

XMM-Newton Technical Note

IDT

XMM-SOC-CAL-TN-0242 Issue 1.0

Sco X-1 out of field observations with XMM-Newton

M. R. Alexandersen¹, I. Valtchanov², F. Fürst², D. Lumb², M. Freyberg³, P. Friedrich³

¹Technical University of Denmark

²XMM-Newton Science Operations Centre

³Max Planck Institute for Extraterrestrial Physics

November 12, 2025

Revision history

Revision number	Date	Revision author	Comments
1.0	Oct 8, 2025	M. R. Alexandersen	first issue



1 Introduction & Background

Continuous calibration of X-ray space telescopes is essential to maintain accurate and reliable measurements and data. Over time, launch vibrations and years in orbit can cause small shifts or wear to the telescope's components, affecting its calibration and response. Recent cross-calibration efforts between XMM-Newton and NuSTAR show that XMM-Newton's EPIC-pn instrument generally measures fluxes about 20% lower than NuSTAR in the 3-12 keV energy range (Fürst, F. 2022), suggesting that its effective area model is too high. A possible explanation is potential misalignments of optical components that were not measured or present during pre-launch ground measurements. By determining the exact misalignments, it is possible to update the ground-measured telescope model and effective area.

Because the telescope is in space, the alignment of its optical components cannot be directly measured like on the ground. Instead, a possible way to estimate it is through off-axis observations of bright X-ray sources just outside the field of view. Here stray light gives rise to so-called single-reflection arcs in the focal plane, mainly produced when off-axis photons miss the primary parabolic mirror and reflect only once from the secondary hyperbolic mirror (Spiga 2015; Buitrago-Casas et al. 2017). Each individual mirror shell produces its own unique single-reflection arc, with the position and intensity of the arcs being highly sensitive to the alignment of the telescope's optical components (Spiga 2015). By comparing actual observations to simulations, it is theoretically possible to gain a better understanding of the true in-orbit alignment of the components.

This work applies the above-mentioned method to XMM-Newton's EPIC-pn instrument using out-of-field single-reflection observations of the X-ray binary Scorpius X-1 (Sco X-1), as well as ray-tracing simulations with ESA's SciSim (Gabriel et al. 2005) software. In theory, by systematically adjusting the simulated mirror and X-ray baffle alignment until the single-reflection arcs approximately match the observations, a better estimate of the true alignment of the components can be obtained, which could then make it possible to update the telescope model and effective area calibration. However, as this TN will show, this is quite a complex task and therefore not trivial to achieve. Furthermore, because of additional complications introduced by the RGA in the optical path, this work will only consider the EPIC-pn camera, while the two MOS cameras will be addressed in future work.

The remainder of this TN is organized as follows: Section 2 describes the observations and data reduction used for the analysis. Section 3 gives a detailed description of the simulation tool and how it was utilized to reproduce the out-of-field observations. Section 4 presents the results of the analysis. Section 5 provides a summary and discussion of the findings and their impact. Section 6 gives a brief outlook and recommendations for future work based on this analysis.



2 Observations & Data reduction

The dataset used in this analysis consists of 28 dedicated off-axis calibration observations of Sco X-1 taken between 2024-03-04 and 2024-03-08 with EPIC-pn in Full Frame mode using the THICK filter under program 093479. The observations are split into four perpendicular directions around the source, each containing seven pointings at different off-axis angles ranging from approximately $34'$ to $79'$. Figure 1 shows the observation pattern used for the dataset, with each circle representing the area observed by each observation.

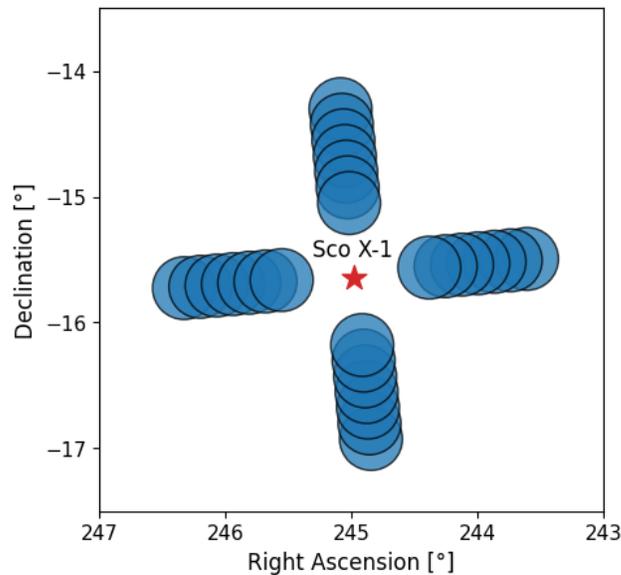


Figure 1: Sky observation pattern for the Sco X-1 off-axis dataset used in this analysis. Each blue circle represents the $\sim 30'$ field of view observed by the EPIC-pn instrument during each pointing, while the red star marks the location of Sco X-1. The dataset consists of 28 observations, split into four perpendicular directions, each containing seven pointings at increasing off-axis angles.

