McGill University

Master's Thesis

Change points in Sgr A*'s X-ray flaring rate: fact or artifact?

Author:

Élie Bouffard

Supervisor and Co-Supervisor:

Prof. Daryl Haggard and

Prof. Nicolas Cowan

A thesis submitted in partial fulfillment of the requirements

for the degree of Master of Science

in the

Department of Physics

McGill University

Montreal, Quebec

Canada

Abstract

Change points in Sgr A*'s X-ray flaring rate: fact or artifact?

by Élie BOUFFARD

An unusual object, G2, had its pericenter passage around Sgr A*, the 4×10^6 $M_\odot$ supermassive black hole in the Galactic Centre, in Summer 2014. Several research teams have reported evidence that following G2's pericenter encounter the rate of Sgr A*'s bright X-ray flares increased significantly. Our analysis carefully treats varying flux contamination from a nearby magnetic neutron star and is free from complications induced by using data from multiple X-ray observatories with different sensitivities and spatial resolutions. We test this scenario using a massive dataset from the *Chandra X-ray Observatory*, the only X-ray instrument that can spatially distinguish between Sgr A* and the nearby Galactic Centre magnetar throughout the full extended period encompassing G2's encounter with Sgr A*. We use X-ray data from the 3 Ms observations of the *Chandra X-ray Visionary Program* (XVP) in 2012 as well as an additional 1.5 Ms of observations up to 2018. We use detected flares to make distributions of flare properties. Using a simulation model of X-ray flares accounting for important factors such as the different $Chandra$ instrument modes, we test the null hypothesis on Sgr A*'s bright (or any flare category) X-ray flaring rate around different potential change points. In contrast to previous studies, our results are consistent with the null hypothesis; the same model parameters produce distributions consistent with the observed ones around any plausible change point.

Résumé

Change points in Sgr A*'s X-ray flaring rate: fact or artifact?

par Élie BOUFFARD

Un objet inhabituel, G2, a atteint son péricentre autour de Sgr A*, le trou noir supermassif de 4×10^6 $M_\odot$ au centre de la galaxie, durant l'été 2014. Certains groupes de recherche ont affirmé que, suite à cela, le taux d'éruptions brillantes a augmenté significativement. Notre analyse traite avec soin la contamination de flux variable provenant du magnétar à proximité et n'a pas de complications induites par l'utilisation de données provenant de plusieurs observatoires à rayons X avec différentes sensibilités et résolutions spatiales. Nous testons ce scénario en utilisant un ensemble de données volumineux du *Chandra X-ray Observatory*, le seul instrument à rayons X capable de distinguer spatialement Sgr A* et le magnétar du Centre Galactique pendant toute la durée de la rencontre entre G2 et Sgr A*. Nous utilisons les données de rayons X issues des observations de 3 Ms du *X-ray Visionary Program* (XVP) de *Chandra* en 2012, ainsi que 1,5 Ms supplémentaires d'observations jusqu'en 2018. Nous utilisons les éruptions détectées afin de créer des distributions de leurs propriétés. À l'aide d'un modèle de simulations d'éruptions rayons X prenant en compte des facteurs importants tels que les différents modes d'instrument de $Chandra$, nous testons l'hypothèse nulle sur le taux d'éruptions rayons X lumineuses (ou de n'importe quelle autre catégorie) de Sgr A* autour de différents points de changement potentiels. Contrairement aux études précédentes, nos résultats sont en accord avec l'hypothèse nulle; les mêmes paramètres du modèle produisent des distributions cohérentes avec les distributions observées autour de tout point de changement plausible.

Acknowledgements

I'd like to thank my supervisor Professor Daryl Haggard for her support, optimism and guidance. Likewise, thanks to my co-supervisor Professor Nicolas Cowan for his useful inputs regarding everything statistics and mathematics. Thanks to Mel Nynka for her help with $Chandra$ data reduction and analysis. Thanks to every member of Professor Haggard's research group for their friendship and help. I'd also like to thank Mike Nowak for housing Daryl Haggard and I at his house for a few days during which he helped me a great deal for my research. Mike also helped me with my research before and after this event. Thanks to Joey Neilsen for his help about Cumulative Distribution Functions. Thanks to Jeffrey Scargle for his help understanding Bayesian Blocks, and thanks to Peter Williams for their public implementation. Thanks to Enmanuelle Mossoux for answering my questions about her work. I'd like also to thank everyone from the McGill Space Institute for their awesome company and for creating such an outstanding learning environment. Finally, thanks to my parents for their encouragement and believing in me.

Contributions

Frederick Baganoff, Sera Markoff, and Mike Nowak are the Co-Principle Investigators (PIs) for the *X-ray Visionary Program* (XVP) observations. Daryl Haggard, Gordon Garmire, Frederick Baganoff, Joey Neilsen and Nanda Rea are the PIs of 15, 10, 7, 4, and 3 of the kept Post-XVP observations, respectively (see Chapter 2). Mel Nynka wrote the *Chandra* data reduction pipeline for the ACIS-S/subarray mode (Section 2.2). She also guided me for the conversion of count rates to luminosity (Section 2.4). Mike Nowak got me started on the data reduction for the ACIS-S/HETG instrument (Section 2.1), but I wrote the code performing the data reduction presented in this thesis. I performed all the work presented in Chapter 3. I used Bayesian Blocks for flare detection (Scargle et al., 2013; Williams et al., 2017), but I modified the implementation of Peter Williams to make it work on any event list (not only *Chandra* events) to use it on simulated data, and I calibrated the prior on Poisson noise (see Chapter 4 and Appendix A). To simulate event lists, I followed Appendix D of Mossoux & Grosso (2017) as a starting point (Section 5.1), but I made the model and simulations presented in the rest of Chapter 5. Joey Neilsen explained to me how to transform a power-law distribution into a Cumulative Distribution Function to draw values from it (see Appendix D). The remainder of the work presented in this thesis is original and will be submitted to a refereed science journal as a culmination of this effort.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Black hole family

Black holes are some of the most unusual objects in the Universe. They are so dense that no force can hold them in equilibrium with gravity. This means that all of their mass is concentrated in a single point of volume zero, called a singularity. A black hole's gravity is so powerful that its escape velocity is greater than the speed of light in vacuum. Black holes were first predicted at the end of the 18th century by English philosopher John Mitchel, who used the term "dark stars". He proposed in 1783 that there could be stars whose gravitational pull would be strong enough to prevent light from escaping their surface. He even predicted that one could find such objects by detecting a star that behaved as if it were in a binary system, but with an invisible companion. This prediction is impressive since we know today of many binary systems containing a black hole.

In the early 20th century, Albert Einstein published his famous Theory of General Relativity which explained gravity as the manifestation of a curvature of spacetime. In the following year, Karl Schwarzchild applied General Relativity to the concept of a black hole and found that its radius R_s (named after him) is linked to its mass M via the famous equation

$$R_s = 2GM/c^2 \tag{1.1}$$

where G is the gravitational constant and c is the speed of light in vacuum.

The term black hole was later popularized by American theoretical physicist John Archibald Wheeler and has been the official term for many decades.

Cygnus X-1 is the first black hole discovered in 1964 (Bowyer et al., 1965) by X-ray observations during a rocket flight above the atmosphere (which blocks X-ray light). Combining these observations with subsequent radio and optical data, it was shown that this system contained a black hole that was accreting mass from its companion star (Oda, 1977).

Following this discovery, scientists have discovered large populations of black holes and we now understand that black holes appear in two categories based on their mass: stellar mass black holes and supermassive black holes (SMBH).

1.1.1 Main categories of black holes

Stellar mass black holes

Stellar mass black holes have masses ranging from $\sim$2.7 $M_\odot$ up to $10^2 M_\odot$. The lower bound comes from the binary neutron star merger GW170817 which was detected by the gravitational wave (GW) detectors LIGO-Virgo on August 17 2017 (Abbott et al., 2017d,e). The remnant of this merger is thought to have initially formed an hyper-massive neutron star which lived for a timescale of the order of milliseconds before collapsing into a black hole (Margalit & Metzger, 2017; Metzger et al., 2018). This is supported by the fact that the remnant's mass ($2.74^{+0.04}_{-0.01} M_\odot$; Abbott et al., 2017d) is above the mass of the most massive known neutron star ($1.97 \pm 0.04 M_\odot$; Demorest et al., 2010) and by the lack of electromagnetic emission consistent with neutron star spin-down (Margutti et al., 2018).

Gravitational waves were predicted by Einstein's Theory of General Relativity and are generated when mass is accelerated, similarly to light being radiated when charged particles accelerate (e.g., bremsstrahlung radiation). These waves also travel at the speed of light. However, instead of being characterized by electric and magnetic fields, they are ripples of spacetime itself. To detect them, LIGO and Virgo use laser interferometry which is accurate enough to probe spatial modulation of one part per 10^{21}. Such a high sensitivity is required to detect astrophysical signals since they originate from 10's of Mpc to Gpc distances, and unlike electromagnetic waves, we can only detect their amplitude, not their energy. Even though their amplitude decreases as $1/r$, their initial amplitude is only comparable to the size of the merger. The first GW detection happened on September 14, 2015 (Willke et al., 2018) with the discovery of the first binary black hole system and merger. As of today, 11 mergers have been observed (Abbott et al., 2016a,b, 2017a,b,c,d).

Formation channels for stellar mass black holes are clearer than those for SMBH. A massive star $M \gtrsim 20 M_\odot$ can reach high enough densities and temperatures at its core to fuse heavier elements than lower mass stars like our Sun. More massive elements are harder to fuse because their nuclei have higher charges, such that their Coulomb barrier is stronger. During fusion, energy is released because the heavier elements produced are more stable (their binding energy to mass energy ratio is higher). However, this trend stops at iron, because fusing iron is an endothermic instead of exothermic process and therefore requires more energy than would be produced. This lowers the amount of radiation pressure that was opposing the force of gravity, resulting in the star's collapse. For massive enough stars, the degeneracy pressure provided by electrons isn't enough to stop the collapse (i.e., the core mass is higher than the Chandrasekhar limit of about 1.4 $M_\odot$). Electrons thus merge with nuclei, forming neutrons. The newly created neutron degeneracy pressure is also not strong enough to stop the collapse of such massive stars and the remnant collapses to a black hole.

Many stellar mass black holes are known in our galaxy often in X-ray binary system

like Cygnus X-1 (see Remillard & McClintock (2006) for a review of known objects). In these systems, X-rays are generated by matter falling from the companion onto the compact object (in this case a black hole, but sometimes it can be a neutron star). Part of the gravitational potential energy of this accreting matter is released as radiation. For excellent reviews discussing accretion, see Abramowicz & Fragile (2013) and Yuan & Narayan (2014).

Supermassive black holes

SMBHs have masses $\geq 10^6 M_\odot$ and are found at the centres of massive galaxies. In their reviews of observational evidence of SMBHs, Ferrarese & Ford (2005) and Graham (2016) explain in detail the diverse detection methods and their historical perspective. One piece of early strong evidence in favour of SMBHs came from quasars, extraordinarily bright objects that were originally thought to be local stars but were later shown to be located at high redsift. For example, the radio galaxy 3C 273 was shown to be extremely bright (absolute magnitude of -26.7; Greenstein & Schmidt, 1964) and at a cosmological distance (redshift of z=0.158; Schmidt, 1963). Yet these two facts are not convincing enough evidence of the existence of a SMBH. Indeed, this behaviour could be explained by a compact group of luminous stars. What finally tipped the balance in favour of a SMBH was the discovery of variability on a 1 year timescale in 3C 273 (Smith & Hoffleit, 1963). This implied that the source had a spatial extent of less than 1 light-year long. The one remaining problem was to explain how can so much power be generated from such a compact source. The idea was suggested (and later widely accepted) that the only reasonable explanation was the conversion of gravitational energy to radiation via accretion onto a massive compact object (Salpeter, 1964; Shakura & Sunyaev, 1973). The presence of a very massive object was also motivated by observations of these objects near the Eddington luminosity, $L_{\rm Edd}$. This limit is defined as the luminosity at which the central object's radiation pressure equates its gravitational pull, preventing new material from accreting (the

central object has to be very massive to be that luminous).

After this discovery, many more similar objects were discovered. Nowadays, one of the strongest evidence for a SMBH comes from our own Galactic Centre. We have observed the influence of a SMBH there, Sagittarius A* (Sgr A* henceforth). Its mass has been measured at $\sim 4 \times 10^6$ $M_\odot$ (Genzel et al., 2010) using the motion of stars in its vicinity (see also the GRAVITY Collaboration Abuter et al. (2018a,b)). Indeed, its relative proximity enables very accurate measurements of the orbits of stars in its vicinity and from the orbital parameters, the mass can be deduced.

Very recently, the Event Horizon Telescope (EHT) Collaboration has released its first look at M87's central black hole, taking a snapshot of the SMBH's shadow. This unprecedented feat provides direct evidence for SMBHs at the centres of massive galaxies (Akiyama et al., 2019). The EHT Collaboration is still analyzing their data for Sgr A*, but we expect to learn a great deal more about its accretion physics in the coming years.

Farther from home, some SMBHs have been found at $z > 6$ with masses $\sim 10^9 M_\odot$ causing theoretical challenges and begging the question of how the Universe grows black holes to these masses within a billion years. A popular idea is that the first generation of stars (Pop. III stars) formed early ($z \sim 20 - 50$) within potential wells from dark matter halos, and collapsed to black hole seeds of masses $\sim 150 M_\odot$ after ~ 2 Myr. These seeds would then grow via accretion and mergers with other seeds (Volonteri, 2010). Another idea is the direct collapse of gas clouds (Begelman et al., 2006) which would form black holes of $10 - 20 M_\odot$ which can grow to $10^6 M_\odot$ by $z = 10 - 20$, and even $10^9 M_\odot$ by $z = 6$ by means of Super-Eddington accretion. Finally, another hypothesis is the formation of primordial black holes during the Big Bang (Sasaki et al., 2018). Hopefully, future *James Webb Space Telescope* and gravitational wave observations will provide enough observational constraints to determine which model -if any- is the correct one.

1.2 Sgr A*, the SMBH in the centre of our galaxy

1.2.1 Quiescence

Our local SMBH's X-ray emission is characterized by: (1) a constant and extremely faint quiescent 2-10 keV unabsorbed luminosity of a few 10^{33} erg/s (Baganoff et al., 2003) and (2) a $\sim$ daily rapid increase of a factor of a few to a few hundred in luminosity that lasts several minutes to hours (Baganoff et al., 2001; Goldwurm et al., 2003; Porquet et al., 2003, 2008; Belanger et al., 2005; Nowak et al., 2012; Neilsen et al., 2013, 2015; Degenaar et al., 2013; Barriere et al., 2014; Mossoux et al., 2016; Ponti et al., 2015; Yuan & Wang, 2015; Yuan et al., 2017; Zhang et al., 2017). In fact, its bolometric luminosity (i.e., integrated across all wavelengths) of $L_{\rm Bol} \sim 10^{-9} L_{\rm Edd}$ is orders of magnitude fainter than that of typical nearby low-luminosity active galactic nuclei (LLAGN) (see Table 9 of Ho, 1999) which have $L_{\rm Bol}$ from $10^{-6} L_{\rm Edd}$ to $10^{-3} L_{\rm Edd}$.

Given such a low luminosity, the best models to describe Sgr A*'s accretion flow are the radiatively inefficient accretion flow (RIAF) models (see Narayan et al. (1998) and Yuan & Narayan (2014) for reviews). The model proposed by Yuan et al. (2003) (see Figure 1.1) shows an effective radiative efficiency $L_{\rm bol} / \left[\dot{M}_{\rm in}(R_B) c^2\right] \approx 2 \times 10^{-5}$, where $\dot{M}_{\rm in}(R_B) \sim 10^{-5} M_\odot$/year (Baganoff et al., 2003) is the accretion rate at the Bondi radius $R_B \sim 10^5 R_s$ at which the thermal energy of the gas is equal to its gravitational potential energy. This accreted material comes from stellar winds of nearby massive stars (Cuadra et al., 2007). This low efficiency is explained by two main factors. Firstly, the vast majority of the accreted mass at the Bondi radius never makes it into the black hole but is instead redirected as an outflow, such that $\dot{M}_{\rm BH} \sim 4 \times 10^{-8}$, where $\dot{M}_{\rm BH}$ is the accretion at R_s. This model is consistent with radio and X-ray observations (Marrone et al., 2006; Wang et al., 2013) showing evidence of a decreasing accretion rate with decreasing radius ($\dot{M}_{\rm BH}$ is constrained between $10^{-9} - 10^{-7} M_\odot$/year).

Secondly, advection dominates over radiation such that the radiative efficiency at R_s becomes $L_{\rm bol}/\left[\dot{M}_{\rm BH}c^2\right] \approx 4 \times 10^{-4}$ (Yuan et al., 2003). Advection dominates in a RIAF because the density is so low that ions don't collide enough with electrons. Such Coulomb collisions are normally how ions cool; they transfer their energy to electrons that can then radiate it away. When this becomes negligible due to low densities, the gas adopts a two-temperature configuration and heat, instead of being radiated away, is advected into the black hole (Quataert, 2003; Yuan & Narayan, 2014). In the case of Sgr A*, the gas density is so low that even electrons do not cool efficiently by radiating. In this regime, both electrons and ions are advection dominated.

The X-ray quiescent emission is well fit by a thermal bremsstrahlung process (Quataert, 2002), in which free electrons in the plasma are deflected by ions, radiating energy. This emission is spatially extended ($\sim$ 1.4") and covers a region around Sgr A* corresponding to the Bondi radius (Baganoff et al., 2003).

At the other end of the spectrum, still following the model of Yuan et al. (2003) of Figure 1.1, the power-law emission at the lowest frequencies is synchrotron radiation produced by a population of non-thermal electrons (because thermalization via collisions is inefficient; Özel et al., 2000). Synchrotron emission from thermal electrons is responsible for the first bump, and the subsequent bumps at higher frequencies are caused by inverse Compton scattering.

1.2.2 Flares

As mentioned in the previous section (and illustrated by the bow ties in Figure 1.1, but see Figure 1.3 for an X-ray image comparing quiescence and flaring states), Sgr A* flares about once a day in the X-ray band. These flares are 10's to 100's of times brighter than quiescence and are characterized by an harder spectra than quiescence (power-law index of $\Gamma = 2$ instead of $\Gamma = 3$ in quiescence) (Nowak et al., 2012). Their spectra are interestingly

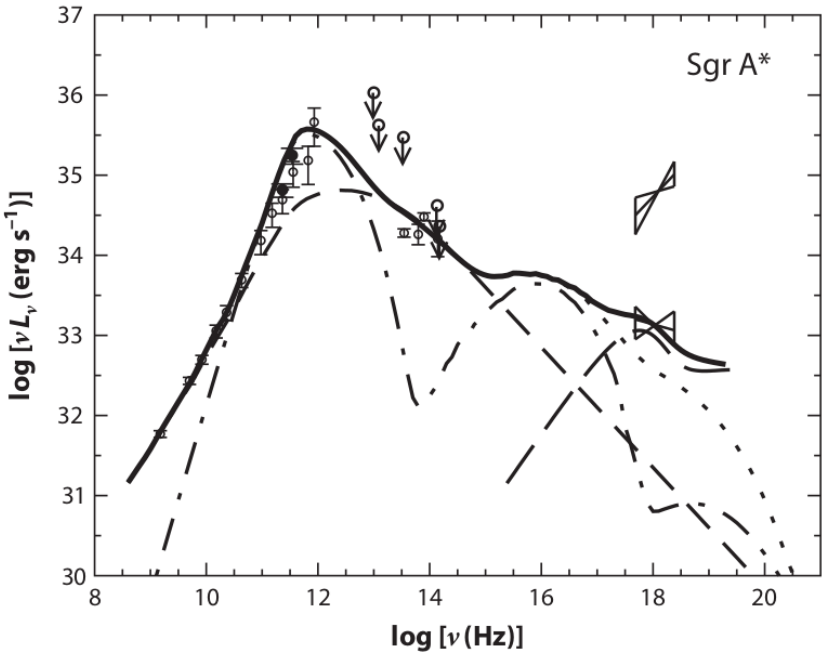

Figure 1.1: Spectral energy distribution of Sgr A* in quiescence (taken from Yuan & Narayan (2014) who used the model proposed by Yuan et al. (2003)). Circles with arrows pointing down are infrared (IR) upper limits and circles with error bars are radio observations. Circles with error bars in the IR region were taken after the model was produced. The lower and upper bow ties represent X-ray measurements of Sgr A* in quiescence and during a flare, respectively. The thick black line is the RIAF model from Yuan et al. (2003), which is the sum of the bremsstrahlung emission from the Bondi radius (long-dashed line), the synchrotron and inverse Compton emission from thermal electrons (dot-dashed line), and the synchrotron emission from non-thermal electrons (short-dashed line). The dotted line represents the sum of the last two components (i.e., without bremsstrahlung).

consistent across all flares (Neilsen et al., 2013; Yuan et al., 2017), suggesting a common origin. They are characterized by timescales of minutes to several hours, further implying an origin physically close to the black hole ($\sim$ 10's of R_s (Quataert, 2003)) unlike the spatially resolved quiescent emission. Recently, the GRAVITY Collaboration confirmed this scale by resolving the orbits of 3 IR flares (Abuter et al., 2018a) by using K-band astrometry measurements of an accuracy of $\sim$ 20-70 microarcseconds (μas). This incredible resolution is achieved via interferometry of four 8 m telescopes of the Very Large Telescope (VLT). The left side of Figure 1.2 shows a simulation of gas in a circular orbit around Sgr A*. The right side of the figure is the equivalent from observations. It shows the projected orbit of

one of the flares, which was observed on July 22 2018 for 30 min. Each point represents a different time, and the positions are the offsets with respect to their median. The orange square shows the position of Sgr A* as determined from the S2 star's orbit (Abuter et al., 2018b), which is consistent with the centroid of the flare's orbit (the pink cross) within the $\pm$ 50 μas uncertainties. The centroid motions of the 3 flares are consistent with a circular orbit of $\sim 10\ R_s$, in agreement with earlier estimations from flare durations.

Despite this knowledge, the physical explanations behind these flares and the X-rays emitted are still illusive. Several mechanisms have been proposed to explain their spectra such as inverse Compton, synchrotron and synchrotron self-Compton (in which electrons emit infrared (IR) photons before inverse Compton scattering them) (Marrone et al., 2008; Dodds-Eden et al., 2009; Eckart et al., 2009; Yusef-Zadeh et al., 2012; Witzel et al., 2012; Nowak et al., 2012; Barriere et al., 2014; Neilsen et al., 2015; Ponti et al., 2017; Zhang et al., 2017; Boyce et al., 2019; Haggard et al., 2019). Many scenarios have been put forward as the physical process leading to the radiation observed including tidal disruption of asteroids, magnetic reconnection (and/or shocks, sometimes combined with gravitational lensing as in Ball et al. (2016, 2018)) leading to particle acceleration (Markoff et al., 2001; Liu & Melia, 2002; Yuan et al., 2003; Liu et al., 2004; Čadež et al., 2008; Kostić et al., 2009; Zubovas et al., 2012; Dibi et al., 2014, 2016; Mezcua, 2017).

