition at the values found for eRASSt J061148.4-662435. Light curves in which three segments had mean fluxes within 4.0% of each other, and the fourth one either had a 69.3% higher, or a 40.9% lower mean flux level, were selected for further analysis with the second condition.

The second condition ensures that the simulated light curves selected to be similar to that of eRASSt J061148.4-662435 have a moderate degree of variability within all segments. To do so, we required that the NIV values in the simulated segments are less than the NEV_b values found for eRASSt J061148.4-662435. We did not want to bias the search by requiring that each segment have a NIV less than or equal to the NEV_b value of the corresponding segment of eRASSt J061148.4-662435. Instead, we required that all four segments have a NIV of less than or equal to 6.0×10^{-3}. That is the largest NEV_b value measured in the observed light curve. Similarly, we also required that at least three segments have a NIV of less than, or equal to 5.5×10^{-3}, at least two segments have a NIV of less than or equal to 1.7×10^{-3}, and at least one segment has an NIV of less than or equal to 1.4×10^{-3}.

We randomly selected the range of fluxes on a linear scale when simulating these light curves. If they are simulated for a smaller range of fluxes, they are more likely to satisfy the requirement of the upper limits on the NIV of the individual segments. However, they are less likely to satisfy the condition on the difference between the mean fluxes of the four segments.

A total of 10^5 light curves were simulated in this way. Out of those, 0.200% satisfied the first requirements on the relation between the mean fluxes in the four segments. And 20.351% satisfied the second requirement of the upper bound on the NIV values in the four segments. However, not a single simulated light curve satisfied both conditions. We detected 453 unique variable sources in the SEP field, but only ≈ 10 of them were sufficiently bright to satisfy the assumptions of these simulations. Therefore, it is very unlikely that a typical AGN evolution would feature the properties observed for eRASSt J061148.4-662435 purely by chance by being observed at the right times. This indicates that this source most likely does not follow a standard AGN luminosity evolution. This applies as long as the simulations can accurately represent AGN variability. It is likely that this source does vary between two distinct flux states within the span of these observations.

Another interesting feature of the light curve of eRASSt J061148.4-662435 is that it appears to be periodic in segments 2 and 3 (Fig. 4.12). To investigate this, we computed the Lomb-Scargle periodograms of the individual segments. These are depicted in Fig.

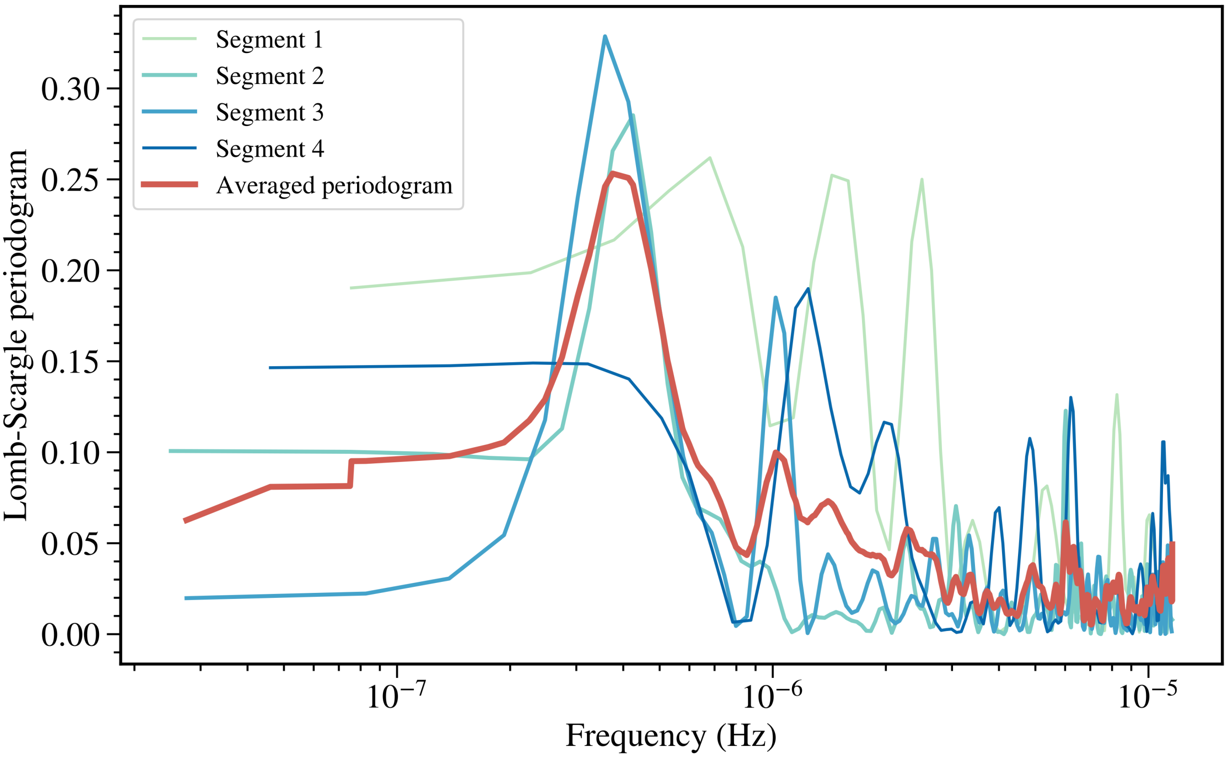

Figure 4.14: Lomb-Scargle periodogram of eRASSt J061148.4-662435, observed by *eROSITA* in eRASS1, 2, and 3. Each of the blue-green curves depict the Lomb-Scargle periodogram of a single segment of eRASS observations, as shown in Fig. 4.12 The mean Lomb-Scargle periodogram of all segments is also shown.

4.14. In segments 2 and 3, they contain peaks at frequencies of 4.25×10^{-7} Hz, and 3.58×10^{-7} Hz, respectively. These correspond to periods of 27.2 days, and 32.4 days, respectively. Segments 1 and 4 were too short to detect periodicities at these frequencies. The mean of the Lomb-Scargle periodograms in all four segments has a peak at a frequency of 3.75×10^{-7} Hz, or a period of 30.9 days.

This potential periodicity is intriguing, as it was detected independently in two segments separated by more than 4 months at a reasonably consistent frequency. The Lomb-Scargle periodogram of segment 3 features a broader peak with a slightly lower peak frequency than the one for segment 2. Nevertheless, the high-frequency tail of the peak is almost identical for both segments.

We do not have sufficient data to confirm the existence of this periodicity. We only observed it twice, each for about 1.5 cycles. However, it is a promising first sign that could be followed up in future.

The light curve of eRASSt J061148.4-662435 (Fig. 4.12) bears some resemblance to the flip-flop phenomenon (see Chapter 5). The light curve can be distinguished into bright and dim states. The count rate observed in the brighter state in segments 1, 3, and 4, remained consistent, and exceeded the count rate in the dim state in segment 2 by

$69.3 \pm 1.8\%$, matching the brightness change observed in Swift J1658.2-4242. Additionally, there is an indication of a periodicity in the light curve, which might be the manifestation of a low-frequency QPO.

BHTs can, in many ways, be considered to be low-mass analogues of AGNs (McHardy et al., 2006). Many phenomena observed in BHTs have subsequently been observed in AGNs as well, at correspondingly longer timescales. The periodograms of BHTs and those of AGNs can often be described by the same broken power laws (Belloni and Hasinger, 1990). Quasi-periodic oscillations have been observed in many BHTs, and are now being detected in an increasing number of AGNs as well (e.g. Gierliński et al., 2008; Agarwal et al., 2021). Flip-flops have, however, only been detected in BHTs, so far. This is because it has so far been impractical to observe AGNs continuously for a sufficiently long time, to be able to detect scaled-up flip-flop timescales. Flip-flops have been observed to feature states that last from between several tens of seconds (Miyamoto et al., 1991) to several kiloseconds (Casella et al., 2004). As we expect these effects to scale linearly with BH mass, the equivalent timescales for AGNs should be in the range of $\approx 10^6 - 10^{12}$ s. The dim state we observed in the light curve of eRASSt J061148.4-662435 lasted for at least 4.3×10^6 s, but could have lasted up to 2.8×10^7 s. That is within the range of timescales we would expect for AGN flip-flops.

However, the transitions between the dim and bright states of eRASSt J061148.4-662435 were not observed. So we do not know how the mean flux level changed. Additionally, the bright flip-flop states observed in BHTs have a harder spectrum than the dim states. In contrast, eRASSt J061148.4-662435 was softer in the bright state. Therefore, although there are some similarities between the light curve of eRASSt J061148.4-662435 and those featuring flip-flops in BHTs, the two phenomena cannot be directly associated.

We observed another likely extragalactic source, eRASSt J054641.8-641524, whose light curve had some similar features to those of eRASSt J061148.4-662435. Its light curve is shown in Fig. 4.15. For this source, we observed a 113% increase in count rate over a course of ≈ 10 days in eRASS1, from 1.982 ± 0.058 cts/s, to 4.223 ± 0.095 cts/s. However, the mean count rate in eRASS2 and 3 was similar to that at the start of eRASS1. The mean count rate in eRASS2 and 3 was 2.347 ± 0.029 cts/s, and 2.285 ± 0.031 cts/s, respectively. We only observed a very small degree of variability in eRASS2 and 3. Nevertheless, the fractional flux increase observed in eRASS1 for eRASSt J054641.8-641524 is even larger than the change observed for eRASSt J061148.4-662435.

The continuous rise in the count rate of eRASSt J054641.8-641524 in eRASS1 could be similar to the transition that occurred in gaps between the segments of observations of eRASSt J061148.4-66243. Both eRASSt J061148.4-66243 and eRASSt J054641.8-641524 may feature a similar flip-flop-like evolution. The differences in their light curves could be merely due to the differing observing intervals, catching the long-term variability at different phases.

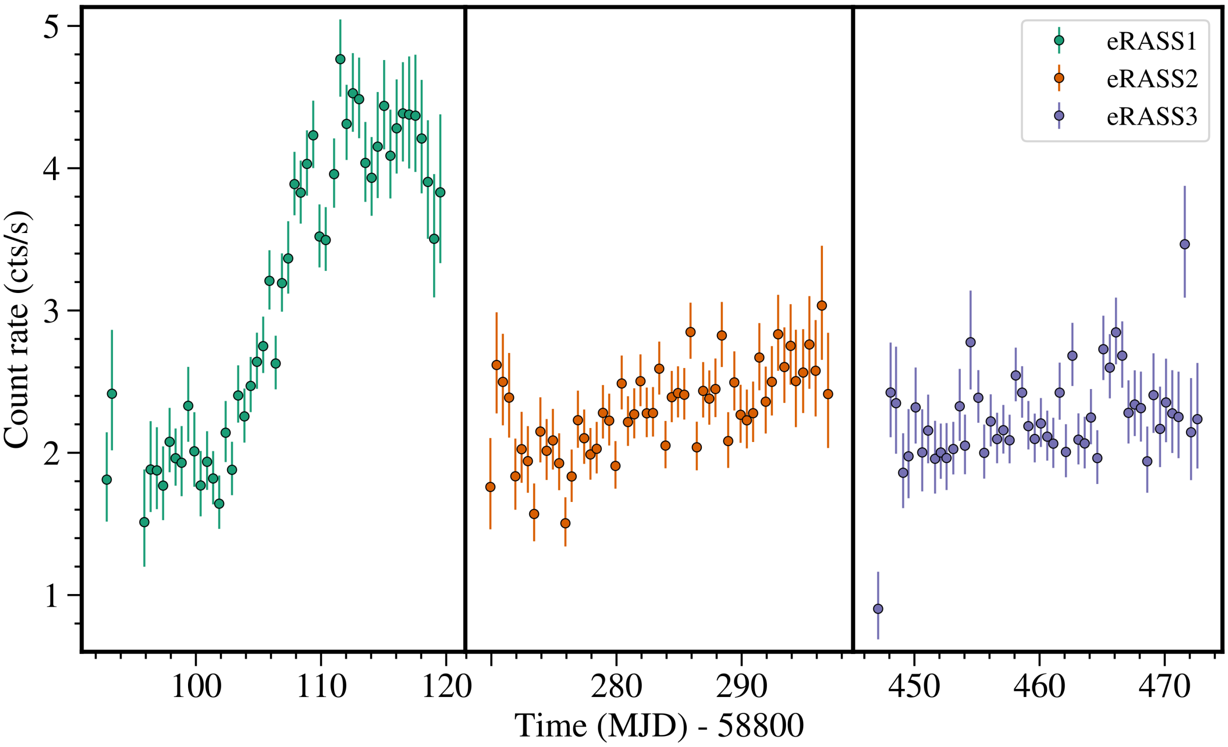

Figure 4.15: Light curve of the likely extragalactic source eRASSt J054641.8-641524, as observed by *eROSITA* in eRASS1, 2, and 3. The light curve has been rebinned by a factor of 3 for display clarity.

4.6.2 Extragalactic transients

The *eROSITA* variable source eRASSt J053942.0-653038 suddenly became significantly brighter on MJD 58869, as is shown in Fig. 4.16. Throughout most of eRASS1, it was detected just barely above the background level, with a mean count rate of $0.0056^{+0.0024}_{-0.0019}$ cts/s. However, in the days after MJD 58869, it was observed with a mean count rate of $0.294^{+0.033}_{-0.030}$ cts/s in the $0.2-5.0$ keV energy band, until it was no longer observed in eRASS1. In this final period of eRASS1 observations, the count rate varied a lot from eroDay to eroDay, but appeared to show a downward trend after the peak. At the end of the eRASS1 observations, it was again observed at a very low count rate. However, that could be a fluke detection, resulting from the low fractional exposures in those bins.

About 150 days later, eRASSt J053942.0-653038 was observed during eRASS2, this time with a mean count rate of $0.1168^{+0.0069}_{-0.0067}$ cts/s in the $0.2-5.0$ keV energy band. The mean count rate also appeared to decrease during the next 30 days, in which it was observed every 4 hours by *eROSITA*. Finally, it was observed again in eRASS3, this time with an even lower mean count rate of merely $0.0680^{+0.0071}_{-0.0064}$ cts/s in the $0.2-5.0$ keV energy band, which is still larger than the source count rate observed for it before the start of this eruption.

We determined its most likely optical counterpart to have a GAIA parallax significance of $p/\sigma_p = 0.884$. It was also observed to have: $W1+1.625\log(F_{0.5-2\ \mathrm{keV}})+6.101 = 5.83$, and $z-W1-0.8(g-r)+1.2 = 0.843$. This source was not matched to the VHS DR5 catalogue, so we could not use the $W1-J$ Galactic-extragalactic distinction. We also observed the spectrum of the most likely optical counterpart of this source with 2dF/AAOmega. However, by the time we observed it, on January 5, 2022, it had become too faint at optical wavelengths to be identified above the LMC continuum. Therefore, we were unable to determine its redshift. Combining these results, we classified it as likely extragalactic with a distinction parameter of $D = 10$. However, the field in which it is located is very crowded, so associating the X-ray detection with a particular optical source is challenging.

The light curve bears some resemblances with the evolution of Tidal Disruption Events (TDEs Rees, 1988; Gomez et al., 2020). It features a sudden large flare that significantly exceeds the pre-flare flux, followed by a gradual decline over hundreds of days. Phinney (1989) determined theoretically, that a TDE should evolve with $R_S \propto T^{-5/3}$, where T is the time since the start of the TDE. However, Lodato et al. (2009) found that the initial drop in luminosity is shallower than $R_S \propto T^{-5/3}$. Lodato and Rossi (2011) described that that relationship is only approximately accurate for about one year after the start of the TDE, in the X-ray band.

We investigated whether the decay in the count rate of eRASSt J053942.0-653038 after the start of the flare could be described by $R_S \propto T^{-5/3}$. The observations immediately after the start of the eruption in eRASS1 were too close to the start to be useful for determining the long-term decay rate. We combined all the observations of segments 2 and 3 (shown in Fig. 4.16), and fitted the decaying count rates with the function: $R_S = A(t-58869)^{-\gamma}$, where t denotes the time of each measurement, expressed in MJD. These observations were best fitted with $\gamma = 0.91 \pm 0.27$ for the full energy band, which is significantly less than

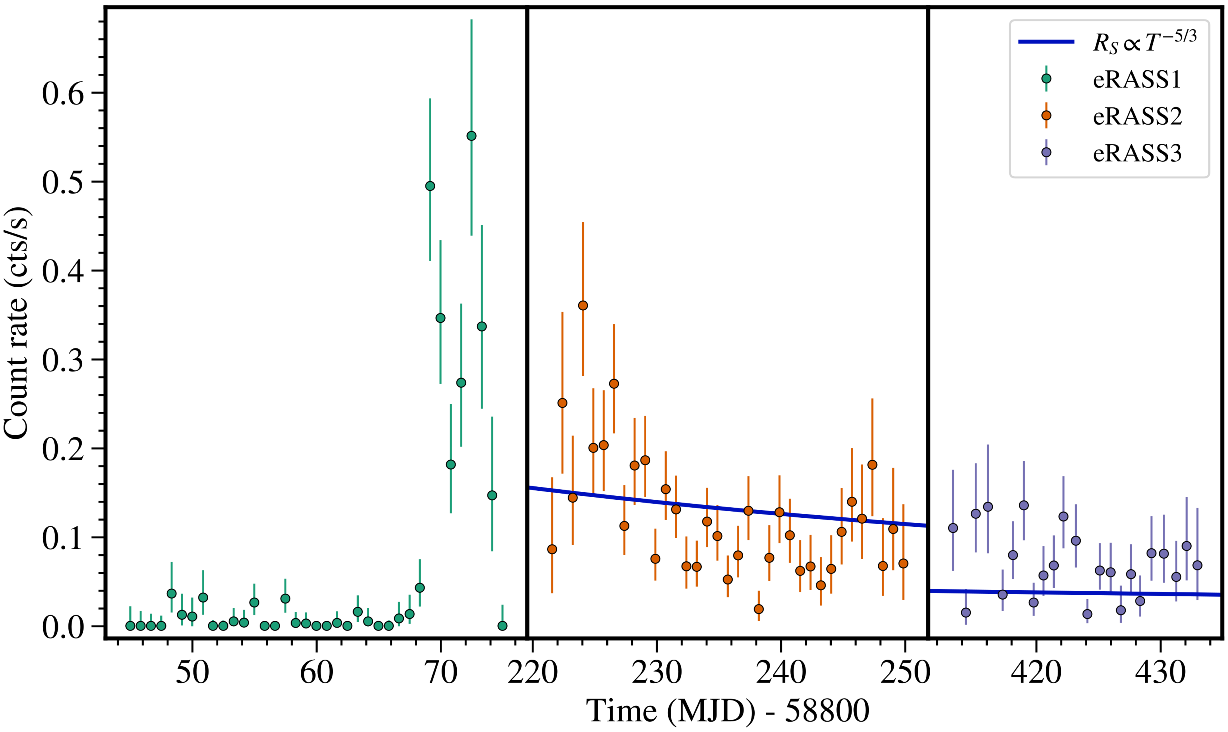

Figure 4.16: Light curve of eRASSt J053942.0-653038, observed by *eROSITA* in eRASS1, 2, and 3. We also indicate the typical TDE luminosity decay $R_S \propto t^{-5/3}$, fitted to the eRASS2 and 3 data. The measurements of the individual eroDays have been rebinned by a factor of 5.

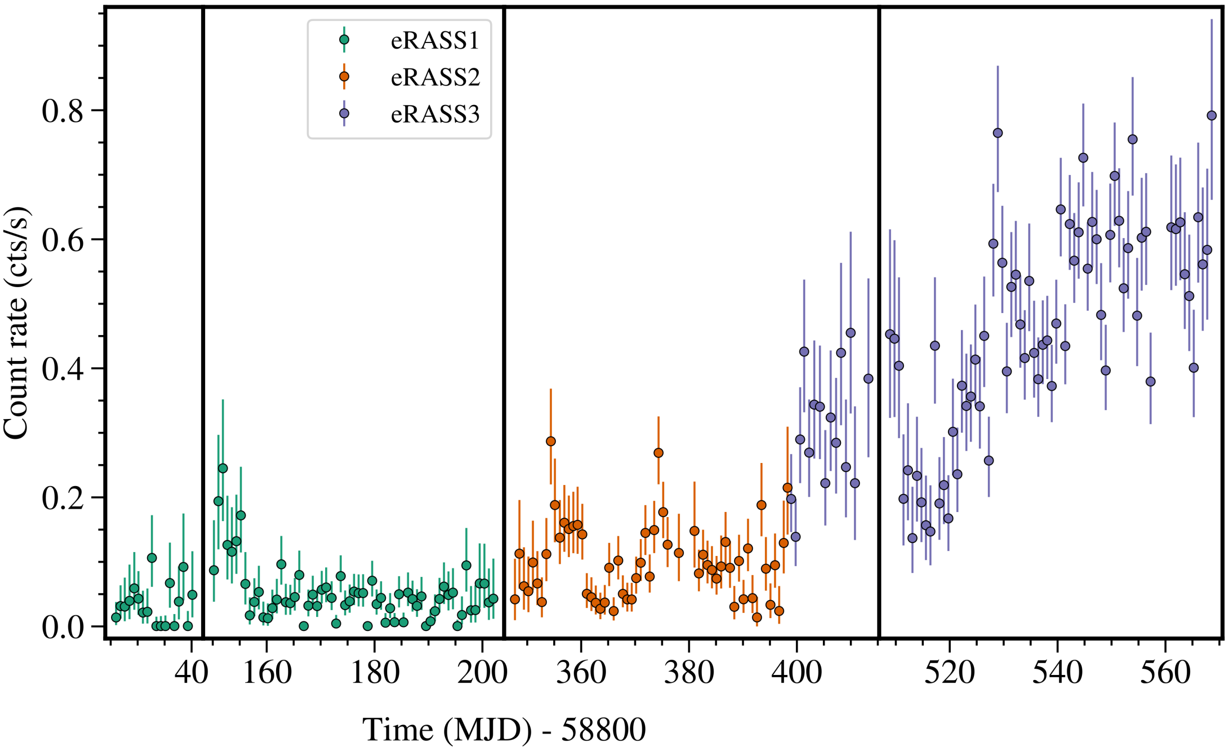

Figure 4.17: The eRASS1, 2, and 3 light curve of eRASSt J055333.8-665751. The light curved was rebinned by a factor of 5, for display clarity.

5/3. Fig. 4.16 also shows the best fit using a power law with $\gamma = 5/3$. This relation is too steep, and underestimates the mean count rate observed in segment 3.

Alternatively, it could be a different type of transient extragalactic source. Without a redshift for it, we are unable to estimate the luminosity at the peak of the eruption. We also lacked follow-up observations of this source during the interval in which it was bright, but noticeably declining. Therefore, the cause of the eruption of eRASSt J053942.0-653038 could not be determined.

We also observed several other likely extragalactic sources, which feature a significant increment or decrement in their observed count rates throughout the observations of them in eRASS1 to 3. Some of these could be AGN ignition and shut-down sources. For example, Fig. 4.17 depicts one such source. The count rate of eRASSt J055333.8-665751 increased by a factor of 12.8. In eRASS1, it had a mean count rate of 0.0468 ± 0.0031 cts/s. In contrast, by the end of eRASS3, after MJD 59345, it had a count rate of 0.599 ± 0.17 cts/s.

We observed the optical spectrum of the most likely counterpart of this source with the AAT. It was well fitted by a quasar model spectrum, at a redshift of $z = 0.121$.

Before the massive increase in the source count rate, we observed this source to feature three flares, each lasting for a few days. These could be episodes of failed ignition of the source.

In other instances, variable sources were not detected in some eRASSs, indicating a possible large change in the source flux to a level below the detection limit. Many more sources in the SEP field might feature such an evolution, but did not satisfy the variability thresholds within any of the three eRASS, and were subsequently not labelled as variable. To detect all of these sources requires a dedicated independent variability search spanning multiple eRASSs. However, the lack of a counterpart of a source in another eRASSs could also indicate a wrong matching of sources, or an issue with the source detection mechanism.

4.6.3 Variable stars

Of the 453 variable sources we identified, 224 were classified as being likely Galactic. Many of these sources have light curves that feature infrequent flares that significantly exceed the non-flaring level.

For instance, the likely Galactic source eRASSt J055930.9-663008 is located merely $4.55'$ away from the SEP. Except for an 8.5 day interval in eRASS3, in which survey mode observations were halted, this source was observed every 4 hours for 556 days. The light curve of eRASSt J055930.9-663008 is shown in Fig. 4.18. It had a steady mean count rate of $0.0404^{+0.0014}_{-0.0007}$ cts/s throughout all 3 eRASSs. However, it was also observed to feature two flares, significantly above its usual, constant flux level. The first, on MJD 59009.38, consisted of a sudden increase to a count rate of 0.6 ± 0.2 cts/s. The second flare occurred on MJD 59236.44, when it became even as bright as 1.7 ± 0.3 cts/s. In both cases, the source dimmed quickly, and returned to its usual flux by the next eroDay. In other sources featuring such flares, these flares can last for several eroDays before the source returns to its pre-flare flux level.

As the source features a low mean source count rate and two prominent, but short flares, it has a large degree of variability. For the first segment of consistent observations, from the start of eRASS1 until the start of the 8.5 day break in observations during eRASS3, we calculated $\mathrm{NEV}_b = 0.26^{+0.19}_{-0.11}$. As the second segment, after the end of the break, did not feature any comparable flare, we measured a significantly lower degree of variability for it, of only $\mathrm{NEV}_b = 0.011^{+0.056}_{-0.011}$.

The star η^2 Dor was observed to have low-frequency, large amplitude, periodic cycles in its light curve, shown in Fig. 4.19. The eRASS1, 2, and 3 light curves all look similar. In all three eRASSs, the source became a lot brighter before dimming down again, with an almost triangular pattern. η^2 Dor had previously been known to have a long duration periodic light curve, based on periods of $21.4-23.3$ days, and 150 days (Tabur et al., 2009).

The similarity of the eRASS1, 2, and 3 light curves indicates the existence of a periodic phenomenon that guides this variability. By cross-correlating the light curves of the three segments with each other, we find that the time between these long-duration flares in eRASS1 and eRASS2 is 178 days. We also find that the eRASS3 light curve can be described as most similar to the eRASS2 light curve, shifted by 188 days. Although these values differ slightly, they indicate that this star features long-term X-ray variability with an approximate period of 183 days, or a fraction of that. This result is inconsistent with the previously detected period of 150 days. The duration of observations of η^2 Dor in each eRASS spans at most 27.8 days, and there is no sign of the cycle occurring more than once

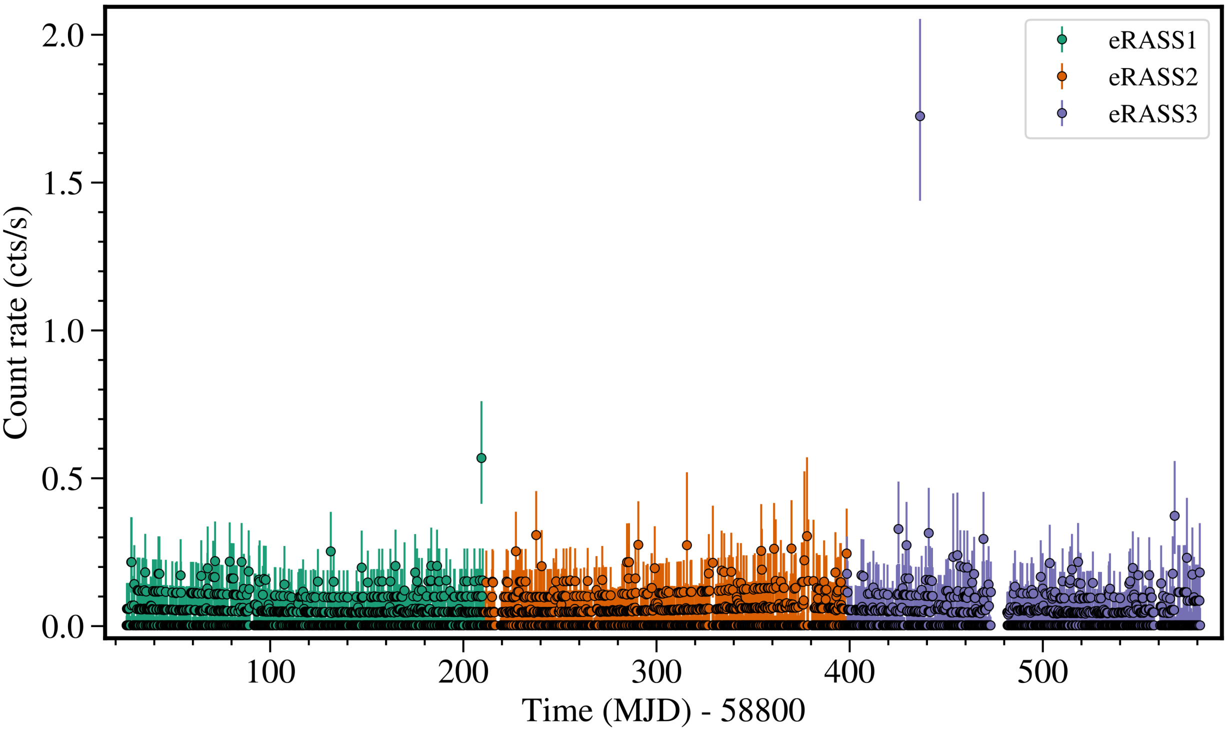

Figure 4.18: Light curve of the likely Galactic source eRASSt J055930.9-663008, as observed by *eROSITA* in eRASS1, 2, and 3. The light curve has not been rebinned.

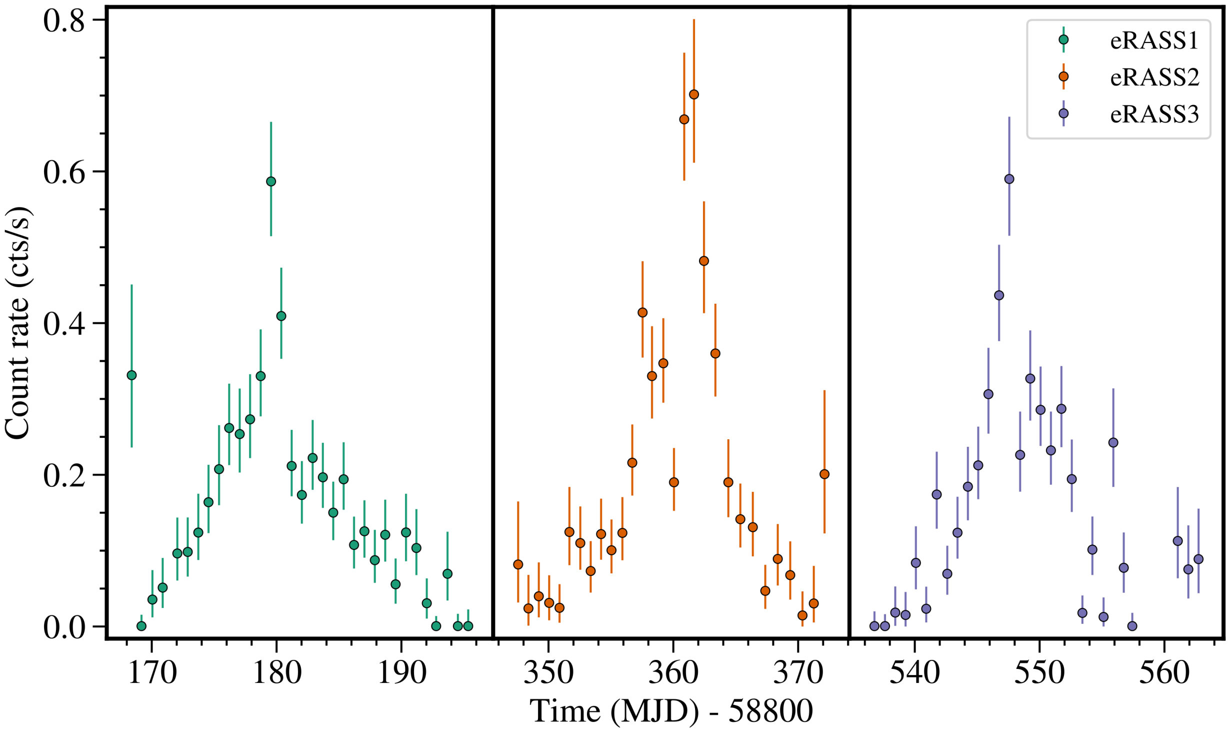

Figure 4.19: *eROSITA* light curve of η^2 Dor in eRASS1, 2, and 3. We rebinned the light curved by a factor of 5, for display clarity.

in that interval. However, we cannot definitively rule out a period of $183/8 = 22.9$ days, which agrees with the findings of Tabur et al. (2009). We could have observed the end and start of previous cycles at the start and end of the observations in each eRASS, without being able to identify them as such.

Chapter 5

Extreme flip-flops of Swift J1658.2-4242 guided by an internal clock

Before the launch of *eROSITA*, we investigated the properties of the BHT Swift J1658.2-4242, and its unusual features, which will be discussed in this chapter. Of greatest interest were the flip-flops detected in several observations during its outburst. As it is a bright source, observed at a constant fractional exposure, there was no need to use updated variability methods discussed in Chapter 3.

5.1 Flip-flops

Flip-flops, first noted by Miyamoto et al. (1991) (hereafter M91), are rare phenomena which appear as top-hat-like flux variations in the light curve of some black hole transients in outburst. Flip-flops are distinguished from absorption dips by their longer duration, their top-hat-like shape, and the positive correlation between flux and hardness during transitions (in contrast, dips have softer spectra at higher fluxes). Flip-flops are also seen to accompany changes in the periodogram, and are considered to be associated with some state transitions. Most state transitions, however, do not involve a flip-flop. Due to the marked and abrupt change in flux which defines the flip-flop, we refer to the higher and lower flux levels as the ‘bright state’ and ‘dim state’ respectively and analyse them separately.

Out of a population of ≈ 60 galactic BHTs (Corral-Santana et al., 2016), only seven systems exhibited properties fitting the description of a flip-flop: GX 339-4 (M91, Nespoli et al. 2003, hereafter N03), GS 1124-683 (Takizawa et al. 1997, hereafter T97), XTE J1550-564 (Homan et al. 2001, hereafter H01, Sriram et al. 2016), XTE J1859+226 (Casella et al. 2004, hereafter C04, Sriram et al. 2013), H1743-322 (Homan et al. 2005, hereafter H05), XTE J1817-330 (Sriram et al. 2012, hereafter S12), and MAXI J1659-152 (Kalamkar et al. 2011, hereafter K11, Kuulkers et al. 2013). In addition, there are other systems in which similar properties have been observed, such as in GRS 1915+105 (Soleri et al., 2008), or MAXI J1535-571 (Huang et al., 2018), but for which an identification as a flip-flop is not certain.

Flip-flops exhibit a large variety of different properties, not only between different systems or different outbursts of the same system but also within a single outburst of a BHT. Some flip-flops are seen at, or very close to, the peak of the outburst (M91, T97, H01, N03, and C04), while others are observed somewhat later, when the flux was dropping back down again (C04, H05, K11, and S12).

The observed flux change during a flip-flop also differs from system to system. Bright states can have an X-ray flux of between 3% and 33% greater than the flux of neighbouring dim states. Transition times have also been observed to span a large range, from fractions of a second, up to more than 1 ks. A large spread of values is also seen in the duration of each state. The time between adjacent flip-flop transitions can be anywhere between a few tens of seconds and several ks. There seems to be a relation between those two parameters, as short-lived flip-flop states have fast transitions (M91, T97, and S12), and long-lived states have slow transitions (N03, H01, C04, and H05).

All flip-flops observed so far have involved a type B QPO in either the bright or the dim state. Therefore, flip-flops seem to require transitions between the SIMS and either the HIMS or the AS.

C04 noticed a hierarchy of QPO states in the flip-flops of XTE J1859+226, with type As occurring at the highest luminosities, type Cs at the lowest luminosities, and type Bs in between. The limiting fluxes separating the different QPO types were found to decrease exponentially with time. All observed flip-flops between QPO types B and C agree with this hierarchy (T97, H01, C04, and K11). However, flip-flops between QPO types A and B disagree with it (M91, N03, and S12) more often than they agree with it (C04, and H05). Despite the fascinating phenomenology of flip-flops, a clear physical interpretation is still lacking.

5.2 Swift J1658.2-4242

5.2.1 Summary of previous observations

At 05:37:23 UTC on February 16, 2018 (MJD 58165), *Swift*/BAT was triggered by an X-ray burst from a point source in the galactic plane, (Barthelmy et al., 2018). This X-ray transient was subsequently called Swift J1658.2-4242. This first detection was quickly followed up with an observation by *Swift*/XRT, by D'Avanzo et al. (2018), and Lien et al. (2018), as well as one observation with the *Mobile Astronomical System of TElescope Robots*, at the Observatorio Astronomico Felix Aguilar (*MASTER-OAFA*; Lipunov et al., 2010), by Lipunov et al. (2018). Grebenev et al. (2018) performed subsequent examinations of data gathered by *Swift*/BAT (Lien et al., 2018), and by *INTEGRAL*, which revealed that the source had started flaring on February 13, 2018 (MJD 58162). Lien et al. (2018) determined a more accurate location of it, at RA = 16h 58m 12.64s, Dec = $-42^\circ\ 41'\ 54.4''$, and found its spectrum to be well described by a highly absorbed power law. Radio observations with *ATCA* strongly supported the notion that Swift J1658.2-4242 is a black hole XRB rather than a neutron star XRB (Russell et al., 2018).

Xu et al. (2018) described the first observation of the source by *NuSTAR*. They fitted the spectrum using a highly absorbed power law with a relativistic reflection component and a Gaussian component to fit an absorption line at 7.03 keV. They found the properties of the source to be consistent with known properties of black hole binaries and therefore described it as a black hole candidate. They also noted a type C QPO, whose centroid frequency increased continuously along with the flux, from 0.14 Hz to 0.21 Hz throughout this observation.

Jin et al. (2019) analysed the spectra detected by *Chandra*, and by *XMM-Newton*. They detected a strong X-ray dust scattering halo (DSH), which significantly affected the observed source spectrum and could be fitted with three dust layers at different distances from the BHT. Multiple methods were used to determine the distance to Swift J1658.2-4242, which was estimated to be $\approx$ 10 kpc.

Xu et al. (2019) examined the first flip-flop transition at the brightest part of the outburst, using data from both *NuSTAR* and *XMM-Newton*. They described a 45% flux decrease in $\approx$ 40 s. The lower flux level contained a type C QPO, but no QPO was detected at the higher flux level. Spectral fitting revealed only small differences in the high and low flux spectra. The strong relativistic reflection effect detected in the LHS by Xu et al. (2018) was no longer present in these observations.

Detailed timing analysis of the HIMS of Swift J1658.2-4242, using observations from *Insight*-HXMT, *NICER*, and *Astrosat* was performed by Xiao et al. (2019). They traced the increase in the QPO frequency alongside a decrease in the rms, determined the independence of QPO frequency and photon energy, and found the phase lag to be close to 0. Their results are consistent with a high inclination for Swift J1658.2-4242.

Jithesh et al. (2019) made a timing and spectral analysis of the three *Astrosat* observations of Swift J1658.2-4242, examining the properties of the detected QPO and its harmonics, their energy dependence, and time lag. The first observation featured a type C QPO which increased in frequency from 1.56 Hz to 1.74 Hz. The second observation featured additional flux variations reminiscent of the one observed by Xu et al. (2019). A 6.6 Hz type C QPO was only observed at low fluxes. By the final observation, the QPO frequency had dropped to 4.0 Hz. The QPO fractional rms was found to increase with photon energy. Positive time lags suggested that the QPO originated in the corona and propagated outwards.

Beri and Altamirano (2018) studied *Swift*/XRT observations of Swift J1658.2-4242 from August 10, and September 28, 2018. They reported that the count rates and spectral shape suggested that Swift J1658.2-4242 had returned to quiescence by the time of these observations.

5.2.2 Observations

Following the initial detection, Swift J1658.2-4242 was observed by *NuSTAR*, *XMM-Newton*, *Astrosat*, *Chandra*, *Swift*, Insight-*HXMT*, *INTEGRAL*, and *NICER* in the X-ray band. *ATCA* also provided coverage in the radio band. We utilised measurements from all these instruments to obtain the best possible understanding of the phenomena occurring during this outburst.

XMM-Newton, *NuSTAR*, *Astrosat*, and *Chandra* provided detailed, high-quality information of the properties of the source at a few distinct parts of the outburst, enabling comprehensive analysis of the light curve, periodogram, and energy spectrum. In addition, *Swift*/XRT and BAT, the *Insight*-HXMT low, medium, and high energy X-ray telescopes (LE, ME, and HE), as well as *NICER* and *INTEGRAL*, provided short snapshots at many different times, allowing an investigation of day-to-day variations, and an overall impression of the entire outburst. These complemented the sparse, yet long-duration observations by other instruments.

Out of the eight *NuSTAR* observations, six were performed simultaneously with *XMM-Newton*, and one simultaneously with *Chandra*. There was also an overlap between the second combined *NuSTAR* and *XMM-Newton* observation analysed here, with the second observation by *Astrosat*. These long-duration observations by *NuSTAR*, *XMM-Newton*, *Astrosat*, and *Chandra* are summarised in Table 5.1.

Swift/XRT performed 78 observations of Swift J1658.2-4242, spanning more than seven months. *Swift*/BAT produces daily average count rates. We used 42 observations by *NICER* to examine the timing evolution of the BHT. *INTEGRAL* obtained 151 observations of the BHT during its outburst, providing good hard X-ray coverage on shorter timescales than *Swift*/BAT.

Swift J1658.2-4242 was also observed with *Insight*-HXMT in 21 individual pointing observations, from MJD 58169 to 58196, with a total net exposure time of $\approx$ 700 ks, at low, medium, and high X-ray energies. Additionally, *ATCA* examined the radio emission from Swift J1658.2-4242 at ten different times within the outburst. These observations are summarised in Table 5.2.

5.3 Data analysis

5.3.1 X-ray data

We produced *Swift*/XRT light curves and spectra using the University of Leicester's online data analysis software[1], using standard input parameters to create one spectrum for each observation. We also applied dead pixel, pileup, and vignetting corrections. The BAT light curve was obtained from the *Swift*/BAT transient monitor website[2] (Krimm et al., 2013), which bins observations made throughout each day into one data point.

---

[1] http://www.swift.ac.uk/user_objects/

[2] https://swift.gsfc.nasa.gov/results/transients/

Instrument	**ObsID**	**Start time (MJD)**	**Exposure (ks)**	**State**	**QPO**
NuSTAR	90401307002	58165.97658	33.337	LHS	C
Astrosat	T02_004T01_9000001910	58169.76510	9.961	HIMS	C
NuSTAR	80301301002	58174.14666	31.537	Flip-flops HIMS↔AS	C and A
XMM-Newton	0802300201	58174.25471	70.526		
NuSTAR	80302302002	58176.75478	24.025	Flip-flops HIMS↔AS	C and A
XMM-Newton	0811213401	58176.79290	60.000		
Astrosat	T02_011T01_9000001940	58180.78589	29.578	Flip-flops HIMS↔AS	C and A
NuSTAR	80302302004	58181.31859	25.076		
XMM-Newton	0805200201	58181.33765	59.800		
NuSTAR	80302302006	58188.16432	28.245	HIMS	C
XMM-Newton	0805200301	58188.21508	51.200		
NuSTAR	80302302008	58192.46058	26.714	HIMS	C
XMM-Newton	0805200401	58192.54771	50.140		
NuSTAR	80302302010	58205.01353	29.102	Late flip-flops HIMS↔SIMS	No QPO
XMM-Newton	0805201301	58205.07581	50.900		
NuSTAR	90401317002	58235.95971	28.493	SIMS	No QPO
Chandra	21083	58236.05065	30.080		

Table 5.1: List of X-ray observations of Swift J1658.2-4242 by *Astrosat*, *Chandra*, *NuSTAR*, and *XMM-Newton* (in Timing mode). Observations with overlapping exposures have been grouped together. Swift J1658.2-4242 was also observed by *Swift*/XRT and BAT, *Insight*-HXMT, *NICER*, *INTEGRAL*, and *ATCA*.