For convenience throughout the analysis, the 28 observations were grouped into four quadrants based on the perpendicular direction in which they were taken relative to the source, with each quadrant thus containing seven observations at different off-axis angles. These quadrants were arbitrarily labeled Q1, Q2, Q3, and Q4, where in Figure 1 the Q1 observations lie above the source, Q2 below, Q3 to the left and Q4 to the right. Furthermore, when referring to specific observations throughout the TN, only the last four digits of the observation ID will be used, as these are the only ones that differ and are unique between each observation.

The data reduction of the observations was strictly carried out in Python within ESA Data-labs (Navarro et al. 2020), where all EPIC-pn data were extracted using the XMM-Newton Science Analysis Software (SAS) (Gulbahar et al. 2025) version 22.1.0 with Current Calibration Files (CCFs) as of 2025-05-14. For each observation, event lists were created and filtered using



PATTERN ≤ 4 , FLAG = 0, and energies between 0.2-12.0 keV, which ensured that only good quality single- and double-pixel events were kept. Additionally, the raw detector pixel coordinates from the event lists were converted to physical camera coordinates using the CCFs, to allow for easier comparisons with simulations later in the analysis. Table 1 summarizes the key properties of all 28 observations.

Table 1: Summary of the 28 Sco X-1 off-axis observations used in this analysis.

Date	ObsID	Exposure [ks]	Off-axis [']	RA [°]	DEC [°]	PA [°]	Quadrant
2024-03-04	0934790131	6.21	80.4	245.087	-14.303	98.721	Q1
2024-03-05	0934790132	7.25	73.0	245.076	-14.428	98.722	Q1
2024-03-05	0934790133	7.25	65.5	245.064	-14.551	98.722	Q1
2024-03-05	0934790134	7.25	58.0	245.053	-14.676	98.722	Q1
2024-03-05	0934790135	7.25	50.5	245.042	-14.801	98.722	Q1
2024-03-05	0934790136	7.25	43.0	245.030	-14.925	98.722	Q1
2024-03-05	0934790137	6.53	35.5	245.019	-15.050	98.722	Q1
2024-03-05	0934791431	7.23	32.6	244.915	-16.179	98.722	Q2
2024-03-05	0934791432	7.25	40.0	244.903	-16.303	98.722	Q2
2024-03-05	0934791433	7.25	47.6	244.891	-16.428	98.723	Q2
2024-03-06	0934791434	7.25	55.1	244.880	-16.553	98.723	Q2
2024-03-06	0934791435	7.25	62.5	244.869	-16.677	98.722	Q2
2024-03-06	0934791436	7.25	70.1	244.857	-16.802	98.723	Q2
2024-03-06	0934791437	15.09	77.6	244.845	-16.926	98.723	Q2
2024-03-07	0934791531	8.13	78.2	246.330	-15.727	98.721	Q3
2024-03-07	0934791532	7.25	70.7	246.200	-15.717	98.722	Q3
2024-03-07	0934791533	7.25	63.2	246.071	-15.706	98.722	Q3
2024-03-07	0934791534	7.25	55.6	245.941	-15.695	98.722	Q3
2024-03-07	0934791535	7.25	48.1	245.811	-15.684	98.722	Q3
2024-03-07	0934791536	7.25	40.7	245.682	-15.674	98.722	Q3
2024-03-07	0934791537	6.53	33.2	245.554	-15.663	98.722	Q3
2024-03-07	0934792831	7.23	34.9	244.382	-15.564	98.722	Q4
2024-03-07	0934792832	7.25	42.4	244.252	-15.552	98.722	Q4
2024-03-07	0934792833	7.25	49.9	244.122	-15.540	98.722	Q4
2024-03-08	0934792834	7.25	57.4	243.993	-15.530	98.722	Q4
2024-03-08	0934792835	7.25	64.9	243.863	-15.518	98.722	Q4
2024-03-08	0934792836	7.25	72.4	243.735	-15.506	98.722	Q4
2024-03-08	0934792837	14.58	79.9	243.606	-15.494	98.722	Q4

Apart from just creating event lists, SAS was also used to create detector images in cts/s for each of the 28 observations, shown in Figure 2. Here each row corresponds to observations taken in a specific quadrant (i.e. Q1, Q2, Q3, or Q4), while each column corresponds to observations taken at approximately the same off-axis angle. As seen in the images, the four observations with the highest off-axis angles show very few or no distinct features, indicating that they are too far off-axis to produce clear single-reflection arcs. At lower off-axis angles, single-reflection arcs are



expected to originate from the innermost mirror shells (e.g., shells 58, 57, 56, etc.), whereas at higher off-axis angles, they are expected to arise from the outermost shells (e.g., shells 1, 2, 3, etc.). Counterintuitively, single-reflection arcs from the larger, outermost mirror shells generally appear as the smaller, innermost arcs in the images, while the smaller, innermost shells appear as the larger, outermost arcs (Spiga 2015; Buitrago-Casas et al. 2017).