These flares have fluences and durations following a power-law distribution with indices of $\Gamma \sim -1.5$ and $\Gamma \sim -0.9$ respectively (Neilsen et al., 2013). Although early work showed a correlation between the flares' fluence and duration (Neilsen et al., 2013), more recent work indicates that this originates from detection biases (Yuan & Wang, 2015; Yuan et al., 2017) (more energetic flares tend to be detected for longer than less energetic ones). At first glance, flares also appear to be uncorrelated in time (following a Poisson flaring rate), but they can also show evidence of clustering e.g., 4 flares are detected within ~ 23

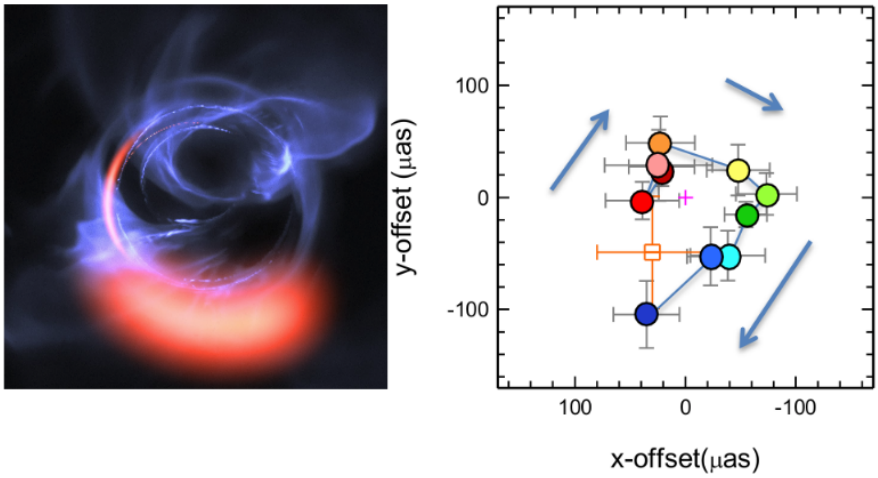

Figure 1.2: *Left*: Simulation of gas moving at 30% of the speed of light in a circular orbit around Sgr A*. This bright gas is actually an IR flare. Figure taken from the GRAVITY press release (https://www.eso.org/public/images/eso1835a/, Credit: ESO/Gravity Consortium/L. Calçada). *Right*: Projected orbit of K-band (2.2 μm) astrometric positions of the centroid of a flare on July 22 2018 observed for 30 min from the GRAVITY collaboration (taken from Abuter et al., 2018a). Each point has an arbitrary color denoting a different point in time and its position corresponds to its offset with the median of all points (small pink cross). The orange square represents Sgr A*'s centroid, as determined by the orbit of the S2 star (Abuter et al., 2018b). The clockwise orbital motion of the gas implies an orbital speed of 30% of the speed of light. Error bars are 1σ and the scale is 10 μas = 1 R_s = 1.2 $\times$ 10^7 km.

ks (ObsID 13854) with an associated probability of 3.5 % (Neilsen et al., 2013, 2015). However, at the moment, there are not enough observations to properly constrain this phenomenon to get a better understanding of the underlying process, but it's been suggested recently (Yuan & Wang, 2015; Yuan et al., 2017) that it might be related to a piecewise-deterministic Markov process (Davis, 1984, used to model earthquakes). These authors notice that flares tentatively deviate from a Poisson process at the $\sim 96\%$ level in one of their two datasets, but only at the $\sim 50\%$ level in the other.

To properly detect flares, it is important to keep contamination from other nearby sources at a minimum, especially in the Galactic Centre and when it comes to the faintest flares. On April 25 2013, this became more difficult even for the high resolution of *Chandra* when a new transient source, an erupting magnetar (SGR J1745-2900), appeared $\sim$ 2.4" from Sgr A* (Kennea et al., 2013; Coti Zelati et al., 2017).

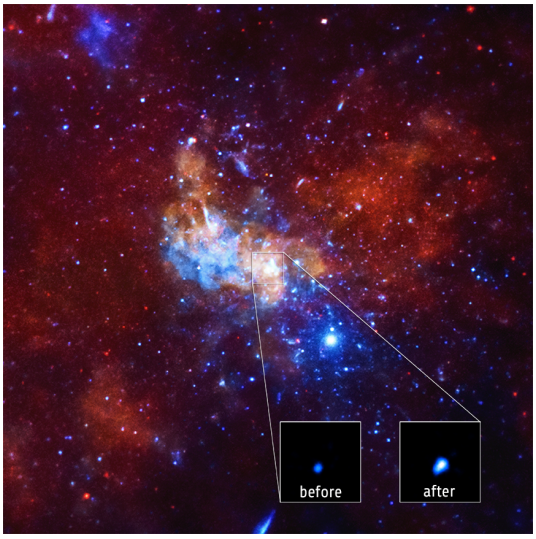

Figure 1.3: Image taken from a *Chandra* press release (http://chandra.harvard.edu/press/15_releases/press_010515.html, Credit: NASA/CXC/Northwestern Univ/D.Haggard et al). Shown are images of Sgr A* in quiescence and during a bright X-ray flare. Analogue of the bow ties in Figure 1.1. Note that the source in the "before" window is the Galactic Centre magnetar (not Sgr A*, since it is much fainter than the magnetar in its quiescent state).

Despite the extra flux leaking into Sgr A*'s extraction region (multiplying its detected quiescence count rate by a factor ~ 2 in the first following months), several authors reported an increase in Sgr A*'s bright/very bright flaring rate at the end of August 2014 following the pericentre passage of an extended object, G2 (Gillessen et al., 2012), around the SMBH (Ponti et al., 2015; Mossoux & Grosso, 2017).

1.2.3 G2

While G2's nature is still up for debate (ideas range from a clumpy gas stream originating from stellar winds collisions (Calderón et al., 2015) to enshrouded stars (Ballone et al., 2016) or even proto-planets (Mapelli & Ripamonti, 2015)), it is known that it is an extended cloud of ~ 3 Earth masses on a very eccentric orbit of $e \sim 0.97$ with a pericentre of $\sim 2000 R_s$ (Gillessen et al., 2012, 2013a,b). Several research groups predicted different

pericentre passage times for G2 depending on their proposed models. For example, Madigan et al. (2016) showed that a pure Keplerian orbit fit to the astrometry and velocity data from Gillessen et al. (2013a,b) and Pfuhl et al. (2015) (from observations performed with the VLT) gives an estimated time between the end of February to mid April 2014 and that the addition of a drag force pushes this date to between the end of May to mid-July. If an inflow is also added from the accretion flow, this prediction becomes from early July to early September. These predicted pericentre passage times are summarized in Figure 1.4. More recently, Gillessen et al. (2018), who had access to recent GRAVITY observations (Abuter et al., 2018a,b), proposed a model involving a drag force without an inflow (the drag forces come from the orbital velocity of G2 through a relatively stationary ambient medium, based on the fact that Sgr A* has both inflow and outflow components). They compared their orbital fit to one using only Keplerian motion and found that the drag force is significant at the 10σ level. They report a pericentre passage time of 2014.58 ± 0.13, earlier but consistent with Madigan et al. (2016).

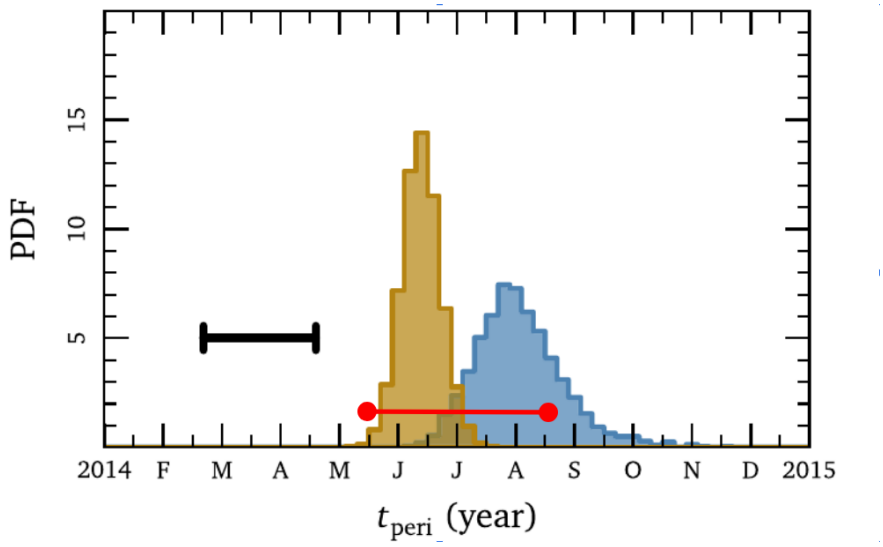

Figure 1.4: G2 pericentre times for different models. Figure taken from Madigan et al. (2016), but with the predicted pericentre time of 2014.58 ± 0.13 from Gillessen et al. (2018) added in red. The black bar is for a Kepler orbit. The derived time for models with inflow is in blue, and the time for models without inflow is in orange.

This cosmic encounter has led to many numerical simulations. Most early predictions

consisted of a brightening in the X-ray (Gillessen et al., 2012) and radio (Narayan et al., 2012; Abarca et al., 2014) which were not observed (Chandler & Sjouwerman, 2014; Haggard et al., 2014; Bower et al., 2015; Borkar et al., 2016). On the other hand, some simulations predicted a delayed brightening years after pericentre (Schartmann et al., 2012) due to G2's angular momentum. More recently, Kawashima et al. (2017), unlike previous works, performed long-term (decade timescale) 3D general relativistic magnetohydrodynamic simulations. They predict an increase in X-ray luminosity 5-10 years after pericentre. However, due to poor understanding of the physical mechanisms behind flares, no direct predictions regarding the flaring rate has been performed to date.

Regarding observations, Sgr A* was closely monitored in Summer and Fall 2014 to detect any increase in activity related to this encounter. Several X-ray observatories performed many observations, namely *Chandra*, *XMM-Newton* and *Swift*. Based on these observations, several research groups reported an increase in the bright flaring rate of Sgr A*.

Ponti et al. (2015) report 80 flares from *Chandra* and *XMM-Newton* in 6.9 Ms from September 1999 to November 2014. They find an increase in the bright/very bright flaring rate (defined as flares with an absorbed fluence greater than 5×10^{-9} erg cm^{-2}) from 0.27 ± 0.04 to 2.5 ± 1.0 a day at 99.9 % confidence after summer 2014 . They also report a decline in the moderate-bright flares from mid-2013 at the 96% level. However, these authors used PIMMS to convert count rates to fluxes. This leads to systematic underestimation of the fluxes because PIMMS doesn't correct the point spread function (PSF) extraction fraction. Ponti et al. (2015) used Bayesian Blocks (Scargle et al., 2013) to detect and characterize flares but they did not calibrate the prior.

Mossoux & Grosso (2017) also use Bayesian Blocks as their detection algorithm, but they calibrate their prior. Using 9.3 Ms of data from 1999 to 2015 from *Chandra*, *XMM-Newton* and *Swift*, they report 107 flares. They also find an increase in the most energetic flaring rate by a factor of 3 following 2014 August 31 and a decay for the faintest flares

by a similar factor. To do so, they use bottom-to-top and top-to-bottom searches on the detected flare fluxes and fluences. However, the way they apply this method removes observing gaps which doesn't properly treat partial flares caught at the beginning or end of an observation (edge flares).

Both of these authors use data from *Chandra* but also from *XMM-Newton* and *Swift* to have as many flares as possible. This complicates things significantly since all these observatories have different sensitivities. Additionally, neither of these authors perform extensive Monte Carlo simulations of X-ray flares. Mossoux & Grosso (2017) do simulate X-ray flares to control their detection bias and they also allow simulated flares to start before the beginning or after the end of a given observation to consider edge effects. However, they do not use their tool to study the impact of other factors such as simulations with multiple flares.

1.3 Outline of Thesis

In this thesis, I use the highest-quality *Chandra* observations of the Galactic Centre in an effort to be as consistent as possible. I use Bayesian Blocks to detect and characterize flares, and I create fluence and duration distributions of flares like Neilsen et al. (2013). By splitting the observations into two datasets around a potential change point related to G2 (and repeating this process for every potential change point), I test the null hypothesis by performing Monte Carlo simulations of each X-ray light curve within each dataset using a model with the same parameters to produce confidence intervals for the flare distributions of both datasets. If there exists an ensemble of parameters producing confidence intervals consistent with the observed distributions before and after each potential change point then the null hypothesis is confirmed, implying no significant change in Sgr A*'s flaring properties due to G2. The motivation is to test previous claims (e.g., Ponti et al. (2015)

and Mossoux & Grosso (2017)) of an increase in the bright flaring rate in late August 2014, corresponding with G2's pericentre passage.

The thesis is organized as follows. In Chapter 2, I explain the $Chandra$ data reduction. In Chapter 3, I describe my handling of potential contamination from the nearby magnetar. In Chapter 4, I explain how I used Bayesian Blocks to detect and characterize flares and show the detected flares. In Chapter 5, I present my Monte Carlo simulation model and look for different possible flare rate variations. Finally, in Chapter 6, I discuss my results and put them into the context of other works before concluding.

Chapter 2

Observations

2.1 X-ray Visionary Program data

Chandra observed Sgr A* for 38 observations during the 2012 *X-ray Visionary Program* (XVP) campaign, collecting a total of ∼3 Ms of data[1]. The observations were all taken with the High Energy Transmission Grating (HETG) Spectrometer with the Advanced CCD Imaging Spectrometer-Spectroscopy (ACIS-S) camera at focus. The HETG disperses some of the photons across the detectors which increases spectral resolution via the grating equation

$$\sin\beta = m\lambda/p \tag{2.1}$$

where β is the diffraction angle, m is an integer representing the diffraction order, λ is the wavelength and p is the grating lines spacing. β is computed from the distance of a given event from the zeroth-order (i.e., undispersed) image. λ is assigned using β (and also a m value deduced from the zeroth-order distance) and the energy E is computed with the usual equation $E\lambda = hc$, where h is Plank's constant and c is the speed of light. m is verified by using the event energy as measured by the ACIS-S camera and the distance from the undispersed image. Since the energy resolution of ACIS-S isn't perfect, I allow

[1] The average exposure time per observation is ∼ 78 ks.

an order sorting tolerance of ± 0.2 (for example, m values found between 0.8 and 1.2 are assigned to $m = 1$).

I perform basic data reductions using the Chandra Interactive Analysis of Observations[2] (`CIAO`) `v.4.9` tools (Fruscione et al., 2006) with the calibration database `CALDB v.4.7.6` and I reprocess the level 2 events file[3] with the `chandra_repro` script before updating the World Coordinate System (WCS) with the `wcs_update` tool. `chandra_repro` is a powerful tool that applies the most recent calibration and makes multiple corrections to the data, including filtering events from bad pixels. `wcs_update` corrects the sources' positions found in the observation in detector coordinates such that it minimizes the difference with respect to known sources from a catalog in RA and DEC. This increases the accuracy of source localisation and is particularly important in the crowded Galactic Centre.

The diffraction order m of each event is determined by the `tg_resolve_events` tool (which automates the steps previously explained) and I keep zeroth and $\pm$first order photons. I use the same extraction region as Nowak et al. (2012), Neilsen et al. (2013) and Mossoux & Grosso (2017) to minimise background. I extract 2-8 keV zeroth-order events from a small 1.25" radius circular region centered around the radio position of Sgr A*[4] and first-order counts (with grating order tolerance of ± 0.2) are extracted from a rectangular box with a 5 pixels (2.5") width also centered on the source (see Figure 2.1). The PHA2 files (FITS files containing events filtered to keep only 0th and $\pm$first order and a spectrum) and the gratings responses (matrices that define the sensibility of each pixel at a given energy) are created with the `tg_extract` and `mktgresp` tools, respectively, while the zeroth order spectra and response files are extracted with the `specextract` tool.

The exposure of each observation is not dead time corrected and is given by the time

[2] Available at `http://cxc.harvard.edu/ciao/`

[3] Level 2 events file refer to events that have gone through extra filtering to keep only good and reliable events

[4] 17:45:40.04909, -29:00:28.118 (Reid & Brunthaler, 2004)

between the first and last event. Dead time refers to the readout time of each event ($\sim$ 0.04104 s) during which no events can be recorded. In this instrument mode, the frame time is 3.1 s, but with the additional readout time it becomes effectively 3.14 s. I refer to this effective time as the frame time from now on.

2.2 Post-XVP data

Chandra observed Sgr A* again in April 2013 and has continued to do so in a series of on-going observations. To create a consistent dataset, I only use observations from the ACIS-S3 1/8th subarray instrument mode (see Figure 2.1). This instrument mode reduces pile-up (see Section 2.3) by only using a 1/8th subarray of the central S3 chip (128-rows), which reduces the frame time to 0.44104 s.

I use a 1.25" radius circular extraction region around Sgr A* to select 2-8 keV events after reprocessing the data with the `chandra_repro` tool and updating the WCS coordinate system with the `wcs_update` tool. I extract X-ray spectra with the `specextract` tool. There is one complication, however; in 2013 April 25, the magnetar SGR J1745-29, at only 2.4" from Sgr A*, went into outburst (Mori et al., 2013; Kennea et al., 2013). Its luminosity was so great that it leaked into Sgr A*'s extraction region. This influences how I handle background contributions for the majority of this dataset (see Chapter 3). Like for the XVP data, I do not perform dead time correction.

I exclude ObsIDs 14944 and 16597 from my analysis because the tool `wcs_match` did not properly correct the WCS coordinate system. This is likely due to their short exposure (20 ks and 18 ks) making the localization of the sources more difficult. Inaccurate positions of the extraction regions would significantly affect the contribution of the magnetar, reducing my ability to correct for the magnetar's contribution to Sgr A*'s quiescent flux (see Chapter 3). Similarly, I exclude ObsIDs 18055 and 18056 because of the presence of a low-mass X-ray binary in outburst, Transient 15 (Ponti et al., 2016) (see Figure 3.2).

My final Post-XVP ACIS-S3 1/8th subarray dataset consists of 1.56 Ms of data across 39 observations.

2.3 Pile-up

When two or more photons hit the same detector region in less than 1 frame time, they are counted as a single count with an energy corresponding to the total energy of the photons. This phenomenon is called pile-up and can happen to zeroth-order events in gratings data or any event in non-grating data. First-order events suffer very little pile-up because they are dispersed to a larger area, making it unlikely that more than one event hit the same detector region in a single frame time.

Pile-up can be described by Equation 2 of Nowak et al. (2012):

$$\Lambda_d = \left[\frac{\left[e^{\alpha \Lambda_i} - 1 \right] e^{-\Lambda_i}}{\alpha \Lambda_i} \right] \Lambda_i \tag{2.2}$$

where Λ_i is the incoming (unpiled) counts per frame, Λ_d is the detected (piled) counts per frame and α is the grade migration parameter (representing the fraction of recorded piled events) which I assume to be $\alpha = 1$.

For gratings data (frame time of 3.14s), I unpile the flaring count rates in the zeroth order of each flaring block (see Section 4.1). I retrieve an average 0th/1st order flaring count rate ratio (quiescence subtracted) across all XVP flares of $\sim$ 1.6, consistent with Nowak et al. (2012).

2.4 X-ray flare energies

To accurately compare between grating and non-grating flare fluences, I must convert counts to energy. Since Sgr A*'s flares have been shown to have very similar X-ray spectra (Nowak et al., 2012; Neilsen et al., 2013), their count rates are directly proportional to

their flux using the same absorption model. I use `XSPEC` (Arnaud, 1996) with a model of `dustscat*tbabs*powerlaw`, the abundances of Wilms et al. (2000) and the cross sections of Verner et al. (1996). I take a power-law index of $\Gamma = 2$, an hydrogen column density of $N_H = 14.3 \times 10^{22}$ cm^{-2}, a dust scattering optical depth $\tau = 0.324(N_H/10^{22}\text{ cm}^{-2})$ and I normalize to the the brightest flare of 2012, which has an absorbed 2-8 keV flux of $F_{\text{2-8}}^{\text{abs}} = 8.5^{+0.9}_{-0.9} \times 10^{-12}$ erg cm^{-2} s^{-1} and an unabsorbed 2-10 keV luminosity of $L_{2-10}^{\text{unabs}} = 19.2^{+7.2}_{-3.7} \times 10^{34}$ erg s^{-1} (Nowak et al., 2012). The unabsorbed fluences reported in Tables 4.1 and 4.2 are pile-up corrected and quiescence subtracted. To normalize to the brightest flare of 2012, I modify the model amplitude until it corresponds to that flare. I save the value $R_0 = 5.07 \times 10^{-4}$ (in units of 10^{-46} cm^{-2}) of this amplitude divided by its unabsorbed luminosity[5]. R_0 is purely model dependent and is the same for gratings and non-gratings observations.

To convert luminosities to count rates using R_0, I compute the normalization factor for each luminosity in Neilsen et al. (2013)'s Table 1 for each of *my* ObsID by multiplying R_0 by each luminosity and extract the count rate[6] associated with the given model. This results in a count rate for each unabsorbed luminosity listed in Neilsen et al. (2013)'s Table 1 for each of my own ObsID's instrument responses. I define R_1 as the ratio between a count rate and its associated unabsorbed luminosity[7]. For each ObsID, I take the mean R_2 of R_1. Finally, I take the median of R_2 across all of my ObsIDs to get a final conversion factor.

With this model, I find a conversion factor between unpiled, quiescence subtracted 2-8 keV flare count rates and their unabsorbed 2-10 keV luminosity of $R_2 = 0.0077$ ct/(10^{34} erg) for gratings flares and $R_2 = 0.013$ ct/(10^{34} erg) for subarray flares. The difference is

[5]This assumes that the unabsorbed 2-10 keV luminosity is linearly proportional to the absorbed 2-8 keV flux. This is justified in my case because the absorption/spectral model is similar for every flare.

[6]Since I do not have a pile-up component in this model, the count rates are incoming (i.e., unpiled) by default.

[7]The only difference between gratings and non-gratings observations in this process is that for gratings data I need to compute the model count rates for the 0th and 1st order independently since they have different spectral files. To get R_1, I use as count rate the sum of the 0th and 1st orders.

due to differing effective areas of these instrument modes. Indeed, ACIS-S/HETG0th+1st orders has an effective area of $\sim$ 200 cm^2 whereas ACIS-S has an effective area of $\sim$ 300 cm^2 (at 5 keV[8]). My conversion factors are consistent with those reported by Yuan & Wang (2015) where the authors report 0.0053 ct/(10^{34} erg) for the 0th order (which translates to $\sim$ 0.008 ct/(10^{34} erg) when adding the $\pm$1st order counts with the $\sim$ 3:2 ratio of Nowak et al. (2012)). For their ACIS-I data (which has similar effective area as ACIS-S3), they get 0.0136 ct/(10^{34} erg).

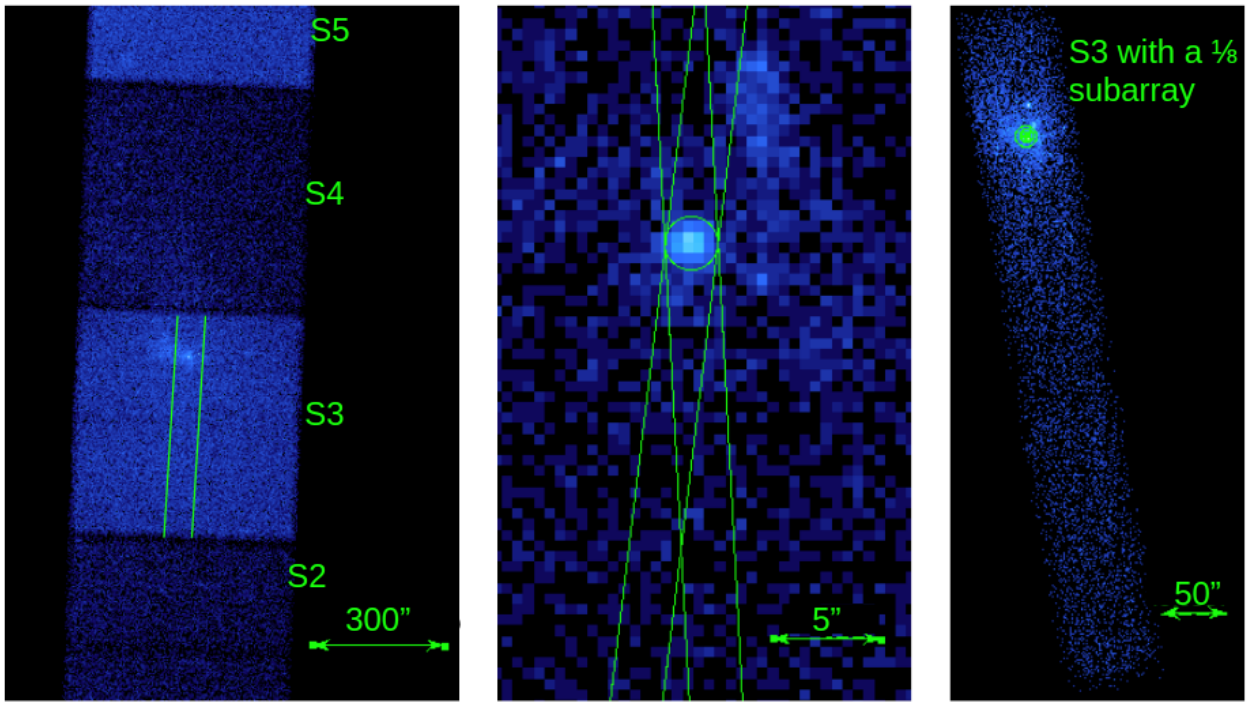

Figure 2.1: Example of observations done with the ACIS-S instrument, with and without the HETG. $Left$: HETG observation (ObsID 14392) with chips S2 to S5 labeled. The green rectangle on the S3 chip represents the 1/8 subarray used in the Post-XVP observations without the HETG. The bright spot on the upper part of this region is Sgr A*. $Middle$: Zoomed in version of the left image around Sgr A*. The green rectangular boxes with a 2.5" width are the $\pm$1st order extraction regions and the circle with a 1.25" radius is the 0th order extraction region. $Right$: ACIS-S3 chip with a 1/8 subarray observation (ObsID 14703). Figure 3.1 shows a zoomed-in version of this image around Sgr A* and the magnetar. 1" $\sim 10^5$ R$_s$ $\sim$ 8000 AU.