Start time (MJD-5800)	**Exposure (ks)**	**Config**	**Frequency (GHz)**	**Flux density (mJy)**
166.83680	23.05	750B	5.50	2.35 ± 0.17
			9.00	2.17 ± 0.15
166.82534	22.36	750B	17.0	2.27 ± 0.09
			19.0	2.30 ± 0.15
171.93565	14.32	750B	5.50	5.95 ± 0.10
			9.00	4.80 ± 0.07
171.92430	14.33	750B	17.0	3.35 ± 0.12
			19.0	3.30 ± 0.20
173.87801	23.01	750B	5.50	50.70 ± 0.60
			9.00	39.50 ± 0.90
173.89942	19.80	750B	17.0	29.60 ± 0.50
			19.0	28.85 ± 0.50
175.62030	16.2	750B	8.40	11.27 ± 0.29
176.59271	24.00	750B	5.50	6.20 ± 0.20
			9.00	4.75 ± 0.20
176.61366	19.77	750B	17.0	3.00 ± 0.15
			19.0	2.90 ± 0.15
182.83993	2.81	EW352	5.50	1.76 ± 0.30
			9.00	1.24 ± 0.10
185.62130	7.43	EW352	5.50	5.30 ± 0.35
			9.00	4.10 ± 0.35
192.93941	11.84	EW352	8.40	0.70 ± 0.20
194.52511	1.21	EW352	5.50	0.95 ± 0.25
			9.00	0.98 ± 0.25
235.88102	8.02	H168	5.50	< 0.3
			9.00	< 0.3

Table 5.2: List of radio observations of Swift J1658.2-4242 by *ATCA*. No radio emission was detected in the last observation, so the 3σ upper limit is quoted instead.

NuSTAR data were reduced using the standard procedure with `nupipeline` and `nuproducts`, using the *NuSTAR* data analysis software (NuSTARDAS) version 1.8.0, version 4.7.8 of the CALibration DataBase (CALDB) from the high energy astrophysics science archive research center (HEASARC), as well as the general functionality of the high energy astrophysics software (HEASoft) version 6.25. In all observations, the lower-left corner of focal plane module B (FPMB) of *NuSTAR* was contaminated by stray light from a nearby source. To prevent this from affecting the extracted source signal, we restricted the size of the source extraction region, as there was some overlap between the source point spread function and the affected region of the CCD. We chose a circular region with a radius of 155" centred at the location of the source, for FPMB. This selection excluded virtually all the stray light seen in the CCD and included almost all the source photons. Focal plane module A (FPMA) did not suffer the same problem, so we chose a much larger source extraction region of a circle with a radius of 270" instead, also centred on the source. This different selection enabled us to test whether our choice of source extraction region for FPMB affected its spectrum. There are slight differences between the spectra of the two instruments at low energies, but they are insignificant to this analysis. Slight normalisation differences between the two instruments are as expected. Background extraction regions were selected manually for each observation and focal plane module, as the area of a circle of radius 30", placed far away from the source, and any other possible contaminants in the field of view. We restricted the *NuSTAR* energy range to 3–79 keV. *NuSTAR* spectra were also rebinned with GRPPHA to contain at least 20 counts per energy bin.

Observations by *XMM-Newton* were taken by the pn, and MOS2 detectors of the European photon imaging camera (EPIC; Strüder et al., 2001). The EPIC-pn camera was set to Timing and Burst mode, in which the pnCCD achieves high temporal resolution but loses spatial resolution along the y direction of the CCD. Thus, instead of a circular region, we selected individual pixel columns. As Swift J1658.2-4242 was very bright throughout all *XMM-Newton* observations, the entire CCD was dominated by source photons. Therefore, we used all observed photons to create the light curves and spectra and did not apply a background correction. The contribution from the background was estimated by comparing the spectra in the strip of 4×200 pixels furthest away from the source, to the spectra of the remaining 60×200 pixels of the CCD. The four-pixel wide column had a comparatively larger soft spectral component than the 60-pixel column. We interpret this as being primarily due to background contamination. As the four-pixel column was still dominated by the source, it would be inadvisable to use this as a background estimate. Instead, we decided to ensure minimal background contamination of the source spectra, by ignoring all energies for which the four-pixel column had a flux equal to at least 5% of the flux in the remaining pixels. This limit was exceeded at all energies below 2.7 keV, but not at any higher energies. Therefore, we restricted the *XMM-Newton* spectra to the energy range 2.7–10 keV.

We used version 16.1.0 of *XMM-Newton*'s Science Analysis System (SAS), implementing `epchain` to create event files, then using the standard procedures to produce spectra and light curves. We used `epiclccorr` to correct the light curves for telemetry dropouts. A few soft proton flares were detected when plotting the *XMM-Newton* light curves at energies exceeding 12 keV. All time intervals for which the count rate at these energies exceeded 4 cts/s were excluded from the light curves. Photons were binned into 6 ms intervals, limited to that value by the frame time for EPIC-pn in Timing mode. We compared the count rates of the observations made in Timing mode with those made in Burst mode shortly before or afterwards, to test whether the *XMM-Newton* observations were suffering from pile up. Such comparisons indicated a negligible level of pileup.

Lien et al. (2018) determined that Swift J1658.2-4242 is heavily obscured, with a hydrogen column density of $N_H \approx 10^{23}$ cm^{-2}. Jin et al. (2019) also detected a strong DSH around the black hole candidate. This provides a challenge for spectral fitting, as the DSH causes a significant, energy-dependent modification of the point spread function of the source. Therefore, pectra extracted from insufficiently large regions are distorted, particularly at low energies. Jin et al. (2017) and Jin et al. (2019) developed the spectral fitting model `dscor` to account for scattering of X-rays by a DSH along the line of sight. However, these models only work for circular source extraction regions and could therefore not be used for the *XMM-Newton* observations in Timing mode. To fit these spectra accurately as well, we developed a new spectral fitting model to describe the effect of the DSH on rectangular extraction regions which are not centred on the source.

We used the *Chandra* Grating-Data Archive and Catalogue webpage[3] to calibrate the data from the *Chandra* high energy transmission grating (HETG) observation of Swift J1658.2-4242. Then we also extracted energy spectra for the positive and negative grating directions.

The two *Astrosat* observations were performed with the LAXPCs. They were first reduced using the `laxpc_make_event` code inside the laxpcsoft package. Then, we generated light curves from the resulting level-2 event files using the task `laxpc_make_lightcurve`. The *Astrosat* observations had a significant background count rate, which was determined by the chain of analysis tasks `laxpc_make_spectra`, `laxpc_make_backspectra`, and finally `laxpc_make_backlightcurve`.

Swift J1658.2-4242 was observed with *Insight*-HXMT in pointed observation mode starting on MJD 58169. There are 21 available individual pointing observations with a total net exposure time of ≈ 700 ks. All the data were reduced following standard procedures using the *Insight*-HXMT data analysis software package HXMTDAS v2.01. The good time intervals (GTIs) were filtered with the screening criteria: (1) Earth elevation angle $>$ 10 degrees; (2) the value of the cutoff rigidity $>$ 8; and (3) the pointing offset angle $<$ 0.04 degrees. The events taken during satellite slews and passages through the South Atlantic Anomaly were also filtered out. For the data detected by the ME, the task `megti` was applied to make corrections to the GTI file.

The *Insight*-HXMT count rate observed by the ME detector was affected by the contribution of nearby sources in the field of view (Xiao et al., 2019). Within the 5.5–10 keV energy range of interest, we determined the contamination from the brightest background source in the field of view, GX 340+0, to be 2.6 cts/s. By comparing the light curves of *Insight*-HXMT ME with those of *NuSTAR*, and *XMM-Newton* for the same energy range,

---

[3] http://tgcat.mit.edu/

we, however, found the best possible agreement for an additional background count rate of 3.9 cts/s. As GX 340+0 is not the only uncorrected background source in the field of view, we used this empirical best fit value to estimate the additional background contribution.

To improve our understanding of the hard X-ray evolution during the outburst, we also created light curves from observations by the IBIS/ISGRI instrument on *INTEGRAL*. We extracted the count rates of individual images, each corresponding to one pointing observation, produced using the ISDC Off-line Scientific Analysis (OSA) software version 11, for five energy bands in the range 20–200 keV.

5.3.2 X-ray timing and spectral analysis

X-ray observations were sorted according to whether flip-flops were detected in them. When we did not detect flip-flops, we generated one periodogram and one energy spectrum for the entire observation. In contrast, when flip-flops were detected, observations were split into individual regions, which were analysed separately. We selected intervals that started or ended at least 100 s before, or after any flip-flop transition, to be used in determining the timing, and spectral properties of each of the flip-flop states.

Medium energy X-ray light curves were extracted for the 5.5–10 keV band. The lower limit of 5.5 keV was chosen because the DSH caused different telescopes to disagree on the flux densities measured at lower energies. The upper limit was chosen because of the energy range of *XMM-Newton* and the desire to maintain consistency across different instruments.

We generated periodograms using the general high-energy aperiodic timing software (GHATS) v.1.1.1 developed by T.M. Belloni, and the Stingray timing software (Huppenkothen et al., 2016). We binned individual photon detections into 10.24 ms bins and generated individual periodograms for intervals of 2^{10} bins. These were then combined to form averaged periodograms.

We created one periodogram for every orbit within each observation by *NuSTAR* and *Astrosat*, as long as the orbit did not feature a flip-flop transition. If it did, the relevant orbit was subdivided further to separate the bright from the dim state. *XMM-Newton* data were binned into 2 ks segments, or shorter ones if a flip-flop transition required it. We excluded all regions with less than 500 s duration, for which low statistics hampered precise spectral fits.

We subtracted the Poisson shot noise from all Leahy-normalised periodograms (Leahy et al., 1983) and subsequently applied the square fractional rms normalisation (Belloni and Hasinger, 1990). The periodograms were rebinned logarithmically and were then exported to XSPEC. We fit them using at least three different Lorentzians, at least one of which was zero-centred, following Belloni et al. (2002). We fit the main QPO peak, the broad-band continuum, and all visible harmonics. We extracted the main QPO properties (ν_0, Q, r, rms) from the fits, whenever a QPO was detected. The rms was computed by integrating fitted periodograms between 0.5 Hz, and 50 Hz, and calculating the associated source and background count rates in the interval. The values of the *NuSTAR* fractional rms in the QPOs were corrected for the effects of the *NuSTAR* deadtime (van der Klis, 1989). We also applied the rms correction to suppress high-frequency power, as described by van der

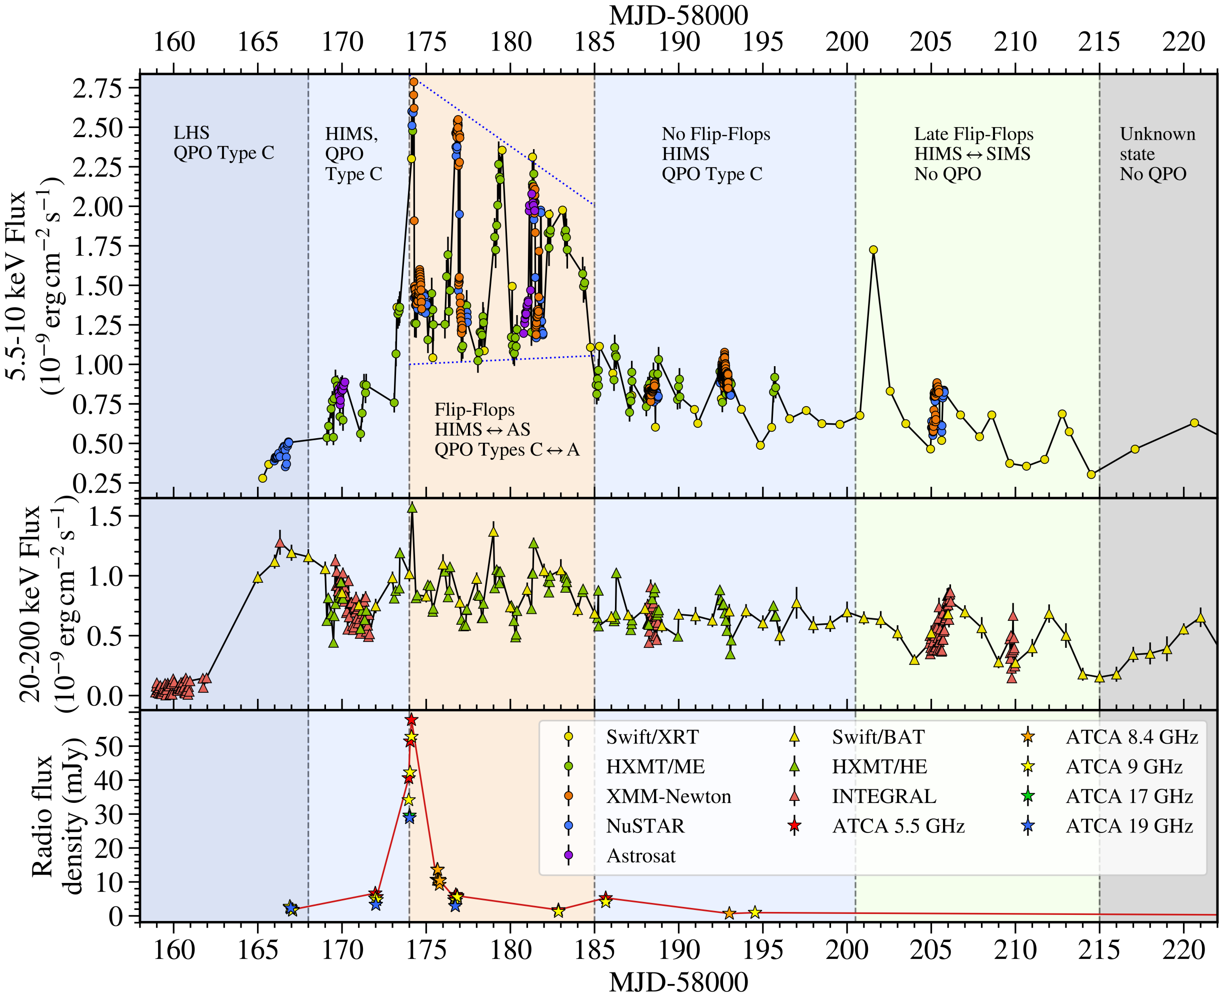

Figure 5.1: Light curve of the 2018 outburst of Swift J1658.2-4242 in three different energy bands. The top panel shows the flux in an intermediate energy X-ray band of 5.5–10 keV, as observed by *Swift*/XRT, *Insight*-HXMT/ME, *XMM-Newton*, *NuSTAR*, and *Astrosat*. In the second panel, we plot the *Swift*/BAT, *Insight*-HXMT/HE, and *INTEGRAL* light curves in the hard X-ray band of 20–200 keV. The radio light curves at 5.5, 8.4, 9, 17, and 19 GHz detected by *ATCA* are plotted in the third panel. We split the data temporally into six regions; the initial flux increase and spectral softening in the LHS and HIMS, the flip-flops, the intermediate HIMS sandwiched between the two flip-flop intervals, the late flip-flops, and finally, another interval for which we lack data to characterise accurately. The rest of the outburst was also observed and analysed, but we primarily focused on the regions shown here. We scaled the count rates of *Astrosat*, *Insight*-HXMT, and *Swift*/BAT to agree with observed fluxes measured by other telescopes.

Klis (1989).

Spectral fits were performed using XSPEC version 12.9.1u. To model the intervening absorption, we used solar abundances from Wilms et al. (2000) and the photoionisation cross-sections from Verner et al. (1996). The spectrum of the very first *NuSTAR* observation, with ObsID 90401307002, was also analysed but is not described here, as it was already thoroughly analysed by Xu et al. (2018), and because it differs greatly from the other *NuSTAR* and *XMM-Newton* spectra we investigated. It features a dominant power law component, and a negligible multicolour disc black body component (Xu et al., 2018).

5.3.3 Radio data

We obtained radio monitoring of Swift J1658.2-4242 with *ATCA*, under project code C3057. Data were obtained at eight epochs from MJD 58166 – 58235. Most radio observations were taken at central frequencies of 5.5 GHz, 9 GHz, 17 GHz, and 19 GHz, where each frequency pair (5.5/9 and 17/19 GHz) was recorded simultaneously. However, our three observations after 2018-03-05 were only taken at 5.5 GHz and 9 GHz. These frequency bands were recorded with a bandwidth of 2 GHz. Additional 8.4 GHz *ATCA* observations were taken on 2018-02-26 and 2018-03-15 as part of observations with the Australian Long Baseline Array (project code V456). These data were recorded with a bandwidth of 64 MHz. The bandpass and flux calibration was performed for all radio observations using PKS 1934-638. Secondary phase calibration was done using the nearby (4.65 degrees away) calibrator J1714-397 for all observations except those taken on 2018-02-26 and 2018-03-15, which used J1713-4257 (2.9 degrees away).

The radio data were flagged and calibrated following standard procedures within the common astronomy software application (CASA, version 4.7.2; McMullin et al., 2007). Imaging was carried out using CLEAN within CASA. Due to a significant amount of diffuse emission in the field, we used a Briggs roust parameter of 0 to help mitigate effects from diffuse emission within the field. Typically, flux densities were determined by fitting for a point source in the image plane. However, due to either short on-source time, or a very compact telescope configuration, flux densities of observations on 2018-03-05, 2018-03-08, and 2018-03-17 (MJD 58182, 58185, and 58194, respectively) were determined by fitting for a point source in the UV-plane using UVMULTIFIT (Martí-Vidal et al., 2014). We also used UVMULTIFIT to search for intra-observational variability within the radio observations, to identify the possibility of radio flaring associated with the X-ray flip-flops.

5.4 Light curves

In Fig. 5.1 we show the light curve of the first 60 days of the 2018 outburst of Swift J1658.2-4242 in three different energy bands. We separated this light curve into six temporal regions based on their different features. In the first region, we detected Swift J1658.2-4242 in the LHS, with the typical flux increase from quiescence. We depict the total outburst in the HIDs of Fig. 5.2 and 5.3. The set of data on the lower right of Fig. 5.3 is from this

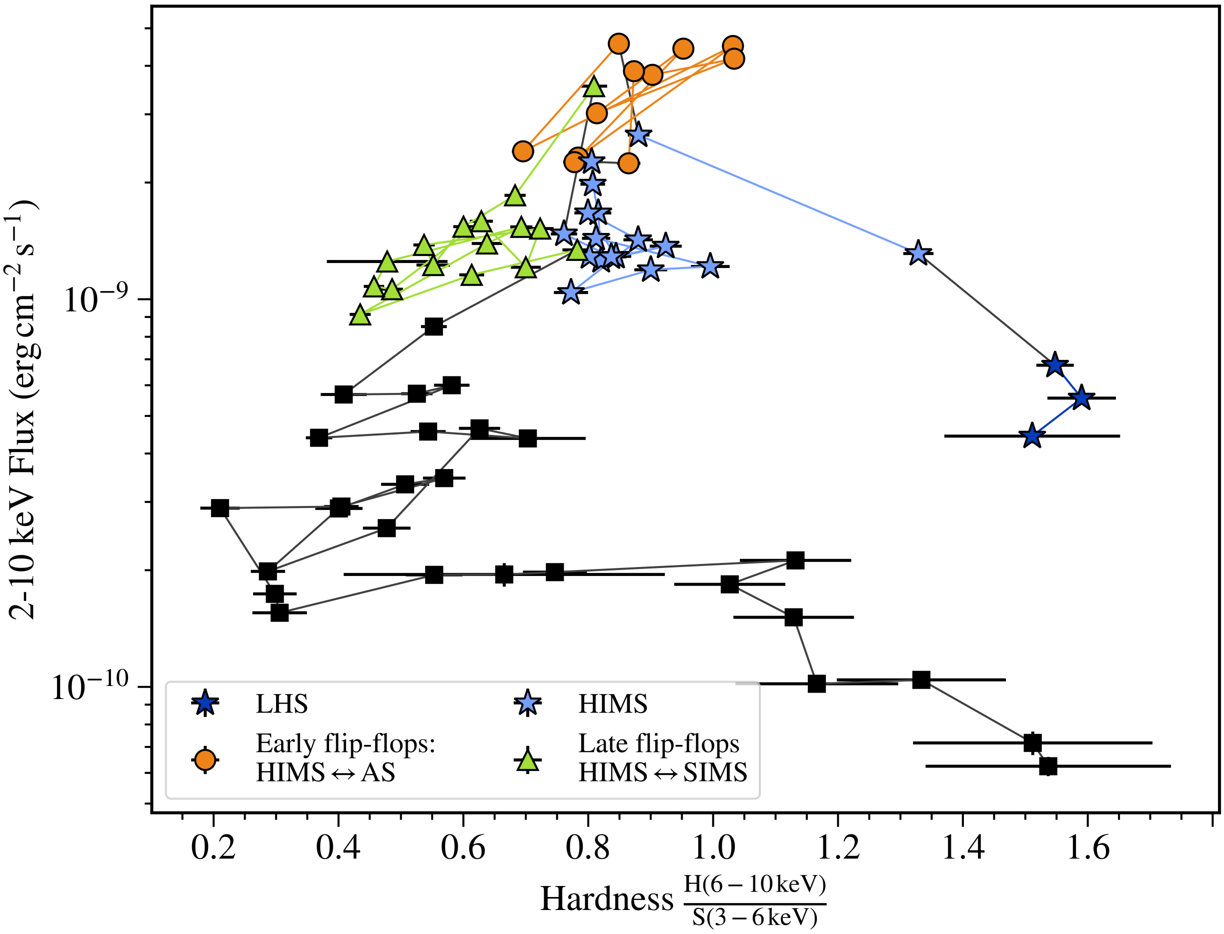

Figure 5.2: HID of all *Swift*/XRT observations. Colours and shapes distinguish the different states of the outburst. Hardness is defined here as the ratio of fluxes in the two energy bands 6–10 keV, and 3–6 keV. Each data point represents one entire observation. Black squares are observations for which a state classification was not possible. The errors in the flux are shown but are almost always smaller than the symbols used for the data points.

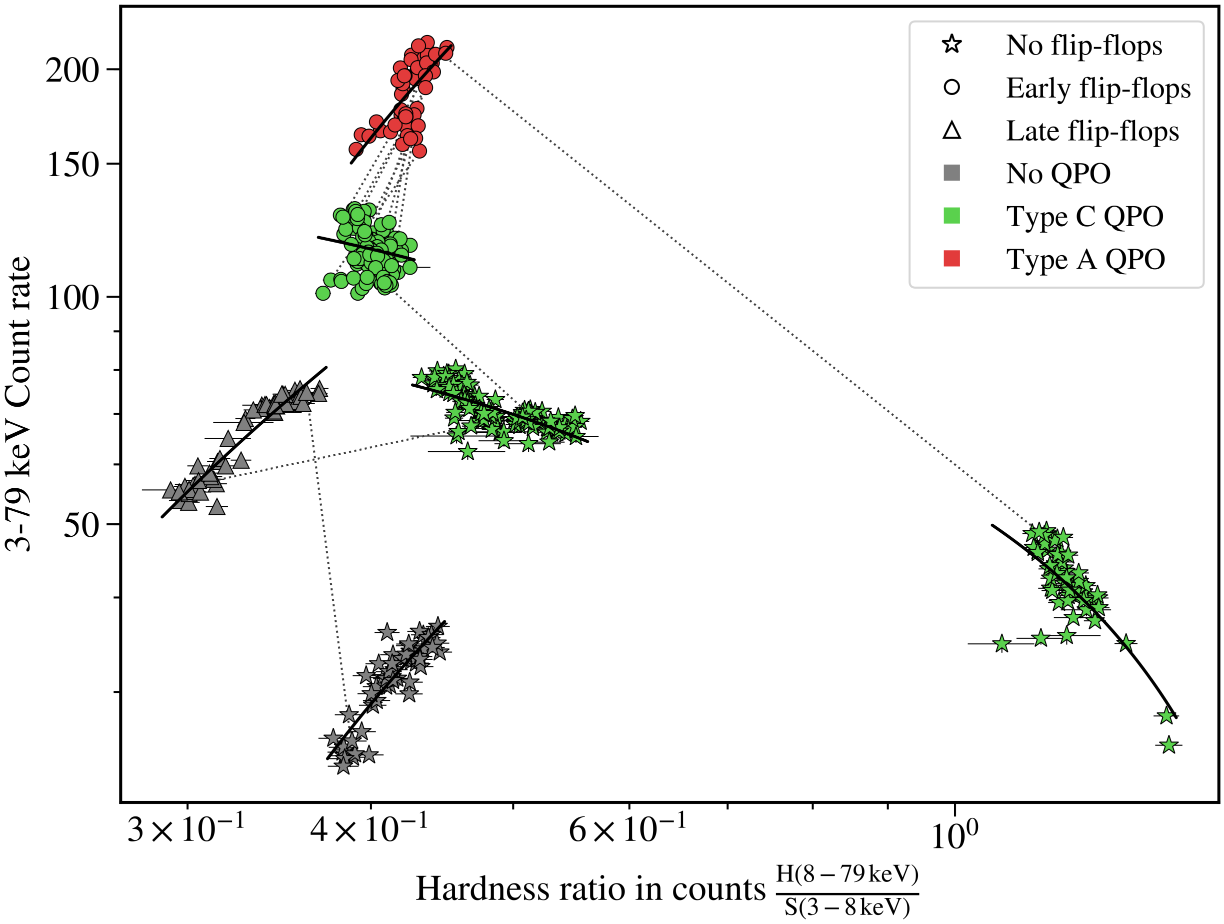

Figure 5.3: *NuSTAR* HID. Colours and shapes distinguish QPO types, and flip-flop detections, respectively. Hardness is defined as the ratio of count rates in the two energy bands 8–79 keV, and 3–8 keV. Data were binned into 200 s intervals. Error bars have been included but are almost always smaller than the symbols used for the data points.

LHS, and shows a negative correlation between intensity and hardness, with a Pearson's correlation coefficient of $r_{xy} = -0.71$, and a two-tailed p-value of $p = 7.5 \times 10^{-11}$.

Next, we observed the HIMS. The time of transition from the LHS to the HIMS is not known, but spectral analyses of *NuSTAR* and *Swift*/XRT observations indicate that it probably occurred around MJD ≈ 58168. Within this state, we detected an increase in the 5.5–10 keV X-ray flux, in parallel with a decrease of the 20–200 keV flux, showing the significant softening of the spectrum. This softening can also be seen in the upper right region of the HIDs of Fig. 5.2 and 5.3.

On MJD 58174, we observed the start of a major radio flare with *ATCA*, which reached flux densities at least 8.7 times brighter than those measured at any other time during the outburst. As the radio flux was still rising when the observation finished, the flare presumably reached even greater fluxes. Merely 200 s after the end of this *ATCA* observation, we measured the highest *NuSTAR* and *XMM-Newton* X-ray flux of the entire outburst, 5.4

times brighter than in the previous *NuSTAR* observation. This can be seen in Fig. 5.4. We also observed the brightest hard X-ray flux of the outburst at the same time, in an observation by *Insight*-HXMT HE in the 20–200 keV band (see Fig. 5.1). When fitting the *NuSTAR* energy spectrum of this brightest state, we measured a total flux in the 3–79 keV energy range, of $F_{3-79} = 9.75 \times 10^{-9}$ erg cm^{-2} s^{-1}. This corresponds to an unabsorbed and unscattered flux of $F_{3-79,0} \approx 1.64 \times 10^{-8}$ erg cm^{-2} s^{-1}. Using the distance estimated by Jin et al. (2019) of $\approx$ 10 kpc, we calculated the unabsorbed and unscattered source luminosity to be $L_{3-79} \approx 2 \times 10^{38}$ erg s^{-1}.

5.4.1 Early flip-flops

Within the same simultaneous observation by *NuSTAR* and *XMM-Newton* in which we measured the greatest intensity of the outburst, there was a sudden flux drop of 43% within 40 s. In subsequent observations by *NuSTAR*, *XMM-Newton*, *Astrosat*, and *Insight*-HXMT/LE and ME, we saw similar flux increases or decreases, with the bright states having a 49–77% higher flux than the dim states (see Fig. 5.4, 5.5, 5.6, and 5.7). Transitions between the two flux levels occurred on timescales of 26–800 s (see Fig. 5.4, 5.5, 5.6, and Table 5.3). Unlike Xu et al. (2019), and Jithesh et al. (2019), we identify these phenomena as flip-flops, as they feature most of the characteristics of flip-flops seen before in other BHTs.

Table 5.3 shows that the greatest ratio of bright to dim state flux between neighbouring states that we observed, occurred during the first observed flip-flop transition and has a value of 1.77. Subsequent flux ratios indicated a decreasing amplitude of fractional flux variation between the two states over time. The second to last of the flip-flops that we observed directly had a flux ratio of merely 1.49. This is reminiscent of a damped oscillation.

In Fig. 5.4, 5.5, and 5.6, we plotted the light curves and hardness ratios of individual *NuSTAR* and simultaneous *XMM-Newton* observations containing flip-flops. 5.4 also features the *ATCA* observation of a radio flare, and 5.6 is complemented by an *Astrosat* observation. These figures show that changes in the total flux occurred coincident with changes in the hardness. Higher flux corresponded to greater hardness, but the fractional increase in hardness was smaller than the fractional increase in flux. We can, therefore, deduce that a flip-flop transition from a dim to a bright state corresponds to a flux density increase for both hard and soft X-ray energies. Nevertheless, the hard X-ray flux must have a larger fractional increase than the soft X-ray flux, to generate the observed change in hardness. However, this trend does not extend to very high energies, as the 20–200 keV light curve of Fig. 5.1 does not feature as significant fractional flux variations within the flip-flop intervals, as the 5.5–10 keV light curve does.

The hardness ratios in Fig. 5.4, 5.5, and 5.6, seem to be correlated with flux variations in the bright states, and anti-correlated with flux variations in the dim states. This is verified in the best fit lines of individual regions within the *NuSTAR* HID of Fig. 5.3. Remarkably, despite being separated by several days, all bright states, and all dim states, lie in the same region of the HID. We combined all the bright state measurements to find

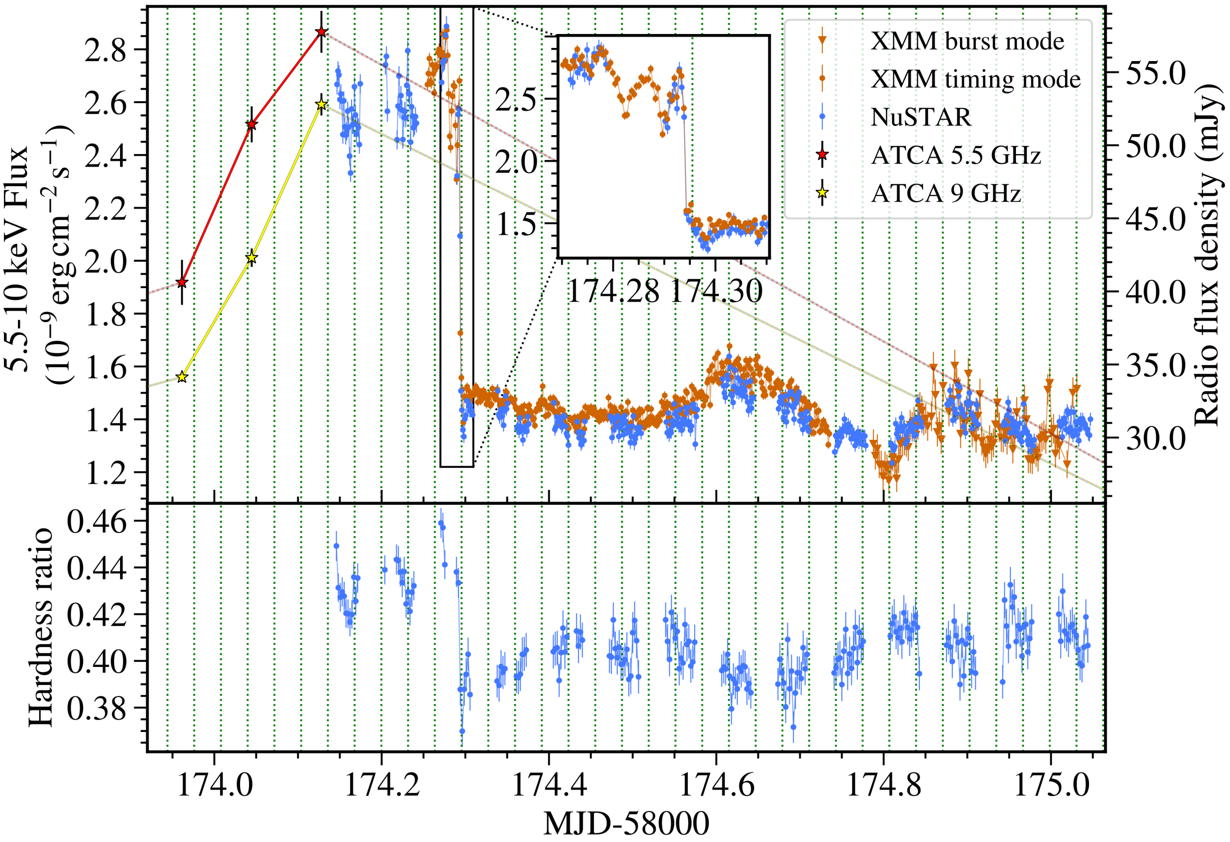

Figure 5.4: Light curve and hardness ratio of the first flip-flop transition observed during the outburst, in *NuSTAR* observation 80301301002, and *XMM-Newton* observation 0802300201. We also show the *ATCA* radio light curve in this interval. *ATCA* detected a flare, and subsequent decay at about the same time as the first flip-flop was observed. Each bin in the light curve covers 100 s intervals, the inset light curve has 40 s bins, and the hardness ratios are determined for 200 s bins. Vertical lines are plotted every 2761 s, for the interval covered in this figure, and also in Fig. 5.5, and 5.6. These lines show that almost all flip-flop transitions occur at integer multiples of a fundamental period.

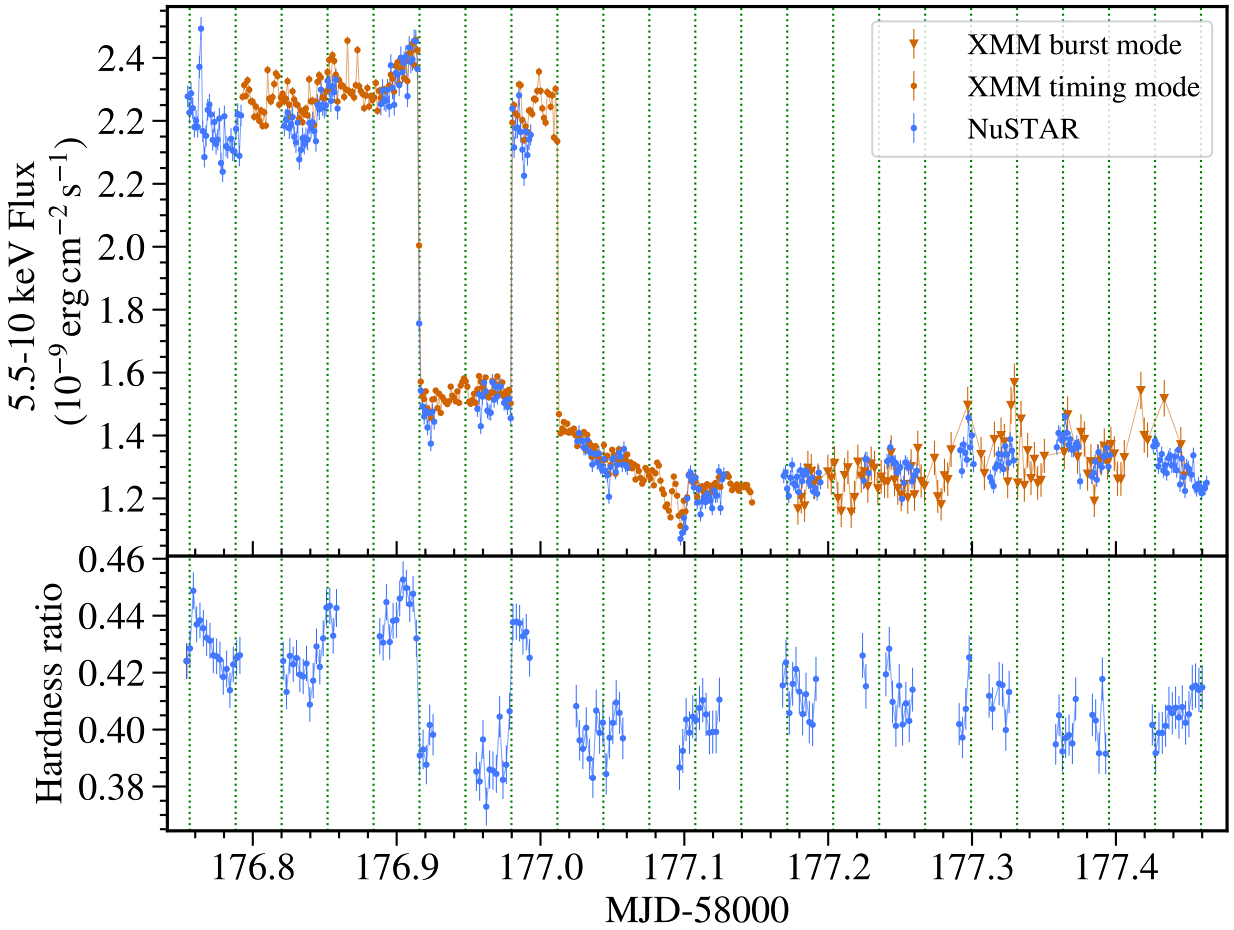

Figure 5.5: Light curve and hardness ratio of the second segment of observations within which we observed flip-flops. This light curve features data from the *NuSTAR* observation 80302302002, and the *XMM-Newton* observation 0811213401. Each bin in the light curve covers 100 s intervals, and the hardness ratios are determined for 200 s bins. Vertical lines are plotted every 2761 s, for the interval covered in this figure, and also in Fig. 5.4, and 5.6. These lines show that almost all flip-flop transitions occur at integer multiples of a fundamental period.

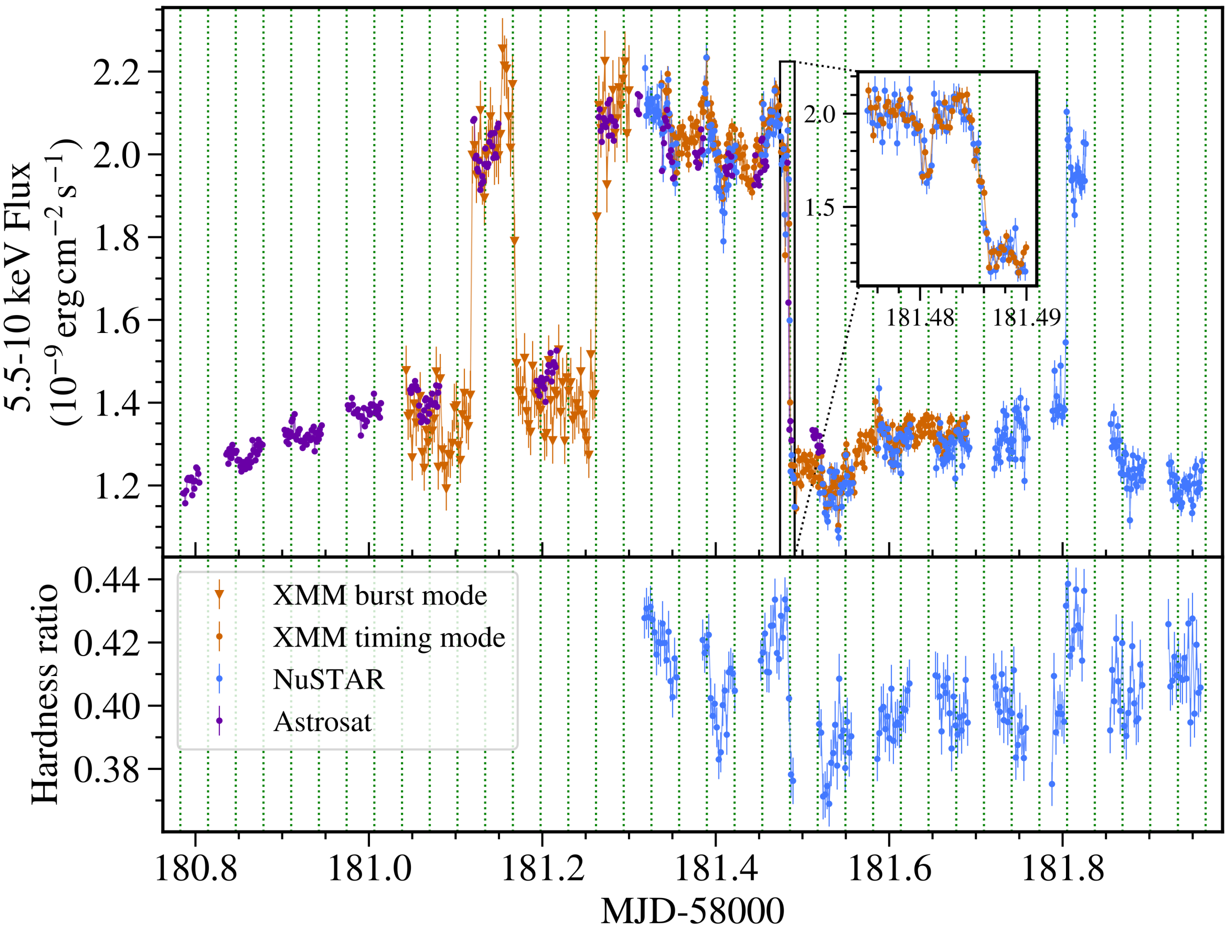

Figure 5.6: Light curve and hardness ratio of the third segment of observations in which flip-flop transitions were observed. This light curve features data from *Astrosat* observation T02_011T01_9000001940, the *NuSTAR* observation 80302302004, and *XMM-Newton* observation 0805200201. Each bin in the light curve covers 100 s intervals, the inset light curve has 20 s bins, and the hardness ratios are determined for 200 s bins. Vertical lines are plotted every 2761 s, for the interval covered in this figure, and also in Fig. 5.4, and 5.5. These lines show that almost all flip-flop transitions occur at integer multiples of a fundamental period.

State	**Flux**	**State duration** (ks)	**Time of centre of transition** (MJD-5800)	**Transition duration** (s)
Bright	9.09	≥ 12.7	174.2942	42 ± 4
Dim	5.15	≥ 65.2		
Bright	8.56	≥ 13.9	176.9158	74 ± 8
Dim	5.47	5.6	176.9895	26 ± 3
Bright	8.41	2.7	$177.0127^{\dagger}$	$62 \pm 10^{\dagger}$
Dim	4.81	≥ 39.0		
Dim		≥ 28.8	181.1200^{*}	$120 \pm 40^{*}$
Bright		4.0	181.1697^{*}	$800 \pm 200^{*}$
Dim		7.6	181.2635^{*}	$320 \pm 80^{*}$
Bright	7.25	19.0	181.4857	160 ± 20
Dim	4.72	27.4	181.8045	36 ± 5
Bright	7.04	≥ 2.0		
Dim	4.59	≥ 9.2		
Dim	2.40	≥ 13.1	205.1654	90 ± 20
Bright	2.95	7.6	$205.2539^{\dagger}$	$110 \pm 30^{\dagger}$
Dim		2.2	$205.2836^{\dagger}$	$700 \pm 100^{\dagger}$
Bright	3.15	18.3	205.5001	50 ± 20
Dim	2.41	≥ 2.0		
Bright	2.91	≥ 0.4	205.5556	50 ± 10
Dim	2.32	7.9	205.6505	100 ± 20
Bright	3.16	≥ 12.2		

Table 5.3: List of flip-flop properties. Each horizontal line denotes a prolonged gap between observations, within which flip-flop properties could not be accurately identified. Two horizontal lines distinguish the early from the late flip-flops. Fluxes are determined in the 3–79 keV range from the *NuSTAR* FPMA spectra, and are expressed in units of 10^{-9} erg cm^{-2} s^{-1}. State durations are calculated based on the first and last measurement performed by *NuSTAR*, *XMM-Newton*, or *Astrosat*, which can be confidently associated with this state. The time of the centre of transition, and its duration, were calculated using Equation 5.1 on the *NuSTAR* light curves, restricted to be within 100 s before the start of the transition, and 100 s after the end of the transition. For the late flip-flops, we restricted the region of the light curve to be fitted to 200 s before the start, and 200 s after the end of the transition. When transitions were not seen by *NuSTAR*, we instead calculated the transition times from the *XMM-Newton* light curves. The symbols † and $*$ denote transition times calculated from the *XMM-Newton* Timing and Burst mode light curves, respectively. Transition times calculated from *XMM-Newton* Burst mode light curves are less reliable estimates and are therefore significantly larger. Blank cells indicate states or transitions for which certain properties could not be determined reliably.

a correlation coefficient of $r_{xy} = 0.66$, and associated two-tailed p-value of $p = 2.5 \times 10^{-7}$, indicating a strong positive correlation between intensity and hardness. In contrast, the combined dim state measurements have correlation statistics of $r_{xy} = -0.17$, and $p = 0.045$.