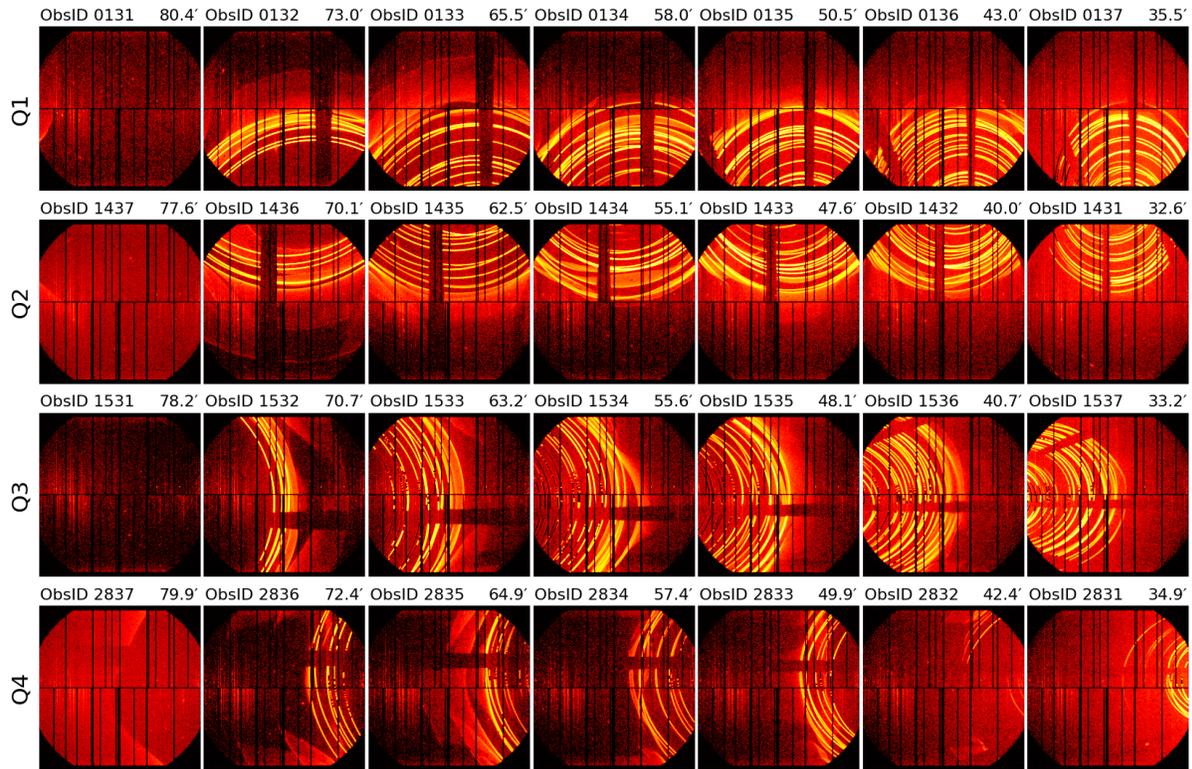


Figure 2: Detector images of the 28 Sco X-1 off-axis observations used in this analysis in cts/s. Clear single-reflection arcs are visible in almost all observations except the highest off-axis pointing in each quadrant. The color indicates the brightness of each pixel in the image, with high-intensity regions shown in yellow and low-intensity regions in dark red. The same color map range was used for all images, in order to better allow for comparisons.

3 SciSim ray-tracing simulations

To calibrate and constrain the alignment of the telescope’s optical components, version 4.0 of the XMM-Newton Science Simulator (SciSim) (Gabriel et al. 2005; SciSim Team 2005) was used to reproduce the Sco X-1 off-axis observations, making it possible to systematically investigate how changes in the alignment of the optical components affect the measured single-reflection arcs. By iteratively adjusting the telescope model within SciSim and comparing the simulated



results to the real observations, the aim was to identify a model that could reproduce the data more accurately, thereby providing a potential improvement to the estimate of the telescope's optical alignment.

Despite the public distribution of SciSim being discontinued on July 19, 2010, it remains fully functional and was installed internally in ESA Datalabs for use in this analysis. While the software is typically operated either through the command line or a graphical user interface, in this work it was executed by calling terminal commands using Python in ESA Datalabs.

3.1 Simulation pipeline

SciSim works by running a series of smaller modules, each simulating a specific part of the overall simulation, such as the satellite pointing, ray generation, mirror reflections, and detector response. The smaller modules are typically run in a pipeline, with the output from one module being used as the input for the next (Gabriel et al. 2005; SciSim Team 2005). Figure 3 illustrates the specific part of the pipeline used in this work, which was run left to right.

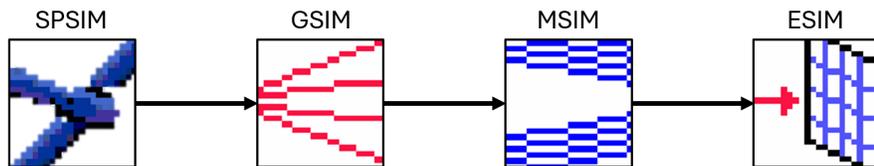


Figure 3: Illustration of the SciSim pipeline used to simulate the off-axis observations in this analysis (adapted from Gabriel et al. 2005). The pipeline runs the four modules SPSIM, GSIM, MSIM and ESIM.