[8] http://cxc.harvard.edu/caldb/prop_plan/pimms/

Chapter 3

Magnetar model and contamination

3.1 The Galactic Centre magnetar SGR J1745-2900

The Galactic Centre is a crowded place. This is why *Chandra*'s high spatial resolution (angular resolution of 0.5") is crucial for observations of individual objects located in the region. Things get even more difficult when new bright sources appear. One such transient showed up on 2013 April 25 when it was detected by *Swift* as a Soft Gamma Repeater (SGR; Mori et al., 2013; Kennea et al., 2013). This magnetar, SGR J1745-2900, is 2.4" away from Sgr A* (Coti Zelati et al., 2017), making *Chandra* the only X-ray observatory capable of spatially resolving the two sources. Figure 3.1 shows a *Chandra* image taken on 2013 June 4 in which the magnetar is much brighter than Sgr A*. In fact, the magnetar only drops to a luminosity comparable to Sgr A*'s quiescence about 3 years after outburst. This means that a bright source is contaminating my Sgr A* extraction region for the majority of my Post-XVP data.

3.2 Estimation of the magnetar's contamination

To correct for the magnetar's contamination, I need to calculate the count rate fraction of the magnetar that overlaps with Sgr A*'s extraction region. To do so, I adopt the background region from Coti Zelati et al. (2017), i.e., an annulus with an inner radius of 5" and

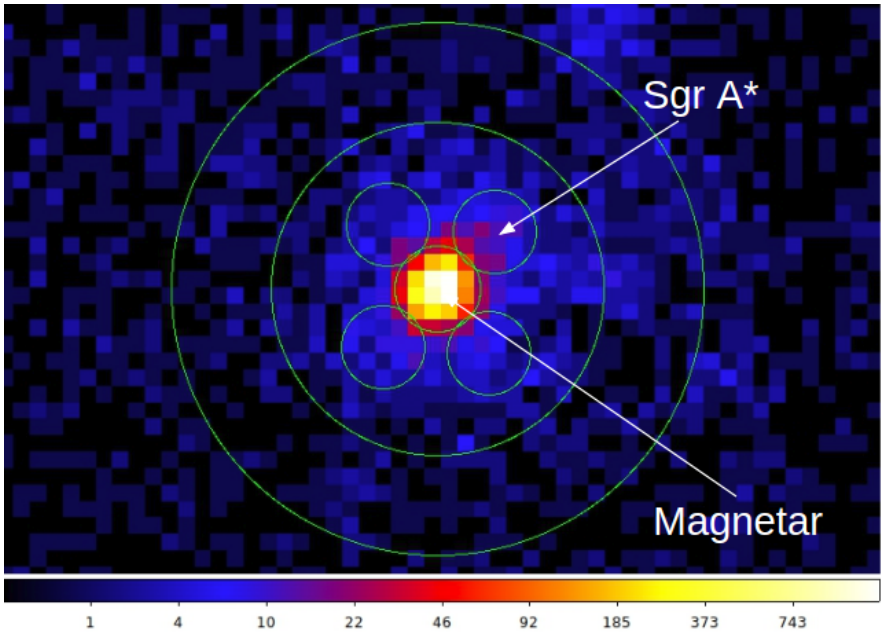

Figure 3.1: Image of ObsID 14703 from *Chandra*. The annulus is the background extraction region (inner radius of 5" and outer radius of 8", center on RA:17:45:40.084, DEC:-29:00:28.70). The brightest spot is the magnetar (extraction region centered on RA:17:45:40.169, DEC:-29:00:29.84 with a radius of 1.3"). The circle towards the upper-right from the magnetar is Sgr A* (extraction region centered on RA:17:45:40.0409, DEC:-29:00:28.118 with a radius of 1.25"). The three other regions have the same radius as Sgr A* and their center's position are : RA:17:45:40.2971, DEC:-29:00:31.57 and RA:17:45:40.0542, DEC:-29:00:31.7615 and RA:17:45:40.2838, DEC:-29:00:27.9185 respectively.

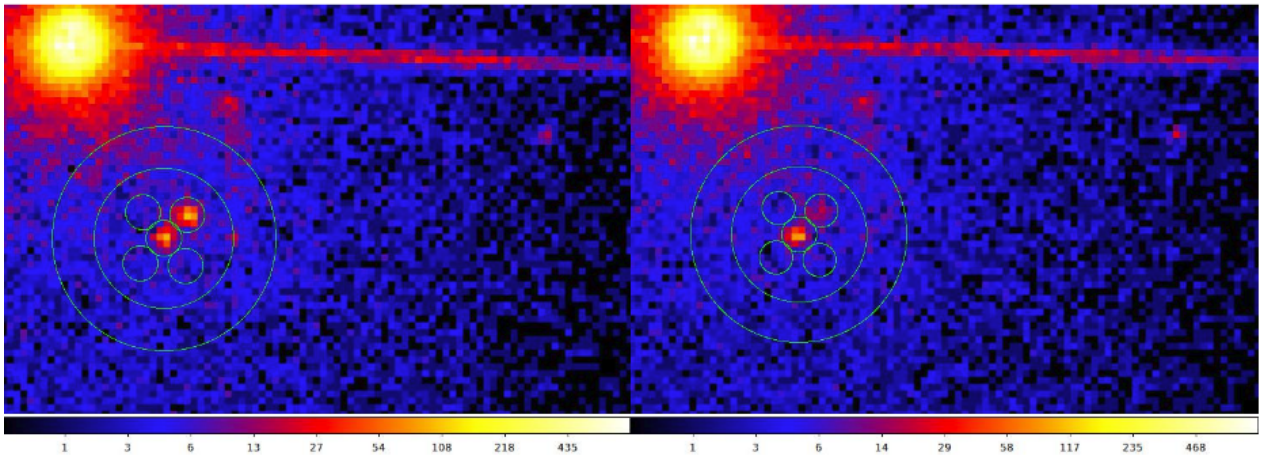

Figure 3.2: Same as figure 3.1, but for ObsIDs 18055 ($left$) and 18056 ($right$). We can clearly see a very bright source that saturates the CCDs (Transient 15). This contaminates significantly my background region. Sgr A* flared in ObsID 18055, which is why it is brighter.

outer radius of 8" (see Figure 3.1). I extract background-subtracted count rates in the 2-8 keV band from 3 regions located at the same radial distance from the magnetar as Sgr A*, with the same size as Sgr A*'s extraction region (radius of 1.25"). I take the mean of the extracted count rates and divide the result by the 2-8 keV background-subtracted count rate of the magnetar's extraction region (radius of 1.3"). This results in the fraction ϵ of the count rate Q_{magn} from the magnetar that leaks out to Sgr A*'s radial distance. In other words, the measured Sgr A* quiescence count rate Q_{eff} (corresponding to the longest block obtained with a Bayesian Blocks algorithm, see Chapter 4) is given by

$$Q_{\text{eff}} = \epsilon \times Q_{\text{mag}} + Q_{\text{sgr}} \tag{3.1}$$

where Q_{sgr} is the *actual* Sgr A* quiescent count rate.

I calculate ϵ for every ObsID of my Post-XVP observations until ObsID 18057 since at that point the magnetar is faint enough that count rates in the testing regions fall to background level. This method makes several assumptions: (1) the PSF is radially symmetric and (2) the 3 regions do not contain any significant X-ray sources.

The resulting plot of ϵ as a function of time since the outburst is shown in Figure 3.3. I get similar results whether or not I take pile-up into account. I correct for pile-up by calculating the predicted incoming count rate (which should be higher than the detected count rate) in the following way. I compute the piled count rates of 10 000 different incoming count rates between 0.001 and 10 counts/s equally separated in logarithmic space using Equation 2.2 with $\alpha = 1$. I identify which incoming count rate corresponded to the detected magnetar count rate for each ObsID and compute the magnetar fraction using this value. I do not consider pile-up in the 3 regions as their count rates are too low for pile-up to be significant. I find $\langle\epsilon\rangle = (1.4 \pm 0.2)\%$ without correcting for pile-up and $\langle\epsilon\rangle = (1.3 \pm 0.2)\%$ if I correct for it. The following analysis uses the ϵ obtained with a pile-up correction.

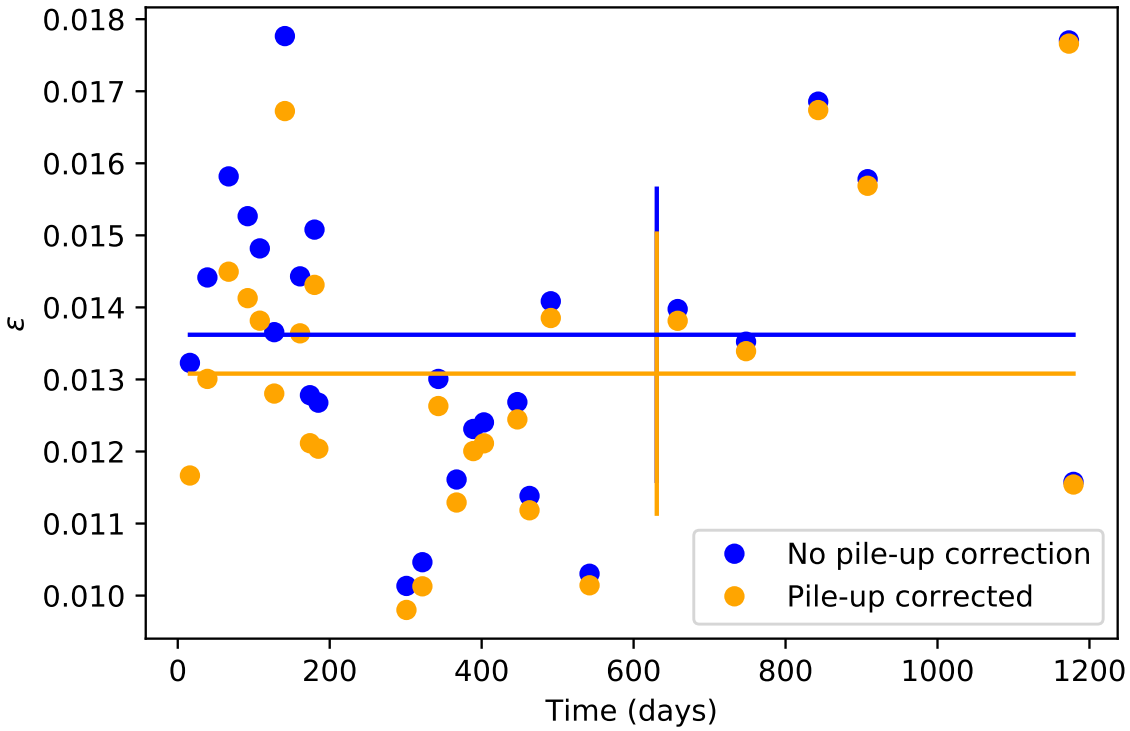

Figure 3.3: Each point corresponds to an ObsID. The horizontal lines are the mean value of their respective dataset (pile-up corrected in orange and not, in blue). Their associated standard deviation is marked with a vertical line. I obtain $\langle\epsilon\rangle = (1.4 \pm 0.2)\%$ without treating pile-up and $\langle\epsilon\rangle = (1.3 \pm 0.2)\%$ if I treat it.

3.3 Implication for Sgr A*'s quiescence count rate

Now that I know what fraction ϵ of the magnetar count rate contaminates Sgr A*'s extraction region, I can estimate its quiescent count rate Q_{sgr}. Following Equation 3.1, I subtract the magnetar contribution $\epsilon \times Q_{\text{magn}}$ from the observed quiescent count rate Q_{eff} of each ObsID (defined as the count rate of the longest block obtained from the Bayesian Blocks algorithm, see Chapter 4). This results in Sgr A*'s quiescent count rate Q_{sgr} for that ObsID. The associated error on each Q_{sgr} is given by

$$\sigma_{Q_{\text{sgr}}} = \sigma_{Q_{\text{eff}}} + \sigma_{\langle\epsilon\rangle} Q_{\text{magn}} \tag{3.2}$$

where $\sigma_{Q_{\text{eff}}}$ is the Poisson noise from the measurement of Q_{eff} from the Bayesian Blocks and $\sigma_{\langle\epsilon\rangle} = 0.2\%$ is the error of $\langle\epsilon\rangle$ when considering pile-up. In general, the Poisson error in a count rate is given by:

$$\sigma_{Poisson} = \frac{\sqrt{N_{cts}}}{t_{block}} \tag{3.3}$$

where N_{cts} is the number of counts within a duration t_{block}. In the present scenario, N_{cts} is the number of counts in the longest block of duration t_{block} from the Bayesian Blocks (see Chapter 4).

To find the mean quiescent count rate from this distribution, I use Monte Carlo simulations. For each point (i.e., for each ObsID), I daw a random point from its associated value and standard deviation. Once I have done this for every point, I save the mean and the standard deviation of the whole set. I repeat this process 10 000 times and take the mean of all the means as well as the mean of all the standard deviations. This results in a value of $\langle Q_{\text{sgr}}\rangle = (0.005 \pm 0.001)$ ct/s for Sgr A*'s average quiescent count rate.

Nowak et al. (2012) find that Sgr A*'s quiescent absorbed flux is $F_{2-8}^{\text{abs}} = 0.147^{+0.004}_{-0.003} \times 10^{-12}$ erg/cm^2/s using an absorbed power-law with index $\Gamma = 3.0^{+0.2}_{-0.2}$. They use an hydrogen column density of $N_H = 12.9^{+0.8}_{-0.8} \times 10^{22}$ cm^{-2}, with the abundances of Wilms et al. (2000) and the cross-sections of Verner et al. (1996). I use XSPEC (Arnaud, 1996) with a model $dustscat * tbabs * powerlaw$ using their parameters. Other parameters include τ, the optical depth to dust scattering at 1 keV and $Halo_{sz}$, the halo size at 1 keV relative to detector beamsize. Still following Nowak et al. (2012), I choose $\tau = 0.324(N_H/10^{22}\text{ cm}^{-2})$ and $Halo_{sz} = 50$.

Using this model, I extract the corresponding count rates Q_{Nowak} for each ObsID using the instrument response files and compare them to the expected value, Q_{sgr}. This is shown in Figure 3.4 and Table 3.1, where we can see that my resulting mean quiescent count rate $\langle Q_{\text{sgr}}\rangle$ is consistent with Nowak et al. (2012)'s prediction. Indeed, I obtain $\langle Q_{\text{sgr}}\rangle = (0.005 \pm 0.001)$ ct/s while $\langle Q_{\text{Nowak}}\rangle = (0.0052 \pm 0.0008)$ ct/s. I obtain the error associated

with their mean count rate by the same method as I use for my own mean count rate, where the uncertainty associated with each ObsID is the Poisson error. The agreement between the two results confirms my treatment of the magnetar's contribution to Sgr A*'s extraction region.

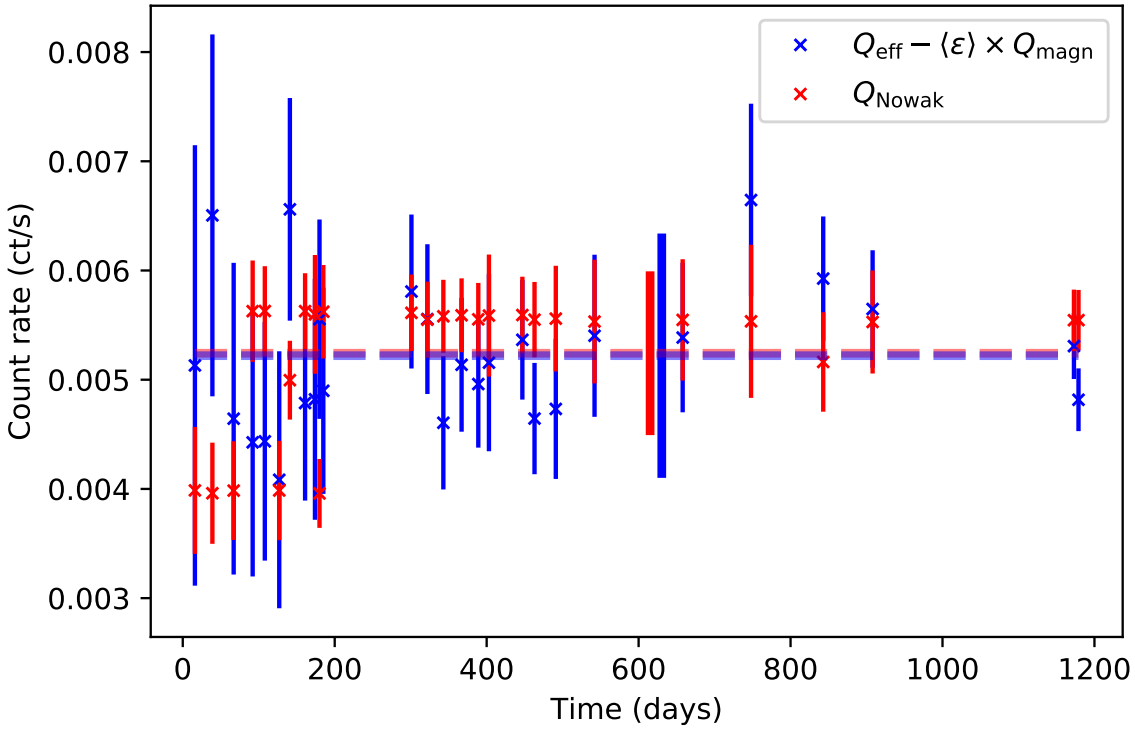

Figure 3.4: This plot shows the difference between Q_{eff} and $\langle\epsilon\rangle \times Q_{\text{magn}}$ in blue as well as the quiescent count rate predicted by Nowak et al. (2012) in red, Q_{Nowak}, at each ObsID present in Figure 3.3. The error bar of each blue point is found by propagating the errors of the parameters of Equation 3.2. The mean count rate $\langle Q_{\text{sgr}} \rangle = (0.005 \pm 0.001)$ ct/s is represented by the blue horizontal line, with its associated 1σ Monte Carlo error bar in the same color. The red line shows $\langle Q_{\text{Nowak}} \rangle = (0.0052 \pm 0.008)$ ct/s. Its error bar is also a 1σ Monte Carlo error bar, and the error bar of each red point is the Poisson error.

Figure 3.5 shows the summary of this analysis. Since my own mean quiescent value is very close to the value measured by Nowak et al. (2012), both choices are reasonable. In the following work, I choose to use my own quiescent value.

3.3.1 A note about the background

To add the quiescent count rates from Nowak et al. (2012) to the magnetar contribution, it makes sense that I previously used background subtracted count rates to compute said fraction ϵ, otherwise I would be adding the background twice. Indeed, Nowak et al. (2012) do not background subtract at all (therefore their quiescence already contains some background). This is also true for my own count rates extracted from the lowest block from the Bayesian blocks. Since those count rates are obtained from event files, they cannot be background subtracted.

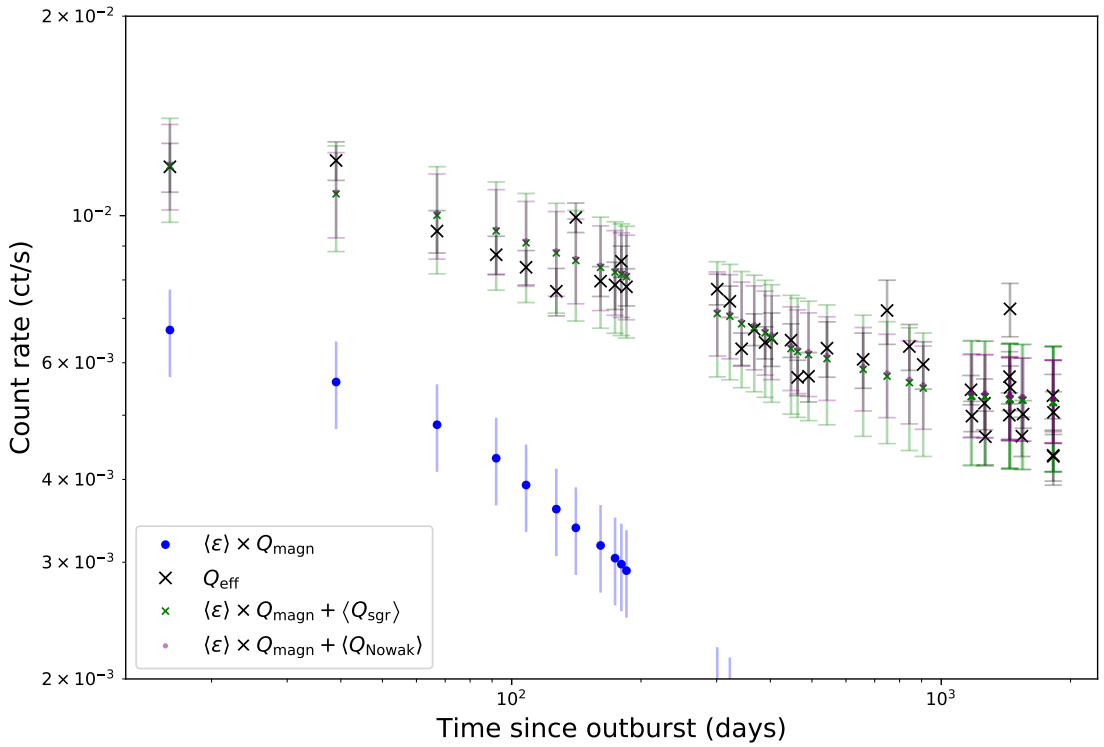

Figure 3.5: Count rates plotted against time since the magnetar outburst. The blue points are the magnetar contribution count rates and the orange points are the count rates of the longest block of each ObsID. The green points show the sum of the count rates from the empirical quiescence and the magnetar contribution, and the purple points represent the sum of the count rates from Nowak et al. (2012)'s quiescence and the magnetar contribution.

In summary, the magnetar SGR J1745-2900, which is located at an angular distance of 2.4" from Sgr A* and was detected on 2013 April 25 by *Swift* and $NuSTAR$, contributes $(1.3 \pm 0.2)\%$ of its count rate to Sgr A*'s extraction region. After correcting for this contamination, we calculate Sgr A*'s quiescent count rate to be (0.005 ± 0.001) ct/s, which is in agreement with the prediction of Nowak et al. (2012) of (0.0052 ± 0.0008) ct/s.