Throughout the observations by *NuSTAR*, *XMM-Newton*, and *Astrosat*, we observed six bright states, seven dim states, and nine transitions between the two states. In Fig. 5.6, we noticed one additional instance of a significant flux difference on either side of a gap within an observation by *NuSTAR*, caused by its low Earth orbit, at about MJD 58181.84. However, the differences in the flux, timing, and spectral properties on either side of the gap, are consistent with the flip-flop behaviour, and we consider that the unobserved region also featured a flip-flop.

In Fig. 5.7, we zoom into the flip-flop region to better illustrate what is happening within it. Thanks to frequent observations by *Swift* and *Insight*-HXMT, we can narrow down the start of the flip-flop interval to MJD 58174. Interestingly, this implies that the start of the flip-flops coincides with the time of the highest intermediate and hard X-ray fluxes, and the major radio flare.

On MJD 58184, following an extended bright state, which appeared to last for more than two days, we observed a final transition to a significantly lower flux. None of our observations over the subsequent 20 days showed any flip-flop activity. Therefore, it seems like this final drop-down transition on MJD 58184–58185 marked the end of this interval of flip-flops, which had lasted for about 11 days.

The *Insight*-HXMT, and *Swift*/XRT light curves during the flip-flop interval revealed the existence of seemingly long (≈ 1 day) dim state periods. These were interspersed with shorter segments involving several flip-flop transitions to and from bright states, which we observed to last for 2.2–19.0 ks (see Table 5.3). However, it is likely that we missed several bright states between individual *Insight*-HXMT and *Swift* observations. In the light curve of Figure 5.7, there appear to be four shorter segments within which bright states and flip-flop transitions were observed. We obtained detailed measurements with *NuSTAR*, *XMM-Newton*, and *Astrosat* of three of these segments, which can be examined in greater detail in Fig. 5.4, 5.5, and 5.6. The duration of these short segments in which one or more bright states were observed, seems to increase from one instance to the next (lasting approximately 0.15, 0.25, 0.5, and 3.3 days, respectively), whereas the time between them, in which no bright state was observed, appears to decrease (lasting approximately 2.5, 2.0, and 1.6 days, respectively).

In the 20–200 keV range, we detected similar up-and-down behaviour in the flux, which corresponds to, but is less pronounced than the light curve evolution observed in the 5.5–10 keV interval. Individual states are harder to identify in the hard X-ray light curve, so we focused on the 5.5–10 keV one. However, as the hardness ratios of Fig. 5.4, 5.5, and 5.6 indicate, another way to single out flip-flop transitions is to monitor the ratio of the disc black body to power law emission.

The flux change during transitions decreased gradually over time, as can be seen by the dotted lines in Fig. 5.1, and 5.7. However, there is only a minimal variation between the flux changes of adjacent transitions. This is illustrated by the blue line in Fig. 5.14, which does not show any discontinuities at the transitions. The gradual flux difference decrease

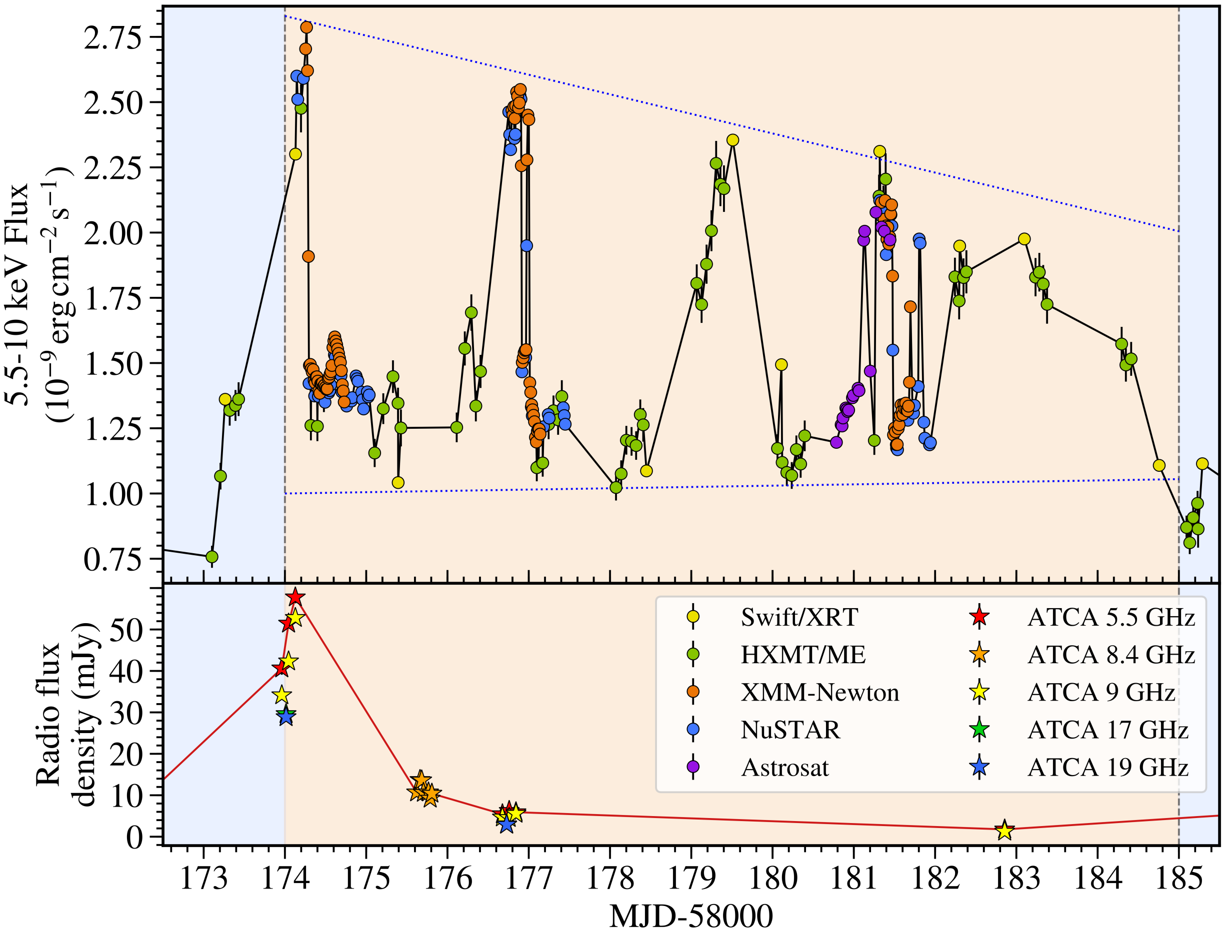

Figure 5.7: Light curve from Fig. 5.1, zoomed into the flip-flop interval. We omitted the 20–200 keV light curve as it does not provide as much useful information as the other two light curves do in this interval.

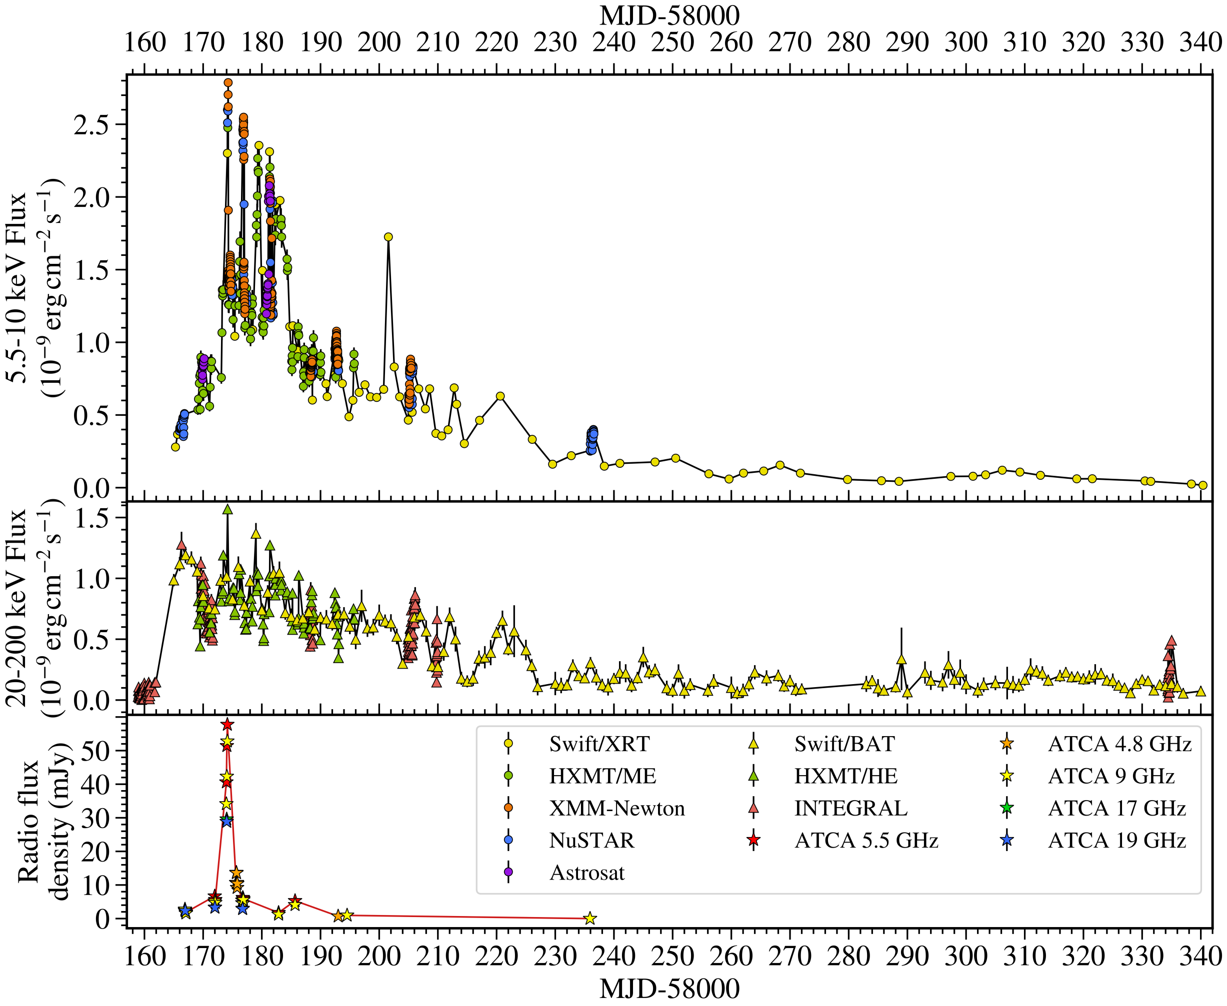

Figure 5.8: Light curve of the entire outburst, for X-rays of 5.5–10 keV, 20–200 keV, and at radio wavelengths.

observed in the flip-flops is reminiscent of the amplitude decay of a damped oscillator.

We determined the time it takes to transition between the two states, by fitting this function:

$$y(t) = \begin{cases} y_0 + m_1(t - t_0) & \text{if } t \leq t_0 \\ y_0 + \frac{y_1 - y_0}{\Delta t}(t - t_0) & \text{if } t_0 < t \leq t_0 + \Delta t \\ y_1 + m_2(t - t_0 - \Delta t) & \text{if } t > t_0 + \Delta t \end{cases} , \tag{5.1}$$

to the region of the light curves from 100 s before the start, to 100 s after the end of each transition. In the above equation, y_0 is the flux just before the transition, and y_1 is the flux just after it; t_0 is the time when the transition starts, and the transition duration is Δt; m_1 and m_2 are the gradients of the light curve just before and after the transition. We found that the transition durations measured from observations by *XMM-Newton* in Burst mode were substantially longer than other measurements, which is a consequence of its larger uncertainties. Transitions from bright to dim states were found to take longer than transitions from dim to bright. Bright to dim transitions lasted for about 42–160 s, whereas dim to bright transitions only lasted for about 26–35 s.

In the inset of Fig. 5.6, we indicate the presence of a short ($\approx$ 90 s) drop in the light curve, with a fractional flux decrease of $\approx 15\%$, just 300 s before the transition from the bright to the dim state. We also see two long ($\approx$ 600 s), almost triangular flux drops of $\approx 18\%$ just before the first flip-flop transition, as can be seen in the inset of Fig. 5.4. These features might be interpreted as failed transitions back down to the dim state.

After the end of the flip-flops, the variability of the light curves on timescales of days was small (see Fig. 5.1). We did not find any evidence of flip-flops in the two *NuSTAR* and simultaneous *XMM-Newton* observations made in this interval, and even frequent monitoring by *Swift* and *Insight*-HXMT did not reveal any significant flux changes. The flux decreased exponentially.

Both *NuSTAR* and *XMM-Newton* observations in this interval lie in almost the same region of the HID (See Fig. 5.3), which is, however, very clearly distinguished from every other observed part of the outburst. We see a negative correlation between hardness and intensity when combining the two observations, with $r_{xy} = -0.75$, and $p = 2.6 \times 10^{-22}$.

5.4.2 Late flip-flops

On MJD 58201, *Swift/XRT* detected an X-ray flare. Five days later, in the following observation by *NuSTAR* and *XMM-Newton*, we detected flip-flops again (See Fig. 5.9). These, however, differ quite significantly from the ones seen previously. They have much smaller flux differences, with flux ratios of between 1.22 and 1.36 between the bright and dim states in the 3–79 keV energy range.

Interestingly, the X-ray flare is placed in the same region of the HID as the early flip-flops were observed in. The late flip-flops observed afterwards, however, fall into a different region, as can be seen in Fig. 5.2.

Unlike the early flip-flops, here we observed a positive correlation between hardness and intensity in both the bright and dim states. The same correlation applies to both, and

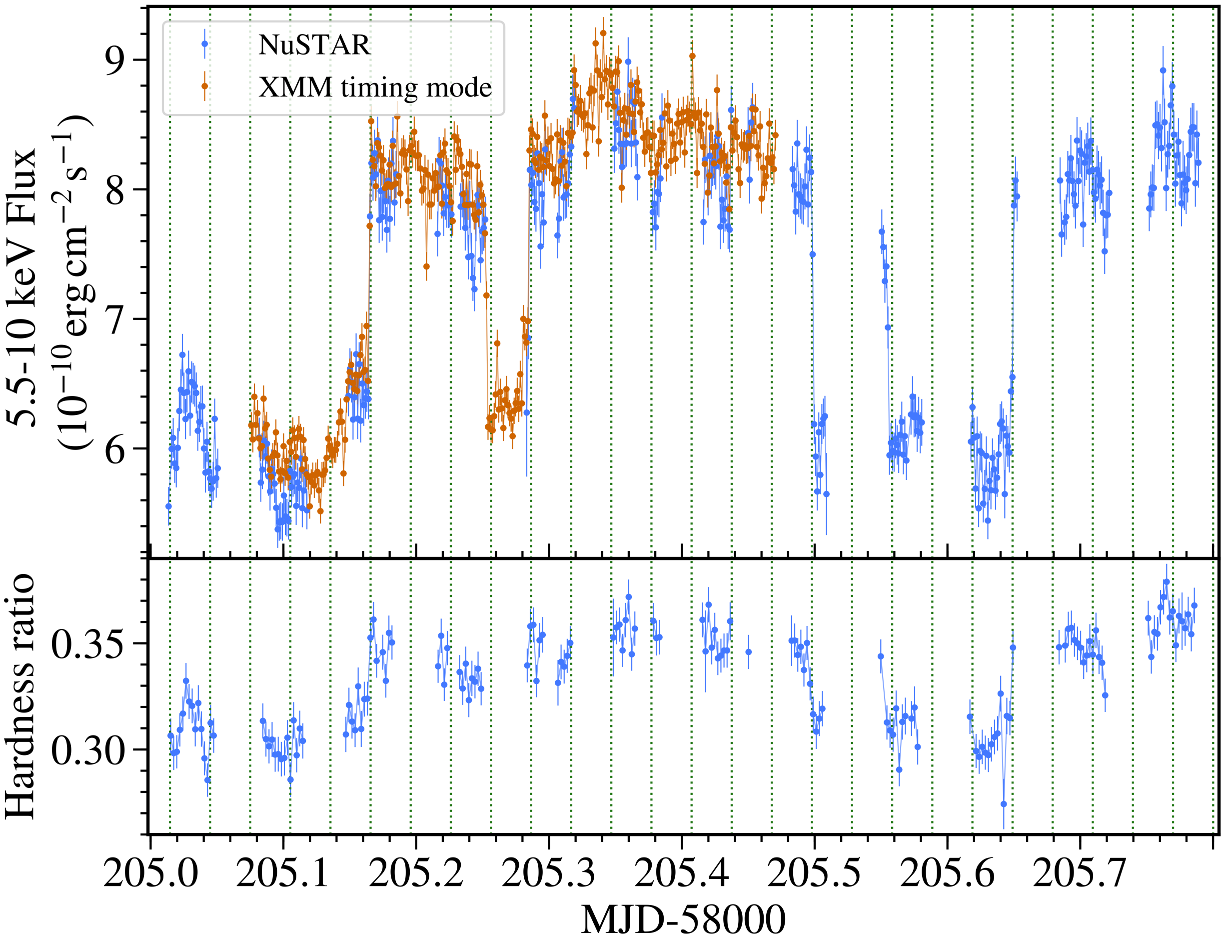

Figure 5.9: Light curves and hardness ratio of the *NuSTAR* and *XMM-Newton* observations of the late flip-flops. We used the same energy ranges and time bins as in Fig. 5.4, 5.5, and 5.6. The Burst mode is not included here, as it was strongly affected by soft proton flares. Vertical lines are drawn every 2610 s, showcasing that the late flip-flops also feature a period underlying the times of transition between the two flip-flop states.

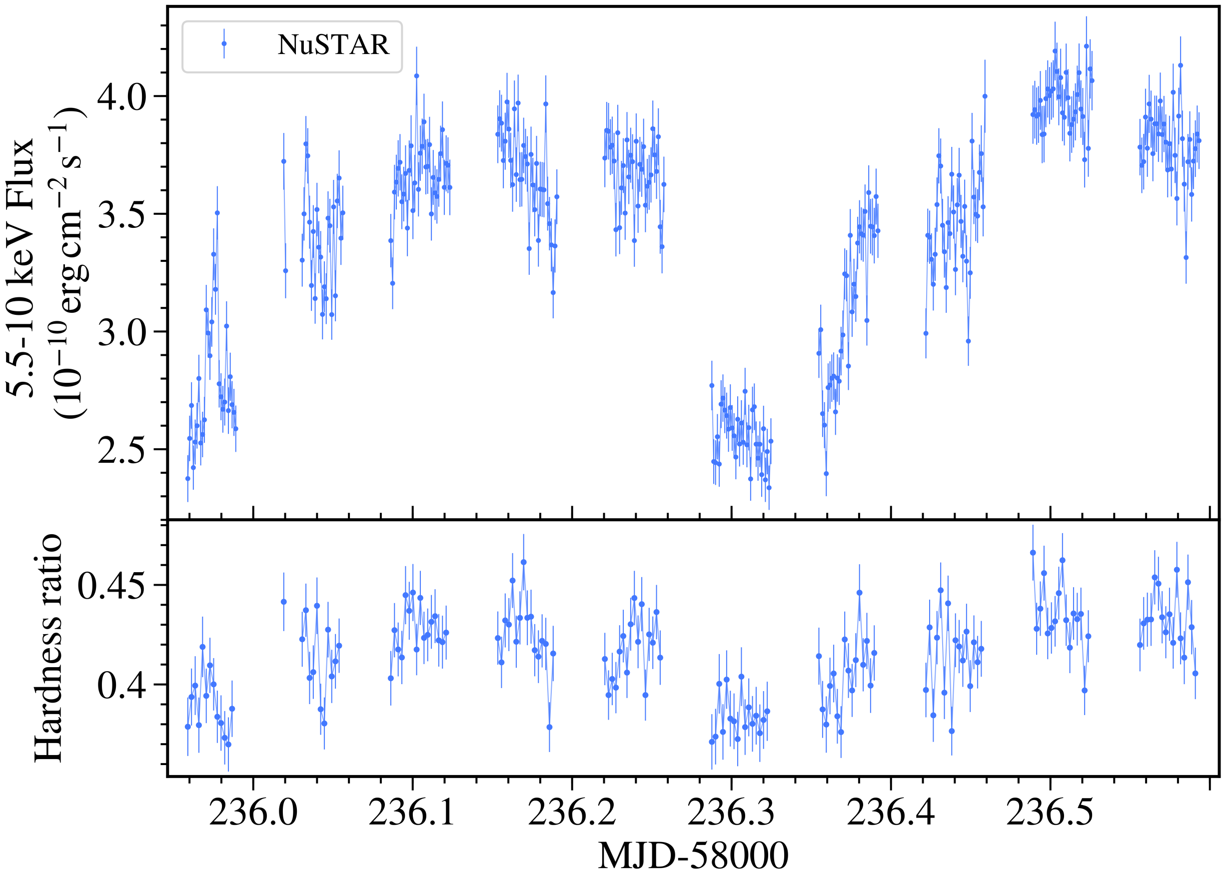

Figure 5.10: Light curve and hardness ratio of the final *NuSTAR* observation, with ObsID 90401317002, starting at MJD 58235.96. We used the same energy ranges and time bins as in Fig. 5.4, 5.5, and 5.6

has a correlation coefficient and p-value of: $r_{xy} = 0.95$, and $p = 4.1 \times 10^{-20}$.

Due to fewer observations of Swift J1658.2-4242 during the late flip-flop interval, we cannot place stringent bounds on its start and end times. The X-ray flare may have marked the start of this interval. In the days that followed the flare, we detected a greater variability between subsequent observations by *Swift*/XRT and BAT than in the 15 days that followed the end of the first flip-flop interval. This variability seemed to end by about MJD 58215, which is our estimate for the end of this late flip-flop interval.

We observed four bright states, four dim states, and six transitions between them. One additional transition was inferred but was not observed. Contrary to the early flip-flops, the late flip-flop transitions from bright to dim states (50–110 s) were faster than the reverse transitions (90–700 s).

The final *NuSTAR* observation, on MJD 58235–58236 contained a drop in the flux of about 27% during a gap in the data, followed by a gradual flux rise up to the previous level (see Fig. 5.10). The hardness ratio correlates with the flux for this observation (with $r_{xy} = 0.89$, and $p = 4.1 \times 10^{-20}$), but the shape of the light curve is vastly different to what was seen during the early or late flip-flops. We, therefore, decided against labelling these flux changes as flip-flops.

In the first *NuSTAR* observation, two absorption dips were seen (Xu et al., 2018). These are distinguished from flip-flops via their short duration, their shape in the light curve, and the increase of the hardness ratio. We did not observe dips in any of the other observations of Swift J1658.2-4242.

The major radio flare seen at the start of the early flip-flop interval died down quickly. During the following observation, ≈ 1.5 days after the detection of the radio flare, the flux had almost entirely returned to its previous level again. We did not detect any additional radio flares at any other point during the outburst.

In Fig. 5.2 we show all *Swift*/XRT observations of the outburst, which continued until September 28, 2018 (MJD 58389). The rest of the outburst features the typical decrease in flux, followed by an increase in hardness, as can be seen in Fig. 5.8. As far as we can tell, Swift J1658.2-4242 never reached the HSS but remained in the intermediate states until it reached the LHS.

The duration of the flip-flop states seems to decrease over time. By this, we define the amount of time for which a particular bright, or dim state was maintained, from the end of the transition leading into it to the start of the transition leading out of it. In the first observation by *NuSTAR*, and *XMM-Newton* within the early flip-flop interval, only one bright and one dim state was seen. In the subsequent two observations, two bright, and two dim states were detected. Moreover, within the late flip-flop interval, a single *NuSTAR* observation spanned four bright, and four dim flip-flop states (See Fig. 5.4, 5.5, and 5.6). In this outburst alone, we recovered some of the variety of flip-flops seen in the literature. We observed a range of flip-flop amplitudes, and durations, and also detected different behaviour in the time domain.

5.5 Discovery of an underlying clock in flip-flop transitions

5.5.1 Constant timescale for early flip-flops

Flip-flop states have vastly different durations. The time between one flip-flop transition and the next was observed to range from anywhere between 2.56 ks to at least 65.2 ks. There does not seem to be a fixed time after which the light curve repeats itself in a periodic manner. However, in the *XMM-Newton* observation 0811213401, we noticed that two of the observed flip-flop states during MJD 58176–58177 had times between transitions which had a ratio very close to 2:1 (See Fig. 5.5, and Table 5.3). The time between the first and second, and between the second and third transition in this observation, is 5590 s and 2785 s, respectively. Based on this finding, we investigated the possibility of an underlying timescale manifesting itself throughout all flip-flops. Starting from one of the transitions of this observation, we extrapolated integer multiples of the shorter of the two durations throughout the early flip-flop interval via the equation:

$$f(n) = Pn + t, \tag{5.2}$$

where t is the centre time of one particular flip-flop transition; P is a period, which we initially estimated to be $P = 2785$ s, and $n \in \mathbb{Z}$. We found that numerous transitions lie very close to one of the $f(n)$ times determined by this formula. Transitions do not happen at every integer n, but when they occur, they lie close to one of the $f(n)$. The light curve does not appear to be periodic, but changes in the light curve occur at seemingly random multiples of a fundamental period, P. Due to the lack of overall periodicity, the standard methods to search for a period over which a cycle repeats itself are not applicable here.

We used the centre times of each transition, which were determined by fitting the light curves around every transition with Equation 5.1. Of the nine transitions we observed, eight are fitted well by Equation 5.2 and a period of about $P = 2785$ s. Often, one of the $f(n)$ times passes through the width of a transition or is just barely shy of it (see Fig. 5.4, 5.5, and 5.6). However, the first transition observed by *XMM-Newton* in Burst mode on MJD 58181 (Fig. 5.6) is not close to any of the times calculated by Equation 5.2.

Excluding this anomalous transition time, and using ≈ 2.78 ks as our initial estimate of the period, we determined a best fit with $P = 2761.0 \pm 0.4$ s, and $t = 58174.2955 \pm 0.0006$ MJD, having a goodness of fit of $\chi^2 = 16.5$, for 6 degrees of freedom. The results of this fit are shown in Fig. 5.11. The vertical lines drawn in Fig. 5.4, 5.5, and 5.6 showcase the best fit results overplotted by the light curve. We also searched for a period larger than 2761.0 s that could fit our observations better than this one, but did not find any. Shorter periods were found, but these turned out to be fractions of the 2761.0 s period.

We next considered the significance of this detection by investigating the likelihood of obtaining at least as good of an agreement from a sample that was not based on a fundamental period. We, therefore, simulated sets of nine random transition times within the window of our observations of the early flip-flop interval, to which we randomly assigned our measured errors in the transition times. The shorter the period is, the greater the likelihood of finding a good agreement between randomised transition times, as the maximum delay between a transition time and an extrapolated period decreases. We, therefore, considered the probability of being able to fit for a period in the way we did, which is at least as long as the one we determined. We fitted all simulated transition times with Equation 5.2, using a given input period, and having set $n = 0$ for the chronologically first transition. For each set of nine simulated transitions and each initial period estimate, we excluded the transition deviating the most from the best fit, leaving only eight transitions to be fitted. In doing so, we ensured that simulated and real data were treated in an identical way.

We fitted the remaining eight transitions again, to determine the closest integer n for each of them, to make sure that the ignored transition did not affect the result. Using the newly found best-fitting integers n, we performed a third fit to find the optimal P and t, and to minimise the χ^2, which we finally compared with the best fit for the actual data. We provided a sample of initial period estimates, starting from 2761.0 s, up to the maximum

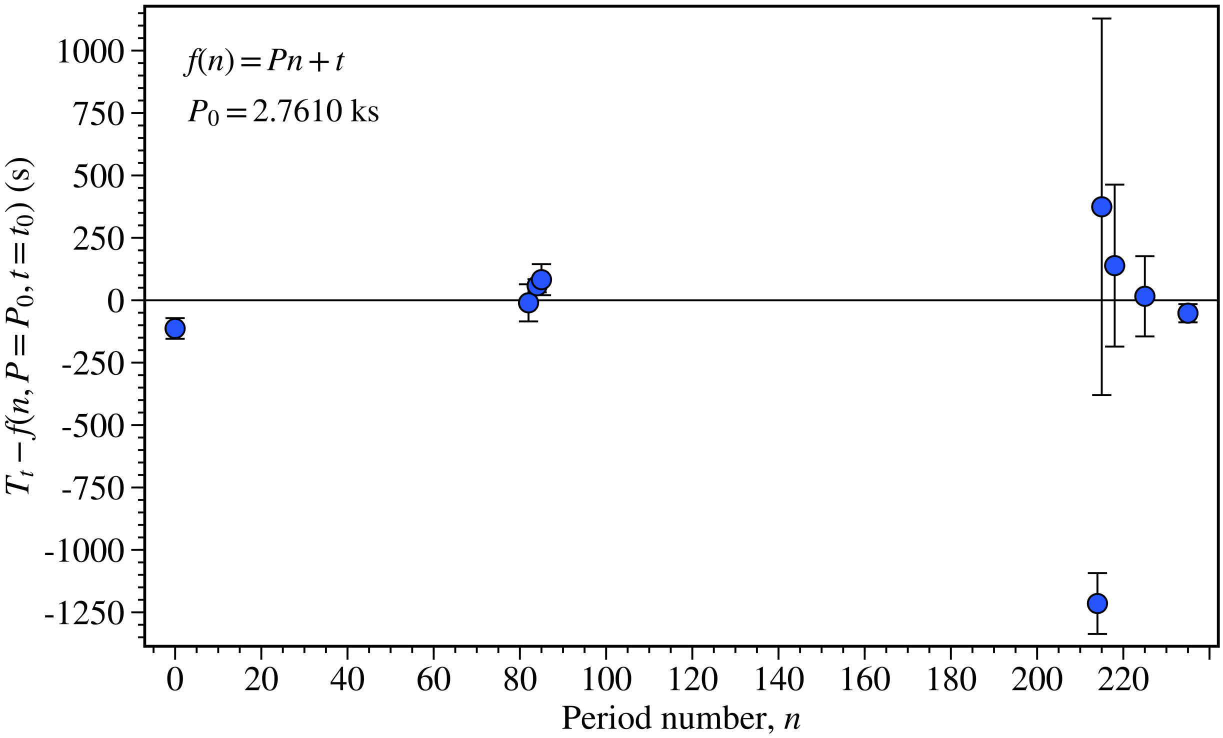

Figure 5.11: The difference between the best fit period and the observed time of the centre of each transition. We ignored the anomalous transition time in the fit, but showed it on the graph nevertheless. The error bars indicate the measured duration of the transition.

period that would allow for nine transitions to occur within our set of observations. The increment in the initial period estimates was chosen in such a way as to correspond to a decrement of one in the number of multiples of the period that could be placed between the start and the end of the early flip-flop interval. In doing so, we sampled the entire range of possible periods.

Every set of simulated transition times was fit with all of these initial period estimates, and we investigated how many simulations had transition times that could be fitted equally well, or better than the actual data, using the same procedure, with at least one period in the specified range. When running 10^5 iterations of this simulation, we found that only 3.38% of randomly generated sets of transition times fitted with a χ^2 of at most 16.5, to a fundamental period at least as large as 2761.0 s, when the worst fitting time was excluded from the fit.

We noticed that the single anomalous transition occurs very close to the halfway point between two of the $f(n)$. This raises the question of whether a period of about half of our first estimate is instead the actual recurrent timescale of the flip-flop transitions. Indeed, a period of $P = 1380.5 \pm 0.2$ s agrees well with all the early flip-flop transition times we observed, fitting with a χ^2 of 18.2 for 7 degrees of freedom. However, we found that in 10^4 simulations, the probability of obtaining at least as good of an agreement between nine randomly generated transition times, for a period at least as long as this one, is 18.2%. Despite the existence of one anomaly, there is a greater significance in the existence of an underlying period of 2761.0 s in the early flip-flops. It is possible that $P = 1380.5 \pm 0.2$ s is the actual underlying period defining the times of transition, but our data are insufficient to distinguish periodicity on these time scales from noise. We, therefore, rely on the agreement obtained for a period of 2761.0 s.

5.5.2 Variable timescale for early and late flip-flops

Next, we turned our attention to the late flip-flops, and the six transitions observed in them. A simple extrapolation of the best fit we had found for the early flip-flops did not fit the transition times of the late flip-flops. Additionally, two transition times were separated by about 2560 s, which suggests that the late flip-flops cannot be fitted well with a period of 2761.0 s and might require a different period. We investigated the possibility of a different linear relation linking the times of the late flip-flop transitions, via Equation 5.2, independent of our results for the early flip-flops. The best fit to these data was obtained with $P = 2610 \pm 20$ s and $t = 58205.165 \pm 0.003$ MJD, and has a goodness of fit of $\chi^2 = 42.5$, for 4 degrees of freedom. These best fit parameters were used to plot the vertical lines in Fig. 5.9. This best fit period underlying the late flip-flop transitions is about 5.5% smaller than the one found for the early flip-flops.

Using a similar simulation as before, after updating the parameters to match the observations of the late flip-flops, we found that for 5×10^5 iterations, the probability that six random transition times within our late flip-flop observations fit with a period at least as large as 2610 s, and a goodness of fit of at most 42.5, is about 7.37%. This higher probability is a consequence of the large χ^2 that we obtained for the best fit to the observed times of transition in this interval. The timing of transitions in the late flip-flops is less consistent than in the early flip-flops.

The detection of similar underlying periods in both intervals is intriguing. We investigated whether all transitions, in both the early and the late flip-flop intervals, could be fitted with one universal equation of the form:

$$f(n) = Pn + An^2 + t, \tag{5.3}$$

where P describes the initial period at $n = 0$; t describes the reference time, corresponding to the time of the transition at $n = 0$; A describes the gradual change in the period, and is equal to $A = \frac{1}{2}\Pi\dot{\Pi}$, where $\Pi(n) = P + 2nA$ is the period at n. We, however, do not measure the periods at any of the n, but rather the time between any two n, such as: $f(n) - f(n-1) = \int_{n-1}^{n} \Pi(n)dn = P + (2n-1)A$.

By setting appropriate initial estimates for P, of the best fitting period found for the early flip-flops, and for A, such that the 5.5% difference in the periods of the early and late flip-flops could be generated within the 32-day range of early to late flip-flop observations, we discovered a fit that achieves a good agreement for all transition times in both intervals, except for the one anomalous transition in the early flip-flops. Excluding that one, the best fit parameters were found to be $P = 2770.0 \pm 0.6$ s, $t = 58174.2963 \pm 0.0008$ MJD, and $A = (-8.70 \pm 0.06) \times 10^{-2}$ s, resulting in an initial rate of period decrease equal to $\dot{\Pi}(0) = (-6.28 \pm 0.04) \times 10^{-5}$. The best fit has 11 degrees of freedom and a goodness of fit of $\chi^2 = 52.6$. We indicate the difference between the measured transition times, and the closest value of Equation 5.3 to each transition time, using these best fitting parameters, in Fig. 5.12. The ability to fit both early and late flip-flop transitions could indicate a strong connection between them.

We once again investigated the significance of this detection, under the assumption that the early and late flip-flops should be considered as one phenomenon rather than two. To do so, we modified the above simulations to instead generate 15 random transition times within our observations of early and late flip-flops. We fitted these times with Equation 5.3 instead, and removed the worst fitting instance. We also had to specify a range of input values for A, which we limited to having an absolute value of less than 0.0870. After 10^4 iterations, we found that the probability of obtaining at least as good of an agreement in randomised data as we had found for our observations, is 0.23%, or 2.8σ.

This analysis is based on the idea that the early and late flip-flops are the product of the same process, which evolves over time, but is not seen for about two weeks. A complementary possibility is that the early and late flip-flops are completely independent realisations of the same process, which are possibly a consequence of the major flares observed in X-ray and radio.

We, therefore, also consider the probability of obtaining a linear fit from Equation 5.2 to both the early and late flip-flop intervals individually, in both instances fitting with a period of at least as large as the one we found in our observations, for equally many transitions as we observed, and having a goodness of fit smaller than the one we detected, and also having periods in both intervals that are separated by less than 5.5%. Using the results

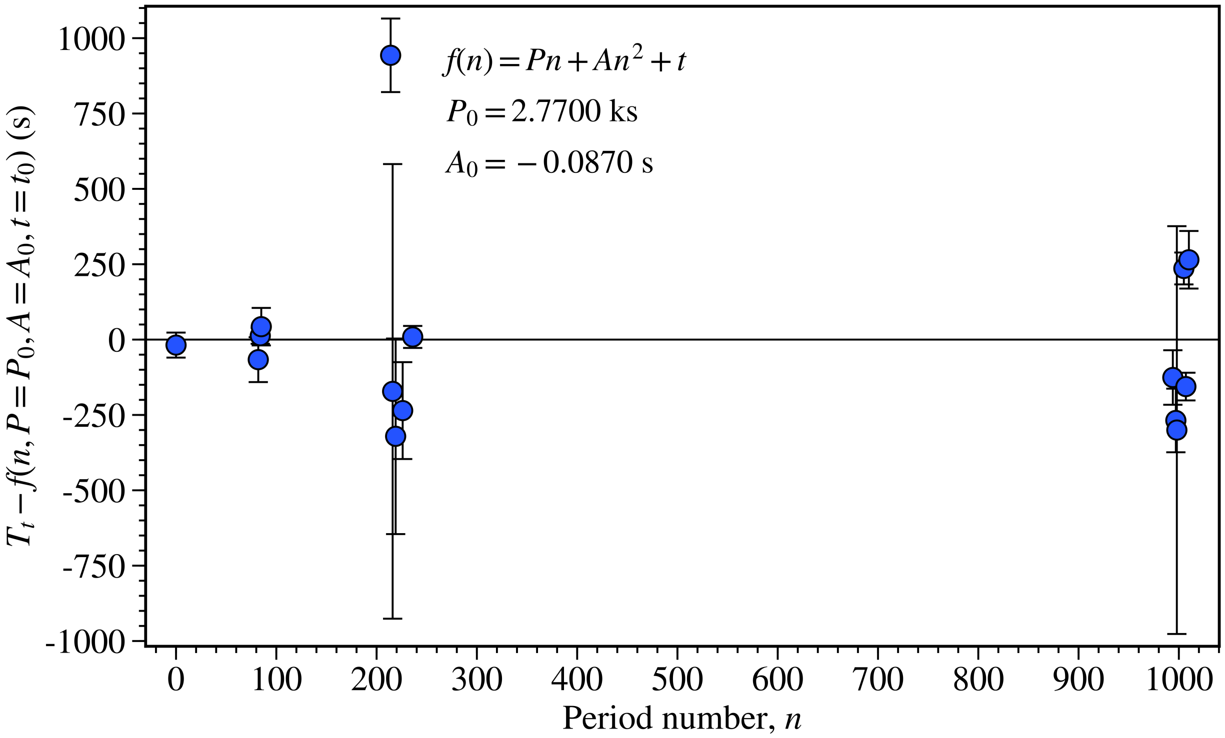

Figure 5.12: The best fit of Equation 5.3 to the times of transition in both the early and late flip-flops. The transition times cannot be fitted with Equation 5.2, assuming a constant, and not too short period. Nevertheless, by adding an additional term to allow for a change in the period, all times of transition could be fitted reasonably well. The anomalous transition time was not used in the fit.

of our simulations for a constant period in the early and late flip-flops described above, we calculated the probability that any randomly chosen combination of an early and late flip-flop period, which managed to meet our fitting criteria applied to simulated data, had a separation of at most 5.5%. This probability was found to be about 26%. Therefore, the probability of randomly generated, and independent early and late flip-flop transition times being able to be fitted with a constant period in each interval, which is greater than or equal to the one we found in the data, which fits at least as well as our measurements did, and have periods separated by at most 5.5%, is $P = 0.0338 \times 0.0737 \times 0.26 = 6.48 \times 10^{-2}\%$, or 3.2σ. This result is based on 10^5 simulations for the early flip-flops, and 5×10^5 simulations for the late flip-flops. In these simulations, we excluded the worst fitting transition to ensure that the data and simulations were handled in the same way. We used the exact same procedure on both the real data and all simulated data. These results, therefore, suggest that this underlying period is significant.

This entire analysis is based on the 15 flip-flop transitions that we observed. There were two more instances when a flip-flop transition was inferred to have occurred in a short gap in the observations, on MJD 58181.8, and MJD 58205.5. We note that in both instances, there was one $f(n)$ time that occurred within the gap and might have coincided with the unobserved transition. Neither of the two unobserved flip-flops contradicts the detection of a period in either the early or the late flip-flop interval. We did not use the gaps as additional measurements of times of flip-flop transition, as the associated error in their measurements is so large that they did not aid the investigation of possible periods.

5.6 Periodogram and QPO analysis

To learn more about the flip-flop transitions, we generated periodograms of individual flip-flop states, and entire observations without flip-flop activity, from the light curves of *Astrosat*, *NuSTAR*, and *XMM-Newton*. In Fig. 5.13 we plotted the periodograms of all the flip-flop states observed by *Astrosat*, as well as the periodogram of the *Astrosat* observation in the HIMS, for comparison.

During the initial rise in flux in the outburst, we found a narrow QPO perched on a strong broad-band continuum featuring a second harmonic, and no subharmonics. It has properties indicative of a type C QPO (Xu et al., 2018). We noticed that the frequency of this QPO increased alongside the flux, rising from 0.13 Hz at the start of the first observation by *NuSTAR*, to 1.9 Hz by the end of the first *Astrosat* observation. This occurred alongside a flux rise from 3.9×10^{-10} erg cm^{-2} s^{-1} to 1.1×10^{-9} erg cm^{-2} s^{-1} in the energy range 3–79 keV.

5.6.1 Flip-flop QPOs

We analysed the periodograms of the flip-flop states individually and noticed a duality of properties, matching the duality of the flux levels. All bright states have similar periodograms, and so do all dim states. However, the periodograms of the bright states differ

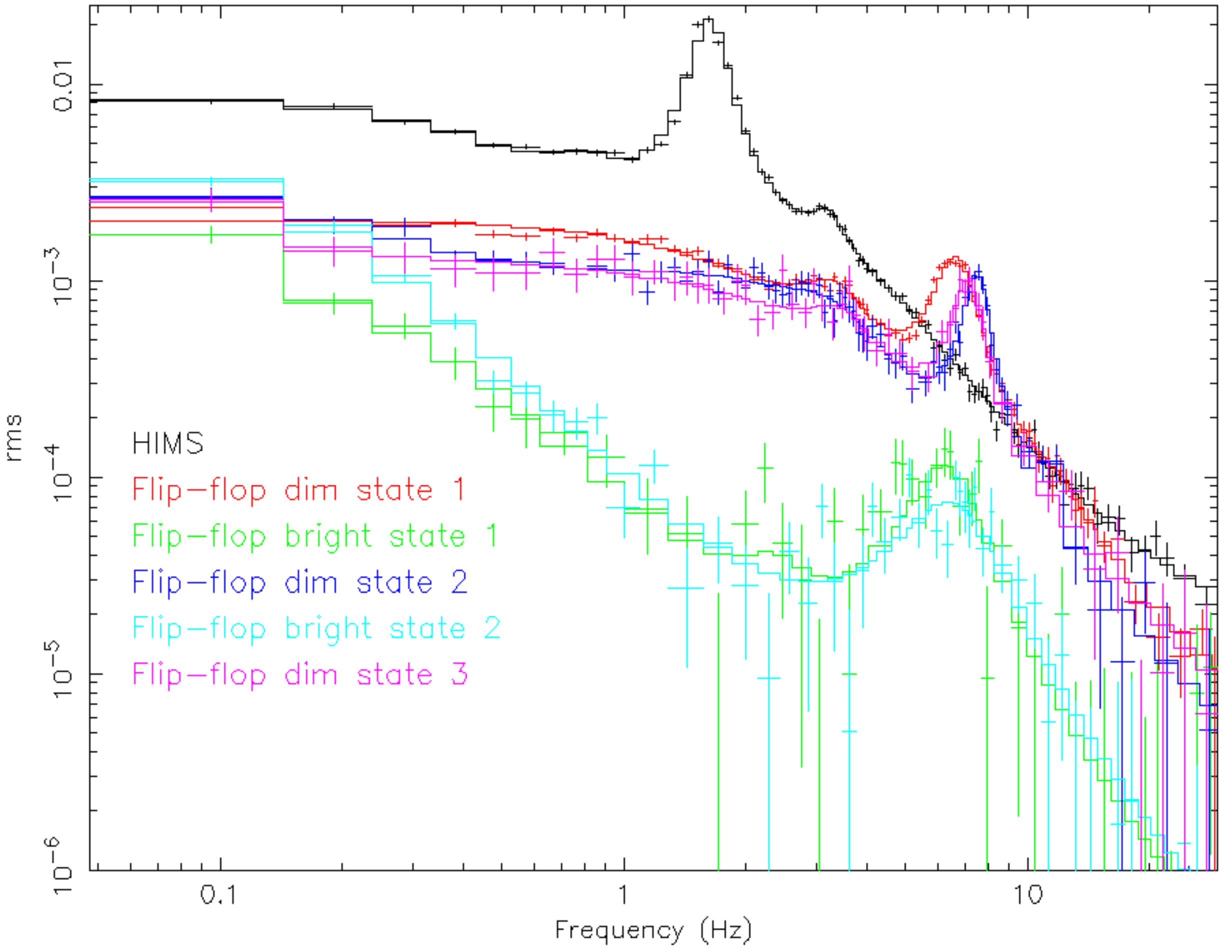

Figure 5.13: Poisson noise subtracted, squared fractional rms normalised *Astrosat* periodograms of different intervals of the outburst. The first *Astrosat* observation in the LHS is marked in black. All other periodograms originate from different regions of the second *Astrosat* observation, which contains flip-flops (See Fig. 5.14). Both the LHS and Flip-flop dim state periodograms have inflated widths due to being generated from a large time interval containing a variable type C QPO frequency. Generating the periodograms of shorter intervals yields a narrower QPO. The same is not true for the type A QPOs in the bright states.