3.1.1 Spacecraft Simulator (SPSIM)

The first stage of the SciSim pipeline was the Spacecraft Simulator (SPSIM) (Siddiqui & Bakker 2005), which generates an attitude history file for the simulated observation. The spacecraft's Right Ascension, Declination, and Position Angle were extracted directly from the XMM-SAS event lists for the 28 Sco X-1 observations, while effects such as jitter and drift were omitted due to the stable pointings of XMM-Newton.

3.1.2 Ray Generator (GSIM)

The second part of the pipeline was the Ray Generator (GSIM) (Beijersbergen & Siddiqui 2005), which takes the attitude history file from SPSIM as input and creates a ray file describing the simulated X-ray photons and their trajectories. Here Sco X-1 was defined as a point source at coordinates RA, DEC (J2000) = 244.9795°, -15.6402°, with no background contribution included. The source was arbitrarily modeled using a simple power-law with a photon index of 1.7 and a flux of 0.1 photons $\text{mm}^{-2} \text{s}^{-1}$ in the 0.5–2.0 keV range. Lastly, an exposure time of 500



seconds was chosen to provide a good balance between high statistics and runtime, producing clear single-reflection arcs while keeping each simulation runtime to around 15 minutes.

3.1.3 Mirror Module Simulator (MSIM)

The next step was the Mirror Module Simulator (MSIM) (Beijersbergen 2005a), which is a ray-tracing tool that takes the ray file from GSIM and simulates the paths and reflections of the photons within the telescope's mirror module, producing an updated ray file containing all remaining photons after passing through it. When using MSIM, a mirror model file must be specified, which contains a detailed description of the mirror module geometry and physical properties, including the positions, shapes, misalignments, and distortions of the mirrors, X-ray baffles, and spiders. The model can be simulated either as an `accurate` model, including effects like component misalignments and distortions, or as an `ideal` model, assuming a perfect Wolter Type I mirror module.

In this analysis, MSIM was run using the `accurate` model to obtain the most realistic results, as mirror distortions and component misalignments were found to significantly affect the position and intensity of the single-reflection arcs. The specific mirror model used was `fm2model_xrb.f7493.2_v1.dat`, which provides a description of the mirror module in front of EPIC-pn based on ground measurements taken prior to launch. Throughout the rest of this TN, this mirror model will be referred to as the `FM2 model`, which is the model that was investigated and used for calibration in this analysis.

3.1.4 European Photon Imaging Camera Simulator (ESIM)

The last module in the pipeline was the European Photon Imaging Camera Simulator (ESIM) (Bakker & Jalota 2005), which uses the MSIM ray file to simulate the EPIC-MOS and EPIC-pn detector responses, producing an event file containing the energy, detection time, and position of each detected photon. In this analysis, the camera was set to the EPIC-pn camera, the filter to `THICK`, and the science mode to `Full Frame`. Furthermore, the detector was shifted +2.6 mm along the Y-axis and -2.0 mm along the Z-axis, as this was found to be required to match the boresight between simulations and observations.

3.2 Additional simulation tools

In addition to the main pipeline modules, SciSim provides several tools that can be applied to intermediate outputs from the different modules before passing them further through the simulation pipeline. In this work, `ssimfilter` (Beijersbergen 2005b) was frequently used on the intermediate ray file output from MSIM to filter and extract rays reflected from specific mirror shells, making it possible to determine exactly which shell each simulated single-reflection arc originated from. Furthermore, `ssimraystat` (Beijersbergen 2005b) was used on the MSIM ray file output to measure the total illuminated effective area before and after modifying the FM2 model, making it possible to quantify the effect of the changes.



3.3 Adjusting mirror model

When adjusting the FM2 model in this analysis, almost every optical component of the telescope could be shifted, tilted, or distorted in various ways. This led to a huge parameter space with thousands of possible parameters and many more possible combinations, which was not feasible to explore. This work therefore only focused on adjusting the tilt of the mirrors and the shift of the X-ray baffles, as these were found to have the largest impact on the position and intensity of the single-reflection arcs. The general method for tilting and shifting these components is described in the following sections, while the specific choices and arcs selected for matching are presented later alongside the results. The coordinate system used in SciSim is furthermore defined such that the X -axis is along the telescope optical axis, the Y -axis is perpendicular to this and runs parallel to the length of the EPIC-pn CCDs, while the Z -axis completes the right-hand coordinate system and runs parallel to the width of the EPIC-pn CCDs.

An overview of the XMM-Newton mirror model used in SciSim is shown in Figure 4. Here the left panel shows a general diagram of the mirror assembly, with a few of the mirror shells, X-ray baffle sieve plates, and spider structure being visible. The right panel furthermore illustrates the exact telescope model implemented in the FM2 model and used in SciSim, which mainly consists of several different hollow cylinders making up the X-ray baffles, spiders, and blocking cylinder, while the 58 mirror shells are implemented as parabolic and hyperbolic shells.