ObsID	Date	Time since outburst	$\langle\epsilon\rangle \times Q_{\rm mag}$	$Q_{\rm eff}$	$Q_{\rm Nowak}$	$Q_{\rm sgr}$
–	–	Days	$\times 10^{-3}$ ct/s	$\times 10^{-3}$ ct/s	$\times 10^{-3}$ ct/s	$\times 10^{-3}$ ct/s
14702	2013-05-12	16	7 ± 1	12 ± 1	4.0 ± 0.6	5 ± 2
14703	2013-06-04	39	5.6 ± 0.8	12.1 ± 0.8	4.0 ± 0.5	6 ± 2
14946	2013-07-02	67	4.8 ± 0.7	9.5 ± 0.7	4.0 ± 0.4	4 ± 1
15041	2013-07-27	92	4.3 ± 0.6	8.7 ± 0.6	5.6 ± 0.5	4 ± 1
15042	2013-08-11	108	3.9 ± 0.6	8.4 ± 0.5	5.6 ± 0.4	4 ± 1
14945	2013-08-31	127	3.6 ± 0.5	7.7 ± 0.6	4.0 ± 0.5	4 ± 1
15043	2013-09-14	141	3.4 ± 0.5	9.9 ± 0.5	5.0 ± 0.4	7 ± 1
15044	2013-10-04	161	3.2 ± 0.5	8.0 ± 0.4	5.6 ± 0.3	4.7 ± 0.9
14943	2013-10-17	174	3.0 ± 0.5	7.9 ± 0.6	5.6 ± 0.5	5 ± 1
14704	2013-10-23	180	3.0 ± 0.4	8.5 ± 0.5	4.0 ± 0.3	5.5 ± 0.9
15045	2013-10-28	185	2.9 ± 0.4	7.8 ± 0.5	5.6 ± 0.4	4.9 ± 0.9
16508	2014-02-21	301	1.9 ± 0.3	7.7 ± 0.4	5.6 ± 0.4	5.8 ± 0.7
16211	2014-03-14	322	1.9 ± 0.3	7.4 ± 0.4	5.5 ± 0.3	5.5 ± 0.7
16212	2014-04-04	343	1.7 ± 0.3	6.3 ± 0.4	5.6 ± 0.3	4.6 ± 0.6
16213	2014-04-28	367	1.6 ± 0.2	6.7 ± 0.4	5.6 ± 0.3	5.1 ± 0.6
16214	2014-05-20	389	1.5 ± 0.2	6.4 ± 0.4	5.6 ± 0.3	5.0 ± 0.6
16210	2014-06-03	403	1.4 ± 0.2	6.2 ± 0.6	5.6 ± 0.5	4.9 ± 0.8
16215	2014-07-16	447	1.1 ± 0.2	6.5 ± 0.4	5.6 ± 0.3	5.3 ± 0.5
16216	2014-08-02	463	1.1 ± 0.2	5.7 ± 0.4	5.5 ± 0.3	4.6 ± 0.5
16217	2014-08-30	491	0.9 ± 0.1	5.7 ± 0.5	5.6 ± 0.5	4.7 ± 0.6
16218	2014-10-20	542	0.9 ± 0.1	6.3 ± 0.6	5.5 ± 0.6	5.4 ± 0.7
16963	2015-02-13	658	0.7 ± 0.1	6.1 ± 0.6	5.5 ± 0.6	5.3 ± 0.7
16966	2015-05-14	748	0.55 ± 0.08	7.2 ± 0.8	5.5 ± 0.7	6.1 ± 0.9
16965	2015-08-17	843	0.42 ± 0.06	6.3 ± 0.5	5.2 ± 0.5	5.9 ± 0.6
16964	2015-10-21	908	0.32 ± 0.05	6.0 ± 0.5	5.5 ± 0.5	5.6 ± 0.5
18731	2016-07-12	1173	0.15 ± 0.02	5.5 ± 0.3	5.5 ± 0.3	5.3 ± 0.3
18732	2016-07-18	1179	0.17 ± 0.03	5.0 ± 0.3	5.5 ± 0.3	4.8 ± 0.3
18057	2016-10-08	1261	0.15 ± 0.02	5.2 ± 0.5	5.5 ± 0.5	–
18058	2016-10-14	1267	0.15 ± 0.02	4.6 ± 0.4	5.6 ± 0.5	–
19726	2017-04-06	1441	0.10 ± 0.01	5.0 ± 0.4	5.5 ± 0.4	–
19727	2017-04-07	1442	0.11 ± 0.02	5.7 ± 0.5	5.5 ± 0.5	–
20041	2017-04-11	1446	0.11 ± 0.02	7.2 ± 0.7	5.5 ± 0.6	–
20040	2017-04-12	1447	0.10 ± 0.02	5.5 ± 0.4	5.5 ± 0.4	–
19703	2017-07-15	1541	0.09 ± 0.01	4.6 ± 0.3	5.5 ± 0.3	–
19704	2017-07-25	1551	0.09 ± 0.01	5.0 ± 0.2	5.5 ± 0.3	–
20344	2018-04-20	1820	0.055 ± 0.008	5.4 ± 0.4	5.5 ± 0.4	–
20345	2018-04-22	1822	0.048 ± 0.007	4.4 ± 0.4	5.5 ± 0.4	–
20346	2018-04-24	1824	0.049 ± 0.007	4.3 ± 0.4	5.5 ± 0.5	–
20347	2018-04-25	1825	0.056 ± 0.008	5.1 ± 0.4	5.5 ± 0.4	–

Table 3.1: Time since the magnetar's outburst (taken to be 2013 April 25 (Kennea et al., 2013)), count rate $\langle\epsilon\rangle \times Q_{\rm mag}$ leaking from the magnetar to Sgr A*'s extraction region, detected quiescent count rate $Q_{\rm eff}$, predicted quiescent count rate $Q_{\rm Nowak}$ according to Nowak et al. (2012) and $Q_{\rm sgr} = Q_{\rm eff} - Q_{\rm mag}$ for each ObsID. From ObsID 18057, the magnetar is faint enough that count rates in the testing regions fall to background level, resulting in $\epsilon < 0$. This is why the corresponding $Q_{\rm sgr}$ values are undefined.

Chapter 4

Detected flares

4.1 Bayesian Blocks and prior calibration

4.1.1 The Bayesian Blocks algorithm

To identify flares in the data, I choose Bayesian Blocks (Scargle et al., 2013). I start from the implementation provided by Williams et al. (2017), which I modify to use event lists from simulations (Chapter 5) and my own prior calibration (section 4.1.2 and Appendix A). The algorithm assumes that the data can be separated into different blocks of constant count rates. If it detects a significant change in count rate at a specific point in the data, it will start a new block. This is called a "change point". The first point is always a change point, such that the number of blocks is given by $N_{blocks} = N_{cp} + 1$, where N_{cp} is the number of non-trivial change points. The goal of the algorithm is to find the optimal way of separating the data. To do so, it needs a prior on the number of change points. This prior should assign a smaller probability for a large number of blocks, otherwise it would lead to overfitting (e.g., making each point its own block in an extreme case). The prior depends on a single parameter, γ, as follows:

$$P(N_{\text{blocks}}) = P_0 \gamma^{N_{blocks}} \tag{4.1}$$

where $P(N_{\text{blocks}})$ denotes the prior probability of having a number of blocks N_{blocks} and P_0 is a normalization factor. Naturally, $0 < \gamma < 1$. As in Scargle et al. (2013), I define the prior on the number of change points as:

$$ncp\_prior = -\log \gamma \tag{4.2}$$

The intuition behind this parameter is that for larger values of ncp_prior, the algorithm is expecting fewer blocks since it translates to a lower value of γ and thus a lower probability for larger N_{blocks}.

The algorithm aims to maximize the log-likelihood of a given block k having a certain duration $T^{(k)}$ and number of events $N^{(k)}$, translating to a count rate $\lambda = N^{(k)}/T^{(k)}$. Mathematically, for the case of X-ray events data, the maximum likelihood is given by:

$$\log L_{\text{max}}^{(k)} + N^{(k)} = N^{(k)} \left(\log N^{(k)} - \log T^{(k)} \right). \tag{4.3}$$

However, this likelihood must take the prior into account. A simple way to implement this is to subtract Equation 4.2 from Equation 4.3.

Since ncp_prior is linked to the expected number of blocks, it can be mathematically connected to the probability of a given change point being a false positive. I refer to this false alarm probability as p_0 (more is explained about p_0 in Section 4.1.2). ncp_prior should also depend on the number of events in a given light curve. Indeed, if one light curve has more events than another, it should be more susceptible to false positives. Therefore, ncp_prior is a function of p_0 and N, the number of events. It must also depend on the underlying statistics of the data. For example, it behaves differently for Gaussian noise than for Poisson noise. As such, the algorithm needs to be properly calibrated on signal-free data before being applied to real data. The next Section explains in detail how I determine the calibration shown in Figure 4.2 that computes an appropriate value of ncp_prior for

$p_0 = 0.05$ for a light curve with a given number of events N.

4.1.2 Prior calibration

To calibrate ncp_prior on Poisson noise, I follow the procedure suggested in Scargle et al. (2013). I choose a number of expected counts N and a constant noise count rate cr, in units of counts/s. I also adopt a frame time of 0.441 s (I tested with 3.14 s, and it did not change anything; the calibration is frame time independent). I find the expected length L of the simulated observation needed to reach N counts, by computing $N/(cr * frame\_time)$. With this length, I create a time array from 0 to the predicted length multiplied by the frame time such that the time step is one frame time. I generate a Poisson light curve of the same length with an expected count rate $cr \times frame\_time$. This light curve dictates how many counts are in each frame time (since a Poisson distribution only generates non-negative integers). In other words, each array index $0 \leq i \leq (L-1)$ in the light curve has an element $n \geq 0$, and I repeat the corresponding ith element in the time array n times.

The resulting event list is sent to the Bayesian Blocks algorithm, which returns the number of change points detected. For each N considered, I search for the value of ncp_prior that correctly results in a given false positive rate p_0. To do so, I evaluate values of ncp_prior between 3 and 15, with a step of 0.1, and generate 1000 random Poisson light curves. p_0 is then the sum of the detected change points divided by the number of light curves (recall that if the Bayesian Blocks algorithm output a single block, then the number of change points is $N_{cp} = 0$). Figure 4.1 shows an example of p_0 as a function of ncp_prior for 3 different N, with a line indicating $p_0 = 0.05$, corresponding to 50 or fewer change points in 1000 light curves. To identify the closest ncp_prior that gives the desired false positive rate (p_0) for a given N, we take the first value of ncp_prior in Fig 4.1 that results in a p_0 value less than the desired p_0.

To create the final calibration of ncp_prior values corresponding to $p_0 = 0.05$, I simulate

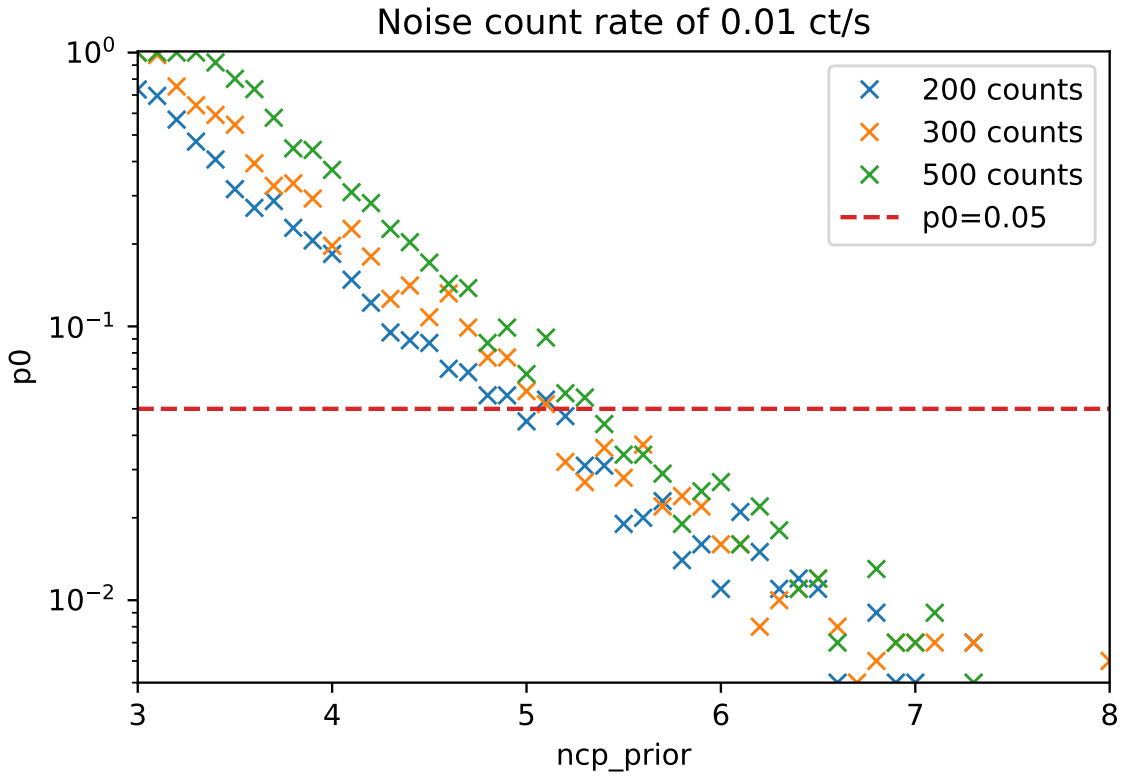

Figure 4.1: Number of change points detected in 1000 random Poisson-generated signal-free light curves divided by 1000 (p_0) as a function of ncp_prior for different number of expected counts in the light curves. The red horizontal dotted line indicates $p_0 = 0.05$

Poisson light curves with the following values of N: 100, 200, 300, 400, 500, 600, 750, 1000, 1250, 1500, 1750, 2000, 2250, 2500 and 3000 expected counts. To accelerate the process, I stop sweeping through ncp_prior as soon as I find the desired value for p_0, and for the next value of N I start looking at values of ncp_prior 0.5 less than the value found for the previous N, because ncp_prior increases with N. For each value of ncp_prior, if the number of change points goes above the number of false positives (50 for $p_0 = 0.05$ and 1000 light curves), I skip it and go to the next value instead of computing the remaining light curves.

Since the value of ncp_prior chosen with my method is a bit noisy (e.g., Figure 4.1), I run the simulation for 10 different noise count rates (0.004, 0.005, 0.006, 0.007, 0.008,

0.009, 0.01, 0.011, 0.012 and 0.1 count/s)[1] and calculate the mean of the different ncp_prior values obtained at each N, as well as their standard deviation. The resulting dependence of ncp_prior on N for the desired value of $p_0 = 0.05$ is displayed in red points in Figure 4.2.

I fit the following function to the data points in Figure 4.2:

$$ncp\_prior(N) = A \log N + B \tag{4.4}$$

where $ncp\_prior(N)$ is the fitted value of the prior for a given expected number of events N, A is the amplitude and B is a constant. Figure 4.2 also shows the corresponding 3σ error region. The error bars on the data points are the standard deviations of the values of ncp_prior across all simulations.

For comparison, I also plot 2 other calibrations used in Scargle et al. (2013). They were derived via similar simulations, but for Gaussian noise, and only considered N values up to 1024. I show the 3σ error region to accommodate the noise and to minimize false positives (recall that higher values of ncp_prior imply that the algorithm is less susceptible to false positives). Roughly, this means that I am 3σ confident that my calibration yields $p_0 \leq 0.05$. My final calibration is the upper bound of the 3σ fit; every time I run Bayesian Blocks (either for simulations or real data), I use this fit and the number of events in the given event list to assign a value of ncp_prior.

4.1.3 Blocks classification

In addition to the Bayesian Blocks implementation, I need a set of criteria to determine which blocks are flaring blocks. This is especially important since I perform extensive Monte Carlo simulations and need to treat simulations and observations consistently.

[1] I choose different count rates to double-check my intuition that changing the count rates does not change the results. Indeed, if the count rate changes, the only modification in the Bayesian Blocks should be the mean count rate of each block (i.e., the probability of finding false change points shouldn't change). The plot with all the different count rates is shown in Appendix A.

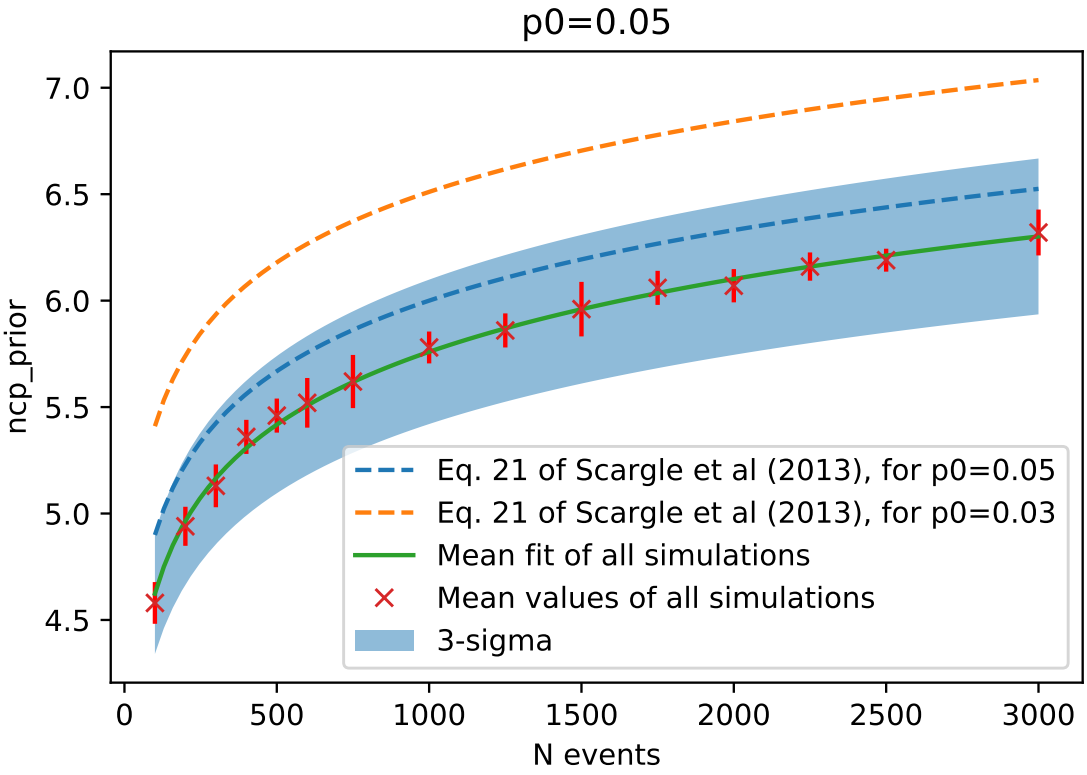

Figure 4.2: Summary plot showing the results of 10 different calibration runs for different noise count rates. Each red point represents the mean value of ncp_prior found across those runs for that given expected number of events N, the error bars are their standard deviation. The green line is the resulting best fit and the blue shaded region is its 3σ error envelope. The blue and orange dotted lines are the calibrations obtained by Scargle et al. (2013) for similar simulations for Gaussian noise for $p_0 = 0.05$ and $p_0 = 0.03$, respectively. They didn't provide error bars.

Since Sgr A* has a low flaring rate ($\sim$ 1.1 flare/day (Neilsen et al., 2013)), its quiescent count rate can be determined from the longest block in each observation. Given that the uncertainty on that block's count rate will be rather small (since it spans a long period), I consider a block as a flaring block if the difference between its count rate and the quiescent count rate is greater than $3\,(\sigma_Q + \sigma_{block})$, where σ_Q and σ_{block} are the quiescent and flaring block's Poisson error on their count rate (defined by Equation 3.3), respectively. Each consecutive flaring block is associated with the same flare. While this criteria may seem conservative, it is necessary otherwise clearly distinct flares separated by a block barely

above quiescence can be considered as a single flare. Figure 4.3 displays two distinct flares that would be grouped together into one flare without this criteria.

Once all the flaring blocks have been found, I unpile them using Equation 2.2 (see Appendix C for details about the XVP data, which is more complex since pile-up is only relevant in the 0th order). I then multiply each block's count rate by their duration, add the resulting number of counts from each flaring block together, and divide it by the total duration of all the blocks in the flare, resulting in the flare's average count rate. Then, I subtract the quiescent count rate from it, and convert the resulting rate to a fluence using the conversion factors presented in Section 2.4.

Since I assume simple Gaussian-shaped flares for my simulations (see Chapter 5), it is possible in principle for two long and bright flares to occur too close together in time to be distinguishable. This is discussed by Yuan & Wang (2015), where the authors find that such a phenomenon is quite unlikely (only 1/3 of their flares showing evidence of substructures have a probability of being made of multiple flares above 5%). I introduce another criterion in an effort to distinguish close flares and test my whole analysis with and without it. The criterion stipulates that if there is a series of consecutive flaring blocks with one of them (not the first nor the last one) having a count rate significantly lower than the others (the count rate difference is larger than 3 error bars added in quadrature between the block and its neighbours), then it is flagged and the flare is considered to be made of two flares. This block is separated in half, with its first half being associated with the left flare and the other half with the right flare. This occurs in only 3 observations, all of which are Post-XVP (ObsIDs 15043, 16218 and 20346). Figure 4.4 shows ObsID 16218 as an example. In Chapter 5, I do not use this criteria unless specified otherwise. I report my flares in Tables 4.1 and 4.2.

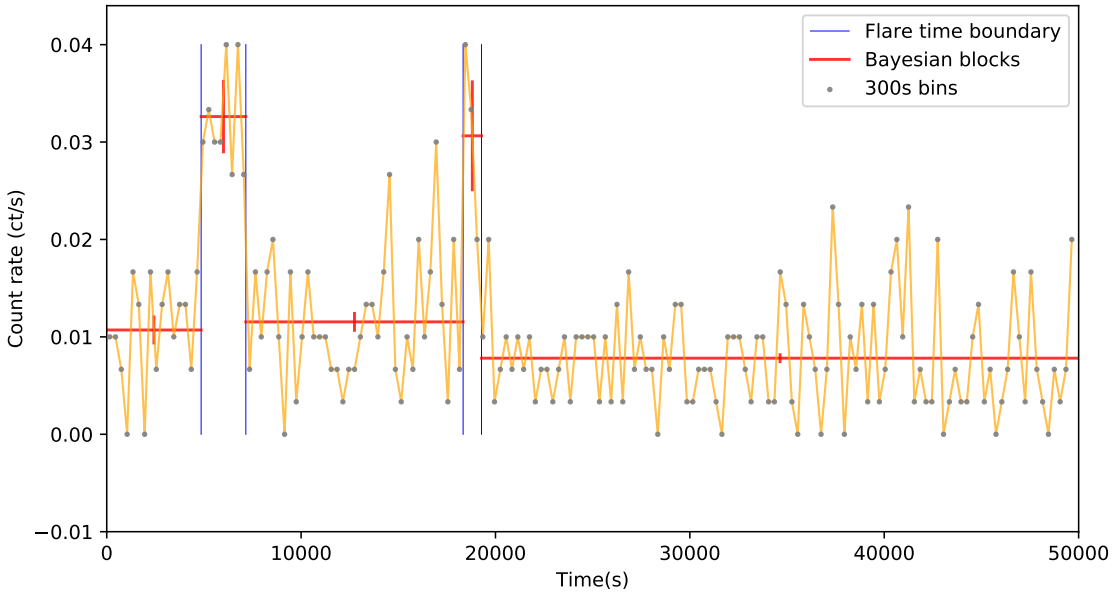

Figure 4.3: Light curve of ObsID 15045. The bin time is 300 s, the Bayesian Blocks are in red with their associated Poisson errors and the blue vertical bars indicate the beginning and end times of the two detected flares with the criteria explained in Section 4.1.3. If one relaxes the constraint, the block between the two flares will be considered a flaring block as well, thus leading to the detection of a single flare instead of 2.