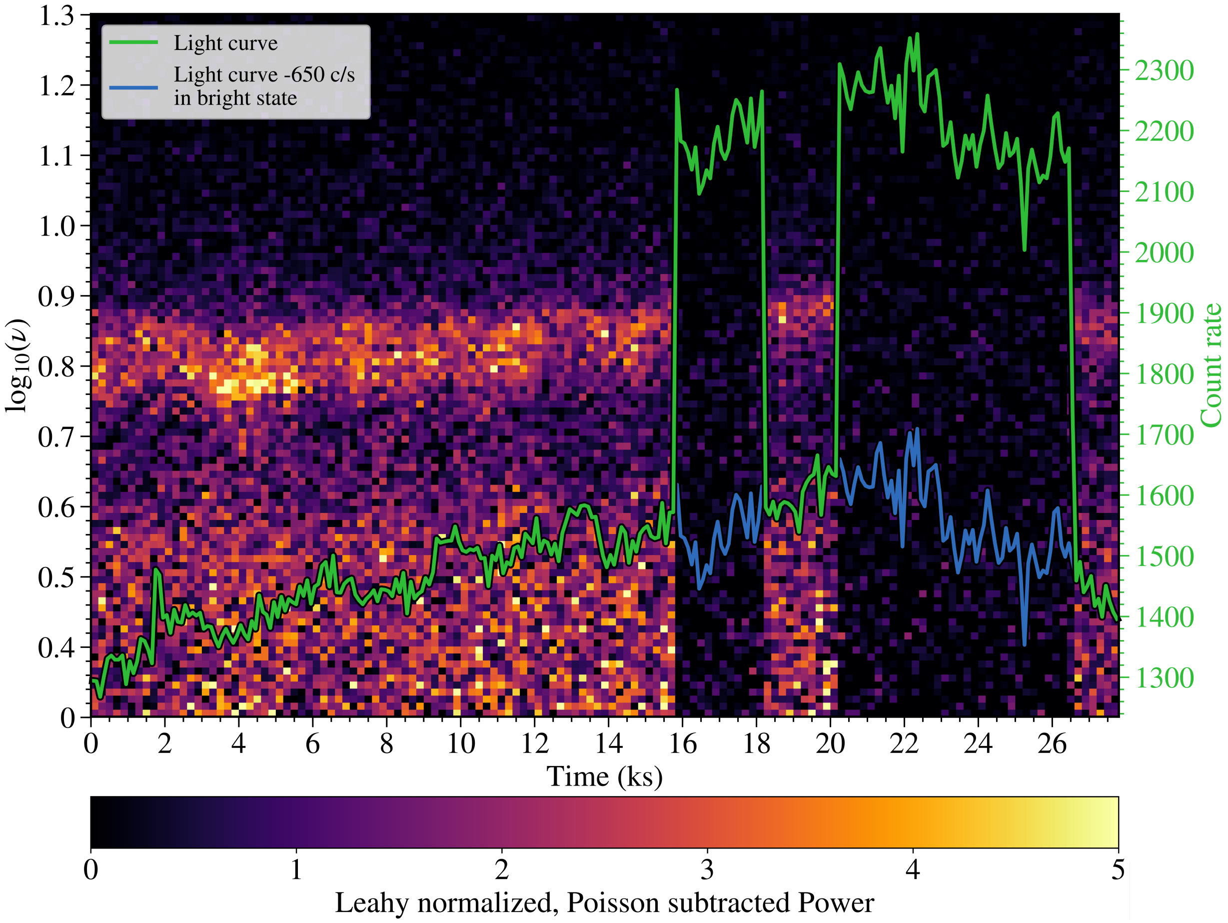

Figure 5.14: Spectrogram of the second *Astrosat* observation, starting on MJD 58180.786, with an overplotted light curve. Colours denote the Leahy normalised, and Poisson noise subtracted power. The light curve is plotted in green. The blue curve represents the same light curve but with a count rate reduced by 650 cts/s in the bright states. Gaps in the observations due to the low Earth orbit of *Astrosat* have been left out for display clarity. The time is measured in seconds of observing time since the start of the observation. The light curve is split into 200 s segments, for which the average count rates and the periodograms were computed. The periodogram frequencies are rebinned on a logarithmic scale. The first three state transitions shown here all lie within gaps in the data. However, we see the entire fourth transition near the end of the observation.

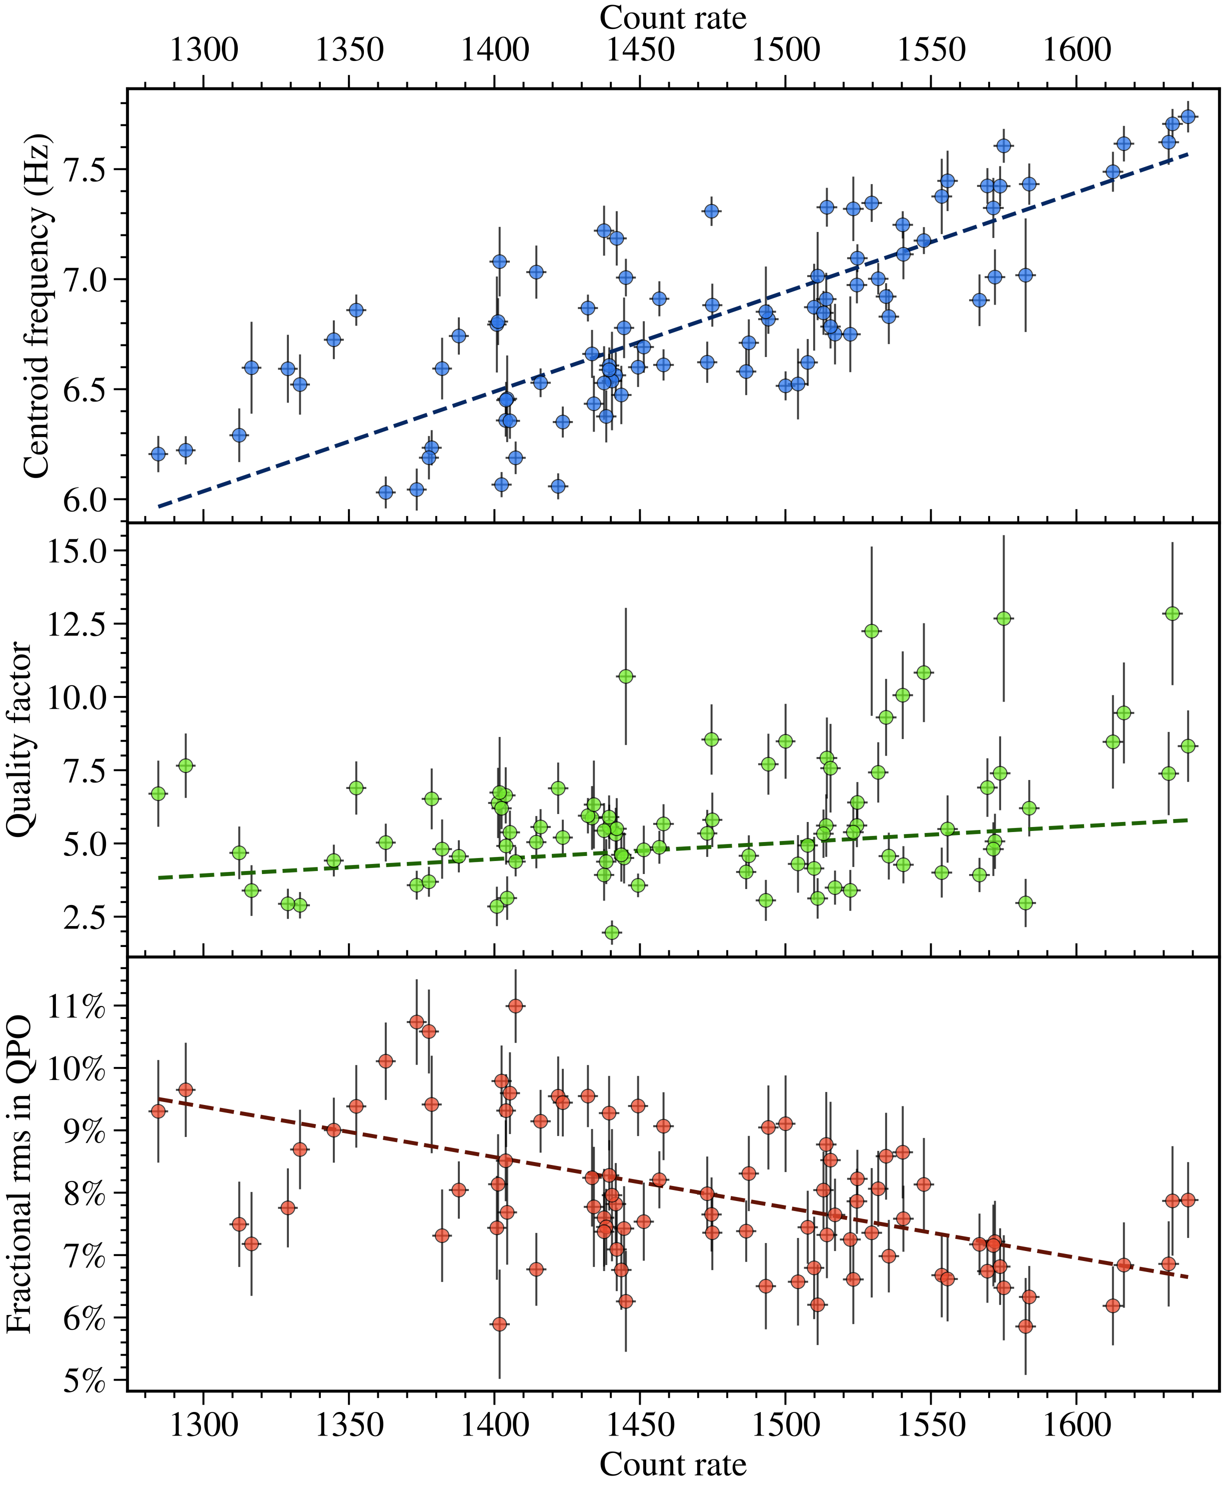

Figure 5.15: Correlation between the best QPO fitting parameters and the count rate, in each of the 200 s intervals of Figure 5.14. We fitted each periodogram between 4 and 10 Hz with a Lorentzian added to a straight line with a constant gradient. The properties of the Lorentzian were extracted to determine the centroid frequency, quality factor, and fractional rms in the QPO.

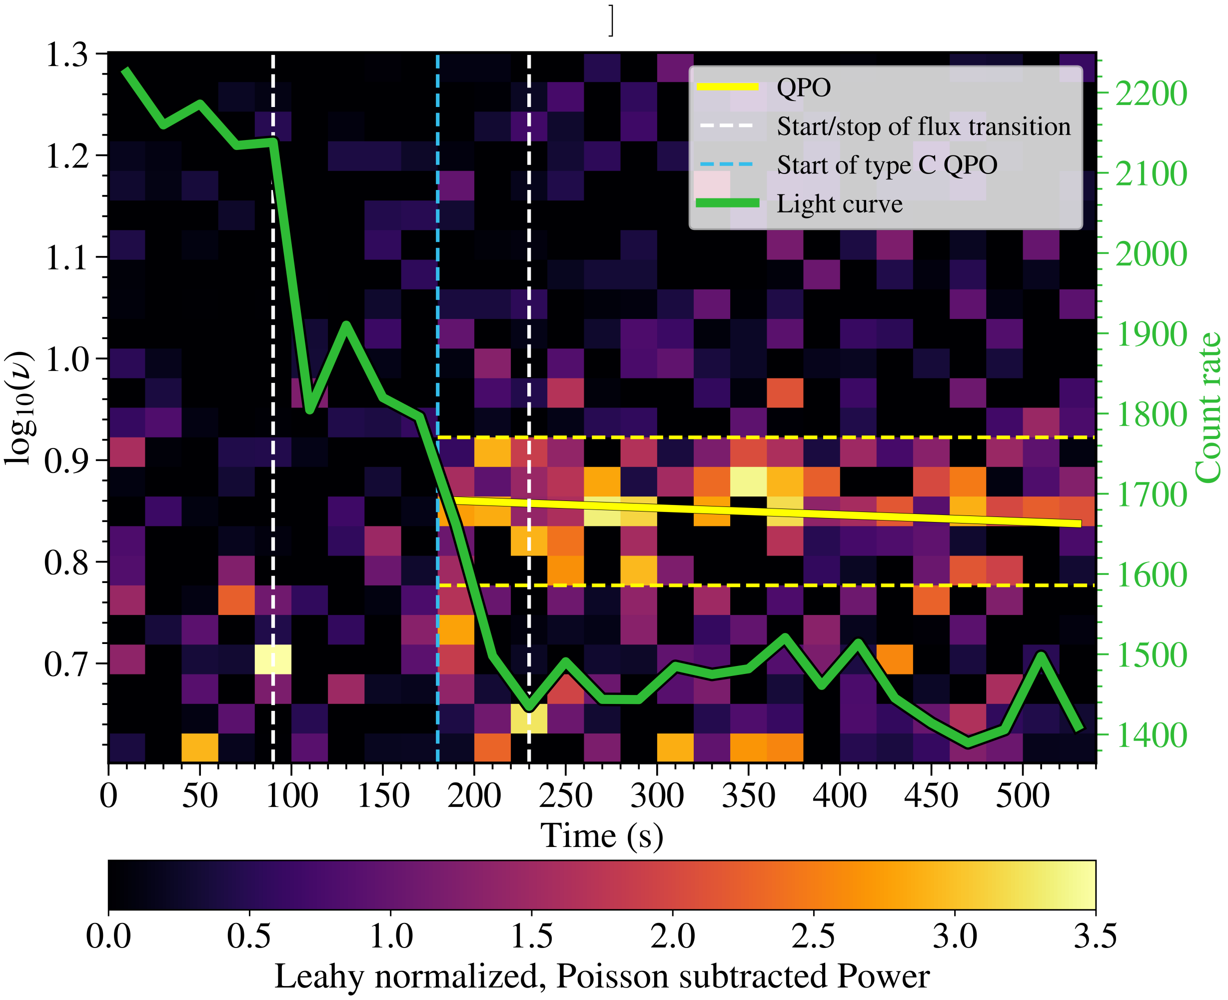

Figure 5.16: Spectrogram of the last flip-flop transition shown in Fig. 5.14, as seen by *Astrosat*. Colours denote the Leahy normalised, Poisson noise subtracted power. The light curve is plotted in green. The time is measured in seconds of observing time since the start of the orbit. We computed the periodogram, and the average count rate, for each of the 20 s intervals into which we divided the data. We plotted a smaller frequency range for this figure due to significant noise at low frequencies, a consequence of the short interval size.

remarkably from those of the dim states (see Fig. 5.13).

The dim flip-flop states detected by *Astrosat*, *NuSTAR*, and *XMM-Newton* all contained a reasonably narrow ($4.7 \leq Q \leq 17$) QPO at frequencies of 6–7.5 Hz (see the purple, red, and dark blue periodograms in Fig. 5.13). They feature a sub-harmonic, but no higher order harmonic. The rms in the QPO is 2.5–8.2%. Together with a strong, flat, broad-band continuum with a break frequency at about 3 Hz, the total rms of the dim states between 0.5 Hz and 50 Hz was found to be 6.4–10.3%. The large effective area of *Astrosat* enabled a more detailed examination of the change in QPO properties on shorter timescales. We noticed that the centroid frequency of the QPO in the dim states of the second *Astrosat* observation scales with the count rate (Fig. 5.14 and 5.15), with $r_{xy} = 0.78$ and $p = 7.3 \times 10^{-19}$. We additionally found that the quality factor increases slightly with increasing count rate ($r_{xy} = 0.38$, $p = 3.2 \times 10^{-4}$). Moreover, the rms in the QPO was found to decrease with increasing count rate ($r_{xy} = -0.51$, $p = 3.8 \times 10^{-7}$), thereby indicating that the rms also correlates negatively with the centroid frequency ($r_{xy} = -0.64$, $p = 2.0 \times 10^{-11}$).

By comparing these results with the distinguishing characteristics of the three different QPO types (Motta, 2016), we found that almost all of these properties suggest that this is a type C QPO. However, this QPO has a large range of quality factors ($4.7 \leq Q \leq 17$) and is usually found at $Q < 10$, which is too wide to fit the standard definition of a type C QPO (Motta, 2016). However, Type C QPOs with smaller quality factors than stated by the standard definition have been seen before, by Motta et al. (2012). This suggests that the lower limit of the quality factor of type C QPOs is ill-defined. Furthermore, the shape of the continuum, the large rms, and the correlations between count rate, centroid frequency, and rms, are all inconsistent with type B and A QPOs.

Whereas the type C QPO of the dim states can easily be seen in the spectrograms of Fig. 5.14 and 5.16, no QPO can be identified during the bright states by eye. So the QPO seems to disappear entirely in dim to bright transitions. This is, however, a consequence of the low amplitude and broad profile of the QPO in the bright state and the short time sampling used in these figures. Generating periodograms over longer time intervals instead revealed the existence of a very low amplitude, broad peak with a centroid frequency of about 5.0–7.2 Hz, a quality factor of 2.0–3.8, and an rms of 1.4–2.3% (see Fig. 5.13). Fig. 5.14 also shows that the broad-band continuum at low frequencies is strongly suppressed in bright states. However, Fig. 5.13, ,indicates that the bright state continuum grows at low frequencies, such that bright and dim states have approximately equal power at 0.1 Hz. The continuum contributes slightly to the total integrated rms in bright states, which is about 2.0–2.7%. These broad, low amplitude QPOs are very difficult to detect and could only be seen in the *Astrosat* periodograms. This QPO type is undoubtedly a type A QPO, because all of these fitting parameters and the shape of the periodogram agree with the standard properties of a type A QPO (N03). We did not detect the type A QPO in the periodograms of *NuSTAR* and *XMM-Newton*. Nevertheless, we found that the shape of their periodograms, and the low rms in the bright states, were consistent with the *Astrosat* result.

Following Lewin et al. (1988), we determined the detection significance of the type A

QPOs. We split the data into individual orbits and only kept data firmly belonging to the bright flip-flop state. We determined the QPO detection significance for each of these intervals and found it to be between $3.2 - 11.6\ \sigma$. The lowest value was found for a very brief interval. Therefore, we conclude that all bright states observed by *Astrosat* featured a type A QPO. Through the similarity in their properties, we infer that a type A QPO was also present in all other bright states of the early flip-flops.

Thanks to the large effective area and short time resolution of *Astrosat*, we were able to compute accurate periodograms over short time intervals as well, enabling the creation of the spectrograms shown in Fig. 5.14, and 5.16. In Fig. 5.16, we zoom into the first-ever detection of a direct transition from a type A to a type C QPO. Other groups (C04, Del Santo et al. 2008) have detected a type C QPO in one observation, and a type A in the next. However, they could not rule out the possibility that the interval between the observations contained a transitional type B QPO. We, however, see multiple direct transitions between types A and C, and can, therefore, determine whether the change between the two QPO types involved a type B QPO. If a type B QPO had occurred during the transition, it would have featured as a very bright spike in the spectrograms, with next to no broad-band continuum, and would presumably have been found at a slightly different frequency than the type C QPO. Visually, we do not see any features matching this description in Fig. 5.16. We also fitted the periodograms during the transition, to search for a type B QPO, but did not find any. We even used intervals as short as 10 s without detecting any type B QPO. The search for QPOs in short time segments becomes increasingly difficult because many cycles are needed to crystallise out a QPO from a strong continuum, and noise. So even though we cannot rule it out entirely, we can rule out the existence of a type B QPO on any timescale longer than about 10 s during this transition.

We note that the change in the periodogram is very sudden. In Fig. 5.16, we see that the type C QPO took less than 20 s to appear. Using smaller time intervals indicates that it took at most ≈ 10 s to get started. There was no indication of the impending change of the QPO in the periodogram before the transition. This is significantly shorter than the flux transition, which took ≈ 140 s instead. Another interesting feature is that the QPO suddenly appears about 100 s after the start of the transition in the light curve, and ≈ 40 s before the flux has finished dropping down to the dim state level.

In Fig. 5.17, we plot the relation between the *Astrosat* total integrated rms and the QPO frequency. In the inset, we plot the deadtime corrected rms in the QPO, in observations by *NuSTAR*, and *XMM-Newton*. Comparing this with the rms as a function of ν plots of other transient BHTs (see e.g. Motta et al., 2011), we can distinguish QPO types, which fall into different regions of this graph. We see the typical negative gradient relation between the rms and centroid frequency of the type C QPOs ($r_{xy} = -0.997$, $p = 1.0 \times 10^{-16}$ for the *Astrosat* observations, and $r_{xy} = -0.91$, $p = 1.0 \times 10^{-20}$ for the combined *NuSTAR*, and *XMM-Newton* observations), and the independence of the rms and centroid frequency of type A QPOs ($r_{xy} = 0.23$, $p = 0.61$). We indicate the region where we would expect the type B QPOs to appear, based on the observations of other BHTs. Compared to similar plots by Motta et al. (2011), we noticed a lack of type B QPOs in this outburst. This also indicates that we only observed QPO types A and C, but no type B QPO.

We also tested whether the change in the timing properties could be explained by the addition of a separate flux component which dilutes the variation seen in the periodogram. We added a component to the dim state *Astrosat* light curve, whose value in every time bin is given by a Poisson distribution with a mean value set equal to the mean counts per bin of this *Astrosat* dim state. This simulates a 100% flux increase by a separate feature with purely Poissonian variability. As expected, this hardly affected the overall shape of the periodogram or the QPO properties. We retained the same broad-band continuum and comparatively narrow QPO. Hence, the increase in flux is not directly responsible for the change in the QPO. Therefore, the source of the type C QPO and the broad-band continuum must either be disrupted or blocked from being viewed, to obtain the observed changes to the periodogram.

In the time between the early and late flip-flops, we once again detected type C QPOs. These were found at frequencies of 3.5–6.1 Hz but were stronger (with an rms in the QPO of $\approx 10\%$), and narrower ($Q \approx 10$) than those observed in the dim flip-flop states.

We did not find any QPOs in the late flip-flops, but we did detect a change to the periodogram in these transitions. *NuSTAR* suffers from a significant and variable deadtime of ≈ 2.5 ms, whose effect can, however, be mostly corrected for by cross-correlating the two light curves to generate a cospectrum (Bachetti et al., 2015). The total integrated rms obtained by this method differs slightly from the one obtained from an ideal non-deadtime affected periodogram. However, substantial changes to the rms are nonetheless significant. Despite the non-detection of a QPO, we still determined the total integrated rms of the *NuSTAR* cospectra in the bright and dim states of the late flip-flops, between 0.01 and 20 Hz. It was found to be about 8–12% in the dim states and 6% in the bright states. We, therefore, still observe a significant change in the periodogram in late flip-flop transitions, which cannot be obtained merely by diluting variations due to the higher count rate of the bright states. The value of the rms in the two states and their position in the HID seems to suggest that the late flip-flops correspond to transitions between the HIMS and the SIMS. However, the bright state, with the lower rms, was found at a greater hardness than the dim state.

The final observation by *NuSTAR* in the SIMS (see Fig. 5.10) also featured some changes in flux. We also did not detect any QPO in this observation. The orbit with the lowest flux level had an rms of 7%, and the orbit with the highest flux had an rms of 10%. So there is some indication of changes to the periodogram in this observation as well. However, the difference is smaller and less significant than the rms differences found in any of the early or late flip-flops. No HFQPOs were detected at any point in the outburst.

5.7 Energy spectra

5.7.1 Dust scattering halo and energy shifts

In Fig. 5.18, 5.19, and 5.20, we plot the *NuSTAR* and *XMM-Newton* spectra of a simultaneous set of intervals from the long dim state observed on MJD 58174. We only include

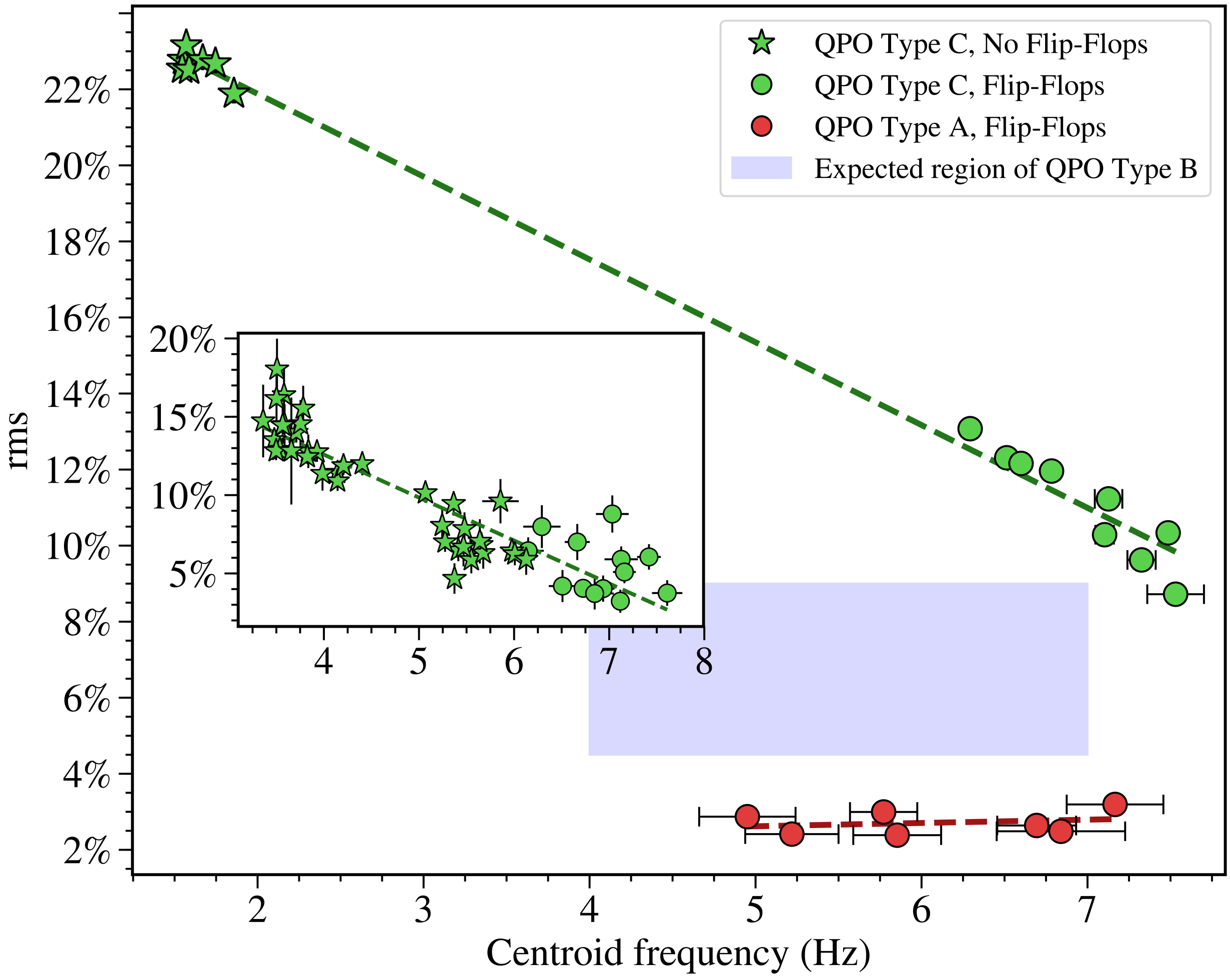

Figure 5.17: Here, we plot the total integrated rms and QPO centroid frequency of each orbit from the two *Astrosat* observations. The inset shows the centroid frequency and the corrected rms in the QPO, detected in individual orbits of *NuSTAR*, and in 2 ks regions of the *XMM-Newton* light curves. Total rms values are computed in the 0.5–50 Hz range. We use colours and shapes to distinguish the type of QPO and whether the measurement was made during the flip-flop period, or outside of it.

times when both instruments were observing Swift J1658.2-4242. As Fig. 5.18 shows, these spectra differ quite a lot, especially at low energies. Therefore, it is not possible to fit them both using the same parameters unless only data at energies exceeding 5.5 keV is considered. This is why we restricted the medium energy X-ray light curve of Fig. 5.1, 5.4, 5.5, 5.6, 5.9, 5.10, 5.7, and 5.8 to the range of 5.5–10 keV. The difference at low energies is caused by the DSH, combined with distinct source extraction methods and regions.

Using the DSH model we developed for the *XMM-Newton* pn CCD in Timing mode, and the standard DSH model by Jin et al. (2017) and Jin et al. (2019) for *NuSTAR*, we managed to remove almost the entire difference between the *NuSTAR* and *XMM-Newton* spectra at low energies, as is shown in Fig. 5.19. The remaining difference between the two spectra at low energies could be due to background contamination or indicate the limit of the usefulness of our DSH model at large hydrogen column densities.

Xu et al. (2018) and Xu et al. (2019) describe an absorption line at an energy of 7.03 keV, very close to the iron K-edge at 7.112 keV. This line can be seen in Fig. 5.19, and might be the signature of an accretion disc wind (Ponti et al., 2012a, 2016; Díaz Trigo and Boirin, 2016). However, the simultaneous *XMM-Newton* spectra feature an emission rather than an absorption line at this energy. There is a contradiction here which needs to be resolved. As the absorption and emission lines are very close to the iron K-edge, we investigated the possibility that a wrong energy calibration at ≈ 7 keV could produce a fake emission or absorption line. By adding a non-zero redshift component to describe a slight shift in energy, having different values for *NuSTAR* and *XMM-Newton*, we achieved a good agreement between the two spectra at these energies.

Therefore, either the *NuSTAR*, the *XMM-Newton* spectra, or both, require an energy recalibration. After the end of the late flip-flop region, we examined the *Chandra* HETG spectrum, which was obtained simultaneously to the last *NuSTAR* observation. As the *Chandra* spectrum is also strongly affected by the DSH at low energies, we restricted the fit to 5.5–10 keV to ensure that the DSH correction would not bias our results. The *NuSTAR* spectrum contains a strong absorption line at ≈ 7 keV, as in all other observations. We fitted the possible absorption or emission line by including the additive `gauss` model in the spectral fit. For *NuSTAR*, this possible absorption line was best fitted with a centroid energy of $7.224^{+0.099}_{-0.056}$ keV, a variance of $(1.47 \pm 0.85) \times 10^{-1}$ keV, and a normalisation of $\left(-2.23^{+0.75}_{-0.53}\right) \times 10^{-4}$ photons cm^{-2} s^{-1}, with the negative sign indicating that this is an absorption line. This line has an associated equivalent width of $-28.5^{+9.6}_{-6.8}$ eV. However, the *Chandra* spectrum was best fitted without the inclusion of an emission or absorption line at this energy. When the additive Gaussian component was included, and the centroid energy allowed to vary, it did not fit anywhere near the value found by *NuSTAR*. When forcing the centroid energy and variance of the Gaussian included in the *Chandra* fit to be equal to their best fitted values in the *NuSTAR* spectrum, we found the best fit Gaussian normalisation to be: $\left(1.9^{+3.2}_{-2.9}\right) \times 10^{-4}$ photons cm^{-2} s^{-1}, corresponding to an equivalent width of $\left(2.4^{+4.0}_{-3.7}\right) \times 10^{-2}$ keV. This does not support the existence of either an absorption or an emission line. Therefore, due to the higher resolution and better calibration of *Chandra* spectra, we decided to use an energy shift in the *NuSTAR* and *XMM-Newton*

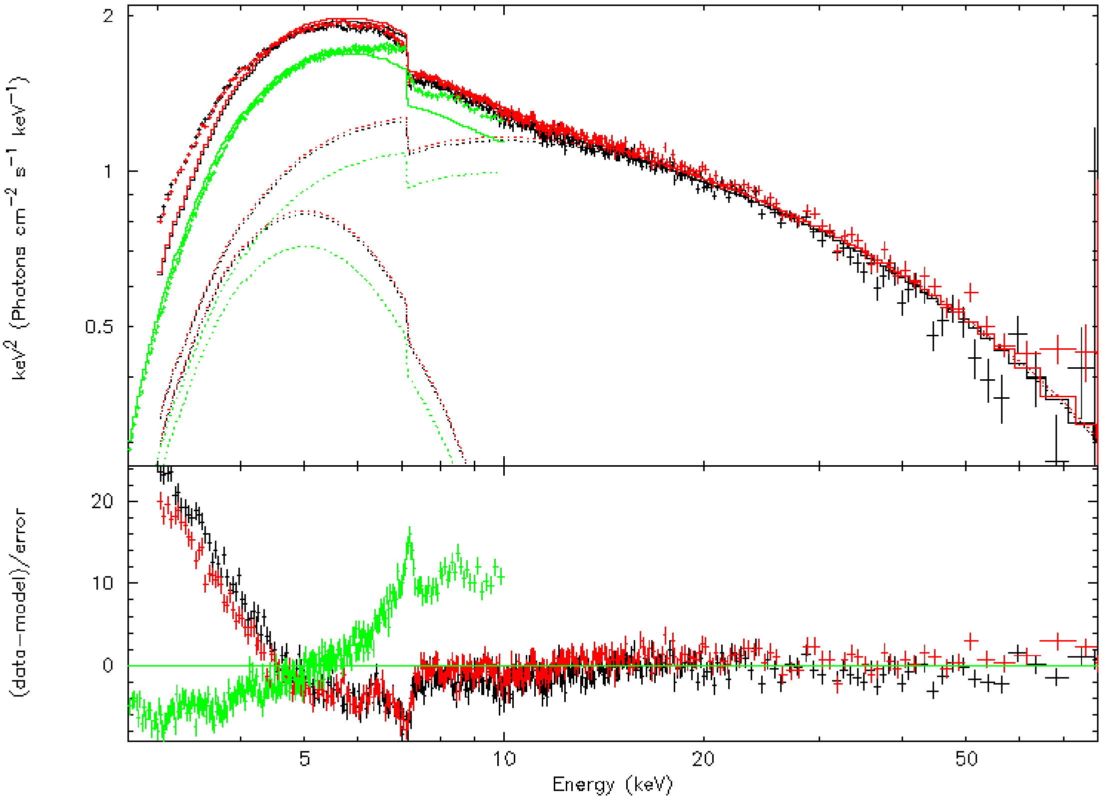

Figure 5.18: Simultaneous *NuSTAR* (black for FPMA, red for FPMB) and *XMM-Newton* (green) spectra of the long dim state observed on MJD 58174, without a DSH, or an energy scale correction.

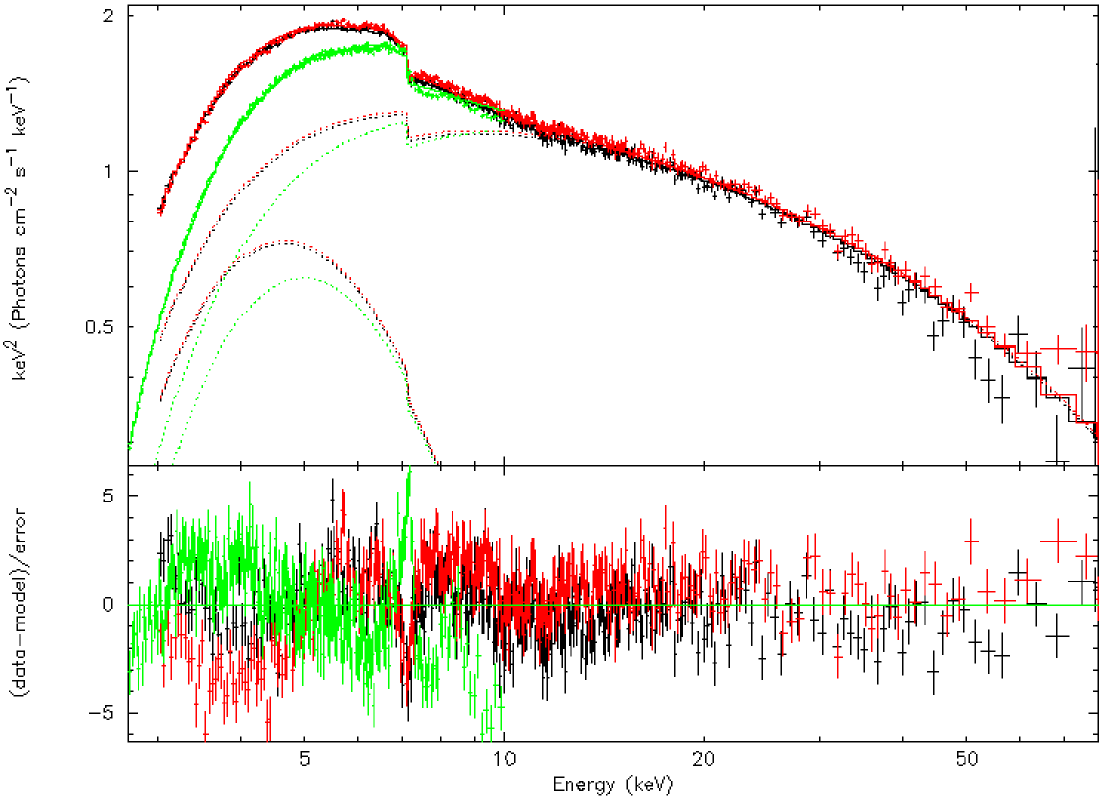

Figure 5.19: Simultaneous *NuSTAR* (black for FPMA, red for FPMB) and *XMM-Newton* (green) spectra of the long dim state observed on MJD 58174, with a DSH, but without an energy scale correction.

spectra to ensure their consistency with the non-detection of an emission or absorption line at ≈ 7 keV by *Chandra*.

We obtained further support for this idea when fitting for the energy of the iron K-edge. We not only found that the best fit energy of the edge differs from the value it should have, of 7.112 keV, but that in using the best fit edge energy, we no longer see absorption or emission lines at these energies.

To measure the energy shift of the iron K-edge, we fitted *NuSTAR* and *XMM-Newton* spectra between 5.5–9 keV using the combination of XSPEC models: `edge*diskbb`. By comparing the fitted edge energy with the value it should have, of 7.112 keV, we can determine a redshift which can be implemented into a model for the entire spectrum, to correct for the energy offset observed at the iron K-edge, and ensure consistency in the spectra of *NuSTAR*, *XMM-Newton*, and *Chandra*.

The best fit results of the shift in iron K-edge energy from its expected value, in the *NuSTAR* and *XMM-Newton* spectra, are shown in Fig. 5.21. As expected, *NuSTAR* spectra require a negative redshift energy correction, and *XMM-Newton* spectra require a positive redshift energy correction. When applying these energy corrections via a redshift component, we obtain the spectra shown in Figure 5.20. We still retain slight differences

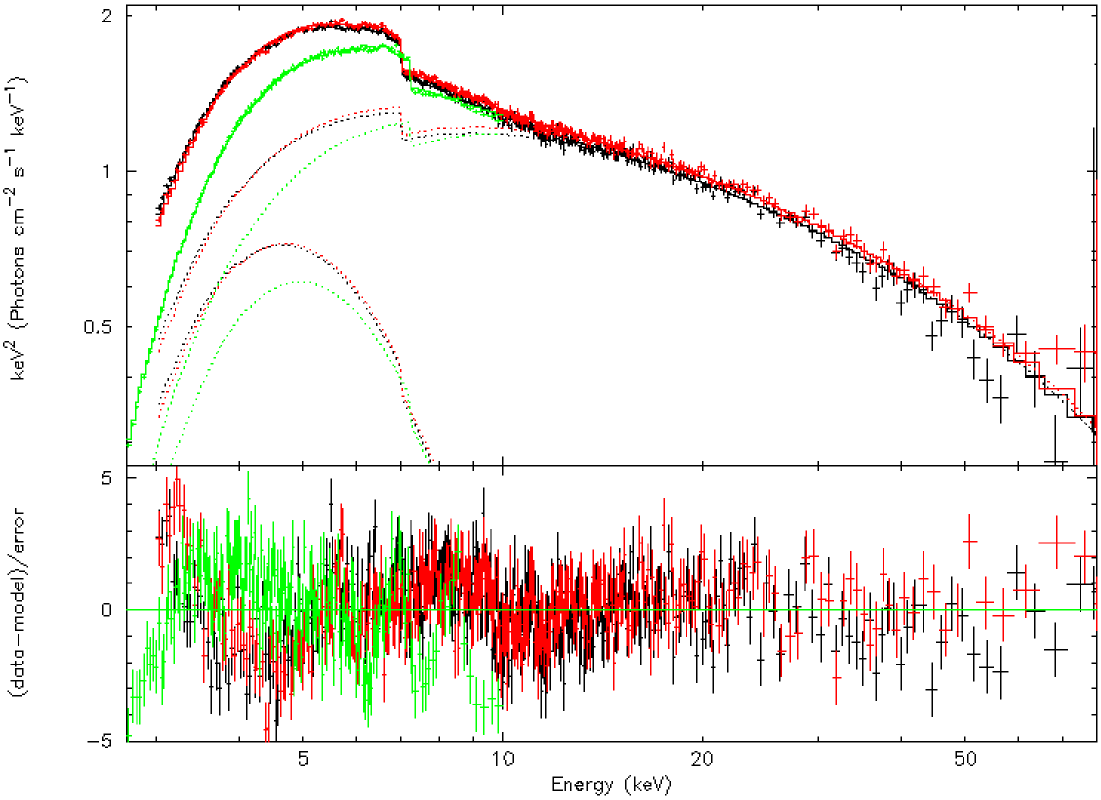

Figure 5.20: Simultaneous *NuSTAR* (black for FPMA, red for FPMB) and *XMM-Newton* (green) spectra of the long dim state observed on MJD 58174, with both a DSH, and an energy scale correction.

between the spectra, but they are now undoubtedly more consistent than without the DSH correction and the energy calibration correction.

5.7.2 Spectral analysis

Comparing the spectra of the bright, dim, and late flip-flop states, we found that they are remarkably similar (see Fig. 5.22). Visually, besides having a different normalisation, these spectra are more similar to each other than to either of the three non-flip-flop spectra shown in this figure. The difference between the bright and dim flip-flop spectra is larger at higher energies. The hardness ratio, therefore, increases simultaneously with the flux, as can be seen in Fig. 5.4, 5.5, 5.6, and 5.9.