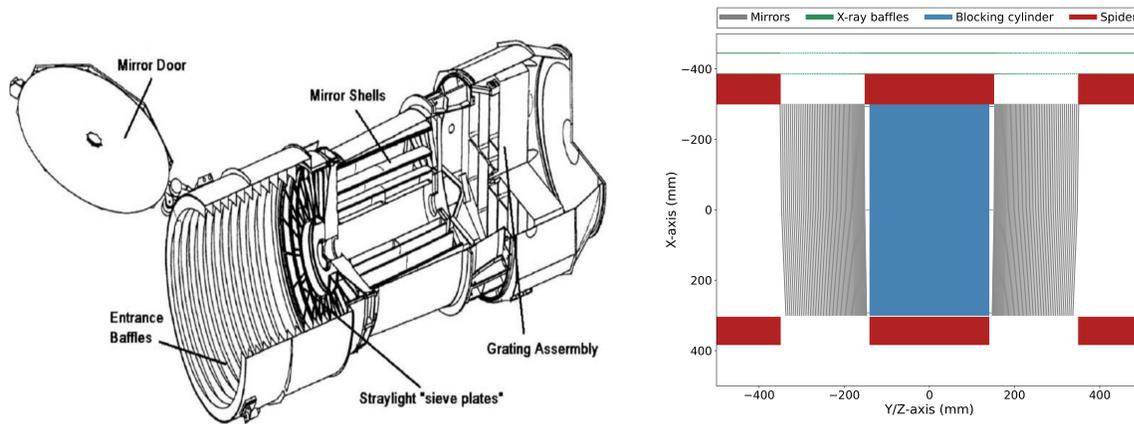


Figure 4: (Left) Diagram of the XMM-Newton mirror assembly (adapted from Lumb et al. 2012), showing the mirror shells, X-ray baffle sieve plates, spider structure, etc. Note that the grating assembly in the image is not present for the mirror module in front of EPIC-pn. (Right) Illustration of the exact telescope model implemented in the FM2 model and used in SciSim, including the 58 mirror shells, two X-ray baffle sieve plates, spiders, and blocking cylinder.



3.3.1 Tilt of the mirror assembly

The simplest parameter in the FM2 model modified in this analysis was the orientation of the entire mirror assembly, which includes the full mirror assembly with X-ray baffles, spiders, and mirrors. This was controlled by the following line in the FM2 model file (Beijersbergen 2005a):

```
Stack XMM_MM
0 0 0 0 0 0 0 -700 +373 ; axis, angle(deg), displ(mm), xrange(mm)
```

In this analysis only the tilt of the mirror assembly was adjusted, which is described in the line by an axis of rotation \vec{n} (red) and a rotation angle θ in degrees (blue). However, tilting the mirror assembly directly in this format was unintuitive and difficult to visualize, and it was instead described as three individual rotations; first around the X -axis by θ_x , then the Y -axis by θ_y , and finally the Z -axis by θ_z . This was implemented using the following rotation matrices:

$$R_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) \\ 0 & \sin(\theta_x) & \cos(\theta_x) \end{bmatrix}, R_y = \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) \\ 0 & 1 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) \end{bmatrix}, R_z = \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

Combining the above into a single rotation matrix using a 1-2-3 Euler angle rotation gave:

$$R = R_z \cdot R_y \cdot R_x \quad (2)$$

To implement this rotation in the FM2 model file, it had to be expressed as a vector \vec{n} describing the rotation axis and a rotation angle θ around this axis. The vector was here simply found as the eigenvector of the rotation matrix R corresponding to the eigenvalue 1, while the rotation angle was calculated as:

$$\theta = \arccos\left(\frac{R_{11} + R_{22} + R_{33} - 1}{2}\right) \quad (3)$$

where R_{ij} is the element in the i th row and j th column of R .

3.3.2 Tilt of individual mirror shells

Another adjustment applied to the FM2 model in this analysis involved the orientation and alignment of the 58 individual mirror shells. Compared to tilting the entire mirror assembly, this introduced a significantly larger parameter space, since each shell was adjusted independently. In the FM2 model, the alignment of each shell was defined by a dedicated line, such as the following example for mirror shell 1 (Beijersbergen 2005a):

```
1 shell id
0 0 0 -0.043 0.0416 ; phi(deg), theta(deg), displ(mm)
```

Similar to the mirror assembly adjustment, this analysis focused exclusively on tilting the individual mirror shells. The first parameter in the model describing this was the azimuthal



direction angle φ of the tilt in the (y, z) -plane (red), while the second parameter defined the rotation angle θ applied along that direction (blue). For the same reasons as with the mirror assembly tilt, the tilt of the individual mirror shells was implemented as a series of sequential rotations around different axes, but in this case using only a rotation about the Y -axis by θ_y , followed by a rotation about the Z -axis by θ_z . Using the rotation matrices in Eq. (1), the combined 2-3 Euler rotation matrix was obtained as:

$$R = R_z \cdot R_y \tag{4}$$

To implement this rotation in the FM2 model, it was assumed that each mirror shell initially pointed along the X -axis with the unit vector $\vec{v} = [1, 0, 0]$. Applying the rotation matrix then gave the new pointing vector $\vec{u} = [\vec{u}_x, \vec{u}_y, \vec{u}_z]$:

$$\vec{u} = R \cdot \vec{v} \tag{5}$$

Using the above, the first parameter required in the FM2 model φ was found as the angle between the positive Y -axis and the rotated unit vector \vec{u} projected onto the (Y, Z) -plane:

$$\varphi = \arctan2(\vec{u}_z, \vec{u}_y) \tag{6}$$

Furthermore, the second parameter required in the FM2 model θ was found as the total rotated angle of the spacecraft, corresponding to the angle between the original and rotated vector. Since both vectors are unit vectors, it follows that $\|\vec{v}\| \cdot \|\vec{u}\| = 1$, and since the original vector was $\hat{v} = [1, 0, 0]$, it follows that $\vec{v} \cdot \vec{u} = \vec{u}_x$. The final angle was therefore determined as:

$$\theta = \arccos\left(\frac{\vec{v} \cdot \vec{u}}{\|\vec{v}\| \cdot \|\vec{u}\|}\right) = \arccos(\vec{u}_x) \tag{7}$$