4.2 Detected flares

I find that the flaring rate is consistent across the datasets as I detect 40 XVP flares in 3 Ms (1.2 ± 0.2 flare/day) and 18 post-XVP flares in 1.56 Ms (1.0 ± 0.2 flare/day)[2]. If I add the flare separation criteria, I have the same number of XVP flares, but I find 21 post-XVP flares giving a rate of 1.2 ± 0.3 flare/day, consistent with Neilsen et al. (2013). The properties of each flare are provided in Tables 4.1 and 4.2. I compare my detected flares in detail with other works (Neilsen et al., 2013; Ponti et al., 2015; Mossoux & Grosso, 2017) in Appendix B. My results are mostly consistent with the different authors, but some differences occur for the faint/short flares, mainly due to what different authors consider

[2]Post-XVP light curves are available at https://www.dropbox.com/sh/gu8zt95h8eukj4n/AAC4ZyuWOjF8A96YzdW1z4fsa?dl=0 and XVP light curves are at https://www.dropbox.com/sh/d2ox22wy6ek0caj/AAAX8UiO3nAPUZWBTqYJyy3ha?dl=0

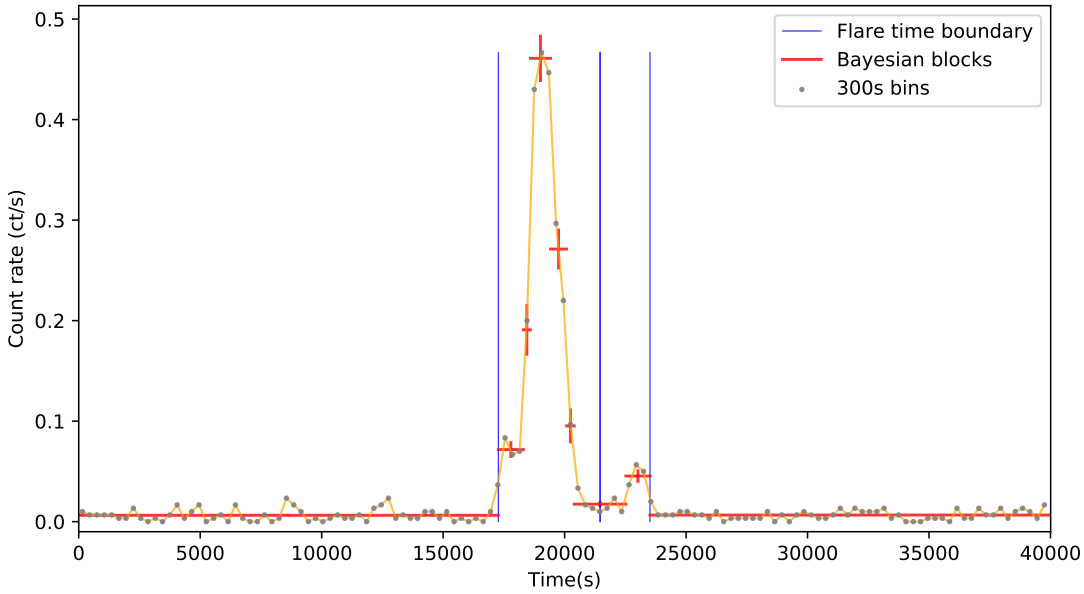

Figure 4.4: Like Figure 4.3 but for ObsID 16218. There is an apparent dip between two flares. However, the block between the two flares is still significantly above quiescence. If I use the flare separation criteria, this block is cut in half and I end up with two flares instead of one.

a flaring block, the Bayesian Blocks calibration and the flare detection method.

ObsID	Exp (ks)	Quie. cr ($\times 10^{-3}$ ct/s)	Fstart (MJD)	Fend (MJD)	Dur (s)	Mean cr (ct/s)	Fluence (2-10 keV, 10^{37} erg)
13850	60.0	5.9 ± 0.3					
14392	59.0	6.9 ± 0.5	55966.433	55966.458	2109	0.021 ± 0.003	3.74
			55966.602	55966.667	5641	0.145 ± 0.004	101.1
14394	17.4	6.8 ± 0.6					
14393	41.1	8.0 ± 0.4					
13856	39.7	5.6 ± 0.4					
13857	39.5	6.7 ± 0.5					
13854	22.7	8 ± 1	56006.487	56006.494	669	0.054 ± 0.009	3.97
			56006.528	56006.538	817	0.055 ± 0.008	4.96
			56006.586	56006.596	944	0.053 ± 0.007	5.45
			56006.682	56006.687	402	0.11 ± 0.02	5.08
14413	14.7	6.5 ± 0.7					
13855	19.9	6.9 ± 0.6					
14414	20.0	5.9 ± 0.5					
13847	153.4	6.3 ± 0.3	56048.511	56048.549	3251	0.024 ± 0.003	7.56
14427	79.9	5.7 ± 0.5	56054.097	56054.151	4679	0.020 ± 0.002	8.71
			56054.468	56054.490	1899	0.020 ± 0.003	3.46
13848	97.8	6.4 ± 0.3					
13849	178.3	6.8 ± 0.3	56058.690	56058.704	1172	0.028 ± 0.005	3.19
			56059.008	56059.064	4891	0.015 ± 0.002	5.36
			56059.315	56059.330	1244	0.025 ± 0.004	2.96
			56060.137	56060.166	2435	0.055 ± 0.005	15.46
13846	55.8	6.1 ± 0.3					
14438	25.3	6.3 ± 0.5					
13845	133.3	5.9 ± 0.2	56066.576	56066.595	1633	0.018 ± 0.003	2.53
			56067.864	56068.026	14006	0.017 ± 0.001	20.38
14460	23.8	5.6 ± 0.7					
13844	19.8	5.9 ± 0.5					
14461	50.6	7.0 ± 0.4					
13853	73.3	5.7 ± 0.3					
13841	44.8	6.2 ± 0.4					
14465	44.2	5.8 ± 0.5	56126.977	56127.034	4911[a]	0.019 ± 0.002	8.25
			56127.177	56127.203	2255	0.018 ± 0.003	3.56
14466	45.0	7.0 ± 0.4	56128.552	56128.555	286[a]	0.08 ± 0.02	2.62
13842	191.6	6.1 ± 0.3	56130.187	56130.225	3246	0.039 ± 0.003	13.84
			56130.908	56130.919	925	0.055 ± 0.008	5.87
			56131.494	56131.585	7865	0.021 ± 0.002	15.63
13839	175.5	6.5 ± 0.2	56132.389	56132.399	856	0.051 ± 0.008	4.95
			56133.999	56134.180	15614	0.026 ± 0.001	40.33
13840	162.2	6.7 ± 0.2	56136.457	56136.507	4271	0.015 ± 0.002	4.42
			56136.627	56136.655	2346	0.019 ± 0.003	3.72
14432	74.2	5.9 ± 0.3	56138.558	56138.609	4482[a]	0.013 ± 0.002	3.88
			56139.374	56139.416	3651[a]	0.053 ± 0.004	22.57
13838	99.1	6.5 ± 0.3	56141.013	56141.049	3067	0.055 ± 0.004	19.48
13852	155.9	6.9 ± 0.3	56143.317	56143.331	1215	0.053 ± 0.007	7.22
			56144.330	56144.350	1724	0.022 ± 0.004	3.45
14439	111.3	6.3 ± 0.2	56147.132	56147.147	1338	0.026 ± 0.004	3.51
14462	133.8	5.9 ± 0.3	56207.180	56207.191	947	0.031 ± 0.006	3.07
			56208.187	56208.220	2822	0.025 ± 0.003	7.02
14463	30.4	6.7 ± 0.6	56216.240	56216.247	525	0.12 ± 0.02	7.89
13851	106.8	5.3 ± 0.3	56217.096	56217.097	152	0.07 ± 0.02	1.34
			56217.814	56217.878	5486	0.074 ± 0.004	51.76
15568	35.9	6.3 ± 0.4					
13843	120.4	6.4 ± 0.3	56223.381	56223.476	8174	0.034 ± 0.002	29.02
15570	68.4	6.0 ± 0.3	56225.233	56225.262	2436	0.032 ± 0.004	8.10
14468	145.9	5.8 ± 0.2	56230.298	56230.338	3504	0.024 ± 0.003	8.38
			56231.567	56231.592	2178	0.030 ± 0.004	6.90

Table 4.1: List of all *Chandra* XVP observations. Listed for each observation are the quiescence count rate, the exposure (taken as the total time spanned by the Bayesian Blocks), the flare start and end times and the duration. The flare mean count rates are pile-up corrected but not quiescence-subtracted and I indicate their Poisson error bars. The fluences are from 2-10 keV (not 2-8 keV like the mean count rates), are pile-up corrected and quiescence-subtracted.

[a] Flare truncated by the beginning or end of the observation.

ObsID	Exposure (ks)	Fstart (MJD)	Fend (MJD)	Dura (s)	Mean cr (ct/s)	Fluence (2-10 keV, 10^{37} erg)
15041	50.0	56500.146	56500.158	1041	0.035 ± 0.006	2.03
		56500.460	56500.470	850	0.033 ± 0.006	1.55
15042	49.2	56516.357	56516.371	1262	0.028 ± 0.005	1.84
		56516.405	56516.441	3109	0.028 ± 0.003	4.60
15043	50.0	56549.085	56549.108[a]	2036	0.75 ± 0.02	127.86
		56549.108[a]	56549.149	3503	0.49 ± 0.01	142.91
15045	50.0	56593.675	56593.702	2300	0.033 ± 0.004	4.31
		56593.831	56593.842	946	0.031 ± 0.006	1.63
16508	47.8	56710.024	56710.048	2066[b]	0.027 ± 0.004	3.02
16217	37.8	56899.488	56899.546	4979	0.019 ± 0.002	5.07
16218	40.0	56950.557	56950.606[a]	4189	0.191 ± 0.007	58.79
		56950.606[a]	56950.630	2055	0.031 ± 0.004	3.81
16963	24.8	57066.252	57066.262	869	0.044 ± 0.007	2.51
16966	24.6	57156.501	57156.537	3131	0.054 ± 0.004	11.14
18731	86.3	57581.947	57581.957	904	0.036 ± 0.006	2.06
18732	84.4	57587.623	57587.653	2564	0.022 ± 0.003	3.36
20041	33.9	57854.352	57854.379	2287	0.103 ± 0.007	16.70
19703	89.0	57949.959	57950.009	4610[b]	0.015 ± 0.002	3.40
		57950.547	57950.561	1210	0.033 ± 0.005	2.66
20346	33.0	58232.206	58232.219[a]	1082	0.12 ± 0.01	9.39
		58232.219[a]	58232.246	2373	0.117 ± 0.007	20.48

Table 4.2: List of all *Chandra* flares from my Post-XVP flare dataset. Listed for each flare are the observation start date, the exposure (taken as the total time spanned by the Bayesian Blocks), the flare start and end times and the duration. The flare mean count rates are pile-up corrected but not quiescence-subtracted and I indicate their Poisson error bars. The fluences are from 2-10 keV (not 2-8 keV like the mean count rates), are pile-up corrected and quiescence-subtracted. They are obtained from the spectral model used by Neilsen et al. (2013) via a method explained in Section 2.4.

[a] Flares showing important substructure, which could indicate the presence of two flares if I use the criteria explained in Section 4.1.3.

[b] Flare truncated by the beginning or end of the observation.

Chapter 5

Simulations

I simulate *Chandra* X-ray light curves by creating event files and running Bayesian Blocks on them. Quiescence count rates are first determined and then flare fluences and durations are drawn from a power-law distribution. Their number is determined assuming a Poisson flaring rate which also yields their occurrence times. For each rendering of a given entire dataset, I simulate every ObsID and save the detected flare fluences and durations. I run these simulations thousands of times with the same parameters to test the null hypothesis by creating confidence intervals for the flare fluence and duration distributions which I use to compare with the real dataset. I aim to determine if there are indeed change points in the flaring behaviour of Sgr A* as reported by Ponti et al. (2015) and Mossoux & Grosso (2017). Instead, I verify the null hypothesis, i.e, I find that the same selection of best-fit parameters match the data before and after each change point. This Chapter explains each step in detail, including the differences that must be considered when simulating gratings versus non-gratings light curves.

5.1 Modeling event files

The foundation for this section is adapted from Appendix D of Mossoux & Grosso (2017). The basis for my simulations are the event files. These are files containing a list of time stamps corresponding to photon arrival times. The minimum resolvable time between

each event corresponds to the frame time (it is also possible to have multiple events in a single frame time). In the case of the ACIS-S3 chip with a 1/8 subarray, the frame time is about 0.441 s, while in HETG gratings observations it is 3.14 s.

To simulate a constant count rate, I use the frame time, the count rate, and the length of the observation. I generate a random light curve obeying Poisson statistics via a Cumulative Distribution Function (CDF). This function outputs, for each time t, the probability that an event occurs before that t. Naturally, it must be 0 at $t = 0$ and 1 at $t = T$, where T is total duration of the light curve. For a constant count rate, the CDF is:

$$CDF_c(t) = t/T \tag{5.1}$$

For a constant count rate r, the expected number of counts is $N_c = rT$. The actual number of counts N follows a Poisson distribution:

$$P(k;\lambda) = \lambda^k \frac{e^{-\lambda}}{k!} \tag{5.2}$$

where in my case $k = N$ is the random variable and $\lambda = N_c$ is the expected value.

After selecting a random number of events N, I draw N numbers uniformly distributed between 0 and 1 and associate them with a t using the inverse of equation 5.1. I show the proof that the events are uniformly distributed in Appendix E. Each t corresponds to an event time, or simulated photon. Since the resolution of the time array is the frame time, if two random t are less than a frame time apart, they will fall at the same t[1].

A more interesting case is when there is a flare in the light curve. To simulate flares (under the simplifying assumption that they are Gaussian as in Neilsen et al. (2013) and Mossoux & Grosso (2017)), I need more parameters, namely the flare peak count rate A (i.e., its amplitude), the flare's standard deviation σ, as well as the time of the flare's peak, t_{peak}. Since the quiescence count rate of Sgr A* is constant, the flare peak count rate is

[1]See Section 2.3 about pile-up.

taken to be the measured peak count rate minus the quiescence count rate. Again, I utilize the CDF; the CDF of a Gaussian is given by (Evans, 1993):

$$CDF_g(t) = \frac{A\sigma}{N_g}\sqrt{\frac{\pi}{2}}\left(\text{erf}\left(\frac{t_{peak}}{\sqrt{2}\sigma}\right) + \text{erf}\left(\frac{t - t_{peak}}{\sqrt{2}\sigma}\right)\right) \tag{5.3}$$

where N_g is the expected number of counts in the Gaussian and erf is the error function. To get N_g, I integrate the Gaussian flare over the observation, i.e.:

$$N_g = \int_0^T Ae^{\frac{(t-t_{peak})^2}{2\sigma^2}}dt \tag{5.4}$$

Numerically, I limit the integration to $\pm 5\sigma$ around t_{peak} to avoid underflow errors. If the flare occurs less than 5σ from the edge of an observation, I integrate until the edge instead.

Since the quiescent emission is always present, I combine the CDF of the quiescence and the flare, via the sum of the normalized CDFs:

$$CDF_{tot}(t) = \frac{N_c}{N_g + N_c}CDF_c(t) + \frac{N_g}{N_g + N_c}CDF_g(t). \tag{5.5}$$

The expected number of events is $N = N_c + N_g$, so I draw a random number of events N_{tot} around the expectation value using Equation 5.2. I again draw N_{tot} random numbers uniformly distributed between 0 and 1 and find the t to which each of them corresponds using the inverse of Equation 5.5. Numerically, I implement this by taking advantage of the fact that the time array and the CDF are both always increasing with t, such that if I sort the N_{tot} random values, I know that the drawn times will be sorted automatically. Thus, I can use a fast function from $numpy$ called $searchsorted$ to associate each of the N_{tot} values with a time t. Due to the numerical imprecision at high t for exponential functions, like erf, I divide each CDF by its maximum value (which is also its last value) to be sure that it is always normalized. A full example of CDFs is shown in Figure 5.1. Note that the CDFs for the Gaussian and the constant count rate are the CDF times their normalization

factors from Equation 5.5, which gives the total, normalized CDF.

An example of a simulated Poisson light curve from the combined CDF is shown in Figure 5.2. Each point corresponds to a 300 s bin. The red lines are the Bayesian Blocks, obtained with $p_0 = 5\%$ and properly calibrated for Poisson noise (see section 4.1.2 for details); the red vertical line is its corresponding Poisson error bar (computed with Equation 3.3). The blue curve represents the actual Gaussian formed from the input parameters (without any noise).

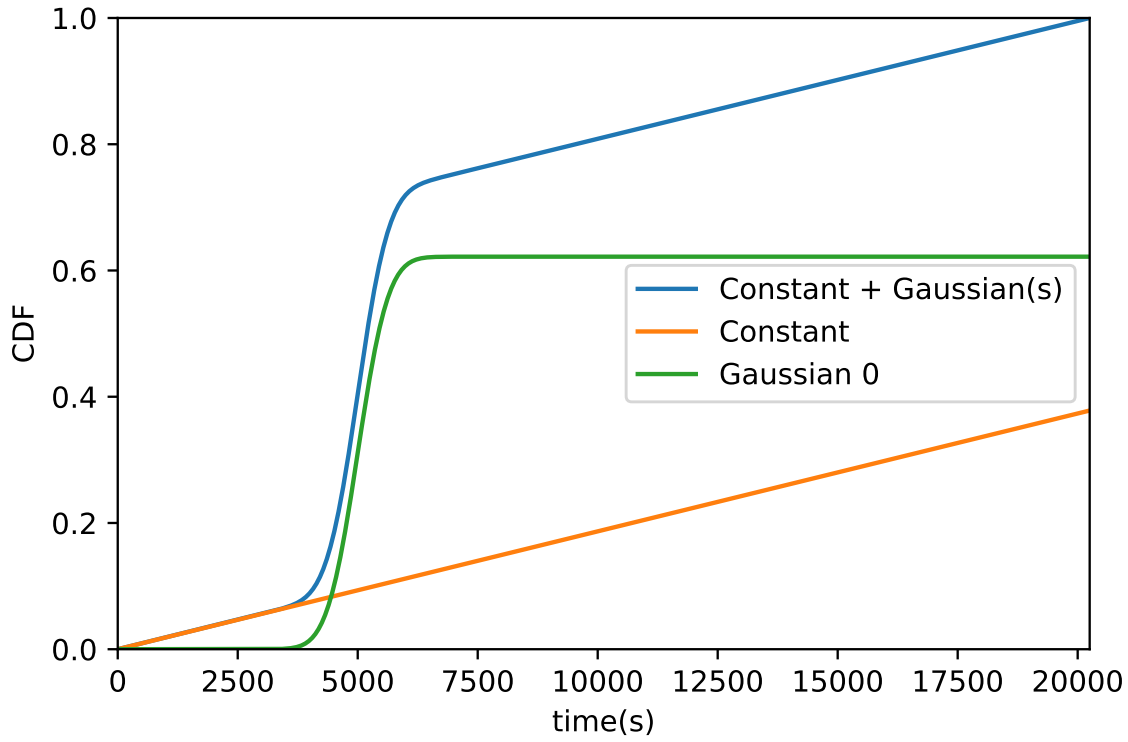

Figure 5.1: Cumulative distribution function (CDF) as a function of time for a simulated light curve made of a quiescent (orange line) and a flaring component (green line). The total CDF is the blue curve and is normalized to its own maximum value.

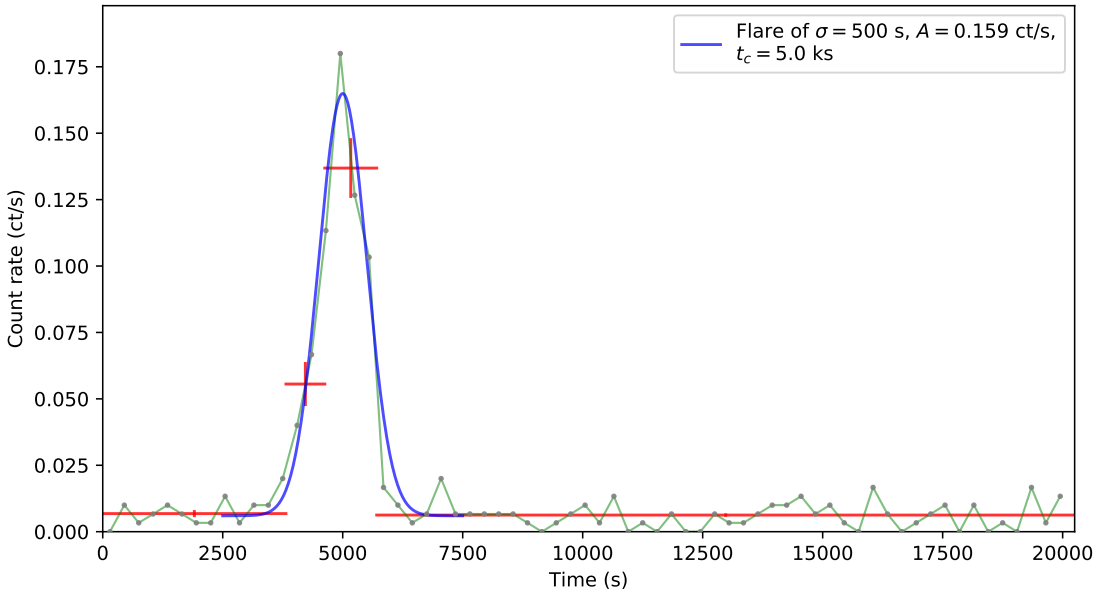

Figure 5.2: A random Poisson light curve generated from the CDF shown in Figure 5.1. The simulated quiescence count rate is 0.006 count/s. There is also a flare centered at 5000 s, with an amplitude of 0.159 count/s and a standard deviation of 500 s. This Gaussian flare is plotted as the blue curve. The grey dots are 300 s bins (the points are plotted at their middle time such that the first point is at 150 s not 0 s). The red horizontal lines are the Bayesian Blocks with their respective Poisson error.

5.2 Simulated quiescence count rates

A crucial part of the simulations is the quiescence count rate. For XVP HETG data, I implement it in a rather simple way. Since I do not detect any trend over long periods in the observed quiescence count rates of this dataset (see Table 4.1), I simulate a random Poisson quiescent count rate around the median of the observed values, 0.0063 ct/s, for each simulated observation. That is, I draw a random Poisson number N of counts around the total expected number of counts from a count rate of 0.0063 ct/s in an observation with a given exposure T. The count rate is then N/T.

For Post-XVP light curves, I need to consider the presence of the magnetar, for which I use my work from Chapter 3. I draw random values from $\langle \epsilon \rangle \times Q_{\rm mag}$ for the corresponding

ObsID and $Q_{\rm sgr}$ and add them together. For ObsIDs after 18732, I draw random Poisson quiescence count rates from $Q_{\rm sgr}$ similar to the XVP case.

5.3 Simulated flares

The total simulated time for a given ObsID is given by the sum of the duration of each of its Bayesian Blocks (given in Tables 4.1 and 4.2 as the exposure) plus an additional 32 ks before and after to allow edge flares, i.e., flares that occur near the edge of a given observation. It is important to consider this since several edge flares were observed in my dataset. (See Figure 5.3 for an example of a simulated light curve with an edge flare.) Flare times are placed randomly in the simulated observation according to a given Poisson flaring rate.