The initial LHS spectrum differs noticeably from the other spectra we measured (see Fig. 5.22). However, we also found that the spectrum in the interval between the two sets of flip-flops, and the spectrum after the completion of the late flip-flops, are visually somewhat similar, besides having a different normalisation. This would indicate that the spectra themselves distinguish the flip-flop intervals from other parts of the outburst. This behaviour is reminiscent of state transitions, which typically are observed to occur at a

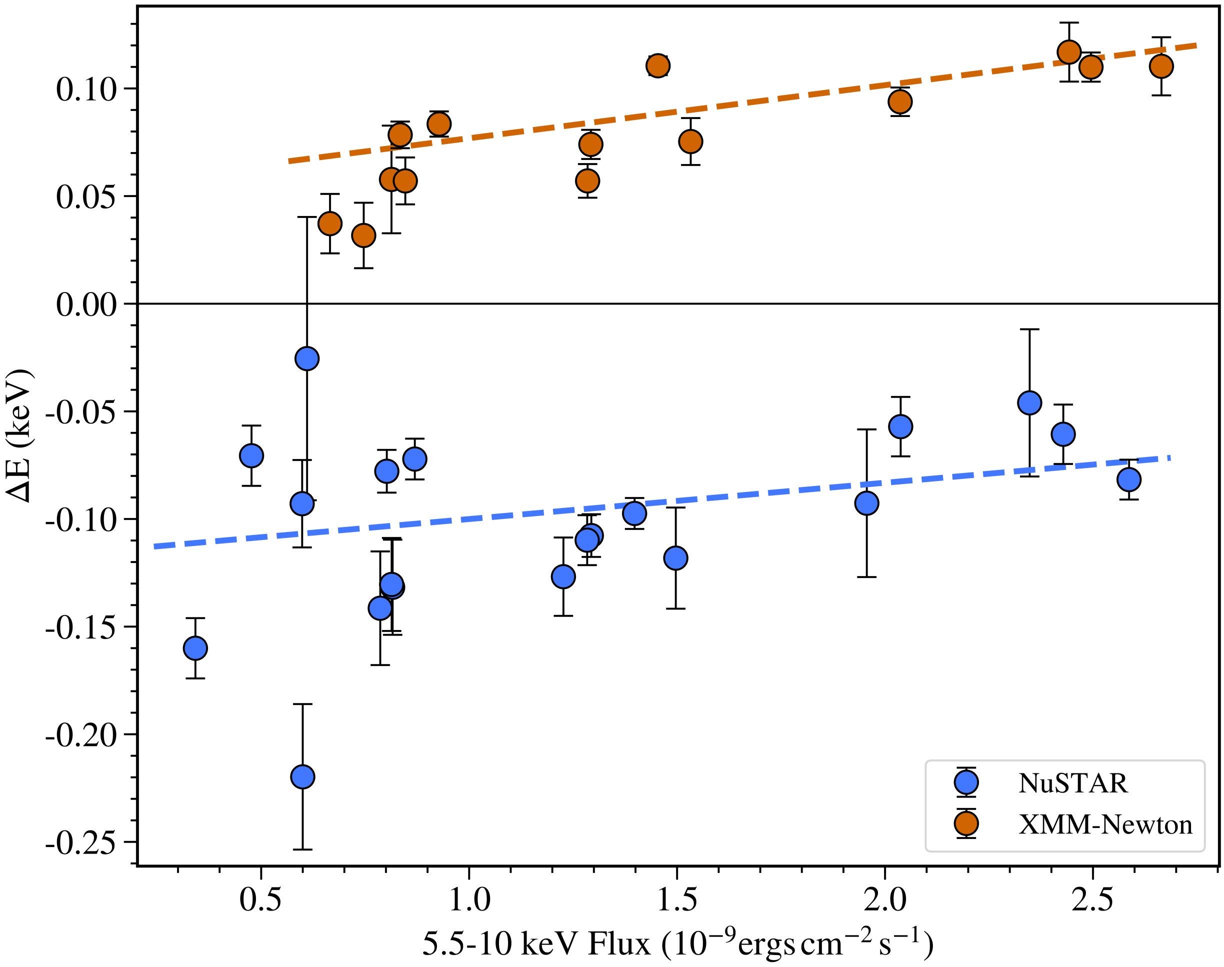

Figure 5.21: The difference between the expected iron K-edge energy of 7.112 keV, and the best fitted edge energy, in *NuSTAR* and *XMM-Newton* spectra. The energy shift was determined by fitting a simple `zedge*diskbb` model to spectra confined to the energy range of 5.5–9 keV.

very similar hardness.

To better understand the differences between the bright and dim states of the flip-flop, and what distinguishes the flip-flop intervals from other parts of the outburst, we fit the spectra of individual flip-flop states and entire observations, which did not feature any flip-flops. We used two sets of XSPEC models in the spectral fitting to obtain two independent estimates of various parameters, which we describe below.

In Model 1, we fit the spectral energy contribution from the multicolour black body component of the accretion disc, using the model `diskbb`. The high energy power law component is represented by the `cutoffpl` model, which includes a cutoff energy, beyond which the power law becomes steeper. We also add a `diskline` component to describe the iron K_α line at 6.4 keV. The effect of absorption by interstellar dust is modelled through the multiplicative `ztbabs` component. We shift the energy of the spectrum to ensure the iron K-edge is located at 7.112 keV by freezing the redshift at the energy shift found previously for each spectrum. We correct for the effects of the DSH through the multiplicative model `dscor`. We also used the `constant` multiplicative model to apply a cross-calibration correction constant between the two focal plane modules of *NuSTAR*. Combining all these elements, our **Model 1** in XSPEC jargon is: `dscor*ztbabs*constant*(diskbb+diskline+cutoffpl)`.

The inclination of the accretion disc around the black hole is unknown. However, the dips seen at the start of the outburst, and the shape of the HID (compare Fig. 5.2 with the HIDs of Muñoz-Darias et al. 2013), suggest that the accretion disc around the black hole has a high inclination relative to the line of sight. To maintain consistency in the spectral fits, we set the inclination equal to 70° for all spectra. The choice of inclination does not affect the other spectral fitting parameters significantly. Unless a close to face-on inclination is chosen, the best fit parameter values agree within their respective errors. The energy of the added `diskline` was set to 6.4 keV. We were unable to fit for the inner and outer disc radii using this model, so we set these parameters equal to 6 GM/c^2, and 1000 GM/c^2, respectively. The remaining spectral parameters were left free, and their values were determined when fitting the data using this Model 1. These fitting parameters are the hydrogen column density (N_H), the inner accretion disc temperature (T_{in}), the emissivity power law index of the iron line (B_{10}), the power law photon index (Γ), the high energy cutoff (E_{cut}), as well as the normalisation of the three additive spectral components, `diskbb`, `diskline`, and `cutoffpl`.

We also fitted the energy spectra using a more physical Comptonisation model by Poutanen and Svensson (1996), `compPS`. As this model describes both the disc black body and the power law component of the spectrum, we replaced both these components of Model 1 by `compPS`, unlike Xu et al. (2019). Therefore, **Model 2** is: `dscor*ztbabs*constant*(diskline+compPS)`. We used the same fixed values for the iron K_α line energy, the inner and outer accretion disc radii, and the inclination of the accretion disc, which were already used in Model 1. We used a spherical corona geometry and specified a multicolour disc black body component by confining the disc temperatures specified in the model to have negative values. The absolute magnitudes of those values then represent the temperatures at the inner edge of the geometrically thin, optically thick accretion disc. We fit the spectra with a free hydrogen column density (N_H), emissivity power law index (B_{10}), coronal electron temperature (T_e), inner accretion disc temperature (T_{in}), coronal optical depth (τ_y), relativistic reflection normalisation (R_r), and normalisation of the `compPS`, and `diskline` components. Spectra fitted with both Models 1 and 2 always had a lower χ^2 at the same number of degrees of freedom when fitted with Model 2. This indicates that the more physical model provides a better description of the data.

We fitted these two models to the *NuSTAR* FPMA and FPMB spectra on their own, as well as to the *NuSTAR* FPMA, FPMB, and *XMM-Newton* spectra in absolutely simultaneous intervals of each observation, or flip-flop state. After correcting for the effects of the DSH and applying an energy correction, there was a good agreement in the best fit parameter values of both sets of fits (see Fig. 5.18, 5.19, and 5.20). Due to the shorter

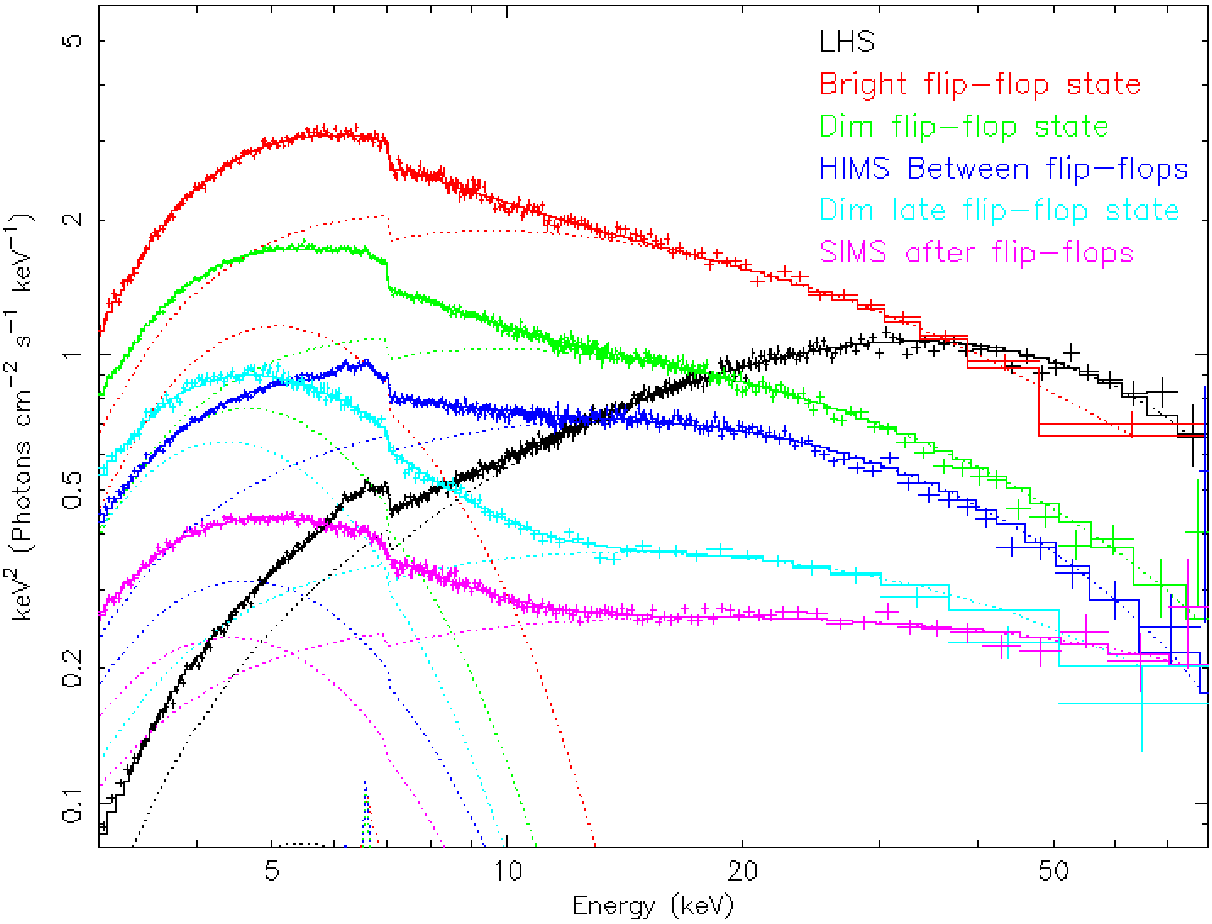

Figure 5.22: Comparison of the *NuSTAR* FPMA spectra and the varying strengths of the individual components they were fitted with, at different times during the outburst. The spectra here are all fitted with Model 1.

XMM-Newton exposure times, the latter selection contained a smaller number of flip-flop states and shorter intervals of simultaneous spectra. To remain consistent, and describe all the observed flip-flops seen by either of the two instruments, we only plot (Fig. 5.23, and 5.24) and list (Tables 5.4, 5.5, 5.6, and 5.7, 5.9, 5.8) the best fit parameter values found when fitting the *NuSTAR* FPMA and FPMB spectra together. In such fits, we only allow the cross-calibration constant to differ between the spectra from the two focal plane modules, along with the difference in the source extractions regions used, an essential element of the DSH model. Errors are quoted at a 90% confidence interval for each parameter.

In Fig. 5.25, we analyse the differences between the bright and dim state spectra. All the spectra in this figure are fitted with Model 2. Initially, the parameters of both the bright and dim state spectra were set equal to the best fit values of the dim state spectrum, in (a). As a result, the dim state spectrum is fitted very well, but the bright state spectrum has huge residuals. Next, we changed the inner disc temperature of the spectral fit describing the bright state to the best fit value that had been found when this spectrum was fitted on its own, in (b). This one change already accounts for most of the difference between the bright and dim state spectra. So, the dominant cause of the change in flux between the dim and bright states is an increase in disc temperature and the effect this has on the power law component.

In plot (c), we also change the `compPS` normalisation of the bright state spectral fit to its best fit value. This removed most of the remaining residuals of the fit, except at low energy, where the greater N_H in the bright state has a noticeable impact. When the N_H itself is changed from the dim to the bright state best fit value, in (d), the bright state spectrum is fitted very well, with only minor differences remaining. By changing all the remaining parameters, we obtain the final graph, (e).

In Fig. 5.23 and 5.24, we distinguished the results of the spectral fits by whether flip-flops were detected in the observations and what QPO type was observed, if any. Interestingly, the observations without a detectable QPO are the least bright, and the bright flip-flop states with a type A QPO are the brightest. All regions in which a type C QPO was observed lie between these two extremes. All identical or comparable spectral parameters involved in the two different fitting models showed similar, though not identical, values and flux dependencies when applied to the same set of data.

In the results of the spectral fits, we found N_H to correlate with flux. When using all the observations shown in Fig. 5.23, and 5.24, we found $r_{xy,1} = 0.93$, $p_1 = 1.0 \times 10^{-8}$, and $r_{xy,2} = 0.83$, $p_2 = 1.2 \times 10^{-5}$ for the two spectral fitting models, respectively. There is, however, a partial degeneracy between N_H and T_{in}, as both parameters push the `diskbb` spectral component to higher energies. We, therefore, analysed whether the observed change in N_H could also be reproduced by a larger change in T_{in}. To do so, we fit all the spectra together, each spectrum with independent fitting parameters as described above. However, the hydrogen column density was tied together for all these spectral fits. Then for comparison, we fit all spectra again, but this time with each spectrum having its own independent hydrogen column density. The second set of fits had 18 fewer degrees of freedom but also had a χ^2 which was smaller by 693. We, therefore, conclude that the observed correlation between hydrogen column density and flux is a real effect. This

Model 1		**80301301002**		**80302302002**			
		bright state	dim state	bright state 1	dim state 1	bright state 2	dim state 2
`ztbabs`	N_H (10^{23} cm^{-2})	14.5 ± 0.6	12.3 ± 0.2	13.8 ± 0.4	13.2 ± 0.7	$13.2^{+0.7}_{-0.9}$	11.9 ± 0.3
	z	0.012	0.014	0.009	0.017	0.007	0.015
`diskbb`	kT_{in} (keV)	1.44 ± 0.03	1.28 ± 0.01	1.46 ± 0.02	$1.29^{+0.04}_{-0.03}$	1.52 ± 0.04	1.27 ± 0.02
	N_{dbb}	67^{+6}_{-5}	74 ± 3	61^{+4}_{-3}	73^{+11}_{-9}	54^{+8}_{-6}	74^{+5}_{-4}
`diskline`	B_{10}	$-2.4^{+0.3}_{-0.2}$	$-1.8^{+0.2}_{-0.1}$	-2.1 ± 0.2	$-1.7^{+0.8}_{-0.4}$	-1.4	-1.8 ± 0.2
	N_{dl} (10^{-3})	4.5 ± 0.1	2.2 ± 0.3	4 ± 1	$2.2^{+1}_{-0.8}$	1^{+17}_{-1}	2.0 ± 0.4
`cutoffpl`	Γ	2.14 ± 0.09	2.06 ± 0.04	2.13 ± 0.07	2.3 ± 0.1	$2.1^{+0.2}_{-0.1}$	2.00 ± 0.06
	E_{cut} (keV)	60^{+20}_{-10}	49^{+5}_{-4}	56^{+11}_{-8}	110^{+220}_{-50}	60^{+30}_{-10}	44^{+5}_{-4}
	N_{cpl}	$3.5^{+0.8}_{-0.7}$	1.6 ± 0.1	3.1 ± 0.5	$2.7^{+0.8}_{-0.7}$	$2.4^{+1.1}_{-0.7}$	$1.3^{+0.2}_{-0.1}$
	χ^2/ν	1803.73/1699	2657.4/2446	2034.98/1863	1312.25/1249	1039.48/1112	2243.91/2073
		= 1.0616	= 1.0864	= 1.0923	= 1.0506	= 0.9424	= 1.0825

Table 5.4: Table of best fit parameters using Model 1, `dscor*ztbabs*constant*(diskbb+diskline+cutoffpl)`, part 1

Model 1		**80302302004**				**80302302006**	**80302302008**
		bright state 1	dim state 1	bright state 2	dim state 2	entire obs.	entire obs.
`ztbabs`	$N_H\,(10^{23}\ \mathrm{cm}^{-2})$	13.1 ± 0.3	11.3 ± 0.3	12.7 ± 0.8	11.3 ± 0.5	$9.4^{+0.3}_{-0.2}$	10.8 ± 0.2
	z	0.008	0.016	0.013	0.018	0.011	0.010
`diskbb`	$kT_{in}\,(\mathrm{keV})$	1.40 ± 0.02	1.31 ± 0.03	1.44 ± 0.05	$1.28^{+0.03}_{-0.04}$	1.37 ± 0.03	1.33 ± 0.02
	N_{dbb}	74 ± 4	70^{+4}_{-3}	56^{+8}_{-6}	65^{+7}_{-6}	20 ± 1	36 ± 2
`diskline`	B_{10}	$-2.0^{+0.3}_{-0.2}$	$-1.6^{+0.3}_{-0.2}$	$-2.1^{+0.7}_{-0.3}$	$-1.8^{+0.4}_{-0.2}$	$-1.80^{+0.10}_{-0.09}$	-1.7 ± 0.1
	$N_{dl}\,(10^{-3})$	$2.8^{+0.9}_{-0.8}$	$1.7^{+0.4}_{-0.3}$	4 ± 2	$2.2^{+0.7}_{-0.6}$	2.4 ± 0.2	1.6 ± 0.2
`cutoffpl`	Γ	1.95 ± 0.08	1.83 ± 0.07	2.0 ± 0.2	$2.0^{+0.10}_{-0.09}$	1.66 ± 0.05	1.75 ± 0.05
	$E_{cut}\,(\mathrm{keV})$	42^{+7}_{-5}	34^{+4}_{-3}	50^{+30}_{-10}	46^{+8}_{-11}	32 ± 2	36 ± 3
	N_{cpl}	1.8 ± 0.3	0.9 ± 0.1	$1.9^{+0.8}_{-0.6}$	$1.2^{+0.3}_{-0.2}$	0.46 ± 0.05	0.54 ± 0.06
	χ^2/ν	1916.61/1840	2103.89/1941	1225.44/1201	1702.75/1572	2823.57/2362	2445.8/2346
		= 1.0416	= 1.0839	= 1.0204	1.0832	= 1.1954	= 1.0425

Table 5.5: Table of best fit parameters using Model 1, `dscor*ztbabs*constant*(diskbb+diskline+cutoffpl), part 2`

Model 1		**80302302010**						**90401317002**
		dim state 1	bright state 1	bright state 2	dim state 2	dim state 3	bright state 3	entire obs.
`ztbabs`	N_H (10^{23} cm^{-2})	$10.3^{+0.5}_{-0.4}$	$11.0^{+0.6}_{-0.5}$	$10.7^{+0.5}_{-0.4}$	10^{+2}_{-1}	$10.1^{+0.6}_{-0.5}$	$11.2^{+0.5}_{-0.4}$	8.2 ± 0.3
	z	0.013	0.020	0.019	0.004	0.026	0.032	0.019
`diskbb`	kT_{in} (keV)	1.21 ± 0.02	1.28 ± 0.03	$1.30^{+0.02}_{-0.03}$	$1.22^{+0.06}_{-0.09}$	1.24 ± 0.03	$1.27^{+0.02}_{-0.03}$	$1.33^{+0.01}_{-0.02}$
	N_{dbb}	80^{+8}_{-7}	66^{+7}_{-5}	57^{+5}_{-3}	70^{+30}_{-10}	71^{+8}_{-7}	62^{+6}_{-5}	$17.0^{+1.0}_{-0.8}$
`diskline`	B_{10}	$-1.7^{+1.4}_{0.57}$	-1.1 ± 1	-2.1	-1.7	-1.5	$-1.8^{+1.0}_{-0.5}$	$-1.1^{+1.0}_{-0.8}$
	N_{dl} (10^{-3})	$0.4^{+0.3}_{-0.2}$	0.5 ± 0.3	$0.3^{+0.6}_{-0.3}$	$0.5^{+1.5}_{-0.5}$	0.4 ± 0.3	0.7 ± 0.4	$0.14^{+0.09}_{-0.06})$
`cutoffpl`	Γ	1.8 ± 0.1	1.8 ± 0.2	2.0 ± 0.1	$2.0^{+0.2}_{-0.3}$	1.6 ± 0.2	2.1 ± 0.1	$1.90^{+0.0.06}_{-0.07}$
	E_{cut} (keV)	60^{+30}_{-10}	40^{+20}_{-10}	70^{+40}_{-20}	130	36^{+15}_{-9}	100^{+100}_{-40}	100^{+90}_{-40}
	N_{cpl}	$0.31^{+0.09}_{-0.07}$	$0.4^{+0.2}_{-0.1}$	$0.6^{+0.2}_{-0.1}$	$0.4^{+0.3}_{-0.2}$	$0.20^{+0.09}_{-0.07}$	0.8 ± 0.2	0.23 ± 0.03
	χ^2/ν	1478.19/1425 $= 1.0373$	1220.84/1272 $= 0.9598$	1569.68/1551 $= 1.0121$	612.87/595 $= 1.03$	1224.99/1169 $= 1.0479$	1531.76/1455 $= 1.0528$	2178.77/2066 $= 1.0546$

Table 5.6: Table of best fit parameters using Model 1, `dscor*ztbabs*constant*(diskbb+diskline+cutoffpl)`, `part 3`

Model 2		**80301301002**		**80302302002**			
		bright state	dim state	bright state 1	dim state 1	bright state 2	dim state 2
`ztbabs`	$N_H\,(10^{23}\ \mathrm{cm}^{-2})$	$12.4^{+0.4}_{-0.3}$	10.8 ± 0.2	11.7 ± 0.3	11.2 ± 0.4	$11.6^{+0.4}_{-0.6}$	10.3 ± 0.2
	z	0.012	0.014	0.009	0.017	0.007	0.015
`diskline`	B_{10}	$-1.6^{+1.6}_{-0.5}$	$-1.3^{+0.4}_{-0.2}$	-1.5 ± 0.3	$-1.4^{+1.0}_{-0.4}$	-0.9 ± 0.9	$-1.2^{+0.6}_{-0.4}$
	$N_{dl}\,(10^{-3})$	2 ± 1	1.5 ± 0.3	$1.9^{+1.0}_{-0.8}$	$1.9^{+0.9}_{-0.7}$	$0.7^{+1.6}_{-0.87}$	1.2 ± 0.3
`compPS`	$kT_e\,(\mathrm{keV})$	72 ± 2	80 ± 1	74 ± 2	89 ± 4	83^{+6}_{-5}	79 ± 2
	$kT_{in}\,(\mathrm{keV})$	-1.45 ± 0.04	-1.31 ± 0.01	-1.45 ± 0.03	-1.32 ± 0.03	$-1.49^{+0.05}_{-0.06}$	-1.31 ± 0.02
	τ_y	0.59 ± 0.04	0.507 ± 0.008	0.57 ± 0.03	0.45 ± 0.02	0.51 ± 0.05	0.51 ± 0.01
	R_r	$0.3^{+0.2}_{-0.1}$	0.29 ± 0.07	0.57 ± 0.03	0 ± 0.1	$0.1^{+0.3}_{-0.1}$	0.4 ± 0.1
	N_{cps}	130 ± 20	117 ± 6	120 ± 10	130^{+20}_{-10}	110^{+2}_{-20}	110^{+9}_{-7}
	χ^2/ν	1774.47/1699 = 1.0444	2616.27/2446 = 1.0696	2004.43/1863 = 1.0759	1315.51/1249 = 1.0533	1039.21/1103 = 0.94216	2211.8/2211 = 1.0670

Table 5.7: Table of best fit parameters using Model 2, `dscor*ztbabs*constant*(diskline+compPS), part 1`

Model 2		**80302302010**						**90401317002**
		dim state 1	bright state 1	bright state 2	dim state 2	dim state 3	bright state 3	entire obs.
`ztbabs`	$N_H\,(10^{23}\ \mathrm{cm}^{-2})$	9.7 ± 0.2	10.4 ± 0.5	$9.7^{+0.2}_{-0.3}$	$8.4^{+1.0}_{-0.9}$	98 ± 0.5	$9.8^{+0.3}_{-0.2}$	7.0 ± 0.2
	z	0.013	0.020	0.019	0.004	0.026	0.032	0.019
`diskline`	B_{10}	-1.5	-1.5	-1.4	-1.5	-1.5	$-1.2^{+1.2}_{-0.6}$	$-0.6^{+0.6}_{-2.0}$
	$N_{dl}\,(10^{-3})$	0.3 ± 0.2	0.5 ± 0.3	$0.1^{+0.3}_{-0.1}$	$0.3^{+0.7}_{-0.3}$	0.4 ± 0.3	0.4 ± 0.3	0.11 ± 0.06
`compPS`	kT_e (keV)	153^{+3}_{-6}	118 ± 7	131^{+3}_{-2}	160^{+10}_{-20}	147^{+4}_{-9}	130 ± 4	150 ± 3
	kT_{in} (keV)	$-1.25^{+0.02}_{-0.01}$	-1.30 ± 0.03	$-1.35^{+0.02}_{-0.01}$	$-1.28^{+0.07}_{-0.04}$	$-1.28^{+0.03}_{-0.02}$	-1.34 ± 0.02	-1.39 ± 0.01
	τ_y	$0.182^{+0.006}_{-0.005}$	0.27 ± 0.01	$0.264^{+0.008}_{-0.007}$	0.18	0.185 ± 0.008	0.272 ± 0.009	0.339 ± 0.007
	R_r	$0^{+0.08}_{-0}$	0.2 ± 0.2	$0^{+0.05}_{-0}$	0	$0.07^{+0.28}_{-0.07}$	$0^{+0.8}_{-0}$	$0^{+0.01}_{-0}$
	N_{cps}	88^{+1}_{-5}	82^{+10}_{-9}	70^{+5}_{-2}	70 ± 20	79^{+6}_{-9}	72 ± 5	22 ± 1
	χ^2/ν	1482.36/1425	1219.41/1273	1574.89/151	614.49/596	1228.62/1169	1542.29/1455	2207.33/2066
		= 1.0403	= 0.9579	= 1.0154	= 1.031	= 1.051	= 1.060	= 1.0684

Table 5.8: Table of best fit parameters using Model 2, `dscor*ztbabs*constant*(diskline+compPS), part 3`

Model 2		**80302302004**				**80302302006**	**80302302008**
		dim state 1	bright state 1	dim state 2	bright state 2	entire obs.	entire obs.
`ztbabs`	$N_H\,(10^{23}\ \mathrm{cm}^{-2})$	11.6 ± 0.3	10.1 ± 0.2	11.0 ± 0.6	9.8 ± 0.4	7.3 ± 0.2	9.3 ± 0.2
	z	0.008	0.016	0.013	0.018	0.011	0.010
`diskline`	B_{10}	$-1.3^{+1.3}_{-0.5}$	-1.4	$-1.5^{+1.5}_{-0.6}$	$-1.4^{+0.8}_{-0.4}$	$-1.0^{+0.5}_{-0.3}$	$-1.0^{+0.9}_{-0.4}$
	$N_{dl}\,(10^{-3})$	$1.1^{+0.8}_{-0.7}$	1.1 ± 0.2	2^{+2}_{-1}	$1.4^{+0.6}_{-0.5}$	1.0 ± 0.2	0.8 ± 0.2
`compPS`	kT_e (keV)	80 ± 2	83 ± 2	80 ± 5	81 ± 3	58 ± 2	80 ± 1
	kT_{in} (keV)	-1.43 ± 0.02	-1.34 ± 0.02	-1.46^{+7}_{-6}	-1.32 ± 0.03	$-1.49^{+0.04}_{-0.03}$	-1.38 ± 0.02
	τ_y	0.48 ± 0.02	0.46 ± 0.01	0.52 ± 0.06	0.51 ± 0.02	$0.91^{+0.06}_{-0.05}$	0.58 ± 0.02
	R_r	0.4 ± 0.1	0.5 ± 0.1	0.3 ± 0.3	0.4 ± 0.2	$1.3 \pm 0.1\pm$	0.7 ± 0.1
	N_{cps}	110^{+11}_{-9}	97^{+7}_{-6}	90 ± 20	100 ± 10	27 ± 3	51^{+4}_{-3}
	χ^2/ν	1890.66/1840	2092.49/1942	1219.77/1201	1687.81/1572	2653.1/2362	2393.55/2346
		$= 1.0275$	$= 1.0775$	$= 1.0156$	$= 1.07367$	$= 1.1232$	$= 1.0203$

Table 5.9: Table of best fit parameters using Model 2, `dscor*ztbabs*constant*(diskline+compPS), part 2`

implies that a significant fraction of the N_H must be local to the system and able to vary within ≈ 2 ks. However, this goes against the assumption of Jin et al. (2019), that all the absorption occurred in the DSH.

We measured a lower average hydrogen column density for the dim flip-flop states than for the bright states. The effect of the different N_H measured in the bright and dim states can be seen in Fig. 5.25 (c), and (d). The low p-value for the null hypothesis, which assumes a constant N_H in both states of $p = 6.4 \times 10^{-4}$, and $p = 1.9 \times 10^{-4}$ for the two models, indicates that the hydrogen column density changes during the flip-flops. Both Xu et al. (2019) and Jithesh et al. (2019) also detected a greater N_H in the bright flip-flop states, though their values differ from the ones we found, as they did not include the DSH model in their spectral fits. As another investigation of whether the N_H changes during the flip-flops, we repeated the above test, comparing the best fit χ^2 of spectral fits of flip-flop spectra, having an identical N_H with the χ^2 of fits with entirely independent values of N_H. We found that the increase in the degrees of freedom by 9, was accompanied by a drop of the total χ^2, of 103. This also supports the notion that the N_H changes during flip-flop transitions.

In both models, we observed an increase in the inner disc temperature with X-ray flux, most notably between the bright and dim flip-flop states. As Fig. 5.25 indicates, this temperature change is the cause of most of the difference between the bright and dim state spectra. This increase is enshrined in the detected correlation between temperature and flux of the early flip-flop bright and dim states, with $r_{xy,1} = 0.92$, $p_1 = 1.6 \times 10^{-4}$, and $r_{xy,2} = 0.91$, $p_2 = 2.4 \times 10^{-4}$ for the two models. Interestingly, the two states by themselves do not show any significant correlation at all, with $p_1 = 0.51$, $p_2 = 0.81$ for the bright states, and $p_1 = 0.98$, $p_2 = 0.59$ for the dim states. At comparable fluxes, the disc has a lower inner temperature in both flip-flop intervals, compared to other parts of the outburst. This is a consequence of the power law component being stronger during the flip-flops than in other regions of the outburst whose spectra we obtained, as shown in Fig. 5.22.

On the two graphs of the dependence of the iron K_α `diskline` normalisation on X-ray flux, we plotted a straight line depicting a constant equivalent width under the assumption of a uniform spectral shape. Fig. 5.22 indicates that even though the spectra have some differences in their spectral shape, this is still a reasonably good assumption. We notice that the equivalent width of the bright and dim flip-flop states is very similar (46 ± 12 eV). This suggests that the disc intercepts a similar solid angle of light emitted by the primary X-ray source in both states. The interval between the early and late flip-flops features a greater equivalent width (83 ± 20 eV), and both the late flip-flops (20 ± 5 eV) and the final *NuSTAR* observation in the SIMS (13 ± 7 eV) have lower equivalent widths than the early flip-flops. the largest equivalent width (166 ± 8 eV) was measured in the first *NuSTAR* observation when the BHT was in the LHS.

The emissivity power law index, B_{10}, does not change significantly throughout the observations we analysed. This was found for both sets of spectral fitting models that we used.

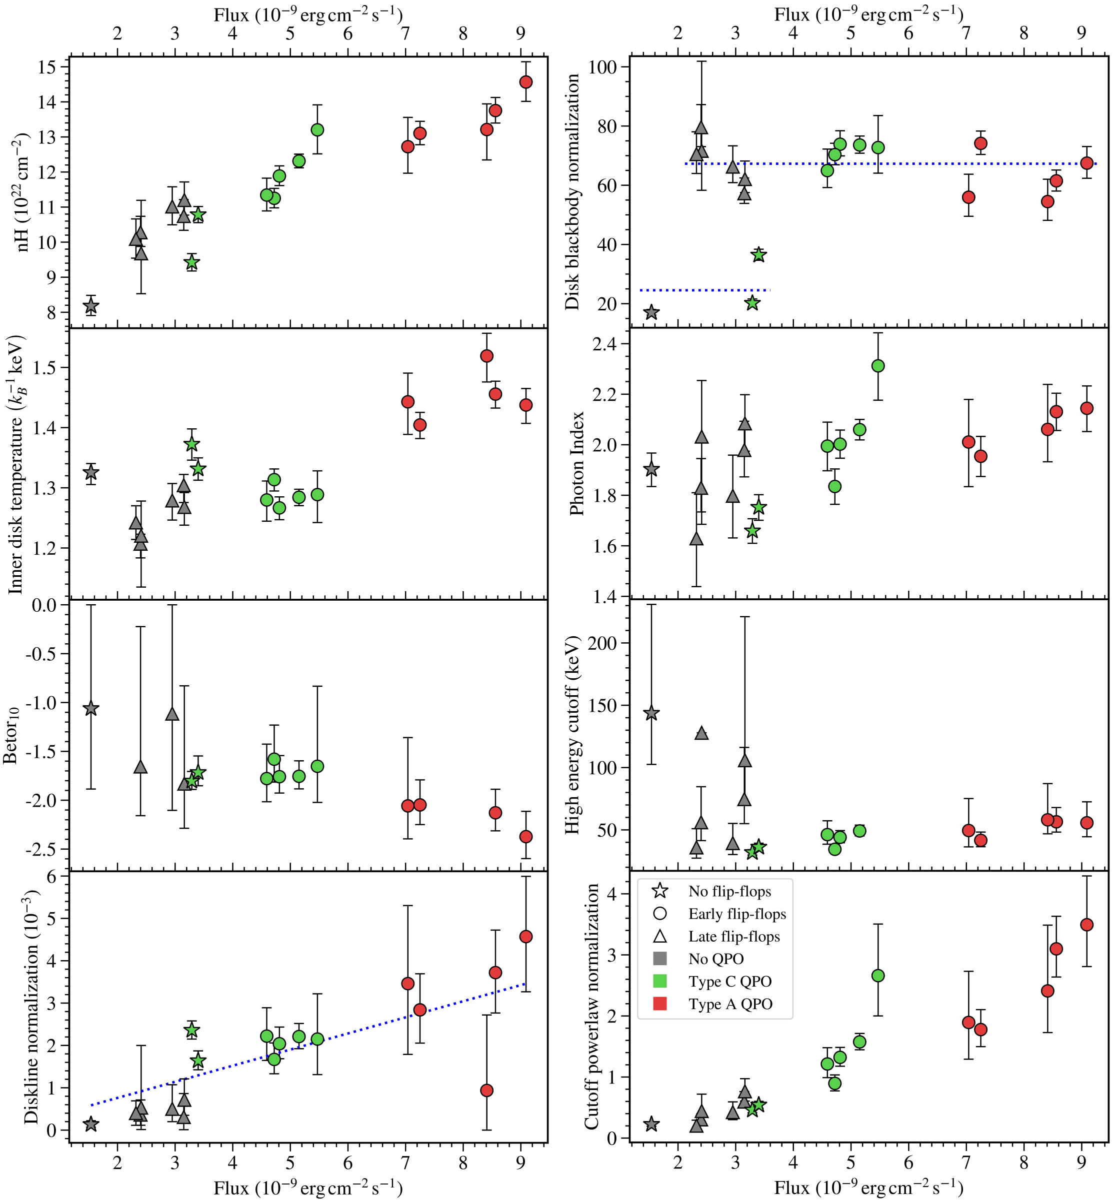

(a) (b)

Figure 5.23: Comparison of the best fit *NuSTAR* spectral parameters of individual observations and flip-flop states fitted with Model 1, `dscor*tbabs*constant*(diskline+diskbb+cutoffpl)`. The spectra are distinguished by their corresponding fluxes, timing properties, and flip-flop state classifications. The first *NuSTAR* observation was excluded, as its spectrum had a negligible black body component and differed too greatly from the spectra shown here. The blue dashed line in the graph of `diskline` normalisation as a function of flux depicts a line of constant equivalent width, assuming no change in spectral shape at an energy of 6.4 keV. The two blue dashed lines in the `diskbb` normalisation graph indicate the average values of this parameter for flip-flop and non-flip-flop spectra. One dim flip-flop observation featuring a type C QPO was omitted from the graph of the high energy cutoff for display clarity, as this parameter was not well constrained in that observation. We also excluded one bright state, and three late flip-flop data points from the graph of Betor$_{10}$, as the fits were insensitive to this parameter, and the errors in the measurements could not be determined.

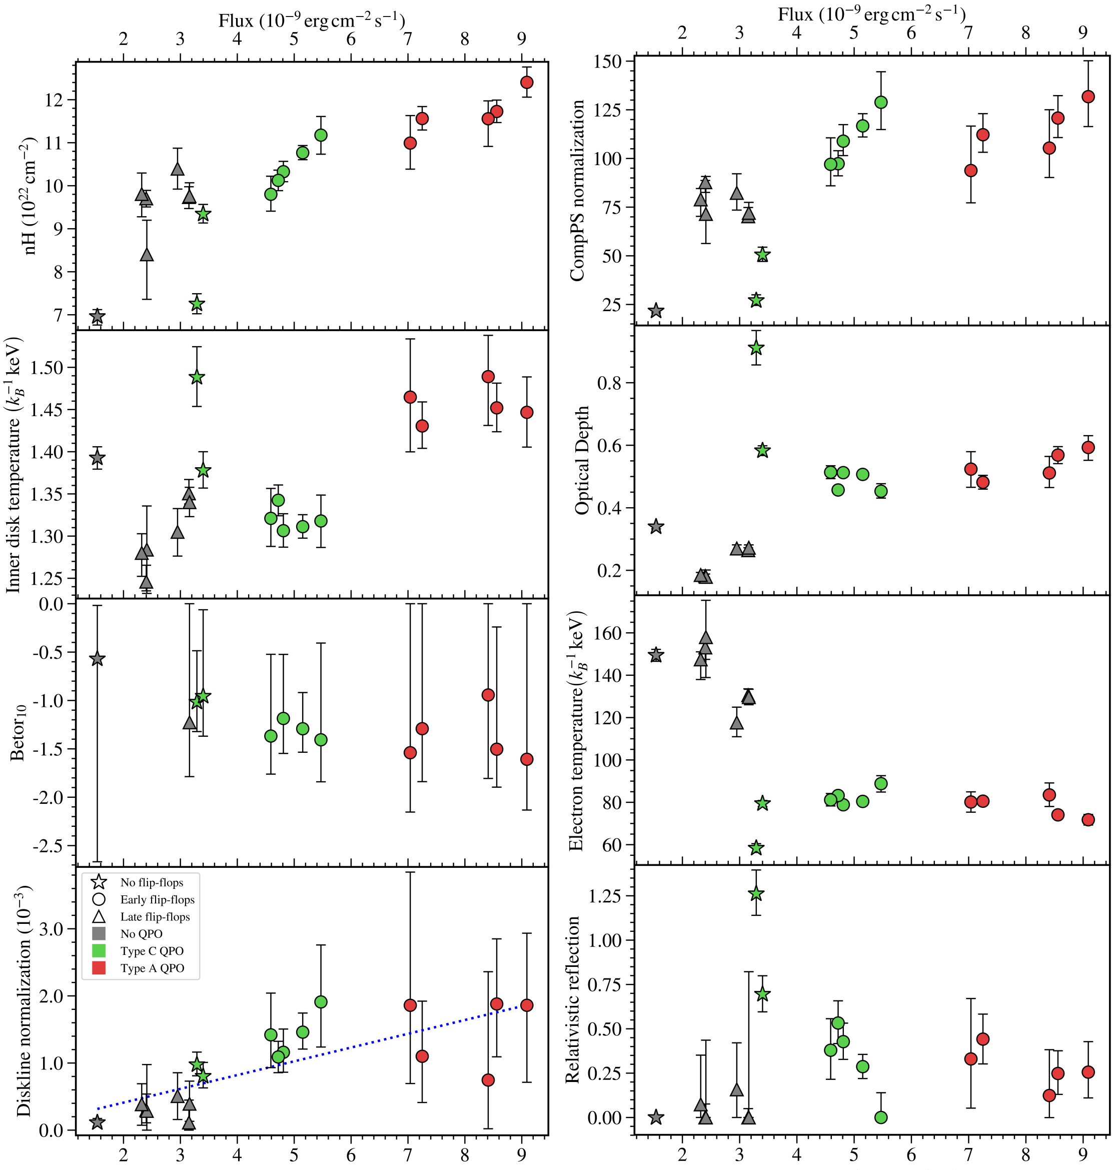

(a) (b)

Figure 5.24: Comparison of the best fit *NuSTAR* spectral parameters of individual observations and flip-flop states fitted with Model 2, `dscor*tbabs*constant*(diskline+compPS)`. We excluded three late flip-flop data points from the graph of Betor$_{10}$, because the fits were insensitive to this parameter, and the errors in the measurements could not be determined.

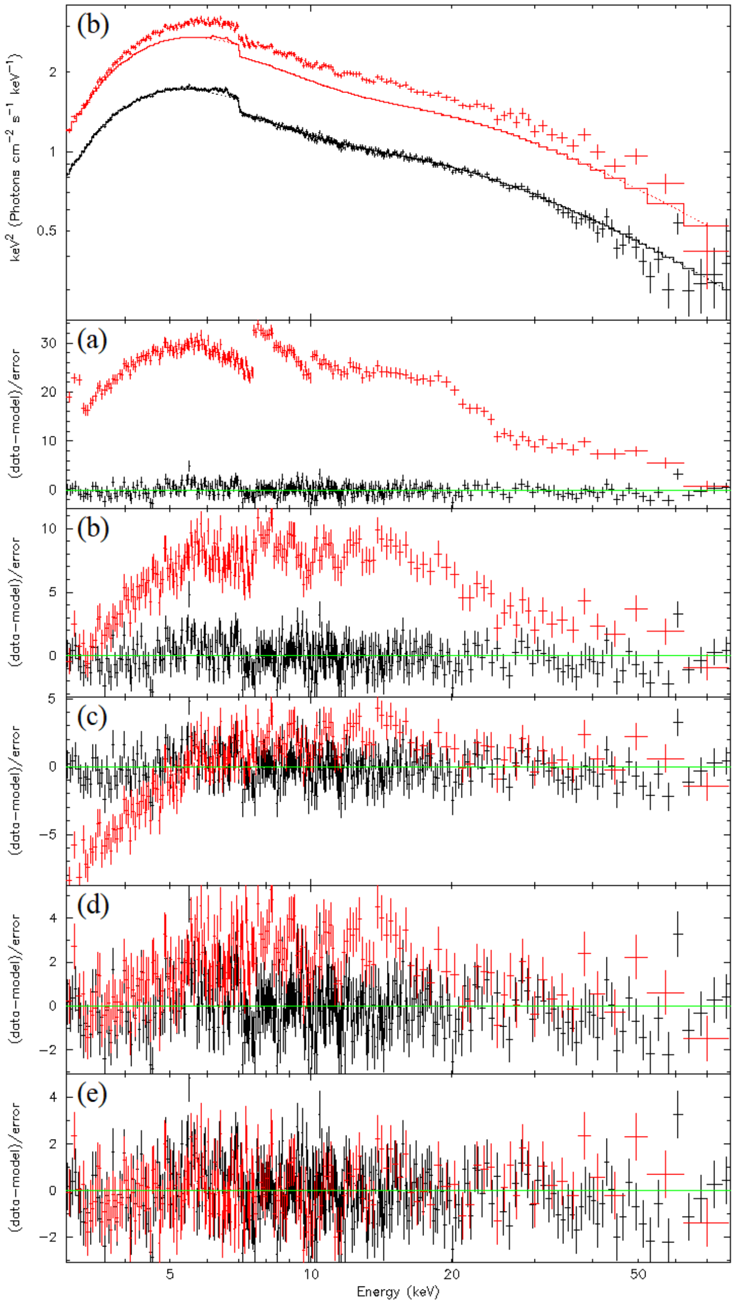

Figure 5.25: The differences between the *NuSTAR* spectra observed in the bright (red), and dim (black) flip-flop states. In panel (a), we show the residuals when fitting the bright state spectrum with Model 2, and setting all parameters to the best fit values found for the dim state. In panels denoted with (b), we show the fit and residuals when using the same fit but changing the bright state inner disc temperature to its best fit value. This removes most of the difference between the two spectra. The `compPS` normalisation of the bright state spectrum is also changed to its best fitting value in panel (c). Most of the remaining difference between the two spectra is corrected by changing N_H, in panel (d). Finally, plot (e) shows the residuals of the best fits to both bright and dim spectra when all parameters of the bright state spectrum fit are set to their best fitting values.