3.3.3 Shift of X-ray baffles

The last relevant components of the FM2 model adjusted in this analysis were the two X-ray baffles, which are designed to block most of the stray light from entering the telescope. Although baffle misalignments do not change the positions of the single-reflection arcs in the focal plane, they can affect their brightness and, in some cases, prevent certain arcs from appearing altogether. Like the mirror assembly, the two baffles could be tilted or shifted along all three axes, controlled by the following lines in the FM2 model (Beijersbergen 2005a):

```
XBaffle OuterSievePlate
0 0 0 0 0 0 0 -445 -444 ; axis, angle(deg), displ(mm), xrange(mm)
XBaffle InnerSievePlate
0 0 0 0 0 0 0 -386 -385 ; axis, angle(deg), displ(mm), xrange(mm)
```

Initial testing showed that tilting and shifting the baffles had essentially the same effect on the single-reflection arcs. Therefore, in order to keep the analysis simple and reduce the parameter space, only shifts were considered. Furthermore, despite the two baffles technically being possible to shift independently of each other, they were for simplicity's sake assumed to be misaligned by the exact same amount, and thus shifted together. The only parameters changed in the two lines above were thus the three displacement parameters (green).



3.3.4 Automatic Matching of Simulated and Observed Arcs

Since each simulation was computationally expensive, manually adjusting the tilt parameters when matching the mirror assembly or individual mirror shells was impractical and time-consuming. To address this, a custom Python function, available internally in the XMM-Newton SOC Bitbucket repository, was developed to compute a loss metric quantifying how closely a given simulated single-reflection arc matches an observed one. By minimizing this loss function, the best-fit tilt values were determined.

The function worked by first manually fitting a circle¹ to the observed single-reflection arc of interest. After this, a simulation was run and the intermediate MSIM ray file was filtered with `ssimfilter` to keep only rays reflected from the mirror shell being matched, which was then also fitted with a circle. Lastly, 100 evenly spaced points were generated along arcs of both the observed and simulated circles within the detector field of view, and the loss metric was computed as the RMS distance between corresponding point pairs.

Although a full grid search of Y - and Z -axis tilts could find the minimum of the loss function relatively easily, it was computationally expensive, taking several days for just two parameters. Instead, the function `gp_minimize` from the `scikit-optimize` python library was used to perform Bayesian optimization, which performed a smart search of the parameter space, with new points being selected by balancing exploration of uncertain regions and improvement of areas with low loss. By doing this, the optimal tilt parameters could be found much more efficiently within a few hours of simulation time.

3.4 Nominal on-axis effective area

The final step in the analysis was to determine how the modifications to the FM2 model affected the telescope's nominal on-axis effective area. This was done by placing a monochromatic point source with a flux of $0.001 \text{ photons mm}^{-2} \text{ s}^{-1}$ at the nominal on-axis in SciSim, such that the source was focused in the focal plane at the nominal target location used in actual observations. The energy of the source was varied from 0.2 to 12.0 keV in 0.1 keV steps, with a 500 s exposure observation being simulated for each energy step. By processing the intermediate MSIM ray file with `ssimraystat`, the effective area was measured for both the unmodified and modified FM2 model, allowing the change in effective area to be determined.

4 Results

4.1 Unmodified simulations

Before modifying the mirror model, all 28 Sco X-1 off-axis observations in the dataset were simulated using the standard FM2 model to test how well the default simulations matched the observations. Figure 5 shows the results, with the event lists from SciSim being converted

¹The arcs are not circular in reality, but this approximation is adequate within the observed detector region.



into images. Like the real observations, the simulations show clear single-reflection arcs in all pointings, except at the highest off-axis angle in each direction, which again appears to be too far off-axis. A closer look also reveals that the simulated arcs are generally simpler than the observed ones, with fewer crossing and scattered features.