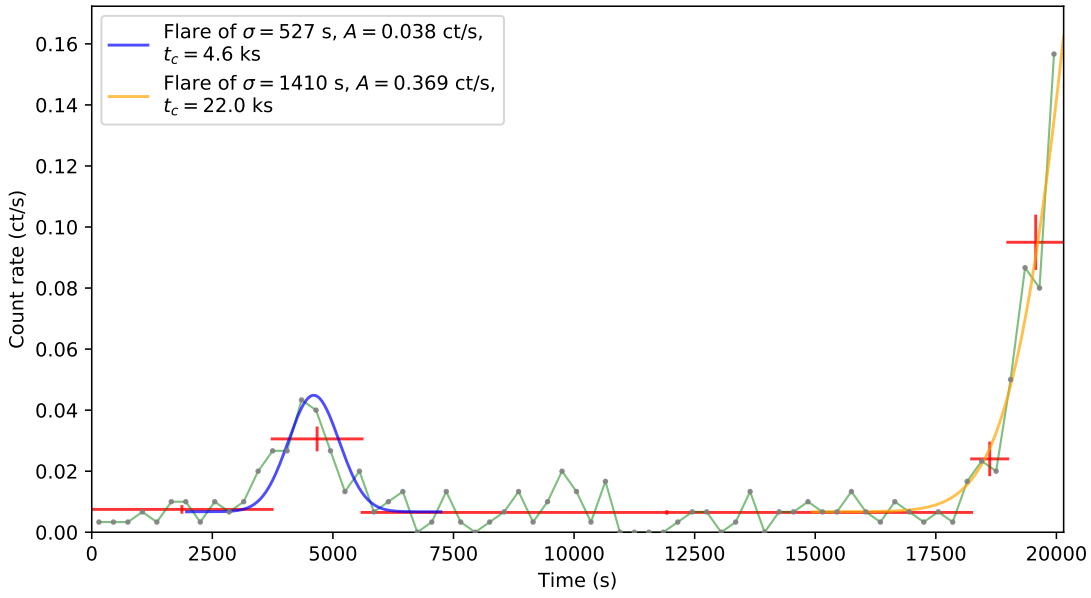

Figure 5.3: Similar to Figure 5.2. Example of a simulated light curve containing a normal flare in blue and an edge flare in orange.

To simulate Gaussian flares at each flare time, I assume that flare fluence[2] and duration distributions follow a power-law (Neilsen et al., 2013) and compute their associated CDF. I explain how to compute the CDF of a power-law distribution in Appendix D. I start by drawing a duration from its CDF. I only consider flares that occur at times t in a given ObsID of exposure T such that $-3\sigma \leq t \leq T + 3\sigma$. (I assume that σ is given by the flare duration divided by 4, as will be explained later in this section). If a given flare occurs outside this range, it is rejected from the simulation since it would not be detectable. If a flare is accepted, I draw a fluence from its CDF. I assume that there is no correlation between fluence and duration as shown in Yuan et al. (2017), where the authors prove that the apparent correlation seen in Neilsen et al. (2013) is caused by detection bias. Indeed, faint flares can only be detected if they are short, otherwise their mean count rate is too low to be detectable.

From these quantities, I find the mean count rate of a given simulated flare by dividing the fluence by the duration and converting the inferred luminosity to a count rate using the conversion factor found in Section 2.4 for the corresponding observing mode. This count rate needs to be piled prior to running Bayesian Blocks. To do so, I use Equation 2.2 and the count rate that is piled is the sum of the mean count rate of the flare and the quiescence. I recover the piled flare mean count rate by subtracting the quiescence from the computed total piled count rate. For XVP HETG data, it is slightly more complicated because pile-up only happens in the 0th order counts. I explain how I handle HETG pile-up in Appendix C.

To find the amplitude A of the Gaussian flare, I suppose that the mean count rate is found by considering a duration of $\pm 2\sigma$ around the time t_{peak}. Knowing that 95.5% of the

[2] I choose to use fluence instead of luminosity since fluence is an integrated quantity and thus contains more signal than luminosity.

Gaussian is contained in this range, I can find its amplitude A as follows:

$$cr_m = \frac{A\int_{t_{peak}-2\sigma}^{t_{peak}+2\sigma} e^{\frac{(t-t_{peak})^2}{2\sigma^2}} dt}{4\sigma} = \frac{0.955A \times \sqrt{2\pi}\sigma}{4\sigma} \tag{5.6}$$

$$A = \frac{2\sqrt{2}cr_m}{0.955\sqrt{\pi}} = 1.67cr_m \tag{5.7}$$

where cr_m is the mean count rate (count/s). In practice, the faintest flares are often detected for less than 4σ (or not detected at all) and the brightest flares can be detected for more than 4σ. I discuss this in the following section.

In summary, for a given flare time, I draw a flare duration and I convert it to the σ of a Gaussian by dividing it by 4. I draw a fluence if the flare is at most 3σ before (after) the observation's start (end). I calculate the luminosity of the flare and convert it to a mean count rate which is then piled. Finally, I compute the amplitude A of the flare and before simulating the event list, I remove the extra 32 ks that were added on each side of the observation.

I run Bayesian Blocks on the simulated event lists and retrieve the detected flares (and the quiescence count rate) in the same way I did for the data in Section 4.1. The value of ncp_prior used by the algorithm is computed from the total number of events N in the event list (using the calibration explained in Section 4.1.2).

5.4 A discussion about my model's response function

At this point, I can test how well my model retrieves flare parameters (i.e., its response function). I show in Figure 5.4 an example of a simulation that uses the detected parameters of the two flares from ObsID 14392 as input. My algorithm retrieves fluence values very close to the originals (3.5 and 100.0 vs 3.74 and 101.1, all in units of 10^{37} erg). The

duration values aren't as good (1300 s and 6300 s whereas the input values were 2109 s and 5641 s, respectively). The relative inaccuracy of the durations are acceptable since my goal is to study the behaviour of Sgr A*'s flaring rate for flares of different fluences, (i.e., I want to determine if the flaring rate of bright/very bright flares is constant or if it changes at some point as claimed by Ponti et al. (2015) and Mossoux & Grosso (2017).)

Figure 5.5 shows the fluence response function obtained by simulating a typical quiescence count rate of 0.006 ct/s in subarray mode. Each point represents one simulation and I show 10 000 in total. Flare durations and fluences are drawn from *uniform* distributions between 500 s and 8000 s and 1.3 to 275 $\times 10^{37}$ ergs, respectively. I use uniform distributions here because this gives a better general understanding of what is happening. To isolate the response function, only 1 flare was simulated at the centre of an observation with an exposure of 50 ks. This removes complicating edge effects, which bias the detected fluences. I plot an orange line to guide the eye around the perfect response. The overall median of the detected/input flare fluences is 1.06. If I look at flares simulated with fluences less than 50 $\times 10^{37}$ ergs, this improves to 1.017. This is interesting since $\sim 94\%$ of simulated flare fluences assuming a power-law index of -1.7 are below this value, meaning that my algorithm does an excellent job for the vast majority of my simulated flares. This power-law index is used in my model to explore change points as explained in the following sections. I performed the same simulations for the grating instrument mode and found a global median of 1.06, which becomes 0.999 for simulated fluences of less than 50 $\times 10^{37}$ ergs. This difference originates from the different conversion factor between count rate and luminosity (the detected/input flare fluences ratio is dragged down because the mean count rates go down, so the flares tend to be detected for a shorter time, hence the lower median).

The detected/input flare fluence ratio tends to increase with fluence because bright flares tend to be detected for a longer time. On the other hand, faint flares tend to be detected for shorter periods. Furthermore, long and bright flares are made of multiple

blocks, so the pile-up correction will be stronger for the highest blocks. This, combined with extra duration, increases the detected fluence with respect to the input fluence because when the flare is injected in the simulation, only the global mean count rate is piled. This is the best I can do when injecting the flare into the simulation since I cannot predict how many blocks will be created during its detection. However, this effect is small. Figure 5.5 indicates that the response function tends to ~ 1.05 for the brightest flares. The dominant effect must thus be the extra detected duration which adds back some of the extra $\sim 4.5\%$ of the Gaussian (Equation 5.6).

Long and faint flares have their fluence generally underestimated more than others because the amplitude of the Gaussian flare goes down, which in turn makes it harder to pick up for the Bayesian Blocks tool. The latter will detect the flare for a shorter amount of time, translating to even lower fluence.

The strong noise in the plot comes from the inherent Poisson counting noise and the fact that durations and fluences were drawn independently (flares with similar fluence can have very different durations).

I provide the code that simulates flares of given durations and fluences in a given instrument mode at `https://github.com/Elie23/X-ray-flare-simulator`.

5.5 Parameters needed for the model

My model needs multiple parameters. I previously described the quiescence, the conversion factor between flare count rate and luminosity, the 0th/1st order ratios in both quiescence and flares for XVP simulations, the exposure time, and the extra time added before and after each simulated light curve. Other parameters include the power-law indices of the duration and the fluence distributions (Γ_{Dura} and Γ_{Fluence}, respectively), the simulated range of duration (from D_{min} to D_{max}) and fluence (from F_{min} to F_{max}), and the flare rate R.

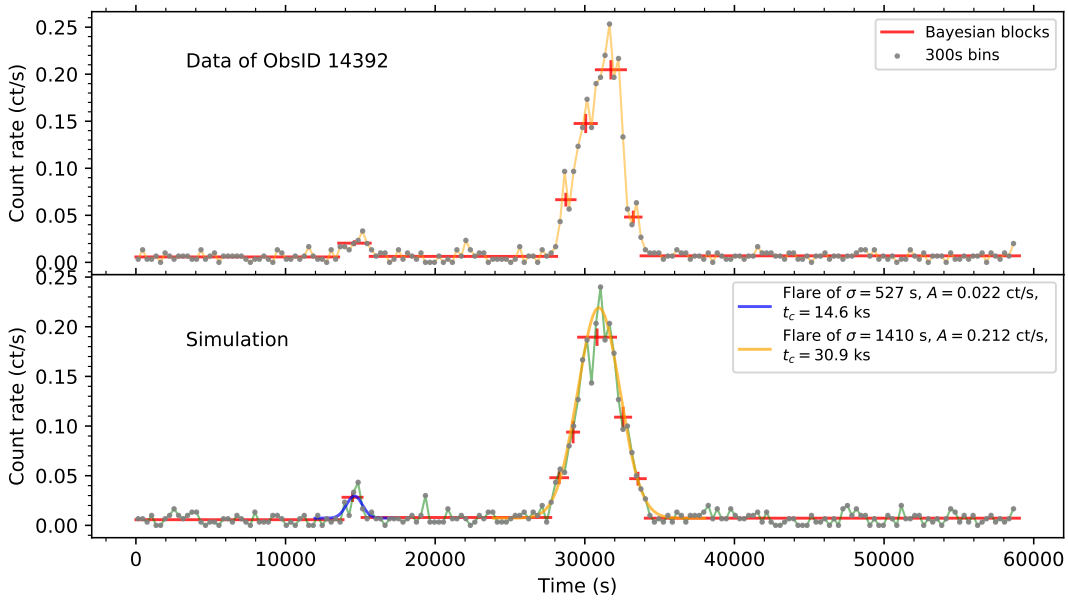

Figure 5.4: Top: 2-8 keV light curve of ObsID 14392 showing the combined zeroth and first order events in 300 s bins. Shown in red are the Bayesian Blocks. Bottom: Simulated light curve using the same quiescence and flare parameters as those measured in ObsID 14392 as input. The two flares detected in ObsID 14392 are represented by the blue and orange Gaussians, respectively. Their parameters are stated in the legend. The Bayesian Blocks algorithm retrieves unabsorbed flare fluences of 3.5×10^{37} erg and 100.0×10^{37} erg (the input values were 3.74×10^{37} erg and 101.1×10^{37} erg).

5.6 Simulating entire datasets

With these tools, I can simulate entire datasets made up of multiple ObsIDs. For each ObsID in a given instrument mode within the dataset of interest, I draw a quiescence count rate as explained in Section 5.2. I convert a given R to a rate per frame which I use to draw a Poisson number for each frame in the ObsID (of duration equal to the corresponding observed ObsID plus an additional 32 ks before and after). This number, effectively always 0 or 1 given the low flaring rate of Sgr A*, is the number of flares in each given frame time. For each flare, I proceed as explained in Section 5.3, and retrieve its properties with the Bayesian Blocks.

I do this for a given dataset and reconstruct the flare fluence distribution to compare

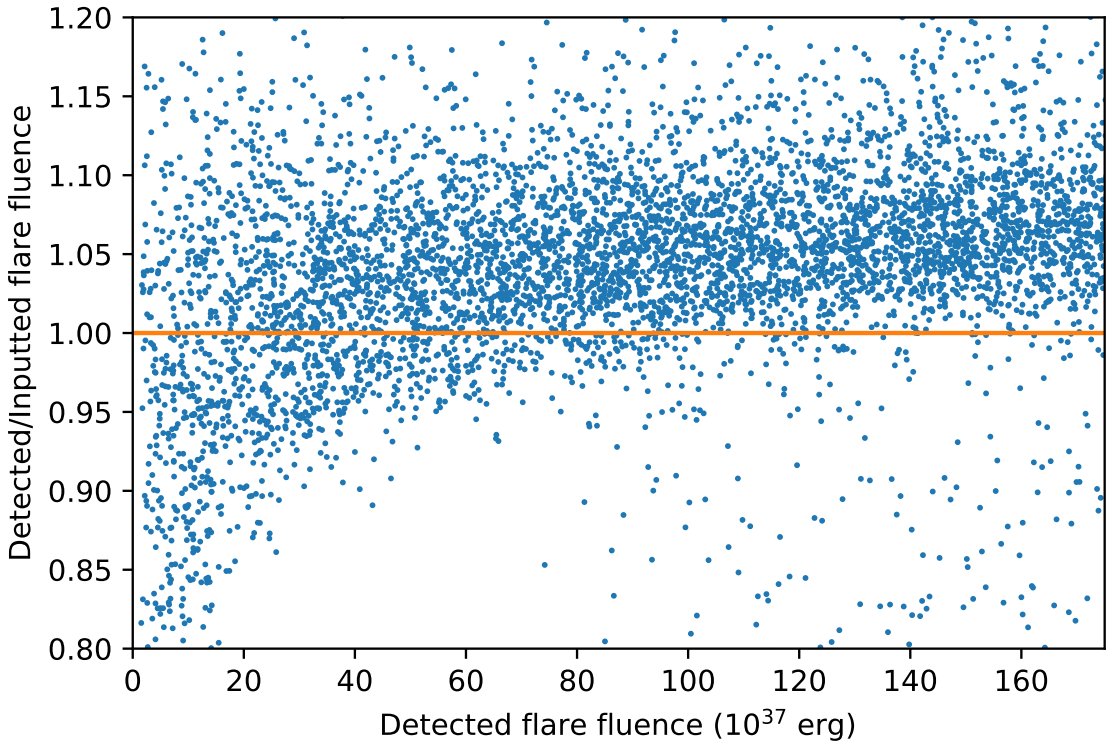

Figure 5.5: Detected/input flare fluence plotted against detected fluence. From a sample of 10 000 points, where each represents a subarray simulation in which a flare was simulated and then detected with my algorithm. Flare fluences were drawn from a uniform distribution between 1.3 and 275 $\times 10^{37}$ ergs. Each fluence was assigned a random duration drawn from a uniform distribution between 500 and 8000 s. The orange line represents a perfect response would be.

it with observations. I split the simulation data and the real data in the same fluence bins to compare them. These are used to produce a binned power-law distribution, where the y-value is equal to the number of flares within that bin divided by the midpoint of the bin. I simulate each given dataset 3000 times[3] to produce 95% (2.5% to 97.5%) confidence shaded regions for each bin. To do so, I save the number of detected flares within each fluence bin for each simulated dataset. After a complete run of 3000 simulations, I get a distribution of the number of flares in each bin and I use it to create 95% confidence

[3]I observe robust bin statistics after $\sim$ 1000 simulations in all of my datasets. I perform 3000 simulations to obtain more accurate error intervals.

intervals. This allows comparison between simulations and real data. The error bar on each real data bin is assumed to be $1 + \sqrt{0.75 + N}$ divided by the midpoint of that bin, where N is the number of flares in that bin (Gehrels, 1986). By using the same model parameters for different datasets, I can test the null hypothesis by determining if the flare fluence distributions from the datasets are consistent with the confidence intervals from the model for each bin. If this is not the case, it would indicate a significant change in the fluence distribution (which would imply a change in the flaring rate of bright or faint flares, for example). Figures 5.6 and 5.7 compare the XVP and Post-XVP datasets as well as the Pre and Post-G2 datasets for different change points (after April 4th 2014 and August 30th 2014), respectively. The confidence intervals were obtained from the same model parameters for each plot.

5.6.1 Simulations of the XVP and Post-XVP observations

Since the XVP and Post-XVP datasets were obtained in different instrument modes with different sensitivities, one might expect to detect different flare fluence distributions. I can test if such a change is physical or due to detection biases since my model takes these factors into consideration automatically. In Figure 5.6, I show the fluence distributions from my XVP (top) and Post-XVP (bottom) datasets over-plotted with 95% confidence regions produced by 3000 Monte Carlo simulations of each datasets. I use the following model parameters to produce confidence regions of the flare fluence distributions of both datasets to test the null hypothesis: $\Gamma_{\text{Dura}} = -0.8$, $\Gamma_{\text{Fluence}} = -1.7$, $D_{\text{min}} = 500$ s, $D_{\text{max}} = 8000$ s, $F_{\text{min}} = 1.3 \times 10^{37}$ erg, $F_{\text{max}} = 275 \times 10^{37}$ erg and a flaring rate of 52 flares per 3 Ms ($R \sim 1.5$ per day). The duration and fluence power-law indexes are consistent with those reported by Neilsen et al. (2013), but my flaring rate is higher since it takes into account detection biases. However, I do not claim that these parameters are the real physical parameters of Sgr A*; they are model parameters and my model isn't a perfect

representation of Sgr A*'s behaviour due to some potential shortcomings explained in Chapter 6.

If a flare duration or fluence is detected outside the simulated range in a given simulation, it is appended to the closest bin and the midpoint of that bin is adjusted accordingly. For XVP simulations, I find no flare fainter than 1.3×10^{37} erg because the detection efficiency rapidly drops below $\sim 2 \times 10^{37}$ erg, but I find at least 1 flare brighter than 275×10^{37} erg in 8.8% of simulations. For Post-XVP simulations, we find at least 1 flare fainter (brighter) than 1.3×10^{37} erg (275×10^{37} erg) in 47% (2.3%) of simulations. This instrument mode has a much better sensitivity, such that its detection efficiency only starts to drop for fluences below $\sim 1.5 \times 10^{37}$ erg. Flares around 1.3×10^{37} erg are much harder to detect in the gratings instrument mode because they translate to even lower count rates and flares of this fluence are already difficult to detect. These flares are more easily detected in subarray simulations, which explains the much higher percentage. Furthermore, as explained in Section 5.4, faint flares tend to be detected for shorter periods, translating to underestimated fluence (some faint flares get pushed to fluences lower than 1.3×10^{37} erg). On the other hand, the higher percentage of simulations containing flares brighter than 275×10^{37} erg in the XVP simulations comes from the fact that the simulated time is higher (3 Ms vs 1.56 Ms) and that each XVP observation is on average longer than each Post-XVP observation (average durations of 78 ks and 40 ks, respectively). This makes flares less likely to happen on the edge of a given observation (and thus it is more likely that bright flares will be detected for longer).

Since I find 40 XVP flares and 18 Post-XVP flares, I use 7 and 5 bins, respectively. Having too many bins risks introducing holes in the distribution while too few do not characterize it robustly. As can be seen from Figure 5.6, the noise is important enough to make each confidence interval produced by simulations using the same model parameters consistent with its respective bin for every bin of the two observed flare fluence distributions, confirming the null hypothesis. Many complex factors contribute to this noise,

namely the random draw from a CDF of a low number of flares, edge effects, overlapping flares, the inherent Poisson counting noise, and the Bayesian Blocks detection process. The fact that scatter is that important for the same model parameters makes a Markov-chain Monte Carlo approach more complicated and unnecessary. I tried this approach but since the same model parameters can produce very different results (and similar results can be obtained from different parameters), the walkers get confused. I thus found the model parameters by trial and error.

If I use an additional criterion to separate flares showing a dip in their center (explained in the last paragraph of Section 4.1.3), I change the maximum fluence to 150×10^{37} erg since the flare in ObsID 15043 becomes two flares. Leaving the other model parameters at their original values, I also find that the two datasets are consistent.

5.7 Looking for change points around G2's pericentre passage

Ponti et al. (2015) and Mossoux & Grosso (2017) report an increase in the bright flaring rate on August 31, 2014 which they argue is caused by the pericentre passage of G2 around Sgr A*. Ponti et al. (2015) also argue that this change could be caused by a noise process, where the flaring rate is constant on average but shows clustering on shorter timescales which was only detected because of the increased monitoring frequency around G2's pericentre passage. These authors also mention that they do not observe this change if they limit themselves to *Chandra* data only. Indeed, most (4 out of 5) of the bright/very bright flares responsible for this shift are from *XMM-Newton*. Mossoux & Grosso (2017)'s results also rely on 2 more recent *Chandra* flares from ObsID 16966 and 17857 (this observation isn't in the present work since it is a gratings observation instead of ACIS-S3 subarray

observation, and I want to keep my datasets as consistent as possible) and a $Swift$ flare from February 2015.

I can test if the inclusion of new $Chandra$ data and my simulations are consistent with such a change point. I tested 4 different points in time; ObsIDs 16212 (April 4th, 2014), 16214 (May 20th, 2014), 16215 (July 16th, 2014) and ObsID 16217 (August 30th 2014). The motivation for these times comes from the different G2 pericentre times from different models. A purely keplerian orbit fit gives an estimated time between the end of February to mid April 2014. The addition of a drag force pushes this date to between the end of May and mid-July, and if an inflow is also added, this prediction shifts to between July and September (Madigan et al., 2016).

I find that the same model parameters presented above produce 95% confidence intervals consistent with all the flare fluence distributions from these different datasets, once again confirming the null hypothesis. Since the only ObsID in which I find flares is 16217, the only variation in the other tested ObsIDs is how the exposure is split between Pre and Post-G2. This difference doesn't change the 95% confidence regions of my simulations until ObsID 16217. Therefore, I only show my results for ObsIDs 16212 and 16217 in Figure 5.7. I observe that despite the removal of the flare in ObsID 16217 from the Post-G2 dataset, the reduced exposure increases the scatter enough to make the entire dataset consistent with the Pre-G2 observations. I also test the change point reported by Mossoux & Grosso (2017) where they find a decrease in the flaring rate of the less energetic flares on 2013 July 27 (ObsID 15041). I do not recover their findings.

As in the previous section, I tested the impact of the additional flare separation criteria and found that the datasets remain consistent.

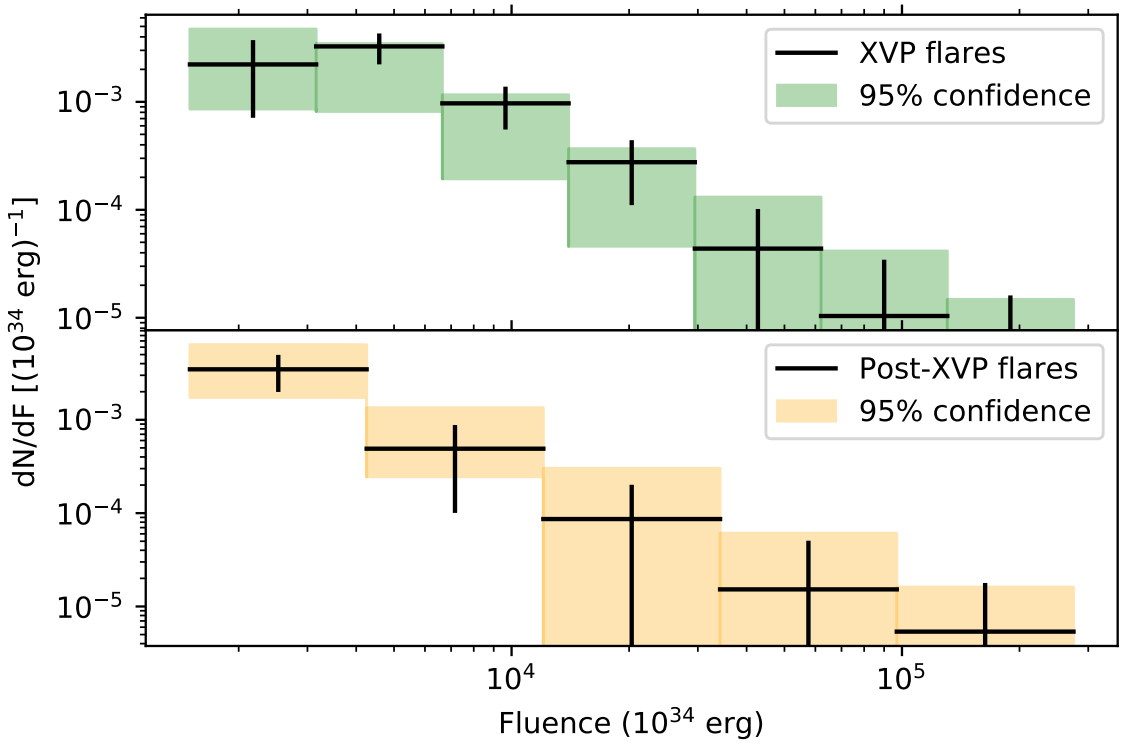

Figure 5.6: Binned unabsorbed fluence differential distribution of XVP (top) and Post-XVP (bottom) flares (see Tables 4.1 and 4.2) in black, with Poisson errors of $1 + \sqrt{0.75 + N}$, where N is the number of flares in that bin. The 95% confidence region obtained from 3000 Monte Carlo simulation is also shown in green (orange) for the (Post-)XVP dataset.