The normalisation of the `diskbb` component, is given by the equation: $N_{dbb} = \left(\frac{R_{in}}{D_{10}}\right)^2 \cos i$. It is proportional to the square of the accretion disc inner truncation radius (R_{in}), and the cosine of the disc inclination (i). The inclination of the accretion disc is not known, and the distance to the binary (D_{10}, measured in units of 10 kpc) has a significant uncertainty, so absolute measurements of the truncation radius from the best fitted values of N_{dbb} are very uncertain. We can, however, determine fractional changes in the inner radius, assuming a constant inclination. In Fig. 5.23, we observe that the `diskbb` normalisation is almost constant in all flip-flop and late flip-flop states, and significantly larger (by a factor of 2–4) than in observations not involving any flip-flops. All the flip-flop observations are consistent with variance around a common mean of 67.3 ± 7.1, having $r_{xy,1} = -0.35$, and $p_1 = 0.18$. In contrast, the observations without flip-flops differ too much to have the same normalisation. Therefore, this suggests that all flip-flops correspond to a similar geometry, which significantly differs from the geometry responsible for the spectra of other observations.

The photon index of the power law spectral component, Γ, does not change much within the spectra we considered. The early flip-flops have an almost constant $\Gamma \approx 2.05 \pm 0.12$, with $p = 0.39$.

The high energy cutoff increases slightly with increasing flux during the early flip-flops, having $p_1 = 0.02$. However, there is insufficient evidence to support a change in the cutoff within the range of bright or dim state observations individually. In the last two *NuSTAR* observations, in which we did not detect any QPO, we found that their spectra have a higher average cutoff (83 ± 40 keV) than any of the spectra of intervals with a measurable type C or A QPO (52 ± 21 keV).

The normalisation of the `cutoffpl` component increases almost linearly with flux for all observations, with no clear distinction between the flip-flop and non-flip-flop states or the nature of the QPO. Using all best fit `cutoffpl` values, we obtain a correlation coefficient of $r_{xy} = 0.94$, and a p-value of $p = 3.6 \times 10^{-9}$.

We noticed an almost constant optical depth for the bright and dim flip-flop states of $\tau_y \approx 0.512 \pm 0.042$. Nevertheless, it features a very gradual rise from dim to bright state, with $r_{xy} = 0.67$, and $p = 0.032$. A greater optical depth of $\tau_y \approx 0.74 \pm 0.16$ was found in the interval between the two sets of flip-flops. The remaining observations of the late flip-flops and the SIMS, for which we did not identify any QPO, had a lower average optical depth of $\tau_y \approx 0.242 \pm 0.057$.

We observed a similar trend of the `compps` electron temperature and the `cutoffpl` high energy cutoff as a function of flux. This was expected, as these two parameters are related. The flip-flops states have $T_e \approx 80.2 \pm 4.5\ \mathrm{k_B}^{-1}$ keV, and only feature a minimal negative dependence of electron temperature on flux, with $r_{xy} = -0.55$, and $p = 0.097$. A lower electron temperature is found between the two sets of flip-flops, with an average of $T_e \approx 68 \pm 11\ \mathrm{k_B}^{-1}$ keV. Observations of the late flip-flops and the SIMS were found to have a higher average electron temperature of $T_e \approx 140 \pm 14\ \mathrm{k_B}^{-1}$ keV.

Finally, the relativistic reflection component of the `compps` model does not indicate a dependence on flux for the early flip-flop states, which have an average of $R_r \approx 0.30 \pm 0.15$.

This also supports the idea that the area of the disc, as seen from the central X-ray source, remains constant during flip-flop transitions. Between the early and late flip-flop intervals, we detected an increase in this component, up to $R_r \approx 0.98 \pm 0.28$. This is followed by a drop-down to close to 0 for the late flip-flops and the SIMS.

To study the temporal evolution of the spectral features and determine whether these change immediately at the transition in the light curve and spectrogram, we fitted the spectra of 2 ks segments on either side of every observed flip-flop. The resulting fits indicated an immediate change of the inner disc temperature and the power law normalisation. Furthermore, there was a slight suggestion of delayed hydrogen column density changes in these spectra. However, changes were small and occurred within the error margin, so a possible early or delayed change in N_H could not be corroborated.

The final *NuSTAR* observation still had a strong power law component, in addition to the multicolour disc black body component. Neither of them could be excluded from the fit or was significantly dimmer than the other. This is indicative of the intermediate states. It does not agree with the HSS definition, in which the disc black body component dominates. The remaining observations by *Swift/XRT* needed to be fitted with both a black body and a power law component. Therefore, it seems like Swift J1658.2-4242 never reached the HSS.

5.8 Discussion

We have investigated the differences between the bright and dim flip-flop states, and the distinctions between the flip-flop intervals, and other parts of the outburst, to enhance our knowledge of the changes during, and the causes of the extreme flip-flops we observed. However, there is a distinct lack of models to describe these phenomena.

The flip-flops we observed seem to be of an extreme and rare variety, which has never been seen before. Previously observed flip-flops had flux ratios of between 1.03 and 1.33 between the bright and dim states, corresponding to transitions between QPO types A and B, or between types B and C. However, the flip-flops of Swift J1658.2-4242 had flux ratios of up to 1.77 and corresponded to transitions between QPO types A and C.

Flip-flops have been observed a few times before, but they are not very common, having only been seen in seven other systems so far, out of ≈ 60 known transient black hole binaries (Corral-Santana et al., 2016). However, it is possible that many flip-flops were missed, due to short exposure times, sparse monitoring, long flip-flop state durations, short intervals of flip-flop activity, or small flip-flop amplitudes.

The duration and frequency of *XMM-Newton*, and *NuSTAR* observations in the critical period of transitions between intermediate states in the outburst of a BHT are unusual, and enabled the discovery of 17 flip-flop transitions, 15 of which were directly observed, and two of which were inferred. This is more than has ever been seen before within a single outburst of a BHT, of a comparable flip-flop duration. Therefore, these observations are ideal for examining these phenomena within a single outburst, without worrying about the differences between different outbursts, and different sources. As flip-flop states of the kind we observed for Swift J1658.2-4242 last for at least 2.7 ks, but can also last for tens of ks, or much longer, short observations risk missing flip-flops entirely. Part of the apparent scarcity of flip-flops in the literature may be merely the product of shorter and fewer observations than would be required to detect flip-flops. Daily averaged count rates seem to indicate the presence of flip-flops in several other outbursts but could not be verified without direct detection through long continuous observations. Flip-flops are possibly much more common than they appear to be.

In the standard definition of spectral-timing states used to describe the outburst of a BHT (Belloni et al., 2011), the dim flip-flop states we observed are classified as HIMS, as they have both strong black body and strong power law spectral components, feature a type C QPO, and have an rms in the range typical of the HIMS. Furthermore, the bright flip-flop states are classified into the AS (Belloni, 2010; Motta et al., 2012) due to their location in the HID. The AS has similar timing properties as the SIMS but is found at a greater hardness and a significantly higher flux than the dim state HIMS. Therefore, the late flip-flops are classified into the HIMS (dim state) and SIMS (bright state) due to their location in the HID, their rms, and the properties of their periodograms.

Interestingly, we always found a negative correlation between count rate and hardness in all observations with a type C QPO (see Fig. 5.3). In contrast, the bright flip-flop states, which feature a type A QPO, and later observations in which no QPO was detected, have a positive correlation between count rate and hardness. This is predominantly a result of the relation between the path traced by a BHT in an outburst in the HID, and the QPO types observed at particular points of the outburst.

Compared to previously observed flip-flops, those in Swift J1658.2-4242 fit into the sub-sample with long-duration states (of order ≈ 0.1–10 ks), and long transition times (of order ≈ 10–100 s). We did not observe any short duration flip-flops, which have durations of order ≈ 10 s, and transition within fractions of a second. This supports the idea that although the duration and transition time of a flip-flop can change quite a lot within one outburst of a system, they remain within about one order of magnitude of each other. No system has been found yet, in which both long and short duration flip-flips were found. This could, however, be the result of an insufficiently large sample of flip-flops.

The non-appearance of a type B QPO is at odds with previous findings by C04, according to which a flip-flop transition always involves a type B QPO. They plot a hierarchy of QPO types as a function of flux, with type A occurring at the greatest fluxes, type C at the lowest fluxes, and type B in between the two. The flux limits separating the three QPO types decayed exponentially with time. Although all flip-flops between types B and C described in the literature follow this hierarchy, several A to B flip-flops contradict it. It was thought that type B QPOs are essential for flip-flops and have to appear in either the bright or the dim state of every flip-flop. We, however, detected the first instance of flip-flops not involving a type B in either of the two states. Therefore, we conclude that flip-flops do not have to involve a type B QPO.

Indeed, all hard-to-soft state transitions were thought to require a type B QPO. Although the bright flip-flop states fit into the properties of the AS rather than the SIMS, the observed change in rms is equivalent to the common distinction between the HIMS and SIMS. We, however, found that these kinds of state transitions do not require a type B QPO.

In Fig. 5.26, we plot a graph distinguishing QPO types in flux and time, analogous to Fig. 14 of C04. Plotting the temporal evolution of the maximum type C flux and the minimum type A flux as decaying exponentials, we also found a clear gap between the two. However, we detected different decay rates of these two exponentials, whereas C04 found parallel lines in this log-linear light curve.

C04 observed type B QPOs in the gap between types A and C, but we did not. However, it should be noted that they detected long-lived bright, or dim flip-flop states in this flux interval, rather than short transition periods, as we did. Perhaps we did not detect any type Bs in this interval because they take longer to switch on than the transition lasts. However, the length of the transitions we found is longer than the observed time it took other type B QPOs in similar systems to appear or disappear. Therefore, it seems as though longer transitions would also not have involved type B QPOs.

It seems possible that there is a QPO phase diagram of parameter values involving more than just the luminosity, which separates the different QPO states of a transient black hole, with transitions between different states happening at specific thresholds. So far, we know of thresholds separating types A and B, and types B and C. We now know of a threshold separating types A and C directly, without a type B in between, in analogy to the solid to gaseous transition in a phase diagram.

The direct transition between QPO types A and C also supports the argument of Motta et al. (2012). They suggest that types A and C originate from the same physical process, but that type Bs have an entirely different origin.

We note that it is possible for the type A QPOs to be present in the PSDs at all times, but remain undetected whenever a type C appears, due to its shallow amplitude and low power compared to the type C QPO continuum and background. One possibility is that the type C QPO replaces the type A QPO. Alternatively, the type A QPO is always maintained in the PSD, and the type C QPO switches on and off alongside changes in the flux. This is possible, as the type A QPO is much weaker than the type C QPO, so the type A QPO might be undetected if a type C QPO is also present. We could not rule out either possibility. An investigation into distinguishing these two scenarios could assist the interpretation of the origin of both flip-flops and QPOs.

We classify the early flip-flops of Swift J1658.2-4242 as transitions between the HIMS and the AS. Transitions between these states have not been described before, and most flip-flops observed so far would be associated with transitions between the HIMS and the SIMS. The extreme flip-flop properties observed in this system are possibly the result of transitions from the HIMS to the AS, rather than to the SIMS.

The state of the system at the time of the jet ejection, and the start of the radio flare, is unknown. The last observation with *Astrosat*, four days before the start of the flare was observed, found the source to be in the HIMS. The first observation after the start of the flare, by *NuSTAR*, was the first detected bright flip-flop state, which is classified into the AS. Even though we did not observe a peak of the radio flux, we can infer that it almost certainly reached its peak during the first flip-flop observation.

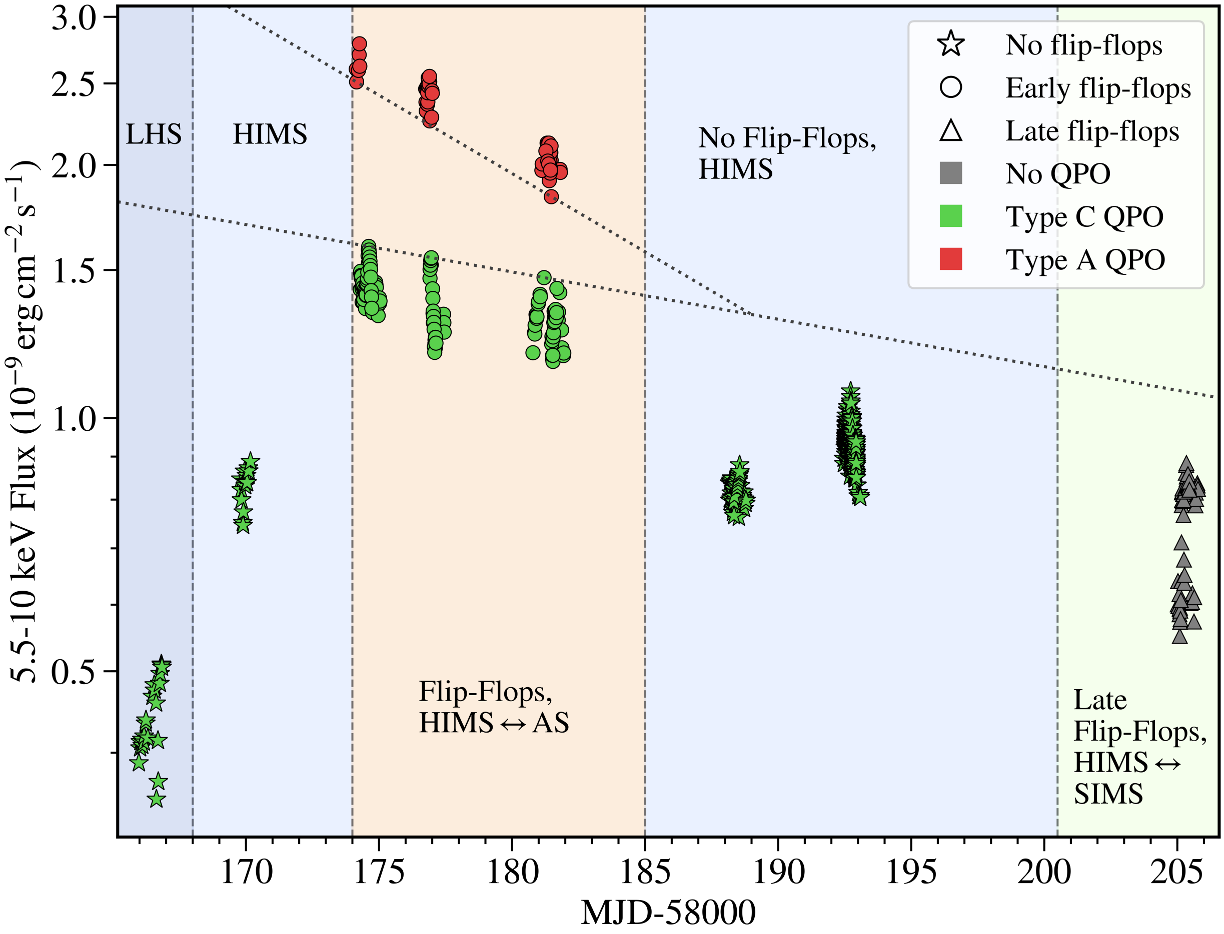

Figure 5.26: Swift J1658.2-4242 light curve, with flux plotted on a logarithmic scale, analogous to Fig. 14 of C04. We also plotted exponentials describing the minimum type A flux, and the maximum type C flux, as a function of time. We only included *NuSTAR*, and *XMM-Newton* data, as their spectra enable an accurate flux determination, and as the QPO type can be unambiguously identified from their light curves. We removed data points during transitions between the two states to show the limiting fluxes of each state when there is no transition.

We note that the power law involved in the flip-flops differs from the standard power law of the LHS. The dominant power law of the LHS is the source of its high variability, with rms values of 10–40% (Motta, 2016). Nevertheless, during the flip-flops, we saw a relative increase in the strength of the power law and an associated sharp drop in the rms variability. This suggests that the corona generating the Comptonised power law spectral feature has different properties during the flip-flop intervals than at the start or end of an outburst. This idea is supported by the high energy cutoff of the power law component in the spectra.

Within the sole detailed observation of a flip-flop transition, we found flux variation to precede changes in the periodogram but take longer to transition. This suggests the existence of two different yet interrelated mechanisms: one for changes in the averaged flux and another for differences in the periodogram.

The simultaneity of the radio flare with the start of the flip-flop interval implies a possible causal connection between the two effects. Jet ejections and subsequent radio flares have been linked with state transitions before (Fender et al., 2004) and are characteristic features of X-ray binaries in outbursts. However, we see no other radio flare at any later time, despite observing numerous flip-flops. We did not even see a flare on MJD 58175 when an *ATCA* observation coincided with a *NuSTAR* and *XMM-Newton* measurement of a bright flip-flop state. If all flip-flop transitions are related to jet ejections, we should have detected a noticeably higher radio flux in this observation.

Due to the lack of additional radio flares, our observations suggest that they are only associated with the first flip-flop transition. This view is supported by the observations described by C04. They detect an ejection episode just before the start of the first X-ray flip-flop interval, and another one within a second flip-flop interval, 11 days later. Furthermore, K11 also found a rapid drop in radio emission, possibly belonging to the end of a radio flare, just before the start of the flip-flops. Therefore, the entire flip-flop behaviour might be caused by the system settling back to an equilibrium state following the disruption caused by the initial jet ejection.

One might expect that a radio flare could also have occurred at the start of the late flip-flop interval. However, we cannot investigate this possibility due to a lack of radio observations in this epoch. We encourage further analysis of similar events, to determine whether there is a connection between the X-ray flip-flops, and the radio flares.

We pose the question of whether the early and late flip-flops are two different instances of the same phenomena, or whether they are governed by entirely distinct mechanisms. The drop in the early flip-flop amplitude and state duration over time seems to continue into the late flip-flops. This is reminiscent of a damped oscillation, which encompasses both the early and late flip-flops. There is a greater similarity between the spectra of the early and late flip-flops, than with the spectra obtained in the interval between the two sets of flip-flops, or after the late flip-flops have stopped. For instance, the `diskbb` normalisation implies that the flip-flops have very similar values of $R_{in}^2 \cos i$, which are significantly larger than corresponding values in observations outside of either of the two flip-flop intervals.

However, there are also some obvious distinctions between the two sets of flip-flops. The late flip-flops have a smaller amplitude, a greater frequency of transitions, and a shorter fundamental period underlying the timing of transitions than the early flip-flops do. All early flip-flops are found at almost the same locations on the HID, with only minor variations over several days. The late flip-flops, however, lie in an entirely different region of the HID and feature different correlations between hardness and intensity. We did not detect any QPOs in any of the late flip-flop states. However, we observed a noticeable difference in the rms of the bright and dim states. On average, the transition times were longer than their earlier equivalents, but the duration of individual flip-flop states was shorter. The late flip-flop spectra also contain a noticeably smaller iron K_α equivalent width than the early flip-flops, suggesting a difference in geometry.

In both the early and the late flip-flop intervals, we observed a period relating the times of transition from one state to the other. The period found in the late flip-flops is slightly smaller than the one detected in the early flip-flops. We showed that it is possible to fit the early and late flip-flop transition times under the assumption of a decreasing period, or as two separate events. The similarity in the periods found in the two intervals suggests that either the late flip-flops could be a continuation of the early flip-flops. Alternatively, they could be separate events, but flip-flops within the same outburst could be constrained to have a limited variation of the transition period. The existence of this period implies a semi-periodic mechanism, by which the system can change from one to the other state only at specific points in time, but does not have to change at every multiple of the period. The time between one transition and the next appears to be a random integer multiple of the underlying period.

The interval between the early and late flip-flops is clearly distinguished from them via its spectral, variability, and timing properties. This indicates that there is a different physical mechanism at play during it. These different properties also suggest that this interval did not feature any unobserved flip-flops. However, of the two *NuSTAR* and *XMM* observations in this interval, the latter observations show a significantly greater similarity to the flip-flop spectral and timing properties. For example, the QPO in the first of these intra-flip-flop observations has a frequency of 3.9 Hz. In contrast, the QPO in the second of these observations has a frequency of 5.7 Hz, much closer to the QPO frequency in the dim flip-flop states observed previously, of 6.2–8.3 Hz. The fitted spectral properties also support this interpretation, with the latter intra-flip-flop observation having similar parameter values as the late flip-flops, as can be seen in Fig. 5.23, and 5.24. This possibly indicates a gradual return to the set of conditions enabling flip-flops to occur.

We notice similarities in the light curves of some, but not all, black hole transients featuring flip-flops. Swift J1658.2-4242, XTE J1859+226 (C04, and Sriram et al. 2013), H1743-322 (H05), possibly also the 2002-2003 outburst of GX 339-4 (N03), all feature an initial flip-flop region at the brightest part of the outburst, which is followed by several days of almost constant flux before a late flip-flop interval is initiated. Flip-flops in the latter interval were only detected and described for Swift J1658.2-4242 and XTE J1859+226. However, comparable regions in the light curves of other black hole transients appear very similar and might also contain flip-flops. None of these systems exhibits a third region that appears to contain flip-flops. These similarities suggest that flip-flops could be linked to a particular type of outburst evolution, which only occurs for some BHTs.

An intriguing distinction between the times when flip-flops occurred and when they did not, can be seen in the fitted truncation radius of the accretion disc. Fig. 5.23 points to a consistent value of the `diskbb` normalisation for all flip-flop states, both within the early and the late flip-flop interval. Other observations within a similar part of the outburst, which do not feature flip-flops, have a range of different values for this normalisation. However, these are all a factor of 2–4 smaller than what was found within the flip-flop intervals. The `diskbb` normalisation is defined as: $N_{dbb} = \left(\frac{R_{in}}{D_{10}}\right)^2 \cos i$, where R_{in} is the inner radius of the accretion disc, D_{10} is the distance to the source, in units of 10 kpc, and i is the inclination. The distance to the system can be assumed to be constant. Therefore, as long as the inclination does not vary, changes of the `diskbb` normalisations are due to a difference in the truncation radii. The spectral fits, therefore, indicate that the disc is truncated at a specific radius whenever flip-flops occur. Following the same reasoning, other times of the outburst, which do not feature flip-flops, have no consistency in their truncation radii.

The fitted `diskbb` normalisation falls into the range of 17–74 within the *NuSTAR* spectra of observations after the initial flux rise. Even though we do not know the exact inclination of the accretion disc relative to the line of sight, we assumed it to be $\approx 70°$. Along with the estimate of the distance to the system, of $\approx$ 10 kpc (Jin et al., 2019), we can approximate the range of `diskbb` normalisations to correspond to accretion disc truncation radii in the range of 12–25 km. We used the conversion by Kubota et al. (1998) to obtain the actual truncation radius from the spectrally fitted one. This correction increases the radius by 19–65%, depending on the choice of κ, and we chose the maximum of $\kappa = 2.0$. The resulting radii are still too small compared with the gravitational radius of a 10 $M_{\odot}$ black hole, of 14.9 km. The small spectrally fitted inner truncation radii could result from an overestimated black hole mass, an underestimated inclination, or an underestimated distance to the black hole. However, these quantities cannot differ much from their estimated values. Tomsick et al. (2009) points out that mixed results were obtained from using the spectrally fitted normalisation of the disc black body component to estimate the inner radius. Additionally, in the relativistic rigid body precession model of the QPO, its frequency is determined by the truncation radius of the accretion disc. A larger truncation radius should yield a lower QPO frequency. However, the highest QPO frequencies were observed in the dim flip-flop states, when the spectrally fitted truncation radius was at its largest. The accretion disc inner radius as a function of time remains uncertain. However, it is seemingly very close to the black hole in the brightest phase of the outburst, especially whenever there is no indication of flip-flop activity.

We found that the major spectral difference between the dim and bright flip-flop states results from an increase of the inner disc temperature at a constant inner radius. An additional increase of the power law normalisation and the hydrogen column density are required to explain most of the changes in the flux and spectral shape between the dim and bright flip-flop states.

The flip-flop interval features a decrease in flip-flop amplitude over time. This decrease seems to continue into the late flip-flop phase, for which both the absolute and the absolute fractional flux difference is substantially smaller than they were in the early flip-flop interval. This suggests an association of flip-flops with a damped oscillation with an irregular period. In this picture, the mechanism responsible gradually loses its amplitude and approaches an equilibrium state. Perhaps the jet ejection triggered an imbalance in the system, which initiated an instability that pushed the system from one unstable state to another, resulting in an oscillation between the two unstable states.

The physical mechanism causing flip-flops to occur remains unknown. The association of flip-flops with alterations of QPOs suggests that understanding the former might require knowledge of the latter. The ongoing discussion about the origin of QPOs complicates the search for a model describing the cause of flip-flops. However, as our findings of transitions between QPO types A and C demonstrated, an understanding of flip-flops could also aid in the search for the origin of QPOs. Below, we present some ideas to provide the reader with a flavour of some ingredients that could play a role in shaping the flip-flop phenomenology.

Flip-flops are described by two timescales; the duration of a state before the next transition and the time the BHT takes to change from one state to the other. For Swift J1658.2-4242, these timescales are of order ≈ 10 ks, and ≈ 100 s, respectively. Other BHTs feature flip-flops with a wide variety of timescales, but the ratio between the two seems to be reasonably similar. This could suggest that the two timescales are related to each other, and to the physical parameters of the disc and the black hole. Indeed, this ratio is consistent with the expected ratio of viscous to thermal timescales within the inner region of a standard accretion disc.

A thermal instability could also account for the temperature variation observed during flip-flop transitions. On the other hand, we note that thermal-viscous instabilities typically result in much larger temperature variations (see e.g. Fragile et al., 2018) than we observed in the flip-flops (of at most 20%). It also remains unclear how this instability could generate the observed change of the QPO or the underlying clock in the times of transition.

A completely different framework posits the existence of an ionised, relativistic, beamed outflow misaligned with the spin of the black hole, therefore precessing with a complex pattern. Our discovery of an underlying clock governing the times of flip-flop transitions indicates the existence of a fundamental complex periodicity in the system, which would be achieved by this model. An increase in flux from dim to bright states would result from bulk Comptonisation of disc and corona emission whenever the outflow points into the line of sight. The bulk Comptonisation of this outflow might explain the observed spectral changes between dim and bright states, namely the increase in temperature and power law normalisation (eventually also the increased hydrogen column density, if the outflow is multi-phase) in the bright states. To explain the variation in the periodogram, we would require this outflow to have a significant radial extent, opening angle, and optical depth, to ensure that a narrow type C QPO would be broadened and weakened into a type A QPO due to the scattering of radiation, and the corresponding time lags across this outflow.

This model provides a straightforward explanation for the near-consistency of the flux differences of adjacent flip-flop transitions. It also immediately justifies why the type C QPO that appears after the end of a bright state is so similar to the type C QPO observed just before the start of the bright state; it was there throughout but was scattered to a type A QPO in the bright state. The pause in flip-flop activity would be explained by the outflow never pointing into the line of sight for a prolonged period. Eventually, the complex system of rotations would point the outflow back into the line of sight once in a while, and the flip-flops would be seen yet again, but this time with a lower amplitude of flip-flop flux variation. Alternatively, the pause in flip-flop activity could result from the outflow being quenched and subsequently re-established.

We tried to model this hypothesis with a precessing conical outflow having a constant opening angle. We defined the system to be in a bright flip-flop state whenever the angle between the line of sight and the outflow centre was less than the opening angle. Larger angles of separation were instead assigned to the dim state. We used this binary model to fit the occurrence of bright and dim states in our observations. We compared our results against simulated flip-flops with the same number of transitions within our observing times. We could not fit our observations well without adding even more complexity, suggesting that this scenario cannot explain the wealth of observational constraints. Using this model, it was possible to fit randomly generated flip-flop states equally well in 45% of our simulations. Therefore, our simulations do not support this interpretation with these specific assumptions. However, it is possible that a more complex system of rotations is required, or that the assumptions we used, particularly of a constant opening angle, are too simplistic. We also note that this model cannot account for the flip-flops observed in other systems, with transitions between different sets of QPO types.

In conclusion, some of the requirements of this model seem unlikely, even though they are still within the realms of physical possibility. A precise, verifiable prediction of this model is a high degree of polarisation in the bright flip-flop states, which future X-ray polarimetric observations by *IXPE* or *eXTP* could investigate. However, at the moment, this model should be regarded as tentative and incomplete. Finally, we note that it is possible for flip-flops to be caused by a combination of several different effects, rather than one single effect, as we have considered here.

Chapter 6

Summary and future outlook

6.1 Conclusions

6.1.1 *eROSITA* variability methodology

We investigated and described how to best perform an eRASS light curve variability analysis. We also improved previously determined variability methodologies, particularly for low count rate sources, and light curves with varying exposure times.

We determined a method for selecting significantly variable eRASS sources, to enable variability analysis to focus on the most intriguing sources of a large sample. We also determined reliable thresholds on the two variability quantifiers SCATT_LO and AMPL_SIG. These thresholds determine the significance of the detected variability and minimise the rate of false positives. They are defined as functions of both the count rate and the number of bins of *eROSITA*-like light curves.

The NIV quantifies the intrinsic degree of variability of a source within the duration of the observations and is usually estimated by the NEV. However, the NEV methodology can only accurately estimate the NIV for light curves with many source counts in every bin. We presented a new method for estimating the NIV that remains accurate regardless of the number of source counts per bin. It also works for light curves with varying exposures. It uses the Bayesian excess variance (bexvar) to estimate the standard deviation of the logarithmic count rate in the light curve. Then it converts that into an estimate of the NIV. This method is based on the assumption of a pink noise PSD but remains more accurate than the NEV for both red, and pink noise PSDs, except at count rates of $\gtrsim 15$ cts/s. We encourage the use of this NEV$_b$ method to estimate the NIV or the rms variability of all but the brightest X-ray sources.

The NIV is offset from the band-limited power by aliasing, and the red noise leak, which leak variability power from higher and lower frequencies into the observed frequency interval. As eRASS light curves consist of short observations separated by much longer breaks, aliasing significantly affects observed periodograms. For power law PSDs, aliasing causes a flattening of the PSD towards the Nyquist frequency.

The NIV of one light curve of a variable source can substantially differ from the NIV in another light curve covering the same frequency interval, even if the variability is governed by the same intrinsic PSD in both. Therefore, the sampling error must be determined, to compare the variability of different sources or the same source at different times. The best way to reduce sampling errors is to split a light curve into shorter segments of equal length, and compute the geometric mean of the NEV_b measurements of individual segments. We determined the value of the sampling error as a function of the number of bins in each segment, the number of segments, and the intrinsic variability of the source in the selected frequency interval, for sources exhibiting pink noise variability. The reduction of the sampling error with an increasing number of segments over which NEV_b is averaged, also depends on whether the individual segments are adjacent, or randomly placed. It can be reduced much more rapidly in the latter case, especially for strongly variable sources.

We also investigated how to analyse the periodograms of variable sources observed by *eROSITA*. If a regularly sampled light curve consists of bins with differing exposure times, its periodogram will feature an additional noise component above the Poisson noise. We determined the dependence of this fractional exposure noise on its mean and variance throughout the light curve.

The code that we developed and used for an accurate variability analysis, particularly for, but not limited to, eRASS light curves, is available at https://github.com/DavidBogensberger/eROSITA_SEP_Variability. It contains further instructions detailing when these methods can be used, how they should be used, how to interpret the results, and when to consider modifying them for use in other areas of the sky, or other instruments.

6.1.2 Variable *eROSITA* X-ray sources in the SEP field

The frequency and consistency of *eROSITA* observations of the SEP field during its eRASSs, enables a deep investigation of the mid to long-term X-ray variability properties of thousands of X-ray sources. However, it comes with the challenges of varying fractional exposures, and a large parameter space in the count rate and exposure time.

Out of the ≈ 8200 X-ray sources observed in each eRASS located in the SEP field, we identified an average of 197 sources which show a significant degree of variability, per eRASS. By matching sources between the three eRASSs, we identified 453 distinct variable sources in this field. We expect about 33 of these variable sources to be false positives.

We used thresholds on the SCATT_LO and AMPL_SIG variability quantifiers to detect variable sources. Of the two methods, SCATT_LO is better suited for detecting variability, as it managed to detect 80.0% of all variable sources identified by either method. Additionally, 39.3% of all variable sources were only identified as such by SCATT_LO. AMPL_SIG might be more susceptible to wrongly classifying non-variable sources as variable. However, AMPL_SIG also identifies many sources as variable, that SCATT_LO missed. Therefore, it is preferable to use both methods to identify variable sources.

Variable sources were distinguished into likely galactic, likely extragalactic, and unknown sources. We observed the optical spectrum of *eROSITA* selected X-ray sources, and determined their source type and redshift. We also used optical photometric catalogues, and previously established relationships between the magnitudes in various bands, to distinguish between likely galactic and likely extragalactic sources. In so doing, the 453 variable sources were classified into 179 likely extragalactic sources, 224 likely galactic sources, 7 XRBs in the LMC, and 43 sources for which we lacked the information to classify them into one of the other categories.

We analysed the periodograms of the 20 significantly variable, likely extragalactic sources with at least 2×10^3 total source counts in the SEP field, in the first three eRASSs. Most of them are well fitted by a single aliased power law, with $\alpha \approx 1$. These results are consistent with previous studies for the frequency range probed by *eROSITA*.

Finally, we investigated some fascinating variable X-ray sources. For instance, the likely AGN eRASSt J061148.4-662435 was observed to have a light curve consisting of two very distinct flux states, that differed by 70%. Despite significant variability between the observed segments, the variability within each segment was relatively moderate. Within two of the observing intervals, the source flux was observed to evolve almost sinusoidally, with a period of ≈ 31 days. The likely extragalactic source eRASSt J054641.8-641524 also potentially featured two distinct flux states. For this source, we observed a transition in which the count rate increased by 113% within ≈ 10 days.

We detected several likely extragalactic sources featuring large, long-term trends. One of them, eRASSt J053942.0-653038 features a sudden flare, followed by a gradual decrease in its source flux throughout the following year.

Many of the 224 likely galactic variable sources we investigated featured infrequent, large, short-duration flares. However, η^2 Dor was instead observed to feature a similar triangular light curve shape in all three eRASSs, indicating a long-term periodicity of 183 days, or a fraction of that.

6.1.3 Flip-flops in Swift J1658.2-4242

We observed the black hole candidate Swift J1658.2-4242 throughout its outburst from February to September 2018, using a suite of X-ray and radio instruments. Swift J1658.2-4242 underwent some extreme changes in its light curve and PDS, which have never been seen before.

Flip-flops with flux ratios of up to 1.77 were observed, which coincided with changes between QPO types A and C. We reported the first detection of a direct transition between these two QPO types. The presence of a type B QPO with a duration of longer than ≈ 10 s is ruled out in the transitions from QPO types A to C. The start of a type C QPO during a bright-to-dim transition was observed to be delayed relative to the start of the transition in the light curve, but the QPO changed significantly faster.

A major radio flare was detected at the same time as the greatest X-ray flux was reached, and the flip-flop interval started, suggesting that these events are related. We found a second interval of flip-flop activity, which featured smaller flux ratios between the two states, a greater frequency of transitions, and no identifiable QPO, but a significant change in total rms between the bright and dim flip-flop states. It started about 16 days after the end of the first flip-flop interval. However, only one radio flare was observed throughout the entire outburst.

We found flip-flop transitions at seemingly random integer multiples of 2.761 ks after the time of the first flip-flop. A slightly lower period of 2.61 ks was observed in the late flip-flops. The existence of an underlying period in both flip-flop intervals, separated by no more than 5.5%, was detected with a significance of 3.2σ. This suggests an inherent base timescale, defining when flip-flop transitions can occur.

Spectral fitting of flip-flop states revealed that the major increase in flux and the slight change in their spectral shape was primarily the result of an increase in the inner accretion disc temperature. The remaining disparity was caused by an additional increase in the power law normalisation and hydrogen column density. In addition, we found that the inner accretion disc radius remained almost constant throughout all flip-flop states and was considerably larger than at times without flip-flops.

We demonstrated the necessity of using a DSH model to describe the spectrum of Swift J1658.2-4242 and other highly obscured sources. Energy corrections were applied to *NuSTAR* and *XMM-Newton* spectra to ensure consistency between them, and the *Chandra* spectrum.

We considered a misaligned, precessing, and beamed outflow as a tentative explanation for the phenomenology we observed. In this model, flip-flops would be generated at semi-periodic times whenever the outflow passes into the line of sight.

Flip-flops have rarely been observed so far, but this might reflect more on the nature of past observations, rather than their absolute rarity, as long continuous observations are required to detect them. We encourage further multiwavelength analysis of similar phenomena in other systems and a continued search for the physical nature of flip-flops.

6.2 Future outlook

Our investigations into the X-ray variability of Swift J1658.2-4242, and of X-ray sources detected by *eROSITA* in the SEP field during its all-sky surveys, have also opened up a lot of new questions, suggesting a direction for future research. In the following sections, we will discuss some of the avenues that could be explored in future, that could help to improve our understanding of the physical processes occurring in the regions close to black holes across the mass range.

6.2.1 *eROSITA* variability methodology

We performed a detailed investigation of how accurately the NIV can be estimated by using different methods. A similar analysis could be done for the periodogram. It would very useful to determine how accurately a measured periodogram can describe the PSD of a variable source, from a set of *eROSITA*-like observations. If this analysis were performed for periodograms of *eROSITA* light curves, it would need to utilise our findings regarding the total Poisson and varying fractional exposure noise. In addition, it would need to account for the effects of aliasing, and possibly also of the red noise leak.

In particular, such an investigation could seek to determine how accurately the power law index of a PSD can be estimated. It would analyse the distribution of measured power law indices from simulated light curves generated from a single power law PSD. The spread and potential offsets of the measured power law indices would be investigated. These could depend on a variety of parameters, including the number of bins per segment, the number of segments in the light curve, the average number of source counts per bin, the average fractional exposure per bin, the input power law index, and the degree of variability of the source. This could enable the error of a measured periodogram power law index to be determined more accurately than before. It could also help identify the sources for which there is the greatest likelihood of measuring a reasonably accurate periodogram.

Furthermore, it would be valuable to investigate how well a PSD consisting of a broken power law can be distinguished from a PSD of a single power law. This analysis would need to investigate what the minimal number of bins and the minimal average source counts per bin has to be, above which an input broken power law model can be accurately fitted to the periodogram of an observed light curve. This would be a more challenging task, as it would need to ensure that the false positive rate of broken power law detection remains small. It would also need to investigate the impact of different values for the two power law indices, and the break frequency. The correlation between the fitting parameters could also be investigated.

The aliasing effect adds a complication to this, which could also be analysed. This could answer questions such as: how does the aliasing effect impact the measurement accuracy of the break frequency as a function of its actual value? Or what is the reliability of a measured high-frequency power low index, for periodograms featuring a break close to the Nyquist frequency?

We generated variability detection thresholds for the parameters SCATT_LO and AMPL_SIG, for light curves between 50 and 1000 bins. These thresholds could be extended to longer light curves, using the same methodology we used in Section 3.4. With a stable and reliable matching of sources between eRASSs, sources located close to the SEP will have eRASS light curves consisting of multiple thousands of bins. By extending the variability thresholds we determined, the variability significance of light curves spanning multiple eRASSs could be determined, even for sources close to the SEP. Many sources are likely to feature larger differences in the source flux in the interval between the last observation of it in one eRASS, and the first observation of it in the next eRASS, than within all observations in the same eRASS. By extending the thresholds, and investigating variability over an interval spanning multiple eRASSs, many additional sources will be identified as variables, at the same false positive probability.

We encourage the further use and investigation of $\overline{\sigma_b}$, and NEV_b as variability measures. The NEV has been widely used so far, but it has several deficiencies. These two variability measures we introduced here more accurately estimate the true source variability, and are, therefore, more useful to any variability analysis. It would be interesting to investigate how $\overline{\sigma_b}$ scales with other source properties, such as the luminosity or the black hole mass.

Both of these parameters have been observed to correlate with the NEV. This could be done by simply using our conversion between the two parameters.

However, such variability analyses should seek to estimate the relationship between the source luminosity, or black hole mass, and the band-limited power. That is the intrinsic, stationary quantity defining the strength of the source variability in general, not just at the time of observation. Estimating the band-limited power includes taking the scatter in the NIV into account. So fitting the band-limited power against the AGN luminosity or black hole mass is likely to yield a tighter correlation, with a smaller degree of scatter around the best fit line, relative to the total error size.

6.2.2 *eROSITA* variability analysis

We have collected a set of significantly variable *eROSITA* sources and analysed the variability properties of a select few in greater detail. However, a lot of science remains to be done for most of the 453 variable X-ray sources we identified. This is a unique data set of long-term variable sources observed almost consistently for up to 18 months. We identified, but did not analyse the variability of XRBs near the SEP. Most of the variable sources were classified as likely galactic sources, which we hardly investigated. For instance, this data set could enable an investigation of the fraction of stars showing significant X-ray variability, as well as a study of the frequency of X-ray flares for different types of stars. Even amongst the AGNs, there remain many whose exact variability properties still need to be analysed in detail.

There is a lot of prospect for further long-term variability analysis of a consistently sampled data set of hundreds of AGNs. For instance, there could be further investigation into the distribution of properties of the low-frequency periodograms observed, and whether these agree with previous models.

There could also be a search for long-term periodicities, to potentially detect low-frequency QPOs in AGNs. *eROSITA* is ideally suited to detect periodicities with periods of days to months. A dedicated search for periodicities would at least be able to place limits on the likelihood of AGNs exhibiting them, within a certain frequency range.

Recently, a few possible QPOs have been identified in various AGNs (Gierliński et al., 2008; Gupta et al., 2018; Agarwal et al., 2021; Ashton and Middleton, 2021). They are usually found to have periods of a few hours. However, the number of AGNs with possible QPOs is still very small. In contrast, QPOs are observed in most BHTs, usually in the low-frequency regime of $\approx 10^{-3} - 30$ Hz (Motta, 2016). When the timescale of these low-frequency QPOs is scaled linearly with BH mass to apply to AGNs, we would expect them to have frequencies in the range of hours to millennia. So far, only the high-frequency limit of these possible AGN QPOs has been probed, and those probably fall into the high-frequency QPO class.

A further avenue to explore could be to investigate what fraction of AGNs detected by *eROSITA* show signs of variability within a selected frequency range. A similar analysis could be done for stars.

6.2.3 Flip-flops

We have yet to determine how frequently flip-flops occur. It is possible that they are much more common than they may currently seem. Their apparent scarcity might be merely a consequence of the typical observing patterns of X-ray telescopes. So far, including this work, flip-flops have only been reported in nine individual outbursts of a BHT. These have occurred in eight individual systems. There are, however, indications of similar phenomena in many other outbursts, that were not identified as flip-flops. For instance, Liu et al. (2022a) observed a sequence of high and low flux states that visually resemble the flip-flops. However, they did not observe any direct transitions between the two states. They also did not identify any QPO in either of the two states. Therefore, the variations in brightness could not definitively be identified as flip-flops, with all of their usual features. Perhaps the standard flip-flop definition is too strict, and should instead include all light curves featuring a top-hat-like flux evolution similar to those shown in Figs. 5.4, 5.5, and 5.6. Perhaps there are different classes of flip-flops. Perhaps only some of them involve strong QPOs, and significant changes to the QPO during transitions.

Long-term *Swift*/XRT monitoring of BHTs in outburst often features intervals during which a significant spread in the flux is observed from one observation to the next. These intervals often occur in the days and weeks after a prominent peak in the X-ray flux. These effects could have been caused by flip-flops and occurred during a part of the outburst in which similar flip-flops have often been found. However, BHTs are intrinsically very variable, so it is impossible to associate large-variable intervals with flip-flops unambiguously.