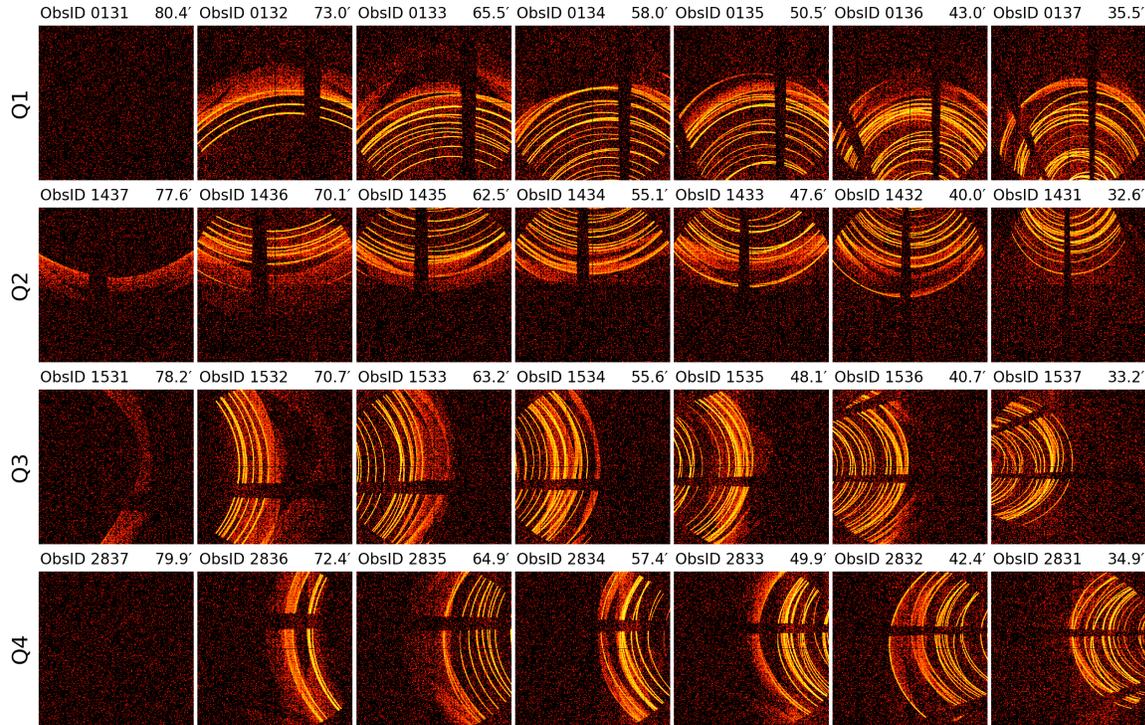


Figure 5: Simulated detector images of the 28 off-axis Sco X-1 pointings using the standard FM2 mirror model. While the simulations reproduce the overall observed arc structure, the exact positions and intensities of the arcs do not match the observations. Furthermore, many of the simulated arcs appear simpler and lack some of the crossing and scattered features present in the real data.

The initial analysis focused on the four Sco X-1 observations with ObsID 0132, 1436, 1532, and 2836, corresponding to the second-highest off-axis angle in each direction. These were chosen since they showed the simplest and fewest single-reflection arcs from the outermost mirror shells of the telescope. A direct comparison between the observations and the unmodified FM2 simulations for these four pointings can be seen in the scatter plot in Figure 6, where the raw event files have been plotted. Looking at the plots, a clear mismatch between observations and simulations can be seen, with the positions of the single-reflection arcs not aligning in any of the four perpendicular directions, suggesting that the FM2 model does not accurately represent the current state of the mirror module in front of EPIC-pn.

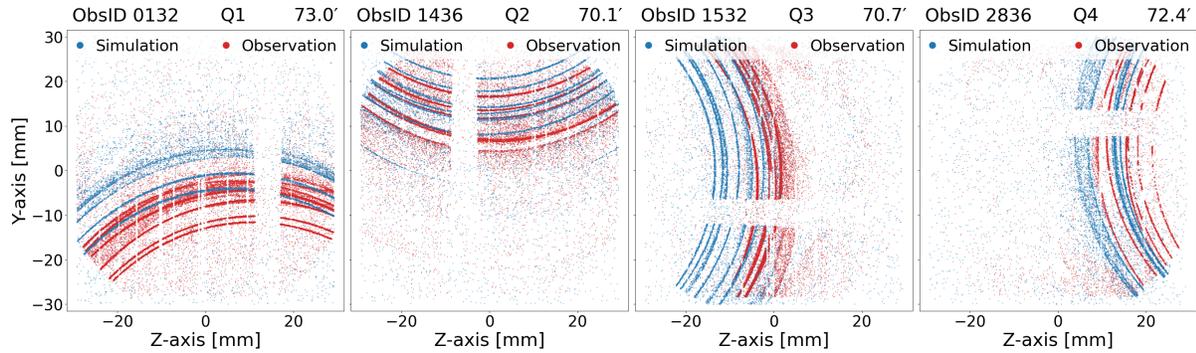


Figure 6: Observations (red) and simulations (blue) for four Sco X-1 pointings using the standard FM2 model. The energy range for the observations were restricted to 3.0–3.2 keV to reduce the number of counts. A clear mismatch between the positions of the observed and simulated single-reflection arcs can be seen.

4.2 Matching single-reflection arcs by tilting the mirror assembly

The simplest way to modify the FM2 model is to tilt the entire mirror assembly, including the mirrors, baffles, and spider. For this, the four pointings mentioned in the previous section were used, where the innermost single-reflection arcs in the simulations were matched to the innermost arcs in the observations. In theory, these arcs should be caused by reflections from the outermost mirror shell 1, however single-reflections from this shell were found to be blocked by the spider and X-ray baffles, and therefore missing completely in the simulations. Because of this, the innermost arc being matched was assumed to be from mirror shell 2.