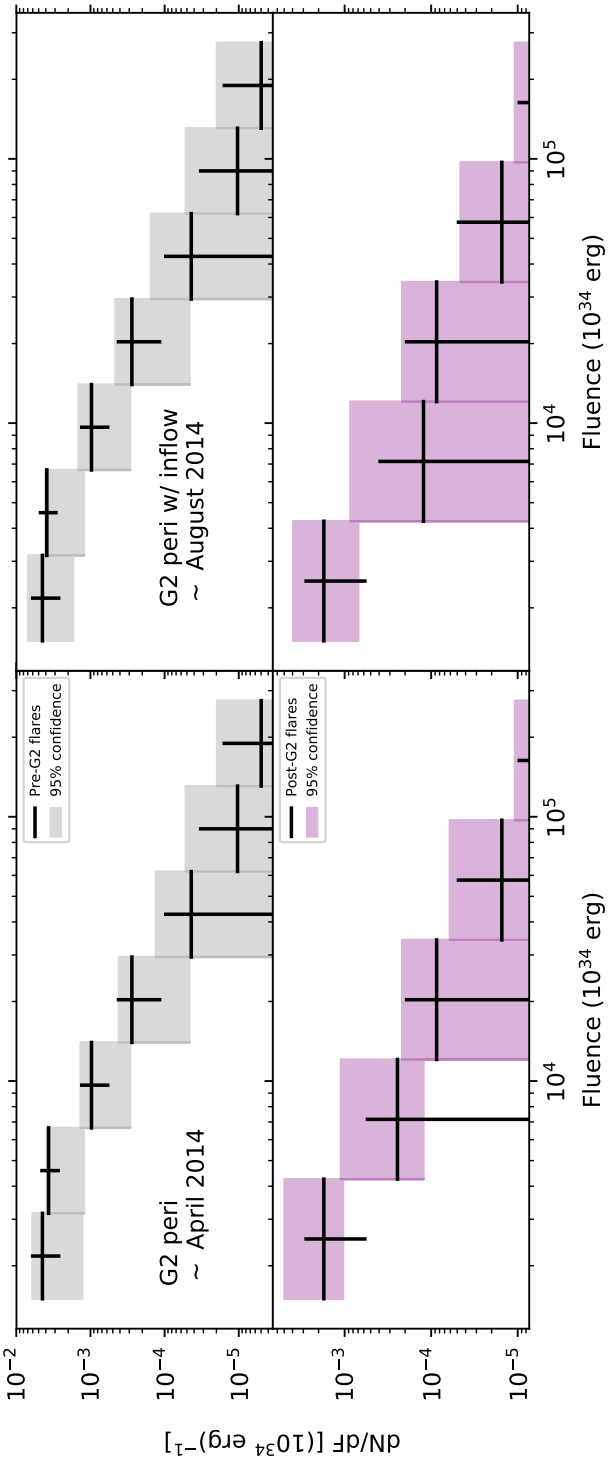

Figure 5.7: Similar to Figure 5.6, but for different change points from the pericentre of G2. The 95% confidence shaded region for the (Post) Pre-G2 dataset is in grey (purple). Left: Test for a change point after ObsID 16212 (April 4th, 2014). Right: Test for a change point after ObsID 16217 (August 30th 2014).

Chapter 6

Discussion and conclusion

Based on the simulations and analysis presented here, I do not report any significant change in Sgr A*'s bright or faint flaring rate at any point in my datasets. This is in contrast with the findings of Ponti et al. (2015) and Mossoux & Grosso (2017) who find that such an increase occurs from 2014 August 31 (after ObsID 16217) for the bright flares (the latter also report a decrease of the faint flaring rate on 2013 July 27 (ObsID 15041), which I do not recover). However, these authors also use data from *XMM-Newton* and $Swift$ which cannot detect weak and moderate flares. Most of the brightest flares responsible for the change point are detected by *XMM-Newton*. On the other hand, $Chandra$ allows proper characterization of all flares, which motivates my choice to analyze data from this single observatory in the present work.

Ponti et al. (2015) use PIMMS to convert counts to energies, which systematically underestimates the energies of each flare since it doesn't correct the PSF extraction fraction. These authors also don't properly calibrate their prior for the Bayesian Blocks which could lead to unreliable uncertainties. For example, in their Section 5.5, they simulate event lists containing a fixed number of flare times (each representing an event) drawn from a uniform distribution and run Bayesian Blocks on them. With a p_0 of 0.3%, they find change points in 0.1% of their lists. Their work also doesn't consider edge effects nor the relative flare detection efficiency of different instruments/observing modes.

Mossoux & Grosso (2017) directly fit their spectra in each observation to get a flux for each flare. However, as shown in Neilsen et al. (2013) and Yuan et al. (2017), Sgr A*'s flares have very similar spectra. Fitting them directly can lead to strong deviations due to poor statistics (the spectrum of some flares can have very few counts per energy bin, leading to unreliable fits). For example, flares #38 and #39 (within the same observation) of Table A.2 of Mossoux & Grosso (2017), have count rates are 0.021 ± 0.004 ct/s and 0.041 ± 0.005 ct/s respectively but their mean 2-10 keV unabsorbed fluxes are 9.20×10^{-12} erg s^{-1} cm^{-2} and 2.08×10^{-12} erg s^{-1} cm^{-2}. This counter-intuitive behaviour[1] happens with many other flares reported by these authors and leads to a strong difference in their ratio of count rates to unabsorbed fluxes between the similar (in terms of effective area) ACIS-S3 subarray and ACIS-I instrument modes (248.2 and 148.1×10^{-12} erg cm^{-2}/ct, respectively). They also neglect to report this ratio for the gratings instrument mode.

Mossoux & Grosso (2017) improve upon Ponti et al. (2015) regarding edge effects. However, in their Section 6, where they look for change points, they perform the search without gaps between observations which doesn't take edge flares into account. They look for a signal by performing a top-to-bottom (bottom-to-top) search, where they run Bayesian Blocks on an event list where each event is a flare. They remove the brightest (faintest) flare and re-run Bayesian Blocks. They repeat this process until a change point is found. However, they do not consider the number of trials in the computation of the significance of the change points found. It is unsurprising that they find similar change points for the bottom-to-top and top-to-bottom searches as the signal in one should be the inverse of the signal in the other. Finally, the fact that these points occur at the edge of their light curve is also not surprising because it requires only one change point (a change in the middle would require two) and such spurious fluctuations are characteristic of the Bayesian Blocks.

[1] I admit that it is possible that this strange behaviour could be explained by large error bars, but they were not provided, nor even mentioned, by these authors.

The present work shares some flaws with these previous studies. Namely, I assume that the flares occur at random Poisson times. However, it's possible that Sgr A*'s flaring behaviour possesses clustering as mentioned in Yuan & Wang (2015) where the authors show that flare clustering on timescales of 20-70 ks is significant at the 96% level (for their gratings dataset, but for their ACIS-I sample this number drops to $\sim 50\%$) and may be described by a piecewise-deterministic Markov process (Davis, 1984). Furthermore, I assume that flare durations and fluences are independent. However, Neilsen et al. (2013) show a moderate correlation ($\rho \sim 0.54$) between the two. My simulations *a-posteriori* recover some of this correlation ($\rho \sim 0.23$) from detection biases. I consider this sufficient as Yuan et al. (2017) report no correlation between those quantities.

Despite these potential shortcomings, the change points reported by Ponti et al. (2015) and Mossoux & Grosso (2017) are unlikely to be indicative of an increase in Sgr A*'s bright flaring activity caused by G2. If the increased bright flaring rate observed by these authors was due to G2, this process should stay active for at least a viscous timescale ($\sim$ 3-10 years at ~ 2000 R_s (Yuan & Narayan, 2014; Ponti et al., 2015; Mossoux et al., 2016) assuming the canonical viscous parameter value $\alpha = 0.1$ (with α as defined as in Shakura & Sunyaev, 1973)). In this scenario, my model should have detected this persistent signal with the additional $Chandra$ exposure of 652 ks across 14 observations from June 2016 to April 2018 (from ObsID 18731 to 20347). In addition to the change points reported in the previous Chapter, I can make an independent and quick sanity check on only these observations. This timescale of a few years after pericentre also probes the estimated timescale before an increased activity begins according to RIAF models (Yuan & Narayan, 2014). I observe 3 bright flares between June 2016 and April 2018 (according to Ponti et al. (2015)'s definition of > 120 counts, corresponding to an unabsorbed fluence of 9.2 $\times 10^{37}$ erg in my case), corresponding to a bright flaring rate of 0.4 $\pm$ 0.2 flare/day. For the whole Post-XVP dataset, I detect 6 bright flares in 1.56 Ms giving a bright rate of 0.3 $\pm$ 0.1 flare/day. For the XVP dataset, I have 10 bright flares in 3 Ms translating to a bright flaring rate of 0.29

$\pm$ 0.09 flare/day. All these rates are consistent within errors. The 3σ upper limit of 3 is 12.7 (Gehrels, 1986) giving a rate of 1.7 flare/day, still lower than the predicted rate of 2.52 ± 0.98 flare/day as found by Ponti et al. (2015). As was determined by my model, this increase can thus be excluded.

Since I do not find an increase, I argue that the signal found by these authors with the inclusion of *XMM-Newton* and $Swift$ data is either an analysis artifact or due to flare clustering as suggested by Ponti et al. (2015). This behaviour would thus have been observed due to the increased monitoring frequency of Sgr A* and not by G2's pericentre passage.

It is also worth noting that another gas cloud, G1, which has nearly identical orbital parameters to G2 (Pfuhl et al., 2015), underwent a pericentre passage near Sgr A* in 2001 and no change in Sgr A*'s quiescence or flaring properties were observed (Yuan & Wang, 2015).

My results are in agreement with simulations from Schartmann et al. (2012) and Kawashima et al. (2017) who predict an increase in activity only a few years and 5-10 years after pericentre, respectively. Whereas the former explain this delay by the strong angular momentum of G2 (which delays accretion), the latter use more careful full 3D general relativistic magnetohydrodynamic simulations to model the evolution of magnetic fields as they interact with the cloud. An instability develops, and a magnetic reconnection event is expected to increase the radio and X-ray luminosity on a dynamo-viscous timescale of 5-10 years. Future observations will determine if this prediction is supported.

In this MSc thesis, I have shown empirically that $(1.3 \pm 0.2)\%$ of the count rate from the transient magnetar SGR J1745-2900 contaminates Sgr A*'s extraction region and quantified it as a function of time since its outburst on April 25 2013. Using Bayesian Blocks calibrated on Poisson noise with $p_0 = 5\%$, and adding flaring blocks together with a systematic method I detected and characterized 58 flares (40 XVP flares and 18 Post-XVP flares, respectively) in 4.5 Ms. I simulated X-ray light curves built from parameters such

as flaring rate, quiescence count rate, flare count rate to luminosity conversion factor, fluence and duration distributions, exposure time and pile-up. I tested the null hypothesis by using those simulated light curves to produce confidence intervals using the same model parameters and comparing them to their respective observed flare fluence bins for different datasets (XVP and Post-XVP or Pre-G2 and Post-G2). My results show that the null hypothesis is confirmed for every tested potential change point; I found no evidence of a change point in the fluence distributions above 95% confidence at any point in time. I conclude that my findings, combined with those of Ponti et al. (2015), Yuan & Wang (2015), Yuan et al. (2017) and Mossoux & Grosso (2017) imply that the increased bright flaring rate observed after the pericentre passage of G2 in previous works was due to an underlying non-stationary noise process, which was only detected due to the increased monitoring frequency.

Recent work on simultaneous IR and X-ray observations (Boyce et al., 2019) has put new constraints on the physical mechanisms for Sgr A*'s variability. Within >100 hours of observation, 4 X-ray flares $\sim$ 4 times quiescence were detected, each having an associated IR flare (though additional IR peaks did not coincide with a detectable X-ray flare). Through a cross-correlation analysis, Boyce et al. (2019) find that X-ray flares lead IR by 10-20 minutes but also point out that their 99.7% confidence region is consistent with no time lag at all. This rules out models in which X-rays lag the IR as would be predicted by inverse-Compton processes but is consistent with models predicting the opposite trend in which X-rays and IR come from synchrotron emission generated following magnetic reconnection by two electron populations, an IR generating one and a relativistic, X-ray generating one Dodds-Eden et al. (2010); Yusef-Zadeh et al. (2012). In their model, the synchrotron cooling time of the IR generating electrons is longer than the injection rate of relativistic X-ray emitting electrons, giving rise to a lag of the IR flares with respect to the X-ray flares. Models predicting simultaneous IR and X-ray flares through synchrotron self-Compton are also consistent with the findings. However, the work of Boyce et al.

(2019) cannot pick a singular model that matches observations. Luckily, there will be joint $Spitzer/Chandra$/GRAVITY observations in Summer 2019 which could help break this degeneracy. These exceptional multiwavelength observations will push our understanding of Sgr A*'s accretion physics further than ever before.

Finally, data from the Event Horizon Telescope Campaigns capable of resolving Sgr A* and M87's black holes (e.g., Akiyama et al., 2019) will put strong constraints on the underlying physical processes behind flares and the accretion flow, and inform better general relativistic magnetohydrodynamic simulations. The next few years will certainly be extremely informative regarding black hole accretion physics.

Appendix A

Prior calibration

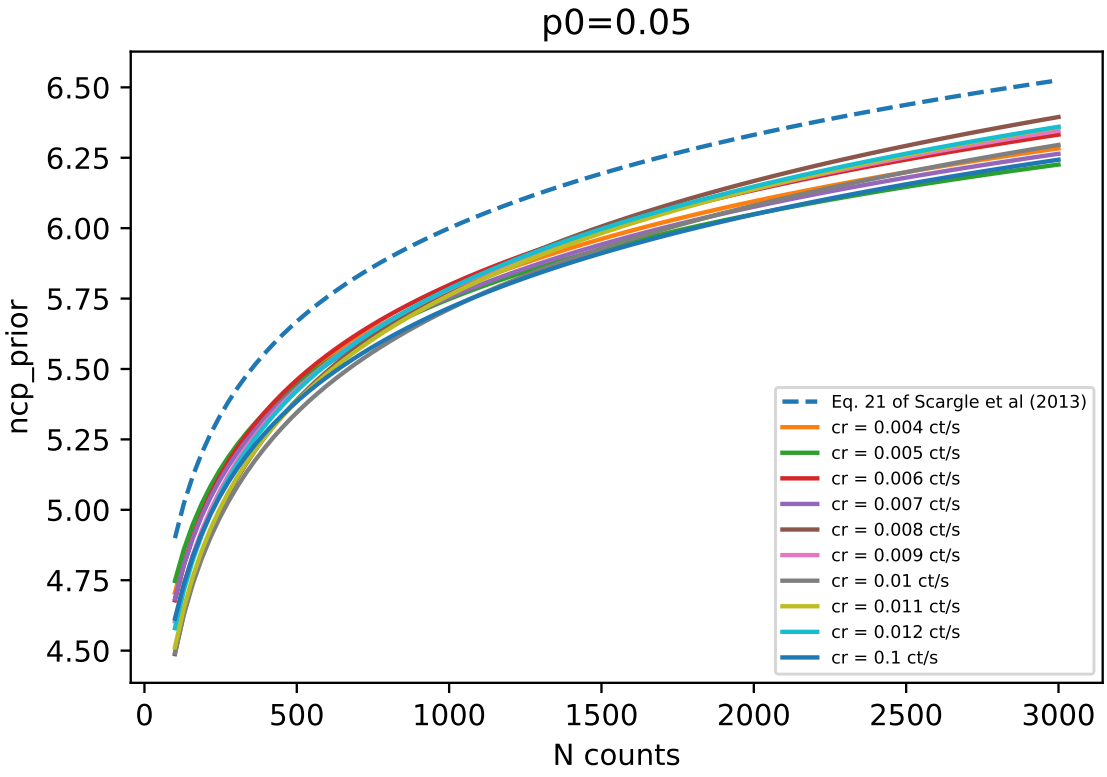

Figure A.1: Behavior of ncp_prior calibrated for $p_0 = 0.05$ for different count rates against different values of expected number of events, N. As I suspected, the behavior is not strongly correlated with the count rate. Instead, it is dominated by random fluctuations. Hence, I choose to take the fit of the mean of these curves, and then use the 3σ overall fit as my calibration for the Bayesian Blocks algorithm.

Appendix B

Comparing with the literature

B.1 Comparison with other works

In this section, I compare my flares with Neilsen et al. (2013), Ponti et al. (2015) and Mossoux & Grosso (2017). To compare fluences with Ponti et al. (2015), I convert from absorbed to unabsorbed values using the scaling between the unabsorbed fluences and absorbed fluxes of Table 1 of Neilsen et al. (2013) assuming a distance of 8 kpc. This should be sufficient for comparison purposes. For the comparisons in this section, I use my flares as obtained without the flare splitting criteria (see Section 4.1.3). Table B.1 shows every flare detected or missed by every work considered and Figures B.2 and B.3 compare the different flares' durations and fluences of each author. From Figure B.3, my work (green '+') shows remarkable agreement with Neilsen et al. (2013) (red triangles) for the XVP dataset. This is expected since I also scale my flares' fluxes to the brightest flare of 2012 from Nowak et al. (2012). Mossoux & Grosso (2017) (blue 'x') instead perform a direct fit, leading to extra scatter and Ponti et al. (2015) (black dots) use PIMMS, disregarding the response matrices and leading to systematic flare fluxes/fluences underestimation.

B.1.1 XVP flares

ObsID	Flare index	This work	N13	P15	M17
14392	1	✓	✓		✓
	2	✓	✓	✓	✓
13857	3				✓
13854	4	✓	✓	✓	✓
	5	✓	✓	✓	✓
	6	✓	✓	✓	✓
	7	✓	✓	✓	✓
13847	8	✓	✓	✓	✓
	9		✓		
14427	10	✓	✓	✓	✓
	11	✓		✓	✓
13849	12	✓	✓	✓	✓
	13	✓	✓	✓	✓
	14	✓	✓		✓
	15	✓	✓	✓	✓
13845	16	✓			✓
	17	✓	✓	✓	✓
13853	18		✓		
14465	19	✓	✓	✓	✓
	20	✓	✓	✓	✓
14466	21	✓	✓	✓	✓
	22				✓
13842	23	✓	✓	✓	✓
	24	✓	✓	✓	✓
	25	✓	✓	✓	✓
13839	26	✓	✓	✓	✓
	27		✓		✓
	28	✓	✓	✓	✓
13840	29	✓		✓	✓
	30	✓			✓
14432	31	✓		✓	✓
	32	✓	✓	✓	✓
13838	33	✓	✓	✓	✓
13852	34	✓	✓	✓	✓
	35			✓	
	36	✓	✓	✓	✓
14439	37	✓	✓	✓	✓
14462	38	✓	✓	✓	✓
	39	✓	✓	✓	✓
14463	40	✓	✓	✓	✓
13851	41	✓	✓		✓
	42	✓	✓	✓	✓
15568	43			✓	✓
13843	44	✓	✓	✓	✓
15570	45	✓	✓	✓	✓
14468	46	✓	✓	✓	✓
	47		✓		
	48	✓	✓	✓	✓
15041	49	✓	–	✓	✓
	50	✓	–		✓
15042	51	✓	–	✓[a]	✓
	52	✓	–	[a]	✓
14945	53		–	✓	
15043	54	✓	–	✓	✓
14943	55		–	✓	
15045	56	✓	–	✓	✓
	57	✓	–	✓	✓
16508	58	✓	–	✓	
16217	59	✓	–	✓	
16218	60	✓[a]	–	✓	✓[a]
	61	[a]	–	✓	[a]
16963	62	✓	–	–	✓
16966	63	✓	–	–	✓

Table B.1: List of all the flares reported in this work, Neilsen et al. (2013) (N13), Ponti et al. (2015) (P15) and Mossoux & Grosso (2017) (M17) within the observations considered in this analysis. I show my results without the flare splitting criteria. I stop at ObsID 16966 since more recent ObsIDs were not available at the time of publishing the other papers.
[a] These authors report these two flares as a single one.

Neilsen et al. (2013) report 39 XVP flares and I find 40. 35 flares are in common. I find 5 flares that they missed in ObsIDs 14427, 13845, 13840 (2 flares in this ObsID) and 14432 (flares # 11, 16, 29, 30 and 31, see Table B.1 and the associated Figures B.2 and B.3). Those flares are rather long (durations above 1600 s) and have low count rates (0.019 ct/s and less). They are thus more likely to be missed by direct Gaussian fitting on the 300 s binned light curves whereas Bayesian Blocks can detect them more easily. However, they report 4 flares that I do not find in ObsIDs 13847, 13853, 13839 and 14468 (flares # 9, 18, 27 and 47). These flares are short (durations from 500 s to 1200 s) and have low fluences (between 8 and 15 counts, or 1.15 and 2 $\times 10^{37}$ ergs). My results agree for the most significant flares, but some differences exist for the weaker flares, which can be explained by the different detection methods used.

Ponti et al. (2015) use Bayesian Blocks like me, but they do not properly calibrate their prior. Despite that, I obtain similar results. For the XVP dataset, they find 37 flares, of which 35 are in common. I find 4 flares that they miss in ObsIDs 14392 (# 1), 13849 (# 14), 13845 (# 16) and 13840 (# 30). All of these flares are faint (unabsorbed fluences below 3.7×10^{37} erg) and are most likely missed because they limit themselves to the 0th order and they don't calibrate ncp_prior. They also report 2 flares that I do not detect in ObsIDs 13852 (# 35) and 15568 (# 43). The flare in ObsID 13852 is not found by Mossoux et al. (2016) nor Neilsen et al. (2013). It is very long (13 ks) and has a very low count rate (0.0049 ct/s) and is therefore consistent with quiescence. The flare at ObsID 15568 is also found by Mossoux et al. (2016) but not by me nor Neilsen et al. (2013). Even though I indeed find a block there, my criterion doesn't consider it significantly above quiescence. Since Mossoux & Grosso (2017) and Ponti et al. (2015) do not explicitly state what they consider

to be a flaring block, it is hard to compare. Either way, this flare is rather faint ($\sim$ 2x quiescence) and its detection is thus sensitive to what different authors consider to be significant.