The lack of flip-flops detections is probably partly due to observing constraints, and the fact that flip-flop states can last for many kiloseconds before transitioning to the bright or dim state. The transition time between states is typically substantially shorter than the duration of individual flip-flop states. So the likelihood of observing a flip-flop transition during a short observation is small. Unless observations span a long, almost continuous interval, lasting significantly longer than the fundamental flip-flop timescale, it is unlikely that flip-flop transitions will be observed, even if many such transitions occur in many BHT outbursts.

Therefore, to improve the chances of detecting flip-flops, we recommend observing for a long, almost continuous interval in the days following the initial peak of the outburst, or any subsequent X-ray flare in the outburst. We also recommend focusing on outbursts of BHTs that have previously exhibited flip-flops.

GX 339-4 has already been observed to feature flip-flops in two separate outbursts, one in 1988 (Miyamoto et al., 1991), and another in 2002 (Nespoli et al., 2003). In its 2021 outburst, there were some indications of flip-flop activity, even though the observed variability could not unambiguously be identified as such Liu et al. (2022a). It is possible that flip-flops do not merely occur randomly, but are more likely to occur in some BHTs, and less likely in others. Particular BHTs might feature flip-flops in most of their outbursts, while other systems might never feature them.

Future systematic searches for flip-flops in outbursts of BHTs could prioritise BHTs that were previously observed to feature flip-flops. This could shed light on whether flip-flops are equally likely to occur in all outbursts of all BHTs, or whether they are more likely to occur for particular sources. It is possible that the occurrence of flip-flops and their various features is related to the properties of the binary system. This could include a dependence on the masses of the black hole and its companion, the inclination of the system, or the spin of the BH. Such an investigation could also help to understand the cause of the flip-flop phenomenon.

Further research is also needed to understand the relationship between flip-flops and the three different QPO types. Casella et al. (2004) depicted a hierarchy of QPO states, in relation to the source brightness. They found that type A QPOs occur when the source is at its brightest. Type C QPOs occur when the source is faint, and type B QPOs occur in an intermediate interval. In many instances, the observed relationship between flip-flop states and the QPO type measured in it agreed with this hierarchy. However, there have been several instances that contradict this model. The causes of the association between QPO type and flip-flop states remain unknown. Further research could also unveil why some systems do not agree with the general picture.

It would also be interesting to investigate how frequently certain types of QPO transitions occur. We currently do not know how common type C to A transitions are, compared to transitions between types C and B, or between types B and A. Questions remain about whether flip-flop transitions involving different pairs of QPOs are the same phenomena, and what mechanisms are responsible for the different QPO transitions observed.

Another possible avenue for future research is to determine whether flip-flop transitions commonly feature flux changes before changes in the periodogram, or vice versa. We observed one flip-flop transition with sufficient sensitivity and temporal resolution to identify that the flux transition started prior to the change in the periodogram, but that the periodogram changed more rapidly. It remains unclear whether the luminosity change triggers a change in the QPO, whether a change of the structure causing the QPO causes the luminosity to vary, or whether a different mechanism affects both parameters, at different times.

The large range of timescales that flip-flop states have been observed to last should also be investigated more. Repeated detections of flip-flops in the same system could indicate that the flip-flop timescale depends on the system properties. In that case, different outbursts of the same source should feature similar flip-flop properties. The possible flip-flops observed in the 1988, 2002, and 2021 outbursts of GX 339-4 seemed to have very different flip-flop durations. If this variation in the flip-flop duration in different outbursts is also seen in other systems, it could instead suggest that flip-flop timescales are not directly related to any of the main properties of the system. We do not know much about what governs the timescale of flip-flops. Nevertheless, we do know that unlike other typical timescales related to BHs, this cannot scale linearly with BH mass. The range of flip-flop timescales observed so far, from several minutes to several hours, spans a much larger logarithmic interval than the logarithmic interval of BH masses expected for these systems. Investigating how the flip-flop timescales relate to various system properties could provide constraints on the model for these systems, and could assist in formulating a universal explanation for these phenomena.

Possibly of greatest interest, is an exploration of whether the timing of flip-flop transitions in the outbursts of other BHTs can also be described to occur at integer multiples of a fundamental flip-flop period. So far, the association of the flip-flop timescale with the flip-flop transitions in the BHT Swfit J1658.2-4242, has not yet satisfied the 5σ significance that is often used to distinguish real effects from unlikely chance occurrences. Investigating the duration of flip-flops in literature and possible future observations will make it possible to determine whether they also occur at multiples of an underlying timescale. If this effect is also seen in other systems, that could provide significant assistance in explaining the physical mechanism of flip-flop transitions.

Finally, it remains to be seen whether flip-flops can occur in AGNs as well. Figs. 4.12, and 4.15 depict light curves of variable likely extragalactic sources, that appear to feature similar long-term variability properties as flip-flops feature in BHTs. However, it was not possible to unambiguously determine whether the observed variability in the two systems can be associated with the properties of the flip-flops observed in BHTs. A detailed long-term variability study over a larger region of the sky, involving a greater number of eRASSs than we used, could try to specifically search for signatures of flip-flop variability.

Bibliography

Abramowicz, M. A. and Fragile, P. C. (2013). Foundations of Black Hole Accretion Disk Theory. *Living Reviews in Relativity*, 16(1):1.

Agarwal, A., Rani, P., Prince, R., Stalin, C. S., Anupama, G. C., and Agrawal, V. (2021). A Possible Quasi-Periodic Oscillation in the X-ray Emission of 3C 120. *Galaxies*, 9(2):20.

Aird, J., Nandra, K., Laird, E. S., Georgakakis, A., Ashby, M. L. N., Barmby, P., Coil, A. L., Huang, J. S., Koekemoer, A. M., Steidel, C. C., and Willmer, C. N. A. (2010). The evolution of the hard X-ray luminosity function of AGN. *"Monthly Notices of the Royal Astronomical Society"*, 401(4):2531–2551.

Allevato, V., Paolillo, M., Papadakis, I., and Pinto, C. (2013). Measuring X-Ray Variability in Faint/Sparsely Sampled Active Galactic Nuclei. *"Astrophysical Journal"*, 771(1):9.

Antia, H. M., Yadav, J. S., Agrawal, P. C., Verdhan Chauhan, J., Manchanda, R. K., Chitnis, V., Paul, B., Dedhia, D., Shah, P., Gujar, V. M., Katoch, T., Kurhade, V. N., Madhwani, P., Manojkumar, T. K., Nikam, V. A., Pandya, A. S., Parmar, J. V., Pawar, D. M., Pahari, M., Misra, R., Navalgund, K. H., Pandiyan, R., Sharma, K. S., and Subbarao, K. (2017). Calibration of the Large Area X-Ray Proportional Counter (LAXPC) Instrument on board AstroSat. *"Astrophysical Journal Supplement"*, 231(1):10.

Antonucci, E., Gabriel, A. H., and Dennis, B. R. (1984). The energetics of chromospheric evaporation in solar flares. *"Astrophysical Journal"*, 287:917–925.

Antonucci, R. (1993). Unified models for active galactic nuclei and quasars. *"Annual Review of Astronomy & Astrophysics"*, 31:473–521.

Ashton, D. I. and Middleton, M. J. (2021). Searching for energy-resolved quasi-periodic oscillations in AGN. *"Monthly Notices of the Royal Astronomical Society"*, 501(4):5478–5499.

Bachetti, M., Harrison, F. A., Cook, R., Tomsick, J., Schmid, C., Grefenstette, B. W., Barret, D., Boggs, S. E., Christensen, F. E., Craig, W. W., Fabian, A. C., Fürst, F., Gandhi, P., Hailey, C. J., Kara, E., Maccarone, T. J., Miller, J. M., Pottschmidt, K., Stern, D., Uttley, P., Walton, D. J., Wilms, J., and Zhang, W. W. (2015). No Time for Dead Time: Timing Analysis of Bright Black Hole Binaries with NuSTAR. *"Astrophysical Journal"*, 800(2):109.

Barthelmy, S. D., Barbier, L. M., Cummings, J. R., Fenimore, E. E., Gehrels, N., Hullinger, D., Krimm, H. A., Markwardt, C. B., Palmer, D. M., Parsons, A., et al. (2005). The burst alert telescope (bat) on the swift midex mission. *Space Science Reviews*, 120(3):143–164.

Barthelmy, S. D., D'Avanzo, P., Deich, A., Gronwall, C., Melandri, A., Page, K. L., and Palmer, D. M. (2018). Swift Trigger 810300 is a probable new Galactic source: Swift J1658.2-4242. *GRB Coordinates Network, Circular Service, No. 22416, #1 (2018/February-0)*, 22416.

Basu, A., Bhattacharya, D., and Joshi, B. C. (2021). Absolute time calibration of LAXPC aboard AstroSat. *Journal of Astrophysics and Astronomy*, 42(2):61.

Belczynski, K., Repetto, S., Holz, D. E., O'Shaughnessy, R., Bulik, T., Berti, E., Fryer, C., and Dominik, M. (2016). Compact Binary Merger Rates: Comparison with LIGO/Virgo Upper Limits. *"Astrophysical Journal"*, 819(2):108.

Belloni, T., editor (2010). *The Jet Paradigm*, volume 794 of *Lecture Notes in Physics, Berlin Springer Verlag*.

Belloni, T. and Hasinger, G. (1990). Variability in the noise properties of Cygnus X-1. *"Astronomy & Astrophysics"*, 227:L33–L36.

Belloni, T., Méndez, M., King, A. R., van der Klis, M., and van Paradijs, J. (1997). A Unified Model for the Spectral Variability in GRS 1915+105. *"Astrophysical Journal Letters"*, 488(2):L109–L112.

Belloni, T., Psaltis, D., and van der Klis, M. (2002). A Unified Description of the Timing Features of Accreting X-Ray Binaries. *"Astrophysical Journal"*, 572(1):392–406.

Belloni, T. M. and Motta, S. E. (2016). *Transient Black Hole Binaries*, volume 440 of *Astrophysics and Space Science Library*, page 61.

Belloni, T. M., Motta, S. E., and Muñoz-Darias, T. (2011). Black hole transients. *Bulletin of the Astronomical Society of India*, 39:409–428.

Belloni, T. M., Sanna, A., and Méndez, M. (2012). High-frequency quasi-periodic oscillations in black hole binaries. *"Monthly Notices of the Royal Astronomical Society"*, 426(3):1701–1709.

Belloni, T. M. and Stella, L. (2014). Fast Variability from Black-Hole Binaries. *"Space Science Reviews"*, 183(1-4):43–60.

Beri, A. and Altamirano, D. (2018). The Low-mass X-ray binary Swift J1658.2-4242 returned to quiescence. *The Astronomer's Telegram*, 12072.

Boller, T., Freyberg, M. J., Trümper, J., Haberl, F., Voges, W., and Nandra, K. (2016). Second ROSAT all-sky survey (2RXS) source catalogue. *"Astronomy & Astrophysics"*, 588:A103.

Brunner, H., Liu, T., Lamer, G., Georgakakis, A., Merloni, A., Brusa, M., Bulbul, E., Dennerl, K., Friedrich, S., Liu, A., Maitra, C., Nandra, K., Ramos-Ceja, M. E., Sanders, J. S., Stewart, I. M., Boller, T., Buchner, J., Clerc, N., Comparat, J., Dwelly, T., Eckert, D., Finoguenov, A., Freyberg, M., Ghirardini, V., Gueguen, A., Haberl, F., Kreykenbohm, I., Krumpe, M., Osterhage, S., Pacaud, F., Predehl, P., Reiprich, T. H., Robrade, J., Salvato, M., Santangelo, A., Schrabback, T., Schwope, A., and Wilms, J. (2022). The eROSITA Final Equatorial Depth Survey (eFEDS). X-ray catalogue. *"Astronomy & Astrophysics"*, 661:A1.

Brunner, H., Liu, T., Lamer, G., Georgakakis, A., Merloni, A., Brusa, M., Bulbul, E., Dennerl, K., Friedrich, S., Liu, A., Maitra, C., Nandra, K., Ramos-Ceja, M. E., Sanders, J. S., Stewart, I. M., Boller, T., Buchner, J., Clerc, N., Comparat, J., Dwelly, T., Eckert, D., Finoguenov, A., Freyberg, M., Ghirardini, V., Gueguen, A., Haberl, F., Kreykenbohm, I., Krumpe, M., Osterhage, S., Pacaud, F., Predehl, P., Reiprich, T. H., Robrade, J., Salvato, M., Santangelo, A., Schrabback, T., Schwope, A., and Wilms, J. (2021). The eROSITA Final Equatorial Depth Survey (eFEDS): The X-ray catalog. *arXiv e-prints*, page arXiv:2106.14517.

Buchner, J. (2014). A statistical test for Nested Sampling algorithms. *arXiv e-prints*, page arXiv:1407.5459.

Buchner, J. (2019). Collaborative Nested Sampling: Big Data versus Complex Physical Models. *"Publications of the Astronomical Society of the Pacific"*, 131(1004):108005.

Buchner, J. (2021). UltraNest - a robust, general purpose Bayesian inference engine. *The Journal of Open Source Software*, 6(60):3001.

Buchner, J., Boller, T., Bogensberger, D., Malyali, A., Nandra, K., Wilms, J., Dwelly, T., and Liu, T. (2022). Systematic evaluation of variability detection methods for eROSITA. *"Astronomy & Astrophysics"*, 661:A18.

Burrows, D. N., Hill, J. E., Nousek, J. A., Kennea, J. A., Wells, A., Osborne, J. P., Abbey, A. F., Beardmore, A., Mukerjee, K., Short, A. D. T., Chincarini, G., Campana, S., Citterio, O., Moretti, A., Pagani, C., Tagliaferri, G., Giommi, P., Capalbi, M., Tamburelli, F., Angelini, L., Cusumano, G., Bräuninger, H. W., Burkert, W., and Hartner, G. D. (2005). The Swift X-Ray Telescope. *"Space Science Reviews"*, 120(3-4):165–195.

Cabanac, C., Henri, G., Petrucci, P. O., Malzac, J., Ferreira, J., and Belloni, T. M. (2010). Variability of X-ray binaries from an oscillating hot corona. *"Monthly Notices of the Royal Astronomical Society"*, 404(2):738–748.

Casella, P., Belloni, T., Homan, J., and Stella, L. (2004). A study of the low-frequency quasi-periodic oscillations in the X-ray light curves of the black hole candidate <ASTROBJ>XTE J1859+226</ASTROBJ>. *"Astronomy & Astrophysics"*, 426:587–600.

Casella, P., Belloni, T., and Stella, L. (2005). The ABC of Low-Frequency Quasi-periodic Oscillations in Black Hole Candidates: Analogies with Z Sources. *"Astrophysical Journal"*, 629(1):403–407.

Chen, W., Shrader, C. R., and Livio, M. (1997). The Properties of X-Ray and Optical Light Curves of X-Ray Novae. *"Astrophysical Journal"*, 491(1):312–338.

Chilingarian, I. V., Katkov, I. Y., Zolotukhin, I. Y., Grishin, K. A., Beletsky, Y., Boutsia, K., and Osip, D. J. (2018). A Population of Bona Fide Intermediate-mass Black Holes Identified as Low-luminosity Active Galactic Nuclei. *"Astrophysical Journal"*, 863(1):1.

Compton, A. H. (1923). A Quantum Theory of the Scattering of X-rays by Light Elements. *Physical Review*, 21(5):483–502.

Corbel, S., Kaaret, P., Jain, R. K., Bailyn, C. D., Fender, R. P., Tomsick, J. A., Kalemci, E., McIntyre, V., Campbell-Wilson, D., Miller, J. M., and McCollough, M. L. (2001). X-Ray States and Radio Emission in the Black Hole Candidate XTE J1550-564. *"Astrophysical Journal"*, 554(1):43–48.

Corral-Santana, J. M., Casares, J., Muñoz-Darias, T., Bauer, F. E., Martínez-Pais, I. G., and Russell, D. M. (2016). BlackCAT: A catalogue of stellar-mass black holes in X-ray transients. *"Astronomy & Astrophysics"*, 587:A61.

Crenshaw, D. M., Rodriguez-Pascual, P. M., Penton, S. V., Edelson, R. A., Alloin, D., Ayres, T. R., Clavel, J., Horne, K., Johnson, W. N., Kaspi, S., Korista, K. T., Kriss, G. A., Krolik, J. H., Malkan, M. A., Maoz, D., Netzer, H., O'Brien, P. T., Peterson, B. M., Reichert, G. A., Shull, J. M., Ulrich, M. H., Wamsteker, W., Warwick, R. S., Yaqoob, T., Balonek, T. J., Barr, P., Bromage, G. E., Carini, M., Carone, T. E., Cheng, F. Z., Chuvaev, K. K., Dietrich, M., Doroshenko, V. T., Dultzin-Hacyan, D., Filippenko, A. V., Gaskell, C. M., Glass, I. S., Goad, M. R., Hutchings, J., Kazanas, D., Kollatschny, W., Koratkar, A. P., Laor, A., Leighly, K., Lyutyi, V. M., MacAlpine, G. M., Malkov, Y. F., Martin, P. G., McCollum, B., Merkulova, N. I., Metik, L., Metlov, V. G., Miller, H. R., Morris, S. L., Oknyanskij, V. L., Penfold, J., Perez, E., Perola, G. C., Pike, G., Pogge, R. W., Pronik, I., Pronik, V. I., Ptak, R. L., Recondo-Gonzalez, M. C., Rodriguez-Espinoza, J. M., Rokaki, E. L., Roland, J., Sadun, A. C., Salamanca, I., Santos-Lleo, M., Sergeev, S. G., Smith, S. M., Snijders, M. A. J., Sparke, L. S., Stirpe, G. M., Stoner, R. E., Sun, W. H., van Groningen, E., Wagner, R. M., Wagner, S., Wanders, I., Welsh, W. F., Weymann, R. J., Wilkes, B. J., and Zheng, W. (1996). Multiwavelength Observations of Short-Timescale Variability in NGC 4151. I. Ultraviolet Observations. *"Astrophysical Journal"*, 470:322.

Czerny, B. (2004). The role of an accretion disk in AGN variability. *arXiv e-prints*, pages astro–ph/0409254.

D'Avanzo, P., Melandri, A., and Evans, P. A. (2018). Swift J1658.2-4242 (trigger 810300): Swift-XRT position. *GRB Coordinates Network*, 22417:1.

Del Santo, M., Malzac, J., Belloni, T., Bazzano, A., Homan, J., Jourdain, E., Méndez, M., and Ubertini, P. (2008). Spectral state transitions in GX 339-4. In *Microquasars and Beyond*, page 91.

Díaz Trigo, M. and Boirin, L. (2016). Accretion disc atmospheres and winds in low-mass X-ray binaries. *Astronomische Nachrichten*, 337(4-5):368.

Edelson, R. A., Krolik, J. H., and Pike, G. F. (1990). Broad-Band Properties of the CfA Seyfert Galaxies. III. Ultraviolet Variability. *"Astrophysical Journal"*, 359:86.

Esin, A. A., McClintock, J. E., and Narayan, R. (1997). Advection-Dominated Accretion and the Spectral States of Black Hole X-Ray Binaries: Application to Nova Muscae 1991. *"Astrophysical Journal"*, 489(2):865–889.

Event Horizon Telescope Collaboration (2019). First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole. *"Astrophysical Journal Letters"*, 875(1):L1.

Event Horizon Telescope Collaboration (2022). First Sagittarius A* Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole in the Center of the Milky Way. *"Astrophysical Journal Letters"*, 930(2):L12.

Fan, X. (2006). Evolution of high-redshift quasars. *"New Astronomy Reviews"*, 50(9-10):665–671.

Fender, R. (2010). 'Disc-Jet' Coupling in Black Hole X-Ray Binaries and Active Galactic Nuclei. In Belloni, T., editor, *Lecture Notes in Physics, Berlin Springer Verlag*, volume 794, page 115.

Fender, R., Corbel, S., Tzioumis, T., McIntyre, V., Campbell-Wilson, D., Nowak, M., Sood, R., Hunstead, R., Harmon, A., Durouchoux, P., and Heindl, W. (1999). Quenching of the Radio Jet during the X-Ray High State of GX 339-4. *"Astrophysical Journal Letters"*, 519(2):L165–L168.

Fender, R. P., Belloni, T. M., and Gallo, E. (2004). Towards a unified model for black hole X-ray binary jets. *"Monthly Notices of the Royal Astronomical Society"*, 355(4):1105–1118.

Fender, R. P., Homan, J., and Belloni, T. M. (2009). Jets from black hole X-ray binaries: testing, refining and extending empirical models for the coupling to X-rays. *"Monthly Notices of the Royal Astronomical Society"*, 396(3):1370–1382.

Ferrarese, L. and Merritt, D. (2000). A Fundamental Relation between Supermassive Black Holes and Their Host Galaxies. *"Astrophysical Journal Letters"*, 539(1):L9–L12.

Fragile, P. C., Etheridge, S. M., Anninos, P., Mishra, B., and Kluźniak, W. (2018). Relativistic, Viscous, Radiation Hydrodynamic Simulations of Geometrically Thin Disks. I. Thermal and Other Instabilities. *"Astrophysical Journal"*, 857(1):1.

Frater, R. H., Brooks, J. W., and Whiteoak, J. B. (1992). The Australia Telescope - Overview. *Journal of Electrical and Electronics Engineering Australia*, 12(2):103–112.

Fryer, C. L. (1999). Mass Limits For Black Hole Formation. *"Astrophysical Journal"*, 522(1):413–418.

Gaia Collaboration (2016). Gaia Data Release 1. Summary of the astrometric, photometric, and survey properties. *"Astronomy & Astrophysics"*, 595:A2.

Gaia Collaboration (2021). Gaia Early Data Release 3. Summary of the contents and survey properties. *"Astronomy & Astrophysics"*, 649:A1.

Geary, R. C. (1935). THE RATIO OF THE MEAN DEVIATION TO THE STANDARD DEVIATION AS A TEST OF NORMALITY. *Biometrika*, 27(3-4):310–332.

Gehrels, N., Chincarini, G., Giommi, P., Mason, K. O., Nousek, J. A., Wells, A. A., White, N. E., Barthelmy, S. D., Burrows, D. N., Cominsky, L. R., Hurley, K. C., Marshall, F. E., Mészáros, P., Roming, P. W. A., Angelini, L., Barbier, L. M., Belloni, T., Campana, S., Caraveo, P. A., Chester, M. M., Citterio, O., Cline, T. L., Cropper, M. S., Cummings, J. R., Dean, A. J., Feigelson, E. D., Fenimore, E. E., Frail, D. A., Fruchter, A. S., Garmire, G. P., Gendreau, K., Ghisellini, G., Greiner, J., Hill, J. E., Hunsberger, S. D., Krimm, H. A., Kulkarni, S. R., Kumar, P., Lebrun, F., Lloyd-Ronning, N. M., Markwardt, C. B., Mattson, B. J., Mushotzky, R. F., Norris, J. P., Osborne, J., Paczynski, B., Palmer, D. M., Park, H. S., Parsons, A. M., Paul, J., Rees, M. J., Reynolds, C. S., Rhoads, J. E., Sasseen, T. P., Schaefer, B. E., Short, A. T., Smale, A. P., Smith, I. A., Stella, L., Tagliaferri, G., Takahashi, T., Tashiro, M., Townsley, L. K., Tueller, J., Turner, M. J. L., Vietri, M., Voges, W., Ward, M. J., Willingale, R., Zerbi, F. M., and Zhang, W. W. (2004). The Swift Gamma-Ray Burst Mission. *"Astrophysical Journal"*, 611(2):1005–1020.

Gendreau, K. C., Arzoumanian, Z., and Okajima, T. (2012). The Neutron star Interior Composition ExploreR (NICER): an Explorer mission of opportunity for soft x-ray timing spectroscopy. In *"Proceedings of the SPIE"*, volume 8443 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 844313.

Gierliński, M., Middleton, M., Ward, M., and Done, C. (2008). A periodicity of ∼1hour in X-ray emission from the active galaxy RE J1034+396. *"Nature"*, 455(7211):369–371.

Gilfanov, M. (2010). *X-Ray Emission from Black-Hole Binaries*, volume 794, page 17.

Gomez, S., Nicholl, M., Short, P., Margutti, R., Alexander, K. D., Blanchard, P. K., Berger, E., Eftekhari, T., Schulze, S., Anderson, J., Arcavi, I., Chornock, R., Cowperthwaite,

P. S., Galbany, L., Herzog, L. J., Hiramatsu, D., Hosseinzadeh, G., Laskar, T., Müller Bravo, T. E., Patton, L., and Terreran, G. (2020). The Tidal Disruption Event AT 2018hyz II: Light-curve modelling of a partially disrupted star. *"Monthly Notices of the Royal Astronomical Society"*, 497(2):1925–1934.

González-Martín, O. and Vaughan, S. (2012). X-ray variability of 104 active galactic nuclei. XMM-Newton power-spectrum density profiles. *"Astronomy & Astrophysics"*, 544:A80.

Gou, L., McClintock, J. E., Steiner, J. F., Narayan, R., Cantrell, A. G., Bailyn, C. D., and Orosz, J. A. (2010). The Spin of the Black Hole in the Soft X-ray Transient A0620-00. *"Astrophysical Journal Letters"*, 718(2):L122–L126.

Grebenev, S. A., Mereminskiy, I. A., Prosvetov, A. V., Ducci, L., Bozzo, E., Savchenko, V., and Ferrigno, C. (2018). INTEGRAL observations of the new X-ray transient Swift J1658.2-4242. *The Astronomer's Telegram*, 11306:1.

Gupta, A. C., Tripathi, A., Wiita, P. J., Gu, M., Bambi, C., and Ho, L. C. (2018). Possible 1 hour quasi-periodic oscillation in narrow-line Seyfert 1 galaxy MCG-06-30-15. *"Astronomy & Astrophysics"*, 616:L6.

Harrison, F. A., Christensen, F. E., Craig, W., Hailey, C., Baumgartner, W., Chen, C. M. H., Chonko, J., Cook, W. R., Koglin, J., Madsen, K.-K., Pivavoroff, M., Boggs, S., and Smith, D. (2005). Development of the HEFT and NuSTAR focusing telescopes. *Experimental Astronomy*, 20(1-3):131–137.

Harrison, F. A., Craig, W. W., Christensen, F. E., Hailey, C. J., Zhang, W. W., Boggs, S. E., Stern, D., Cook, W. R., Forster, K., Giommi, P., Grefenstette, B. W., Kim, Y., Kitaguchi, T., Koglin, J. E., Madsen, K. K., Mao, P. H., Miyasaka, H., Mori, K., Perri, M., Pivovaroff, M. J., Puccetti, S., Rana, V. R., Westergaard, N. J., Willis, J., Zoglauer, A., An, H., Bachetti, M., Barrière, N. M., Bellm, E. C., Bhalerao, V., Brejnholt, N. F., Fuerst, F., Liebe, C. C., Markwardt, C. B., Nynka, M., Vogel, J. K., Walton, D. J., Wik, D. R., Alexander, D. M., Cominsky, L. R., Hornschemeier, A. E., Hornstrup, A., Kaspi, V. M., Madejski, G. M., Matt, G., Molendi, S., Smith, D. M., Tomsick, J. A., Ajello, M., Ballantyne, D. R., Baloković, M., Barret, D., Bauer, F. E., Blandford, R. D., Brandt, W. N., Brenneman, L. W., Chiang, J., Chakrabarty, D., Chenevez, J., Comastri, A., Dufour, F., Elvis, M., Fabian, A. C., Farrah, D., Fryer, C. L., Gotthelf, E. V., Grindlay, J. E., Helfand, D. J., Krivonos, R., Meier, D. L., Miller, J. M., Natalucci, L., Ogle, P., Ofek, E. O., Ptak, A., Reynolds, S. P., Rigby, J. R., Tagliaferri, G., Thorsett, S. E., Treister, E., and Urry, C. M. (2013). The Nuclear Spectroscopic Telescope Array (NuSTAR) High-energy X-Ray Mission. *"Astrophysical Journal"*, 770(2):103.

Heckman, T. M., Miley, G. K., van Breugel, W. J. M., and Butcher, H. R. (1981). Emission-line profiles and kinematics of the narrow-line region in Seyfert and radio galaxies. *"Astrophysical Journal"*, 247:403–418.

Hill, J. E., Burrows, D. N., Nousek, J. A., Abbey, A. F., Ambrosi, R. M., Bräuninger, H. W., Burkert, W., Campana, S., Cheruvu, C., Cusumano, G., Freyberg, M. J., Hartner, G. D., Klar, R., Mangels, C., Moretti, A., Mori, K., Morris, D. C., Short, A. D. T., Tagliaferri, G., Watson, D. J., Wood, P., and Wells, A. A. (2004). Readout modes and automated operation of the Swift X-ray Telescope. In Flanagan, K. A. and Siegmund, O. H. W., editors, *X-Ray and Gamma-Ray Instrumentation for Astronomy XIII*, volume 5165 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 217–231.

Homan, J. and Belloni, T. (2005). The Evolution of Black Hole States. *"Astrophysics and Space Science"*, 300(1-3):107–117.

Homan, J., Miller, J. M., Wijnands, R., van der Klis, M., Belloni, T., Steeghs, D., and Lewin, W. H. G. (2005). High- and Low-Frequency Quasi-periodic Oscillations in the X-Ray Light Curves of the Black Hole Transient H1743-322. *"Astrophysical Journal"*, 623(1):383–391.

Homan, J., Wijnands, R., van der Klis, M., Belloni, T., van Paradijs, J., Klein-Wolt, M., Fender, R., and Méndez, M. (2001). Correlated X-Ray Spectral and Timing Behavior of the Black Hole Candidate XTE J1550-564: A New Interpretation of Black Hole States. *"Astrophysical Journal Supplement"*, 132(2):377–402.

Hopkins, A. M., De Silva, G. M., Springob, C. M., Ryder, S. D., Watson, F. G., and Colless, M. M. (2012). The Telescopes and Processes of the Australian Astronomical Observatory. In *Organizations*, volume 1, pages 27–38.

Huang, Y., Qu, J. L., Zhang, S. N., Bu, Q. C., Chen, Y. P., Tao, L., Zhang, S., Lu, F. J., Li, T. P., Song, L. M., Xu, Y. P., Cao, X. L., Chen, Y., Liu, C. Z., Chang, H.-K., Yu, W. F., Weng, S. S., Hou, X., Kong, A. K. H., Xie, F. G., Zhang, G. B., ZHOU, J. F., Chang, Z., Chen, G., Chen, L., Chen, T. X., Chen, Y. B., Cui, W., Cui, W. W., Deng, J. K., Dong, Y. W., Du, Y. Y., Fu, M. X., Gao, G. H., Gao, H., Gao, M., Ge, M. Y., Gu, Y. D., Guan, J., Gungor, C., Guo, C. C., Han, D. W., Hu, W., Huo, J., Ji, J. F., Jia, S. M., Jiang, L. H., Jiang, W. C., Jin, J., Jin, Y. J., Li, B., Li, C. K., Li, G., Li, M. S., Li, W., Li, X., Li, X. B., Li, X. F., Li, Y. G., Li, Z. J., Li, Z. W., Liang, X. H., Liao, J. Y., Liu, G. Q., Liu, H. W., Liu, S. Z., Liu, X. J., Liu, Y., Liu, Y. N., Lu, B., Lu, X. F., Luo, T., Ma, X., Meng, B., Nang, Y., Nie, J. Y., Ou, G., Sai, N., Shang, R. C., Sun, L., Tan, Y., Tao, W., Tuo, Y. L., Wang, G. F., Wang, H. Y., Wang, J., Wang, W. S., Wang, Y. S., Wen, X. Y., Wu, B. B., Wu, M., Xiao, G. C., Xiong, S. L., Xu, H., Yan, L. L., Yang, J. W., Yang, S., Yang, Y. J., Zhang, A. M., Zhang, C. L., Zhang, C. M., Zhang, F., Zhang, H. M., Zhang, J., Zhang, Q., Zhang, T., Zhang, W., Zhang, W. C., Zhang, W. Z., Zhang, Y., Zhang, Y., Zhang, Y. F., Zhang, Y. J., Zhang, Z., Zhang, Z., Zhang, Z. L., Zhao, H. S., Zhao, J. L., Zhao, X. F., Zheng, S. J., Zhu, Y., Zhu, Y. X., Zou, C. L., and Insight-HXMT Collaboration (2018). INSIGHT-HXMT Observations of the New Black Hole Candidate MAXI J1535-571: Timing Analysis. *"Astrophysical Journal"*, 866:122.

Huppenkothen, D., Bachetti, M., Stevens, A. L., Migliari, S., and Balm, P. (2016). Stingray: Spectral-timing software.

Ingram, A., Done, C., and Fragile, P. C. (2009). Low-frequency quasi-periodic oscillations spectra and Lense-Thirring precession. *"Monthly Notices of the Royal Astronomical Society"*, 397(1):L101–L105.

Ingram, A. and van der Klis, M. (2013). An exact analytic treatment of propagating mass accretion rate fluctuations in X-ray binaries. *"Monthly Notices of the Royal Astronomical Society"*, 434(2):1476–1485.

Ingram, A., van der Klis, M., Middleton, M., Done, C., Altamirano, D., Heil, L., Uttley, P., and Axelsson, M. (2016). A quasi-periodic modulation of the iron line centroid energy in the black hole binary H1743-322. *"Monthly Notices of the Royal Astronomical Society"*, 461(2):1967–1980.

Israel, W. (1967). Event Horizons in Static Vacuum Space-Times. *Physical Review*, 164(5):1776–1779.

Ivanova, N., Justham, S., Chen, X., De Marco, O., Fryer, C. L., Gaburov, E., Ge, H., Glebbeek, E., Han, Z., Li, X. D., Lu, G., Marsh, T., Podsiadlowski, P., Potter, A., Soker, N., Taam, R., Tauris, T. M., van den Heuvel, E. P. J., and Webbink, R. F. (2013). Common envelope evolution: where we stand and how we can move forward. *"Astronomy & Astrophysics"*, 21:59.

Janka, H.-T. (2013). Natal kicks of stellar mass black holes by asymmetric mass ejection in fallback supernovae. *"Monthly Notices of the Royal Astronomical Society"*, 434(2):1355–1361.

Jansen, F., Lumb, D., Altieri, B., Clavel, J., Ehle, M., Erd, C., Gabriel, C., Guainazzi, M., Gondoin, P., Much, R., Munoz, R., Santos, M., Schartel, N., Texier, D., and Vacanti, G. (2001). XMM-Newton observatory. I. The spacecraft and operations. *"Astronomy & Astrophysics"*, 365:L1–L6.

Jentzsch, F. and Nähring, E. (1929). Reflexion von Röntgenstrahlen. *Naturwissenschaften*, 17(50):980–980.

Jin, C., Ponti, G., Haberl, F., and Smith, R. a. (2017). Probing the interstellar dust towards the Galactic Centre: dust-scattering halo around AX J1745.6-2901. *"Monthly Notices of the Royal Astronomical Society"*, 468(3):2532–2551.

Jin, C., Ponti, G., Li, G., and Bogensberger, D. (2019). Exploring the Interstellar Medium Using an Asymmetric X-Ray Dust Scattering Halo. *"Astrophysical Journal"*, 875(2):157.

Jithesh, V., Maqbool, B., Misra, R., T, A. R., Mall, G., and James, M. (2019). Spectral and Timing Properties of the Galactic X-Ray Transient Swift J1658.2-4242 Using Astrosat Observations. *"Astrophysical Journal"*, 887(1):101.

Kalamkar, M., Homan, J., Altamirano, D., van der Klis, M., Casella, P., and Linares, M. (2011). The Identification of MAXI J1659-152 as a Black Hole Candidate. *"Astrophysical Journal Letters"*, 731(1):L2.

Karolinski, N. and Zhu, W. (2020). Detecting isolated stellar-mass black holes in the absence of microlensing parallax effect. *"Monthly Notices of the Royal Astronomical Society"*, 498(1):L25–L30.

Kasen, D. and Woosley, S. E. (2009). Type II Supernovae: Model Light Curves and Standard Candle Relationships. *"Astrophysical Journal"*, 703(2):2205–2216.

Kirchner, J. W. (2005). Aliasing in 1/ f^{α} noise spectra: Origins, consequences, and remedies. *"Physical Review E"*, 71(6):066110.

Knoetig, M. L. (2014). Signal Discovery, Limits, and Uncertainties with Sparse On/Off Measurements: An Objective Bayesian Analysis. *"Astrophysical Journal"*, 790(2):106.

Kormendy, J. and Ho, L. C. (2013). Coevolution (Or Not) of Supermassive Black Holes and Host Galaxies. *"Annual Review of Astronomy & Astrophysics"*, 51(1):511–653.

Kovács, A. and Szapudi, I. (2015). Star-galaxy separation strategies for WISE-2MASS all-sky infrared galaxy catalogues. *"Monthly Notices of the Royal Astronomical Society"*, 448(2):1305–1313.

Kovács, Z., Gergely, L., and Biermann, P. L. (2011). Maximal spin and energy conversion efficiency in a symbiotic system of black hole, disc and jet. *"Monthly Notices of the Royal Astronomical Society"*, 416(2):991–1009.

Kraft, R. P., Burrows, D. N., and Nousek, J. A. (1991). Determination of Confidence Limits for Experiments with Low Numbers of Counts. *"Astrophysical Journal"*, 374:344.

Krimm, H. A., Holland, S. T., Corbet, R. H. D., Pearlman, A. B., Romano, P., Kennea, J. A., Bloom, J. S., Barthelmy, S. D., Baumgartner, W. H., Cummings, J. R., Gehrels, N., Lien, A. Y., Markwardt, C. B., Palmer, D. M., Sakamoto, T., Stamatikos, M., and Ukwatta, T. N. (2013). The Swift/BAT Hard X-Ray Transient Monitor. *"Astrophysical Journal Supplement"*, 209(1):14.

Kubota, A., Tanaka, Y., Makishima, K., Ueda, Y., Dotani, T., Inoue, H., and Yamaoka, K. (1998). Evidence for a Black Hole in the X-Ray Transient GRS 1009-45. *"Publications of the Astronomical Society of Japan"*, 50:667–673.

Kuulkers, E., Kouveliotou, C., Belloni, T., Cadolle Bel, M., Chenevez, J., Díaz Trigo, M., Homan, J., Ibarra, A., Kennea, J. A., Muñoz-Darias, T., Ness, J. U., Parmar, A. N., Pollock, A. M. T., van den Heuvel, E. P. J., and van der Horst, A. J. (2013). MAXI J1659-152: the shortest orbital period black-hole transient in outburst. *"Astronomy & Astrophysics"*, 552:A32.

Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Sutherland, P. G., Kahn, S., and Grindlay, J. E. (1983). On searches for pulsed emission with application to four globular cluster X-ray sources : NGC 1851, 6441, 6624 and 6712. *"Astrophysical Journal"*, 266:160–170.

Lewin, W. H. G., van Paradijs, J., and van der Klis, M. (1988). A review of quasi-periodic oscillations in low-mass X-ray binaries. *"Space Science Reviews"*, 46(3-4):273–377.

Lewis, I. J., Cannon, R. D., Taylor, K., Glazebrook, K., Bailey, J. A., Baldry, I. K., Barton, J. R., Bridges, T. J., Dalton, G. B., Farrell, T. J., Gray, P. M., Lankshear, A., McCowage, C., Parry, I. R., Sharples, R. M., Shortridge, K., Smith, G. A., Stevenson, J., Straede, J. O., Waller, L. G., Whittard, J. D., Wilcox, J. K., and Willis, K. C. (2002). The Anglo-Australian Observatory 2dF facility. *"Monthly Notices of the Royal Astronomical Society"*, 333(2):279–299.

Lien, A. Y., Kennea, J. A., Barthelmy, S. D., Beardmore, A. P., Evans, P. A., Krimm, H. A., Malesani, D., Markwardt, C. B., Palmer, D. M., Page, K. L., Sakamoto, T., Stamatikos, M., and Ukwatta, T. N. (2018). Swift J1658.2-4242: Swift BAT and XRT refined analysis. *GRB Coordinates Network*, 22419:1.

LIGO Scientific Collaboration and Virgo Collaboration (2016). Observation of Gravitational Waves from a Binary Black Hole Merger. *"Physical Review Letters"*, 116(6):061102.

LIGO Scientific Collaboration and Virgo Collaboration (2020). GW190521: A Binary Black Hole Merger with a Total Mass of 150 $M_{\odot}$. *"Physical Review Letters"*, 125(10):101102.

Lintott, C. J., Schawinski, K., Keel, W., van Arkel, H., Bennert, N., Edmondson, E., Thomas, D., Smith, D. J. B., Herbert, P. D., Jarvis, M. J., Virani, S., Andreescu, D., Bamford, S. P., Land, K., Murray, P., Nichol, R. C., Raddick, M. J., Slosar, A., Szalay, A., and Vandenberg, J. (2009). Galaxy Zoo: ‘Hanny’s Voorwerp’, a quasar light echo? *"Monthly Notices of the Royal Astronomical Society"*, 399(1):129–140.

Lipunov, V., Gorbovskoy, E., Ershova, O., Tyurina, N., Kornilov, V., Zimnukhov, D., Krylov, A., Gorbunov, I., Balanutsa, P., Kuznetsov, A., Chazov, V., Kuvshinov, D., Yurkov, V., Gabovich, A., Sergienko, Y., Podesta, R., Lopez, C., Podesta, F., Levato, H., Saffe, C., Juan, S., Buckley, D., Potter, S., Rebolo, R., Serra, M., Gres, O., Budnev, N. M., Ishmuhametova, Y., Tlatov, A., Senik, V., and Dormidontov, D. (2018). Swift trigger 810300: MASTER-OAFA robotic report. *GRB Coordinates Network*, 22418:1.

Lipunov, V., Kornilov, V., Gorbovskoy, E., Shatskij, N., Kuvshinov, D., Tyurina, N., Belinski, A., Krylov, A., Balanutsa, P., Chazov, V., Kuznetsov, A., Kortunov, P., Sankovich, A., Tlatov, A., Parkhomenko, A., Krushinsky, V., Zalozhnyh, I., Popov, A., Kopytova, T., Ivanov, K., Yazev, S., and Yurkov, V. (2010). Master Robotic Net. *Advances in Astronomy*, 2010:349171.

Liu, H., Jiang, J., Zhang, Z., Bambi, C., Ji, L., Kong, L., and Zhang, S. (2022a). Rapidly alternating flux states of GX 339-4 during its 2021 outburst captured by Insight-HXMT. *"Monthly Notices of the Royal Astronomical Society"*, 513(3):4308–4317.

Liu, J., Zhang, H., Howard, A. W., Bai, Z., Lu, Y., Soria, R., Justham, S., Li, X., Zheng, Z., Wang, T., Belczynski, K., Casares, J., Zhang, W., Yuan, H., Dong, Y., Lei, Y., Isaacson, H., Wang, S., Bai, Y., Shao, Y., Gao, Q., Wang, Y., Niu, Z., Cui, K., Zheng, C., Mu, X., Zhang, L., Wang, W., Heger, A., Qi, Z., Liao, S., Lattanzi, M., Gu, W.-M., Wang, J., Wu, J., Shao, L., Shen, R., Wang, X., Bregman, J., Di Stefano, R., Liu, Q., Han, Z., Zhang, T., Wang, H., Ren, J., Zhang, J., Zhang, J., Wang, X., Cabrera-Lavers, A., Corradi, R., Rebolo, R., Zhao, Y., Zhao, G., Chu, Y., and Cui, X. (2019). A wide star-black-hole binary system from radial-velocity measurements. *"Nature"*, 575(7784):618–621.

Liu, T., Merloni, A., Comparat, J., Nandra, K., Sanders, J., Lamer, G., Buchner, J., Dwelly, T., Freyberg, M., Malyali, A., Georgakakis, A., Salvato, M., Brunner, H., Brusa, M., Klein, M., Ghirardini, V., Clerc, N., Pacaud, F., Bulbul, E., Liu, A., Schwope, A., Robrade, J., Wilms, J., Dauser, T., Ramos-Ceja, M. E., Reiprich, T. H., Boller, T., and Wolf, J. (2022b). Establishing the X-ray source detection strategy for eROSITA with simulations. *"Astronomy & Astrophysics"*, 661:A27.

Liu, T., Merloni, A., Comparat, J., Nandra, K., Sanders, J. S., Lamer, G., Buchner, J., Dwelly, T., Freyberg, M., Malyali, A., Georgakakis, A., Salvato, M., Brunner, H., Brusa, M., Klein, M., Ghirardini, V., Clerc, N., Pacaud, F., Bulbul, E., Liu, A., Schwope, A., Robrade, J., Wilms, J., Dauser, T., Ramos-Ceja, M. E., Reiprich, T. H., Boller, T., and Wolf, J. (2021). Establishing the X-ray Source Detection Strategy for eROSITA with Simulations. *arXiv e-prints*, page arXiv:2106.14528.

Lodato, G., King, A. R., and Pringle, J. E. (2009). Stellar disruption by a supermassive black hole: is the light curve really proportional to $t^{-5/3}$? *"Monthly Notices of the Royal Astronomical Society"*, 392(1):332–340.

Lodato, G. and Rossi, E. M. (2011). Multiband light curves of tidal disruption events. *"Monthly Notices of the Royal Astronomical Society"*, 410(1):359–367.

Lomb, N. R. (1976). Least-Squares Frequency Analysis of Unequally Spaced Data. *"Astrophysics and Space Science"*, 39(2):447–462.

Lu, Y. and Yu, Q. (2001). The relationship between X-ray variability and the central black hole mass. *"Monthly Notices of the Royal Astronomical Society"*, 324(3):653–658.

Magorrian, J., Tremaine, S., Richstone, D., Bender, R., Bower, G., Dressler, A., Faber, S. M., Gebhardt, K., Green, R., Grillmair, C., Kormendy, J., and Lauer, T. (1998). The Demography of Massive Dark Objects in Galaxy Centers. *"The Astronomical Journal"*, 115(6):2285–2305.

Mainzer, A., Bauer, J., Cutri, R. M., Grav, T., Masiero, J., Beck, R., Clarkson, P., Conrow, T., Dailey, J., Eisenhardt, P., Fabinsky, B., Fajardo-Acosta, S., Fowler, J., Gelino, C., Grillmair, C., Heinrichsen, I., Kendall, M., Kirkpatrick, J. D., Liu, F., Masci, F., McCallon, H., Nugent, C. R., Papin, M., Rice, E., Royer, D., Ryan, T., Sevilla, P., Sonnett, S., Stevenson, R., Thompson, D. B., Wheelock, S., Wiemer, D., Wittman, M., Wright, E., and Yan, L. (2014). Initial Performance of the NEOWISE Reactivation Mission. *"Astrophysical Journal"*, 792(1):30.

Mainzer, A., Grav, T., Bauer, J., Masiero, J., McMillan, R. S., Cutri, R. M., Walker, R., Wright, E., Eisenhardt, P., Tholen, D. J., Spahr, T., Jedicke, R., Denneau, L., DeBaun, E., Elsbury, D., Gautier, T., Gomillion, S., Hand, E., Mo, W., Watkins, J., Wilkins, A., Bryngelson, G. L., Del Pino Molina, A., Desai, S., Gómez Camus, M., Hidalgo, S. L., Konstantopoulos, I., Larsen, J. A., Maleszewski, C., Malkan, M. A., Mauduit, J. C., Mullan, B. L., Olszewski, E. W., Pforr, J., Saro, A., Scotti, J. V., and Wasserman, L. H. (2011). NEOWISE Observations of Near-Earth Objects: Preliminary Results. *"Astrophysical Journal"*, 743(2):156.

Marconi, A., Risaliti, G., Gilli, R., Hunt, L. K., Maiolino, R., and Salvati, M. (2004). Local supermassive black holes, relics of active galactic nuclei and the X-ray background. *"Monthly Notices of the Royal Astronomical Society"*, 351(1):169–185.

Marocco, F., Eisenhardt, P. R. M., Fowler, J. W., Kirkpatrick, J. D., Meisner, A. M., Schlafly, E. F., Stanford, S. A., Garcia, N., Caselden, D., Cushing, M. C., Cutri, R. M., Faherty, J. K., Gelino, C. R., Gonzalez, A. H., Jarrett, T. H., Koontz, R., Mainzer, A., Marchese, E. J., Mobasher, B., Schlegel, D. J., Stern, D., Teplitz, H. I., and Wright, E. L. (2021). The CatWISE2020 Catalog. *"Astrophysical Journal Supplement"*, 253(1):8.

Martí-Vidal, I., Vlemmings, W. H. T., Muller, S., and Casey, S. (2014). UVMULTIFIT: A versatile tool for fitting astronomical radio interferometric data. *"Astronomy & Astrophysics"*, 563:A136.

Masetti, N., Mason, E., Morelli, L., Cellone, S. A., McBride, V. A., Palazzi, E., Bassani, L., Bazzano, A., Bird, A. J., Charles, P. A., Dean, A. J., Galaz, G., Gehrels, N., Landi, R., Malizia, A., Minniti, D., Panessa, F., Romero, G. E., Stephen, J. B., Ubertini, P., and Walter, R. (2008). Unveiling the nature of INTEGRAL objects through optical spectroscopy. VI. A multi-observatory identification campaign. *"Astronomy & Astrophysics"*, 482(1):113–132.

Masoura, V. A., Mountrichas, G., Georgantopoulos, I., and Plionis, M. (2021). Relation between AGN type and host galaxy properties. *"Astronomy & Astrophysics"*, 646:A167.

Matthews, T. A. and Sandage, A. R. (1963). Optical Identification of 3C 48, 3C 196, and 3C 286 with Stellar Objects. *"Astrophysical Journal"*, 138:30.

McHardy, I. (1988). EXOSAT observations of variability in active galactic nuclei. *"Memorie della Societa Astronomica Italiana"*, 59:239–260.

McHardy, I. M., Koerding, E., Knigge, C., Uttley, P., and Fender, R. P. (2006). Active galactic nuclei as scaled-up Galactic black holes. *"Nature"*, 444(7120):730–732.

McMahon, R. G., Banerji, M., Gonzalez, E., Koposov, S. E., Bejar, V. J., Lodieu, N., Rebolo, R., and VHS Collaboration (2013). First Scientific Results from the VISTA Hemisphere Survey (VHS). *The Messenger*, 154:35–37.

McMahon, R. G., Banerji, M., Gonzalez, E., Koposov, S. E., Bejar, V. J., Lodieu, N., Rebolo, R., and VHS Collaboration (2021). VizieR Online Data Catalog: The VISTA Hemisphere Survey (VHS) catalog DR5 (McMahon+, 2020). *VizieR Online Data Catalog*, page II/367.

McMullin, J. P., Waters, B., Schiebel, D., Young, W., and Golap, K. (2007). CASA Architecture and Applications. In Shaw, R. A., Hill, F., and Bell, D. J., editors, *Astronomical Data Analysis Software and Systems XVI*, volume 376 of *Astronomical Society of the Pacific Conference Series*, page 127.

Merloni, A., Predehl, P., Becker, W., Böhringer, H., Boller, T., Brunner, H., Brusa, M., Dennerl, K., Freyberg, M., Friedrich, P., Georgakakis, A., Haberl, F., Hasinger, G., Meidinger, N., Mohr, J., Nandra, K., Rau, A., Reiprich, T. H., Robrade, J., Salvato, M., Santangelo, A., Sasaki, M., Schwope, A., Wilms, J., and German eROSITA Consortium, t. (2012). eROSITA Science Book: Mapping the Structure of the Energetic Universe. *arXiv e-prints*, page arXiv:1209.3114.

Miyamoto, S., Kimura, K., Kitamoto, S., Dotani, T., and Ebisawa, K. (1991). X-Ray Variability of GX 339-4 in Its Very High State. *"Astrophysical Journal"*, 383:784.

Motta, S., Homan, J., Muñoz Darias, T., Casella, P., Belloni, T. M., Hiemstra, B., and Méndez, M. (2012). Discovery of two simultaneous non-harmonically related quasi-periodic oscillations in the 2005 outburst of the black hole binary GRO J1655-40. *"Monthly Notices of the Royal Astronomical Society"*, 427(1):595–606.

Motta, S., Muñoz-Darias, T., Casella, P., Belloni, T., and Homan, J. (2011). Low-frequency oscillations in black holes: a spectral-timing approach to the case of GX 339-4. *"Monthly Notices of the Royal Astronomical Society"*, 418(4):2292–2307.

Motta, S. E. (2016). Quasi periodic oscillations in black hole binaries. *Astronomische Nachrichten*, 337(4-5):398.

Muñoz-Darias, T., Coriat, M., Plant, D. S., Ponti, G., Fender, R. P., and Dunn, R. J. H. (2013). Inclination and relativistic effects in the outburst evolution of black hole transients. *"Monthly Notices of the Royal Astronomical Society"*, 432(2):1330–1337.

Murray, S. S., Austin, G. K., Chappell, J. H., Gomes, J. J., Kenter, A. T., Kraft, R. P., Meehan, G. R., Zombeck, M. V., Fraser, G. W., and Serio, S. (2000). In-flight performance of the Chandra high-resolution camera. In Truemper, J. E. and Aschenbach, B., editors, *X-Ray Optics, Instruments, and Missions III*, volume 4012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 68–80.

Nandra, K., George, I. M., Mushotzky, R. F., Turner, T. J., and Yaqoob, T. (1997). ASCA Observations of Seyfert 1 Galaxies. I. Data Analysis, Imaging, and Timing. *"Astrophysical Journal"*, 476(1):70–82.

Narayan, R., McClintock, J. E., and Yi, I. (1996). A New Model for Black Hole Soft X-Ray Transients in Quiescence. *"Astrophysical Journal"*, 457:821.

Nenkova, M., Sirocky, M. M., Nikutta, R., Ivezić, Ž., and Elitzur, M. (2008). AGN Dusty Tori. II. Observational Implications of Clumpiness. *"Astrophysical Journal"*, 685(1):160–180.

Nespoli, E., Belloni, T., Homan, J., Miller, J. M., Lewin, W. H. G., Méndez, M., and van der Klis, M. (2003). A transient variable 6 Hz QPO from GX 339-4. *"Astronomy & Astrophysics"*, 412:235–240.

Nidever, D. L., Dey, A., Fasbender, K., Juneau, S., Meisner, A. M., Wishart, J., Scott, A., Matt, K., Nikutta, R., and Pucha, R. (2021). Second Data Release of the All-sky NOIRLab Source Catalog. *"The Astronomical Journal"*, 161(4):192.

Okajima, T., Soong, Y., Balsamo, E. R., Enoto, T., Olsen, L., Koenecke, R., Lozipone, L., Kearney, J., Fitzsimmons, S., Numata, A., Kenyon, S. J., Arzoumanian, Z., and Gendreau, K. (2016). Performance of NICER flight x-ray concentrator. In den Herder, J.-W. A., Takahashi, T., and Bautz, M., editors, *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*, volume 9905 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99054X.

O'Neill, P. M., Nandra, K., Papadakis, I. E., and Turner, T. J. (2005). The relationship between X-ray variability amplitude and black hole mass in active galactic nuclei. *"Monthly Notices of the Royal Astronomical Society"*, 358(4):1405–1416.

Oppenheimer, J. R. and Snyder, H. (1939). On Continued Gravitational Contraction. *Physical Review*, 56(5):455–459.

Oppenheimer, J. R. and Volkoff, G. M. (1939). On Massive Neutron Cores. *Physical Review*, 55(4):374–381.

Padovani, P., Alexander, D. M., Assef, R. J., De Marco, B., Giommi, P., Hickox, R. C., Richards, G. T., Smolčić, V., Hatziminaoglou, E., Mainieri, V., and Salvato, M. (2017). Active galactic nuclei: what's in a name? *"Astronomy & Astrophysics"*, 25(1):2.

Paolillo, M., Papadakis, I., Brandt, W. N., Luo, B., Xue, Y. Q., Tozzi, P., Shemmer, O., Allevato, V., Bauer, F. E., Comastri, A., Gilli, R., Koekemoer, A. M., Liu, T., Vignali, C., Vito, F., Yang, G., Wang, J. X., and Zheng, X. C. (2017). Tracing the accretion history of supermassive black holes through X-ray variability: results from the ChandraDeep Field-South. *"Monthly Notices of the Royal Astronomical Society"*, 471(4):4398–4411.

Paolillo, M., Schreier, E. J., Giacconi, R., Koekemoer, A. M., and Grogin, N. A. (2004). Prevalence of X-Ray Variability in the Chandra Deep Field-South. *"Astrophysical Journal"*, 611(1):93–106.

Papadakis, I. E. (2004). The scaling of the X-ray variability with black hole mass in active galactic nuclei. *"Monthly Notices of the Royal Astronomical Society"*, 348(1):207–213.

Papadakis, I. E., Brinkmann, W., Negoro, H., and Gliozzi, M. (2002). Detection of a high frequency break in the X-ray power spectrum of Ark 564. *"Astronomy & Astrophysics"*, 382:L1–L4.

Pauli, W. (1925). Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren. *Zeitschrift fur Physik*, 31(1):765–783.

Pavlinsky, M., Tkachenko, A., Levin, V., Alexandrovich, N., Arefiev, V., Babyshkin, V., Batanov, O., Bodnar, Y., Bogomolov, A., Bubnov, A., Buntov, M., Burenin, R., Chelovekov, I., Chen, C. T., Drozdova, T., Ehlert, S., Filippova, E., Frolov, S., Gamkov, D., Garanin, S., Garin, M., Glushenko, A., Gorelov, A., Grebenev, S., Grigorovich, S., Gureev, P., Gurova, E., Ilkaev, R., Katasonov, I., Krivchenko, A., Krivonos, R., Korotkov, F., Kudelin, M., Kuznetsova, M., Lazarchuk, V., Lomakin, I., Lapshov, I., Lipilin, V., Lutovinov, A., Mereminskiy, I., Molkov, S., Nazarov, V., Oleinikov, V., Pikalov, E., Ramsey, B. D., Roiz, I., Rotin, A., Ryadov, A., Sankin, E., Sazonov, S., Sedov, D., Semena, A., Semena, N., Serbinov, D., Shirshakov, A., Shtykovsky, A., Shvetsov, A., Sunyaev, R., Swartz, D. A., Tambov, V., Voron, V., and Yaskovich, A. (2021). The ART-XC telescope on board the SRG observatory. *"Astronomy & Astrophysics"*, 650:A42.

Penrose, R. (1965). Gravitational Collapse and Space-Time Singularities. *"Physical Review Letters"*, 14(3):57–59.

Peterson, B. M. (2006). The Broad-Line Region in Active Galactic Nuclei. In Alloin, D., editor, *Physics of Active Galactic Nuclei at all Scales*, volume 693, page 77.

Phinney, E. S. (1989). Manifestations of a Massive Black Hole in the Galactic Center. In Morris, M., editor, *The Center of the Galaxy*, volume 136, page 543.

Ponti, G., Bianchi, S., Muñoz-Darias, T., De, K., Fender, R., and Merloni, A. (2016). High ionisation absorption in low mass X-ray binaries. *Astronomische Nachrichten*, 337(4-5):512–517.

Ponti, G., Fender, R. P., Begelman, M. C., Dunn, R. J. H., Neilsen, J., and Coriat, M. (2012a). Ubiquitous equatorial accretion disc winds in black hole soft states. *"Monthly Notices of the Royal Astronomical Society"*, 422(1):L11–L15.

Ponti, G., Papadakis, I., Bianchi, S., Guainazzi, M., Matt, G., Uttley, P., and Bonilla, N. F. (2012b). CAIXA: a catalogue of AGN in the XMM-Newton archive. III. Excess variance analysis. *"Astronomy & Astrophysics"*, 542:A83.

Poutanen, J., Krolik, J. H., and Ryde, F. (1997). The nature of spectral transitions in accreting black holes: the case of CYG X-1. *"Monthly Notices of the Royal Astronomical Society"*, 292(1):L21–L25.

Poutanen, J. and Svensson, R. (1996). The Two-Phase Pair Corona Model for Active Galactic Nuclei and X-Ray Binaries: How to Obtain Exact Solutions. *"Astrophysical Journal"*, 470:249.

Predehl, P., Andritschke, R., Arefiev, V., Babyshkin, V., Batanov, O., Becker, W., Böhringer, H., Bogomolov, A., Boller, T., Borm, K., Bornemann, W., Bräuninger, H., Brüggen, M., Brunner, H., Brusa, M., Bulbul, E., Buntov, M., Burwitz, V., Burkert, W., Clerc, N., Churazov, E., Coutinho, D., Dauser, T., Dennerl, K., Doroshenko, V., Eder, J., Emberger, V., Eraerds, T., Finoguenov, A., Freyberg, M., Friedrich, P., Friedrich, S., Fürmetz, M., Georgakakis, A., Gilfanov, M., Granato, S., Grossberger, C., Gueguen, A., Gureev, P., Haberl, F., Hälker, O., Hartner, G., Hasinger, G., Huber, H., Ji, L., Kienlin, A. v., Kink, W., Korotkov, F., Kreykenbohm, I., Lamer, G., Lomakin, I., Lapshov, I., Liu, T., Maitra, C., Meidinger, N., Menz, B., Merloni, A., Mernik, T., Mican, B., Mohr, J., Müller, S., Nandra, K., Nazarov, V., Pacaud, F., Pavlinsky, M., Perinati, E., Pfeffermann, E., Pietschner, D., Ramos-Ceja, M. E., Rau, A., Reiffers, J., Reiprich, T. H., Robrade, J., Salvato, M., Sanders, J., Santangelo, A., Sasaki, M., Scheuerle, H., Schmid, C., Schmitt, J., Schwope, A., Shirshakov, A., Steinmetz, M., Stewart, I., Strüder, L., Sunyaev, R., Tenzer, C., Tiedemann, L., Trümper, J., Voron, V., Weber, P., Wilms, J., and Yaroshenko, V. (2021). The eROSITA X-ray telescope on SRG. *"Astronomy & Astrophysics"*, 647:A1.

Rees, M. J. (1988). Tidal disruption of stars by black holes of 10^6-10^8 solar masses in nearby galaxies. *"Nature"*, 333(6173):523–528.

Rees, M. J., Begelman, M. C., Blandford, R. D., and Phinney, E. S. (1982). Ion-supported tori and the origin of radio jets. *"Nature"*, 295(5844):17–21.

Remillard, R. A. and McClintock, J. E. (2006). X-Ray Properties of Black-Hole Binaries. *"Annual Review of Astronomy & Astrophysics"*, 44(1):49–92.

Russell, T. D., Miller-Jones, J. C. A., Sivakoff, G. R., and Tetarenko, A. J. (2018). ATCA radio detection of the new X-ray transient Swift J1658.2-4242. *The Astronomer's Telegram*, 11322:1.

Russell, T. D., Soria, R., Miller-Jones, J. C. A., Curran, P. A., Markoff, S., Russell, D. M., and Sivakoff, G. R. (2014). The accretion-ejection coupling in the black hole candidate X-ray binary MAXI J1836-194. *"Monthly Notices of the Royal Astronomical Society"*, 439(2):1390–1402.

Russell, T. D., Tetarenko, A. J., Miller-Jones, J. C. A., Sivakoff, G. R., Parikh, A. S., Rapisarda, S., Wijnand s, R., Corbel, S., Tremou, E., Altamirano, D., Baglio, M. C., Ceccobello, C., Degenaar, N., van den Eijnden, J., Fender, R., Heywood, I., Krimm, H. A., Lucchini, M., Markoff, S., Russell, D. M., Soria, R., and Woudt, P. A. (2019). Disk-Jet Coupling in the 2017/2018 Outburst of the Galactic Black Hole Candidate X-Ray Binary MAXI J1535-571. *"Astrophysical Journal"*, 883(2):198.

Sahu, K. C., Anderson, J., Casertano, S., Bond, H. E., Udalski, A., Dominik, M., Calamida, A., Bellini, A., and Brown, T. M. (2022). An Isolated Stellar-Mass Black Hole Detected Through Astrometric Microlensing. *arXiv e-prints*, page arXiv:2201.13296.

Salvato, M., Buchner, J., Budavári, T., Dwelly, T., Merloni, A., Brusa, M., Rau, A., Fotopoulou, S., and Nandra, K. (2018). Finding counterparts for all-sky X-ray surveys with NWAY: a Bayesian algorithm for cross-matching multiple catalogues. *"Monthly Notices of the Royal Astronomical Society"*, 473(4):4937–4955.

Salvato, M., Wolf, J., Dwelly, T., Georgakakis, A., Brusa, M., Merloni, A., Liu, T., Toba, Y., Nandra, K., Lamer, G., Buchner, J., Schneider, C., Freund, S., Rau, A., Schwope, A., Nishizawa, A., Klein, M., Arcodia, R., Comparat, J., Musiimenta, B., Nagao, T., Brunner, H., Malyali, A., Finoguenov, A., Anderson, S., Shen, Y., Ibarra-Mendel, H., Trump, J., Brandt, W. N., Urry, C. M., Rivera, C., Krumpe, M., Urrutia, T., Miyaji, T., Ichikawa, K., Schneider, D. P., Fresco, A., Wilms, J., Boller, T., Haase, J., Brownstein, J., Lane, R. R., Bizyaev, D., and Nitschelm, C. (2021). The eROSITA Final Equatorial-Depth Survey (eFEDS): Identification and characterization of the counterparts to the point-like sources. *arXiv e-prints*, page arXiv:2106.14520.

Sartori, L. F., Schawinski, K., Trakhtenbrot, B., Caplar, N., Treister, E., Koss, M. J., Urry, C. M., and Zhang, C. E. (2018). A model for AGN variability on multiple time-scales. *"Monthly Notices of the Royal Astronomical Society"*, 476(1):L34–L38.

Saunders, W., Bridges, T., Gillingham, P., Haynes, R., Smith, G. A., Whittard, J. D., Churilov, V., Lankshear, A., Croom, S., Jones, D., and Boshuizen, C. (2004). AAOmega: a scientific and optical overview. In Moorwood, A. F. M. and Iye, M., editors, *Ground-based Instrumentation for Astronomy*, volume 5492 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 389–400.

Savonije, G. J. (1979). Roche-lobe overflow and massive X-ray binary systems. *"Astronomy & Astrophysics"*, 71:352–358.

Scargle, J. D. (1982). Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *"Astrophysical Journal"*, 263:835–853.

Schwarzschild, K. (1916). Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pages 189–196.

Shakura, N. I. and Sunyaev, R. A. (1973). Black holes in binary systems. Observational appearance. *"Astronomy & Astrophysics"*, 24:337–355.

Shapiro, S. L. and Teukolsky, S. A. (1983). *Black holes, white dwarfs, and neutron stars : the physics of compact objects.*

Sharp, R., Saunders, W., Smith, G., Churilov, V., Correll, D., Dawson, J., Farrel, T., Frost, G., Haynes, R., Heald, R., Lankshear, A., Mayfield, D., Waller, L., and Whittard, D. (2006). Performance of AAOmega: the AAT multi-purpose fiber-fed spectrograph. In McLean, I. S. and Iye, M., editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6269 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 62690G.

Shemmer, O., Netzer, H., Maiolino, R., Oliva, E., Croom, S., Corbett, E., and di Fabrizio, L. (2004). Near-Infrared Spectroscopy of High-Redshift Active Galactic Nuclei. I. A Metallicity-Accretion Rate Relationship. *"Astrophysical Journal"*, 614(2):547–557.

Singh, K. P., Tandon, S. N., Agrawal, P. C., Antia, H. M., Manchanda, R. K., Yadav, J. S., Seetha, S., Ramadevi, M. C., Rao, A. R., Bhattacharya, D., Paul, B., Sreekumar, P., Bhattacharyya, S., Stewart, G. C., Hutchings, J., Annapurni, S. A., Ghosh, S. K., Murthy, J., Pati, A., Rao, N. K., Stalin, C. S., Girish, V., Sankarasubramanian, K., Vadawale, S., Bhalerao, V. B., Dewangan, G. C., Dedhia, D. K., Hingar, M. K., Katoch, T. B., Kothare, A. T., Mirza, I., Mukerjee, K., Shah, H., Shah, P., Mohan, R., Sangal, A. K., Nagabhusana, S., Sriram, S., Malkar, J. P., Sreekumar, S., Abbey, A. F., Hansford, G. M., Beardmore, A. P., Sharma, M. R., Murthy, S., Kulkarni, R., Meena, G., Babu, V. C., and Postma, J. (2014). ASTROSAT mission. In *"Proceedings of the SPIE"*, volume 9144 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91441S.

Soleri, P., Belloni, T., and Casella, P. (2008). A transient low-frequency quasi-periodic oscillation from the black hole binary GRS 1915+105. *"Monthly Notices of the Royal Astronomical Society"*, 383(3):1089–1102.

Spinoglio, L. and Fernández-Ontiveros, J. A. (2021). AGN types and unification model. In Pović, M., Marziani, P., Masegosa, J., Netzer, H., Negu, S. H., and Tessema, S. B., editors, *Nuclear Activity in Galaxies Across Cosmic Time*, volume 356, pages 29–43.

Sriram, K., Rao, A. R., and Choi, C. S. (2012). Fast transition of type-B quasi-periodic oscillation in the black hole transient XTE J1817-330. *"Astronomy & Astrophysics"*, 541:A6.

Sriram, K., Rao, A. R., and Choi, C. S. (2013). A Spectral Study of the Rapid Transitions of Type-B Quasi-periodic Oscillations in the Black Hole Transient XTE J1859+226. *"Astrophysical Journal"*, 775(1):28.

Sriram, K., Rao, A. R., and Choi, C. S. (2016). Study of a Sudden QPO Transition Event in the Black Hole Source XTE J1550-564. *"Astrophysical Journal"*, 823(1):67.

Stella, L. and Vietri, M. (1998). Lense-Thirring Precession and Quasi-periodic Oscillations in Low-Mass X-Ray Binaries. *"Astrophysical Journal Letters"*, 492(1):L59–L62.

Stella, L., Vietri, M., and Morsink, S. M. (1999). Correlations in the Quasi-periodic Oscillation Frequencies of Low-Mass X-Ray Binaries and the Relativistic Precession Model. *"Astrophysical Journal Letters"*, 524(1):L63–L66.

Strüder, L., Briel, U., Dennerl, K., Hartmann, R., Kendziorra, E., Meidinger, N., Pfeffermann, E., Reppin, C., Aschenbach, B., Bornemann, W., Bräuninger, H., Burkert, W., Elender, M., Freyberg, M., Haberl, F., Hartner, G., Heuschmann, F., Hippmann, H., Kastelic, E., Kemmer, S., Kettenring, G., Kink, W., Krause, N., Müller, S., Oppitz, A., Pietsch, W., Popp, M., Predehl, P., Read, A., Stephan, K. H., Stötter, D., Trümper, J., Holl, P., Kemmer, J., Soltau, H., Stötter, R., Weber, U., Weichert, U., von Zanthier, C., Carathanassis, D., Lutz, G., Richter, R. H., Solc, P., Böttcher, H., Kuster, M., Staubert, R., Abbey, A., Holland, A., Turner, M., Balasini, M., Bignami, G. F., La Palombara, N., Villa, G., Buttler, W., Gianini, F., Lainé, R., Lumb, D., and Dhez, P. (2001). The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *"Astronomy & Astrophysics"*, 365:L18–L26.

Sunyaev, R., Arefiev, V., Babyshkin, V., Bogomolov, A., Borisov, K., Buntov, M., Brunner, H., Burenin, R., Churazov, E., Coutinho, D., Eder, J., Eismont, N., Freyberg, M., Gilfanov, M., Gureyev, P., Hasinger, G., Khabibullin, I., Kolmykov, V., Komovkin, S., Krivonos, R., Lapshov, I., Levin, V., Lomakin, I., Lutovinov, A., Medvedev, P., Merloni, A., Mernik, T., Mikhailov, E., Molodzov, V., Mzhelsky, P., Mueller, S., Nandra, K., Nazarov, V., Pavlinsky, M., Poghodin, A., Predehl, P., Robrade, J., Sazonov, S., Scheuerle, H., Shirshakov, A., Tkachenko, A., and Voron, V. (2021). The SRG X-ray orbital observatory, its telescopes and first scientific results. *arXiv e-prints*, page arXiv:2104.13267.

Sutherland, W., Emerson, J., Dalton, G., Atad-Ettedgui, E., Beard, S., Bennett, R., Bezawada, N., Born, A., Caldwell, M., Clark, P., Craig, S., Henry, D., Jeffers, P., Little, B., McPherson, A., Murray, J., Stewart, M., Stobie, B., Terrett, D., Ward, K., Whalley, M., and Woodhouse, G. (2015). The Visible and Infrared Survey Telescope for Astronomy (VISTA): Design, technical overview, and performance. *"Astronomy & Astrophysics"*, 575:A25.

Tabur, V., Bedding, T. R., Kiss, L. L., Moon, T. T., Szeidl, B., and Kjeldsen, H. (2009). Long-term photometry and periods for 261 nearby pulsating M giants. *"Monthly Notices of the Royal Astronomical Society"*, 400(4):1945–1961.

Tagger, M. and Pellat, R. (1999). An accretion-ejection instability in magnetized disks. *"Astronomy & Astrophysics"*, 349:1003–1016.

Takizawa, M., Dotani, T., Mitsuda, K., Matsuba, E., Ogawa, M., Aoki, T., Asai, K., Ebisawa, K., Makishima, K., Miyamoto, S., Iga, S., Vaughan, B., Rutledge, R. E., and Lewin, W. H. G. (1997). Spectral and Temporal Variability in the X-Ray Flux of GS 1124-683, Nova Muscae 1991. *"Astrophysical Journal"*, 489(1):272–283.

Tennant, A. F. and Mushotzky, R. F. (1983). The absence of rapid X-ray variability in active galaxies. *"Astrophysical Journal"*, 264:92–104.

The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration (2021). GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. *arXiv e-prints*, page arXiv:2111.03606.

Thompson, T. A., Kochanek, C. S., Stanek, K. Z., Badenes, C., Post, R. S., Jayasinghe, T., Latham, D. W., Bieryla, A., Esquerdo, G. A., Berlind, P., Calkins, M. L., Tayar, J., Lindegren, L., Johnson, J. A., Holoien, T. W. S., Auchettl, K., and Covey, K. (2019). A noninteracting low-mass black hole-giant star binary system. *Science*, 366(6465):637–640.

Timmer, J. and Koenig, M. (1995). On generating power law noise. *"Astronomy & Astrophysics"*, 300:707.

Titarchuk, L. and Fiorito, R. (2004). Spectral Index and Quasi-Periodic Oscillation Frequency Correlation in Black Hole Sources: Observational Evidence of Two Phases and Phase Transition in Black Holes. *"Astrophysical Journal"*, 612(2):988–999.

Tolman, R. C. (1939). Static Solutions of Einstein's Field Equations for Spheres of Fluid. *Physical Review*, 55(4):364–373.

Tomsick, J. A., Yamaoka, K., Corbel, S., Kaaret, P., Kalemci, E., and Migliari, S. (2009). Truncation of the Inner Accretion Disk Around a Black Hole at Low Luminosity. *"Astrophysical Journal Letters"*, 707(1):L87–L91.

Truemper, J. (1982). The ROSAT mission. *Advances in Space Research*, 2(4):241–249.

Turner, M. J. L., Abbey, A., Arnaud, M., Balasini, M., Barbera, M., Belsole, E., Bennie, P. J., Bernard, J. P., Bignami, G. F., Boer, M., Briel, U., Butler, I., Cara, C., Chabaud, C., Cole, R., Collura, A., Conte, M., Cros, A., Denby, M., Dhez, P., Di Coco, G., Dowson, J., Ferrando, P., Ghizzardi, S., Gianotti, F., Goodall, C. V., Gretton, L., Griffiths, R. G., Hainaut, O., Hochedez, J. F., Holland, A. D., Jourdain, E., Kendziorra, E., Lagostina, A., Laine, R., La Palombara, N., Lortholary, M., Lumb, D., Marty, P., Molendi, S., Pigot, C., Poindron, E., Pounds, K. A., Reeves, J. N., Reppin, C., Rothenflug, R., Salvetat, P., Sauvageot, J. L., Schmitt, D., Sembay, S., Short, A. D. T., Spragg, J., Stephen, J., Strüder, L., Tiengo, A., Trifoglio, M., Trümper, J., Vercellone, S., Vigroux, L., Villa, G., Ward, M. J., Whitehead, S., and Zonca, E. (2001). The European Photon Imaging Camera on XMM-Newton: The MOS cameras. *"Astronomy & Astrophysics"*, 365:L27–L35.

Turner, T. J., George, I. M., Nandra, K., and Turcan, D. (1999). On X-Ray Variability in Seyfert Galaxies. *"Astrophysical Journal"*, 524(2):667–673.

Ubertini, P., Lebrun, F., Di Cocco, G., Bazzano, A., Bird, A. J., Broenstad, K., Goldwurm, A., La Rosa, G., Labanti, C., Laurent, P., Mirabel, I. F., Quadrini, E. M., Ramsey, B., Reglero, V., Sabau, L., Sacco, B., Staubert, R., Vigroux, L., Weisskopf, M. C., and Zdziarski, A. A. (2003). IBIS: The Imager on-board INTEGRAL. *"Astronomy & Astrophysics"*, 411:L131–L139.

Ulrich, M.-H., Maraschi, L., and Urry, C. M. (1997). Variability of Active Galactic Nuclei. *"Annual Review of Astronomy & Astrophysics"*, 35:445–502.

van der Klis, M. (1989). Fourier techniques in X-ray timing. In Ögelman, H. and van den Heuvel, E. P. J., editors, *Timing Neutron Stars*, volume 262 of *NATO Advanced Study Institute (ASI) Series C*, page 27.

Varnière, P., Rodriguez, J., and Tagger, M. (2002). Accretion-ejection instability and QPO in black-hole binaries. II. Relativistic effects. *"Astronomy & Astrophysics"*, 387:497–506.

Varnière, P. and Tagger, M. (2002). Accretion-Ejection Instability in magnetized disks: Feeding the corona with Alfvén waves. *"Astronomy & Astrophysics"*, 394:329–338.

Varniere, P. and Vincent, F. H. (2016). Impact of inclination on quasi-periodic oscillations from spiral structures. *"Astronomy & Astrophysics"*, 591:A36.

Vaughan, S. (2010). A Bayesian test for periodic signals in red noise. *"Monthly Notices of the Royal Astronomical Society"*, 402(1):307–320.

Vaughan, S., Edelson, R., Warwick, R. S., and Uttley, P. (2003). On characterizing the variability properties of X-ray light curves from active galaxies. *"Monthly Notices of the Royal Astronomical Society"*, 345(4):1271–1284.

Vedrenne, G., Roques, J. P., Schönfelder, V., Mandrou, P., Lichti, G. G., von Kienlin, A., Cordier, B., Schanne, S., Knödlseder, J., Skinner, G., Jean, P., Sanchez, F., Caraveo, P., Teegarden, B., von Ballmoos, P., Bouchet, L., Paul, P., Matteson, J., Boggs, S., Wunderer, C., Leleux, P., Weidenspointner, G., Durouchoux, P., Diehl, R., Strong, A., Cassé, M., Clair, M. A., and André, Y. (2003). SPI: The spectrometer aboard INTEGRAL. *"Astronomy & Astrophysics"*, 411:L63–L70.

Venemans, B. P., Verdoes Kleijn, G. A., Mwebaze, J., Valentijn, E. A., Bañados, E., Decarli, R., de Jong, J. T. A., Findlay, J. R., Kuijken, K. H., La Barbera, F., McFarland, J. P., McMahon, R. G., Napolitano, N., Sikkema, G., and Sutherland, W. J. (2015). First discoveries of z $\sim$ 6 quasars with the Kilo-Degree Survey and VISTA Kilo-Degree Infrared Galaxy survey. *"Monthly Notices of the Royal Astronomical Society"*, 453(3):2259–2266.

Verner, D. A., Ferland, G. J., Korista, K. T., and Yakovlev, D. G. (1996). Atomic Data for Astrophysics. II. New Analytic FITS for Photoionization Cross Sections of Atoms and Ions. *"Astrophysical Journal"*, 465:487.

Weisskopf, M. C., Tananbaum, H. D., Van Speybroeck, L. P., and O'Dell, S. L. (2000). Chandra X-ray Observatory (CXO): overview. In Truemper, J. E. and Aschenbach, B., editors, *"Proceedings of the SPIE"*, volume 4012 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 2–16.

Wijnands, R., Homan, J., and van der Klis, M. (1999). The Complex Phase-Lag Behavior of the 3-12 HZ Quasi-Periodic Oscillations during the Very High State of XTE J1550-564. *"Astrophysical Journal Letters"*, 526(1):L33–L36.

Wijnands, R. and van der Klis, M. (1999). The Broadband Power Spectra of X-Ray Binaries. *"Astrophysical Journal"*, 514(2):939–944.

Wilms, J., Allen, A., and McCray, R. (2000). On the Absorption of X-Rays in the Interstellar Medium. *"Astrophysical Journal"*, 542(2):914–924.

Wilson, W. E., Ferris, R. H., Axtens, P., Brown, A., Davis, E., Hampson, G., Leach, M., Roberts, P., Saunders, S., Koribalski, B. S., Caswell, J. L., Lenc, E., Stevens, J., Voronkov, M. A., Wieringa, M. H., Brooks, K., Edwards, P. G., Ekers, R. D., Emonts, B., Hindson, L., Johnston, S., Maddison, S. T., Mahony, E. K., Malu, S. S., Massardi, M., Mao, M. Y., McConnell, D., Norris, R. P., Schnitzeler, D., Subrahmanyan, R., Urquhart, J. S., Thompson, M. A., and Wark, R. M. (2011). The Australia Telescope Compact Array Broad-band Backend: description and first results. *"Monthly Notices of the Royal Astronomical Society"*, 416(2):832–856.

Winkler, C., Courvoisier, T. J. L., Di Cocco, G., Gehrels, N., Giménez, A., Grebenev, S., Hermsen, W., Mas-Hesse, J. M., Lebrun, F., Lund, N., Palumbo, G. G. C., Paul, J., Roques, J. P., Schnopper, H., Schönfelder, V., Sunyaev, R., Teegarden, B., Ubertini, P., Vedrenne, G., and Dean, A. J. (2003). The INTEGRAL mission. *"Astronomy & Astrophysics"*, 411:L1–L6.

Wolf, J., Nandra, K., Salvato, M., Liu, T., Buchner, J., Brusa, M., Hoang, D. N., Moss, V., Arcodia, R., Brüggen, M., Comparat, J., de Gasperin, F., Georgakakis, A., Hotan, A., Lamer, G., Merloni, A., Rau, A., Rottgering, H. J. A., Shimwell, T. W., Urrutia, T., Whiting, M., and Williams, W. L. (2021). First constraints on the AGN X-ray luminosity function at z 6 from an eROSITA-detected quasar. *"Astronomy & Astrophysics"*, 647:A5.

Wolter, H. (1952). Spiegelsysteme streifenden Einfalls als abbildende Optiken für Röntgenstrahlen. *Annalen der Physik*, 445(1):94–114.

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., Ressler, M. E., Cutri, R. M., Jarrett, T., Kirkpatrick, J. D., Padgett, D., McMillan, R. S., Skrutskie, M., Stanford, S. A., Cohen, M., Walker, R. G., Mather, J. C., Leisawitz, D., Gautier, Thomas N., I., McLean, I., Benford, D., Lonsdale, C. J., Blain, A., Mendez, B., Irace, W. R., Duval, V., Liu, F., Royer, D., Heinrichsen, I., Howard, J., Shannon, M., Kendall, M., Walsh, A. L., Larsen, M., Cardon, J. G., Schick, S., Schwalm, M., Abid, M., Fabinsky, B., Naes, L., and Tsai, C.-W. (2010). The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance. *"The Astronomical Journal"*, 140(6):1868–1881.

Xiao, G. C., Lu, Y., Ma, X., Ge, M. Y., Yan, L. L., Li, Z. J., Tuo, Y. L., Zhang, Y., Zhang, W., Liu, H. X., Zhou, D. K., Zhang, L., Bu, Q. C., Cao, X. L., Jiang, W. C., Chen, Y. P., Zhang, S., Lu, F. J., Chen, L., Qu, J. L., Song, L. M., Zhang, S. N., Zhuang, R. L., Shang, R. C., and Jin, Y. J. (2019). Timing analysis of Swift J1658.2-4242's outburst in 2018 with Insight-HXMT, NICER and AstroSat. *Journal of High Energy Astrophysics*, 24:30–40.

Xu, Y., Harrison, F. A., Kennea, J. A., Walton, D. J., Tomsick, J. A., Miller, J. M., Barret, D., Fabian, A. C., Forster, K., Fürst, F., Gandhi, P., and García, J. A. (2018). The Hard State of the Highly Absorbed High Inclination Black Hole Binary Candidate Swift J1658.2-4242 Observed by NuSTAR and Swift. *"Astrophysical Journal"*, 865(1):18.

Xu, Y., Harrison, F. A., Tomsick, J. A., Barret, D., Gandhi, P., García, J. A., Miller, J. M., Uttley, P., and Walton, D. J. (2019). Broadband X-Ray Spectral and Timing Analyses of the Black Hole Binary Candidate Swift J1658.2-4242: Rapid Flux Variation and the Turn-on of a Transient QPO. *"Astrophysical Journal"*, 879(2):93.

Yamaguchi, M. S., Kawanaka, N., Bulik, T., and Piran, T. (2018). Detecting Black Hole Binaries by Gaia. *"Astrophysical Journal"*, 861(1):21.

Zhang, S., Zhang, S. N., Lu, F. J., Li, T. P., Song, L. M., Xu, Y. P., Wang, H. Y., Qu, J. L., Liu, C. Z., Chen, Y., Cao, X. L., Zhang, F., Xiong, S. L., Ge, M. Y., Chen, Y. P., Liao, J. Y., Nie, J. Y., Zhao, H. S., Jia, S. M., Li, X. B., Guan, J., Li, C. K., Zhang, J., Jin, J., Wang, G. F., Zheng, S. J., Ma, X., Tao, L., and Huang, Y. (2018). The insight-HXMT mission and its recent progresses. In *Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray*, volume 10699 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 106991U.

Zhao, X., Gou, L., Dong, Y., Zheng, X., Steiner, J. F., Miller-Jones, J. C. A., Bahramian, A., Orosz, J. A., and Feng, Y. (2021). Re-estimating the Spin Parameter of the Black Hole in Cygnus X-1. *"Astrophysical Journal"*, 908(2):117.

Zheng, X. C., Xue, Y. Q., Brandt, W. N., Li, J. Y., Paolillo, M., Yang, G., Zhu, S. F., Luo, B., Sun, M. Y., Hughes, T. M., Bauer, F. E., Vito, F., Wang, J. X., Liu, T., Vignali, C., and Shu, X. W. (2017). Deepest View of AGN X-Ray Variability with the 7 Ms Chandra Deep Field-South Survey. *"Astrophysical Journal"*, 849(2):127.

Zhu, S. F. and Xue, Y. Q. (2016). Using Leaked Power to Measure Intrinsic AGN Power Spectra of Red-noise Time Series. *"Astrophysical Journal"*, 825(1):56.

Acknowledgements

While we find ourselves here on the final page, the culmination of more than 4 years of work, it is worth remembering that none of this would have been possible without any of the following amazing people. Therefore, I would like to express my deepest gratitude to:

Prof. Kirpal Nandra, for the amazing supervision, guidance, perseverance, and inspiration. For being the voice of reason, always pushing me in the right direction and dreaming big.

Dr. Gabriele Ponti, for the amazing supervision, assistance, and passion. For the invaluable support at the very beginning, for frequently checking on how I was doing, and for the great choice of BHT to observe (even though it didn't feature any winds).

Dr. Mara Salvato, for the consistent support, and for being an invaluable guide in the hectic days preceding AAT observations.

Dr. Johannes Buchner, for always taking the time to listen to my thoughts in detail, and providing equally detailed replies, teaching me lots of finer details of statistics and computing, and writing bexvar.

Dr. Teng Liu, for his extensive knowledge of *eROSITA* data products.

Prof. Scott Croom, for his patience, while I chased down errors in my AAT fld files, and for connecting me to the 2dF/AAOmega controls (illegally).

Dr. Tomaso Belloni, for the invaluable knowledge of BHTs, and a keen eye for evaluating the significance of certain discoveries.

Julien Wolf, for being a great office mate, and for turning my COVID joke into a reality.

The MPE high energy group, the eROAGN group, and the eROTDA groups, for their friendly and supportive environments.

Dr. Chichuan Jin, Dr. Haiwu Pan, Dr. Thomas D. Russell, Prof. James Miller-Jones, Dr. Teo Muñoz-Darias, Pavan Vynatheya, Dr. Federico Vincentelli, Hattie Starck, Dr. Ider Chitham, Dr. Chandreyee Maitra, and Dr. Mirko Krumpe, for providing additional support for this work, and for proofreading it at various stages of completeness.

And to Wolfgang, Michaela, Stefan, Fiona, and Luisa for the constant emotional support. For all the great time we spent together during the pandemic, for supporting my endeavours, however far away they might be (physically and thematically), but nevertheless always being there to fall back upon.