Mossoux & Grosso (2017) find 44 XVP flares, of which I have 40 in common. The 4 flares that I do not find are in ObsIDs 13857, 14466, 13839, and 15568 (flares # 3, 22, 27 and 43). Even though I indeed find blocks in ObsIDs 13857, 14466, and 15568, they are not considered flares because of my significance criterion. The reported flares at ObsIDs 13857 and 14466 are short (durations of 471 s and 380 s) and the one at ObsID 15568 lasts >4907 s but is very faint (0.006 ct/s above quiescence). This translates to rather high error bars, which explains why my method doesn't consider them flares. The flare at ObsID 13839 is undetected by my Bayesian Blocks and also by Ponti et al. (2015), but it is on the short end (750 s according to Neilsen et al. (2013) and 1411 s according to Mossoux & Grosso (2017)). This difference can be explained by the way I calibrated ncp_prior by taking the 3σ upper fit, which is motivated by the noise of the process.

Overall, my results are consistent in both duration and fluence with these authors and only differ for faint/short flares.

B.1.2 Post-XVP flares

Ponti et al. (2015) report 10 flares for the Post-XVP observations that I share. There are substantial differences with what I find. For ObsID 15041, Mossoux & Grosso (2017) and I find two flares (flares # 49 and 50), but Ponti et al. (2015) only find the first one. That second flare is faint (unabsorbed fluence of 1.55×10^{37} erg). Likewise, I find two flares in ObsID 15042 (flares # 51 and 52), but Ponti et al. (2015) report this flare as one and Mossoux & Grosso (2017) do not find any. I argue that there are two distinct flares since we can see in Figure B.1 that the two flares are separated by a non-flaring block. Ponti et al. (2015) also report 2 faint flares (unabsorbed fluence of 0.6 and 0.7 $\times 10^{37}$ erg) in ObsIDs 14945

and 14943 (flares # 53 and 55), but Mossoux & Grosso (2017) and I do not find them. As explained in Mossoux & Grosso (2017), this is likely the result of the prior calibration. Ponti et al. (2015) also report a flare in ObsID 16508 (# 58) and another one in ObsID 16217 (# 59) that Mossoux & Grosso (2017) do not find, but I do. Briefly, my results are a combination of these authors and the differences occur based on the criterion that selects what is considered a flaring block and how the Bayesian Blocks are calibrated.

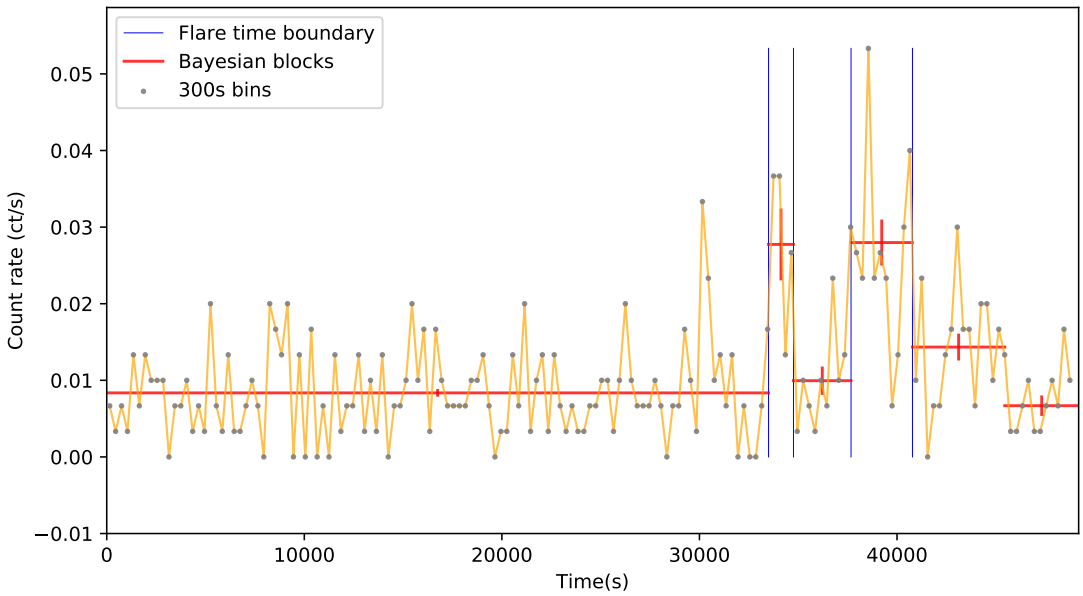

Figure B.1: Light curve of ObsID 15042. Notice the block separating the two flares.

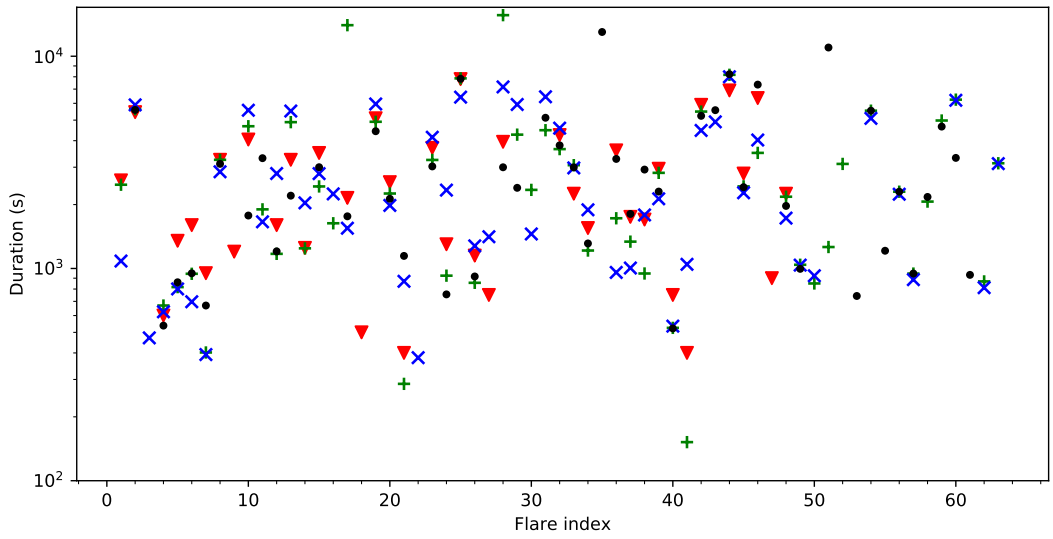

Figure B.2: Durations of the flares detected in this work (green '+'), Neilsen et al. (2013) (red triangles), Mossoux & Grosso (2017) (blue 'x') and Ponti et al. (2015) (black dots). See the associated Table B.1.

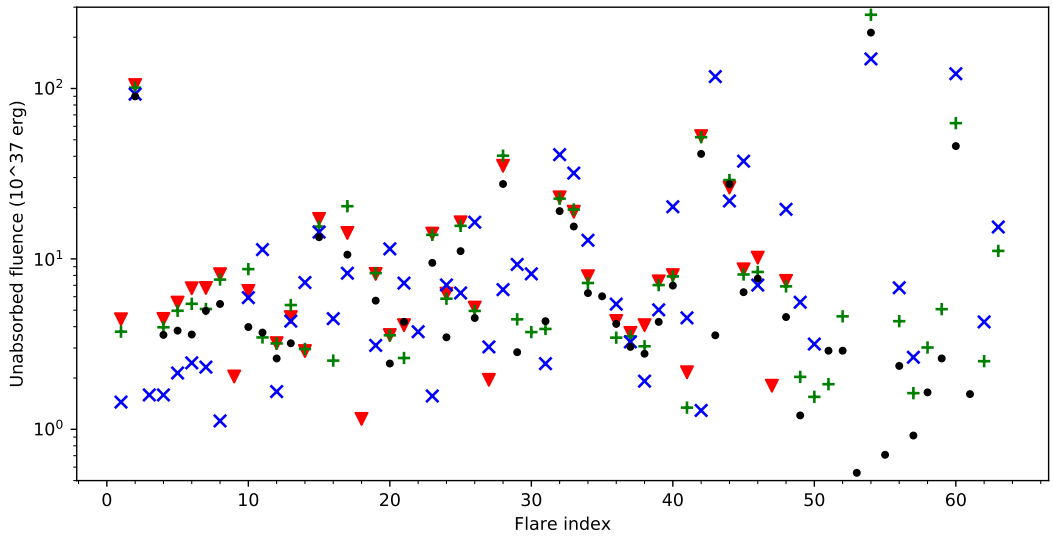

Figure B.3: Like Figure B.2 but for unabsorbed fluences. For Ponti et al. (2015), I converted from absorbed to unabsorbed values using the scaling between the unabsorbed fluences and absorbed fluxes of Table 1 of Neilsen et al. (2013) assuming a distance of 8 kpc. This should be sufficient for comparison purposes.

Appendix C

Pile-up for gratings data and simulations

For HETG data, obtaining the unpiled, quiescence subtracted count rate of each flare to compute the fluence is more complex than simply using Equation 2.2. After running Bayesian Blocks on the gratings event list, I save the number of 0th and 1st order counts in the quiescence block and compute the 0th/1st order count ratio. I save the number of 0th and 1st order counts of each flaring block and unpile the 0th order with Equation 2.2. Using the 0th and 1st order quiescence count rates of that observation, I perform quiescence subtraction to each unpiled flaring block and save the flare 0th/1st unpiled and quiescence subtracted count rate ratio.

I obtain a median value across all observations of quiescence 0th/1st order count rate of 0.45 and a median value of unpiled and quiescence subtracted 0th/1st order flare count rate of 1.6.

For simulations, I assume that these ratios are fixed to those median values for each simulated observation. To pile flare mean count rates, I thus find the total 0th order count rate using these ratios (0th order count rate from the mean flare count rate + 0th order quiescence count rate) and pile it using the pile-up equation. The simulated flare piled mean count rate is then that piled count rate minus the quiescence count rate.

The unpiling process in simulations requires more attention since the (piled) 0th/1st order flare ratio changes due to pile-up in the 0th order. This phenomenon can be taken into account via a calibration that computes the piled 0th/1st order flare ratio for different

unpiled count rates. I generate incoming count rates, apply pile-up in their 0th order assuming a quiescence count rate corresponding to the median of the ones observed (0.0063 ct/s) and using my 0th/1st order ratios found in the real data. I compute the new 0th/1st order ratio in the flare count rate after pile-up and save that value. When the Bayesian Blocks detect a flaring block mean count rate, my algorithm looks for the corresponding piled 0th/1st order flare ratio which itself is linked to the unpiled count rate.

Appendix D

Finding the CDF of a power-law distribution

Suppose we have a power-law distribution of some quantity L, with index Γ. If this quantity is defined between L_{min} and L_{max}, then we can easily normalize it to get its probability density function (PDF):

$$C = \int_{L_{min}}^{L_{max}} L^{-\Gamma} dL = \frac{L_{min}^{-\Gamma+1} - L_{max}^{-\Gamma+1}}{\Gamma - 1} \tag{D.1}$$

The PDF is then obtained directly:

$$PDF(L) = \frac{L^{-\Gamma}}{C} \tag{D.2}$$

Since the CDF represents the probability of L being less than a given value, it can be computed from the PDF by integrating it up to said value, i.e.:

$$CDF(L) = \int_{L_{min}}^{L} \frac{L^{-\Gamma}}{C} dL \tag{D.3}$$

$$CDF(L) = \frac{1}{C} \frac{L^{-\Gamma+1} - L_{min}^{-\Gamma+1}}{-\Gamma + 1} \tag{D.4}$$

Appendix E

Time distribution of Poisson events

Let's assume that we are given a Poisson distribution with mean rate λ and that we know that there is exactly one event before time T. Knowing that a Poisson process has no memory (i.e., each draw is independent), we could be inclined to think that the probability of the event happening within any interval within $[0, T]$ of the same duration should be the same. Indeed, it can be shown that the probability of the occurrence time T_1 being within $[0, t]$ with $t \leq T$ is

$$\begin{aligned} \Pr\left(T_1 \leq t | N(T) = 1\right) &= \frac{\Pr\left(1 \text{ event in } (0, t]\right) \Pr\left(0 \text{ events in } (t, T]\right)}{\Pr\left(N(T) = 1\right)} \\ &= \frac{\lambda t \exp\left[-\lambda t\right] \exp\left[-\lambda\left(T - t\right)\right]}{\lambda T \exp\left[-\lambda T\right]} \\ &= \frac{t}{T} \end{aligned} \tag{E.1}$$

This shows that the time of the event is characterized by an uniform distribution over the interval $[0, T]$. This result can be generalized to N events[1] and is the basis behind the method used in Section 5.1 to create event lists containing a known number of Poisson events.

[1] See Section 1.5 of Sigman (2007)

References

Abarca, D., Sądowski, A., & Sironi, L. 2014, Monthly Notices of the Royal Astronomical Society, 440, 1125

Abbott, B., Abbott, R., Abbott, T., et al. 2017a, Physical Review Letters, 118, 221101

Abbott, B. P., Abbott, R., Abbott, T., et al. 2016a, Physical Review X, 6, 041015

—. 2016b, Physical review letters, 116, 241103

—. 2017b, The Astrophysical Journal Letters, 851, L35

—. 2017c, Physical review letters, 119, 141101

—. 2017d, Physical Review Letters, 119, 161101

—. 2017e, Astrophys. J. Lett, 848, L12

Abramowicz, M. A., & Fragile, P. C. 2013, Living Reviews in Relativity, 16, 1

Abuter, R., Amorim, A., Bauböck, M., et al. 2018a, Astronomy & Astrophysics, 618, L10

Abuter, R., Amorim, A., Anugu, N., et al. 2018b, Astronomy & Astrophysics, 615, L15

Akiyama, K., Alberdi, A., Alef, W., et al. 2019, The Astrophysical Journal, 875, L1

Arnaud, K. 1996, in Astronomical Data Analysis Software and Systems V, Vol. 101, 17

Baganoff, F., Bautz, M., Brandt, W. N., et al. 2001, Nature, 413, 45

Baganoff, F. K., Maeda, Y., Morris, M., et al. 2003, The Astrophysical Journal, 591, 891

Ball, D., Özel, F., Psaltis, D., & Chan, C.-k. 2016, The Astrophysical Journal, 826, 77

Ball, D., Özel, F., Psaltis, D., Chan, C.-K., & Sironi, L. 2018, The Astrophysical Journal, 853, 184

Ballone, A., Schartmann, M., Burkert, A., et al. 2016, The Astrophysical Journal Letters, 819, L28

Barriere, N. M., Tomsick, J. A., Baganoff, F. K., et al. 2014, The Astrophysical Journal, 786, 46

Begelman, M. C., Volonteri, M., & Rees, M. J. 2006, Monthly Notices of the Royal Astronomical Society, 370, 289

Belanger, G., Goldwurm, A., Melia, F., et al. 2005, The Astrophysical Journal, 635, 1095

Borkar, A., Eckart, A., Straubmeier, C., et al. 2016, Monthly Notices of the Royal Astronomical Society, 458, 2336

Bower, G. C., Markoff, S., Dexter, J., et al. 2015, The Astrophysical Journal, 802, 69

Bowyer, S., Byram, E., Chubb, T., & Friedman, H. 1965, Science, 147, 394

Boyce, H., Haggard, D., Witzel, G., et al. 2019, The Astrophysical Journal, 871, 161

Čadež, A., Calvani, M., & Kostić, U. 2008, Astronomy & Astrophysics, 487, 527

Calderón, D., Ballone, A., Cuadra, J., et al. 2015, Monthly Notices of the Royal Astronomical Society, 455, 4388

Chandler, C. J., & Sjouwerman, L. O. 2014, The Astronomer's Telegram, 6017

Coti Zelati, F., Rea, N., Turolla, R., et al. 2017, MNRAS, 471, 1819

Cuadra, J., Nayakshin, S., & Martins, F. 2007, Monthly Notices of the Royal Astronomical Society, 383, 458

Davis, M. H. 1984, Journal of the Royal Statistical Society: Series B (Methodological), 46, 353

Degenaar, N., Miller, J., Kennea, J., et al. 2013, The Astrophysical Journal, 769, 155

Demorest, P., Pennucci, T., Ransom, S., Roberts, M., & Hessels, J. 2010, nature, 467, 1081

Dibi, S., Markoff, S., Belmont, R., et al. 2014, Monthly Notices of the Royal Astronomical Society, 441, 1005

—. 2016, Monthly Notices of the Royal Astronomical Society, 461, 552

Dodds-Eden, K., Sharma, P., Quataert, E., et al. 2010, The Astrophysical Journal, 725, 450

Dodds-Eden, K., Porquet, D., Trap, G., et al. 2009, The Astrophysical Journal, 698, 676

Eckart, A., Baganoff, F., Morris, M., et al. 2009, Astronomy & Astrophysics, 500, 935

Evans, M. 1993, N. hastings, and B. Peacock, Statistical Distributions, John Wiley & Sons

Ferrarese, L., & Ford, H. 2005, Space Science Reviews, 116, 523

Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, in SPIE, Vol. 6270, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62701V

Gehrels, N. 1986, The Astrophysical Journal, 303, 336

Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, Reviews of Modern Physics, 82, 3121

Gillessen, S., Genzel, R., Fritz, T., et al. 2012, Nature, 481, 51

—. 2013a, The Astrophysical Journal, 763, 78

Gillessen, S., Genzel, R., Fritz, T. K., et al. 2013b, The Astrophysical Journal, 774, 44

Gillessen, S., Plewa, P., Widmann, F., et al. 2018, arXiv preprint arXiv:1812.01416

Goldwurm, A., Brion, E., Goldoni, P., et al. 2003, The Astrophysical Journal, 584, 751

Graham, A. W. 2016, in Galactic Bulges (Springer), 263–313

Greenstein, J. L., & Schmidt, M. 1964, The Astrophysical Journal, 140, 1

Haggard, D., Baganoff, F. K., Rea, N., et al. 2014, The astronomer's telegram, 6242

Haggard, D., et al. 2019, in preparation

Ho, L. C. 1999, The Astrophysical Journal, 516, 672

Kawashima, T., Matsumoto, Y., & Matsumoto, R. 2017, Publications of the Astronomical Society of Japan, 69, 43

Kennea, J., Burrows, D., Kouveliotou, C., et al. 2013, The Astrophysical Journal Letters, 770, L24

Kostić, U., Čadež, A., Calvani, M., & Gomboc, A. 2009, Astronomy & Astrophysics, 496, 307

Liu, S., & Melia, F. 2002, The Astrophysical Journal Letters, 566, L77

Liu, S., Petrosian, V., & Melia, F. 2004, The Astrophysical Journal Letters, 611, L101

Madigan, A.-M., McCourt, M., & O'Leary, R. M. 2016, Monthly Notices of the Royal Astronomical Society, stw2815

Mapelli, M., & Ripamonti, E. 2015, The Astrophysical Journal, 806, 197

Margalit, B., & Metzger, B. D. 2017, The Astrophysical Journal Letters, 850, L19

Margutti, R., Alexander, K., Xie, X., et al. 2018, The Astrophysical Journal Letters, 856, L18

Markoff, S., Falcke, H., Yuan, F., & Biermann, P. L. 2001, Astronomy & Astrophysics, 379, L13

Marrone, D. P., Moran, J. M., Zhao, J.-H., & Rao, R. 2006, The Astrophysical Journal Letters, 654, L57

Marrone, D. P., Baganoff, F., Morris, M., et al. 2008, The Astrophysical Journal, 682, 373

Metzger, B. D., Thompson, T. A., & Quataert, E. 2018, The Astrophysical Journal, 856, 101

Mezcua, M. 2017, International Journal of Modern Physics D, 26, 1730021

Mori, K., Gotthelf, E. V., Zhang, S., et al. 2013, The Astrophysical Journal Letters, 770, L23

Mossoux, E., & Grosso, N. 2017, Astronomy & Astrophysics, 604, A85

Mossoux, E., Grosso, N., Bushouse, H., et al. 2016, Astronomy & Astrophysics, 589, A116

Narayan, R., Mahadevan, R., Grindlay, J. E., Popham, R. G., & Gammie, C. 1998, The Astrophysical Journal, 492, 554

Narayan, R., Özel, F., & Sironi, L. 2012, The Astrophysical Journal Letters, 757, L20

Neilsen, J., Nowak, M., Gammie, C., et al. 2013, The Astrophysical Journal, 774, 42

Neilsen, J., Markoff, S., Nowak, M., et al. 2015, The Astrophysical Journal, 799, 199

Nowak, M., Neilsen, J., Markoff, S., et al. 2012, The Astrophysical Journal, 759, 95

Oda, M. 1977, Space Science Reviews, 20, 757

Özel, F., Psaltis, D., & Narayan, R. 2000, The Astrophysical Journal, 541, 234

Pfuhl, O., Gillessen, S., Eisenhauer, F., et al. 2015, The Astrophysical Journal, 798, 111

Ponti, G., De Marco, B., Morris, M., et al. 2015, Monthly Notices of the Royal Astronomical Society, 454, 1525

Ponti, G., Jin, C., De Marco, B., et al. 2016, Monthly Notices of the Royal Astronomical Society, 461, 2688

Ponti, G., George, E., Scaringi, S., et al. 2017, Monthly Notices of the Royal Astronomical Society, 468, 2447

Porquet, D., Predehl, P., Aschenbach, B., et al. 2003, Astronomy & Astrophysics, 407, L17

Porquet, D., Grosso, N., Predehl, P., et al. 2008, Astronomy & Astrophysics, 488, 549

Quataert, E. 2002, The Astrophysical Journal, 575, 855

—. 2003, Astronomische Nachrichten: Astronomical Notes, 324, 435

Reid, M. J., & Brunthaler, A. 2004, The Astrophysical Journal, 616, 872

Remillard, R. A., & McClintock, J. E. 2006, Annu. Rev. Astron. Astrophys., 44, 49

Salpeter, E. 1964, The Astrophysical Journal, 140, 796

Sasaki, M., Suyama, T., Tanaka, T., & Yokoyama, S. 2018, Classical and Quantum Gravity, 35, 063001

Scargle, J. D., Norris, J. P., Jackson, B., & Chiang, J. 2013, The Astrophysical Journal, 764, 167

Schartmann, M., Burkert, A., Alig, C., et al. 2012, The Astrophysical Journal, 755, 155

Schmidt, M. 1963, A Century of Nature, Twenty One Discoveries That Changed Science and the World

Shakura, N. I., & Sunyaev, R. A. 1973, Astronomy and Astrophysics, 24, 337

Sigman, K. 2007, Poisson processes and Compound (batch) Poisson processes, Available at http://www.columbia.edu/~ks20/4703-Sigman/4703-07-Notes-PP-NSPP.pdf, Columbia University, accessed on May 30 2019

Smith, H. J., & Hoffleit, D. 1963, Nature, 198, 650

Verner, D., Ferland, G. J., Korista, K., & Yakovlev, D. 1996, arXiv preprint astro-ph/9601009

Volonteri, M. 2010, The Astronomy and Astrophysics Review, 18, 279

Wang, Q., Nowak, M., Markoff, S., et al. 2013, Science, 341, 981

Williams, P. K., Clavel, M., Newton, E., & Ryzhkov, D. 2017, Astrophysics Source Code Library

Willke, B., Collaboration, L. S., et al. 2018, Symmetry: Culture and Science, 29, 257

Wilms, J., Allen, A., & McCray, R. 2000, The Astrophysical Journal, 542, 914

Witzel, G., Eckart, A., Bremer, M., et al. 2012, The Astrophysical Journal Supplement Series, 203, 18

Yuan, F., & Narayan, R. 2014, Annual Review of Astronomy and Astrophysics, 52, 529

Yuan, F., Quataert, E., & Narayan, R. 2003, The Astrophysical Journal, 598, 301

Yuan, Q., & Wang, Q. D. 2015, Monthly Notices of the Royal Astronomical Society, 456, 1438

Yuan, Q., Wang, Q. D., Liu, S., & Wu, K. 2017, Monthly Notices of the Royal Astronomical Society, 473, 306

Yusef-Zadeh, F., Wardle, M., Dodds-Eden, K., et al. 2012, The Astronomical Journal, 144, 1

Zhang, S., Baganoff, F. K., Ponti, G., et al. 2017, The Astrophysical Journal, 843, 96

Zubovas, K., Nayakshin, S., & Markoff, S. 2012, Monthly Notices of the Royal Astronomical Society, 421, 1315