When matching the arcs, the mirror assembly in the FM2 model was tilted independently around the Y - and Z -axis for each of the four observations using the automatic fitting algorithm. X -axis tilts were not considered in this analysis, since the shadow in the single-reflection arcs, caused by one of the spider spokes blocking photons, already matched well between observations and simulations. Figure 7 compares the resulting simulations (blue) and observations (red), where the innermost single-reflection arc now matches in all four observations, while the remaining arcs do not. This suggests that tilting the entire mirror assembly is not sufficient, and that individual mirror shells are likely misaligned as well. The required tilts also differ significantly between the four observations, ranging from 0.5' to 3.0' for the Y -axis and 0.9' to 2.5' for the Z -axis. This was unexpected, as matching the arcs on one side of the telescope should in theory also reproduce the ones on the other side.

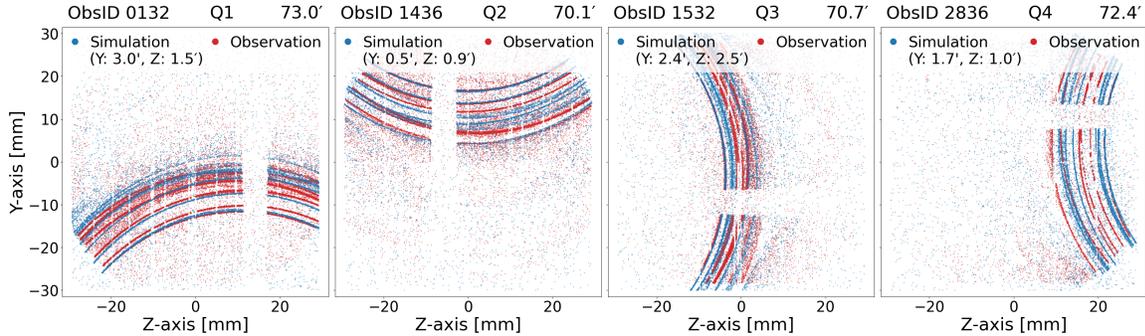


Figure 7: Observations (red) and simulations (blue) for four Sco X-1 pointings after tilting the mirror assembly in the FM2 model to match the innermost single-reflection arc. The applied tilt angles for each pointing are shown in the legends, and are found to be different for each of the four observations.

To further investigate the tilt parameters that best match the innermost single-reflection arc, a two-dimensional grid of Y- and Z-axis tilt values between $-1'$ and $4'$ was generated for each of the four Sco X-1 observations. For each grid point, the loss score between the innermost simulated and observed arc was computed using the custom Python function described in Section 3. The results are shown in Figure 8, where the minima and best-fit tilt parameters are clearly identifiable, but occur at different tilt parameters for each observation, further illustrating that a single mirror assembly tilt cannot simultaneously reproduce all four pointings.

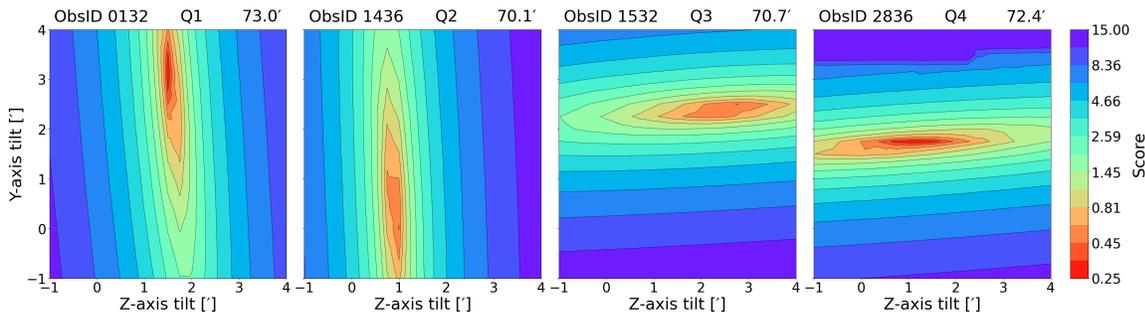


Figure 8: Loss score evaluated over a grid of Y- and Z-axis mirror assembly tilt values for four Sco X-1 pointings. Blue colors indicate higher loss values and a poor fit, while red colors indicate lower values and a good fit. The different minima between observations demonstrate that a single set of best-fit mirror assembly tilt parameters cannot match all pointings simultaneously.

Despite not yielding a perfect match for all single-reflection arcs or observations at once, the effects of the applied mirror assembly tilts on the nominal on-axis effective area were also investigated and shown in Figure 9. The largest effective area decrease occurred when matching the single-reflection arc in the Q1 ObsID 0132 observation, resulting in a $\sim 1\text{--}2\%$ drop, while matching the arcs in the Q2 ObsID 1436 and Q4 ObsID 2836 instead led to a small $\sim 2\text{--}4\%$



increase. Thus, while adjustments to the mirror assembly do affect the effective area, they are unlikely to account for the full discrepancy between XMM-Newton and NuSTAR.

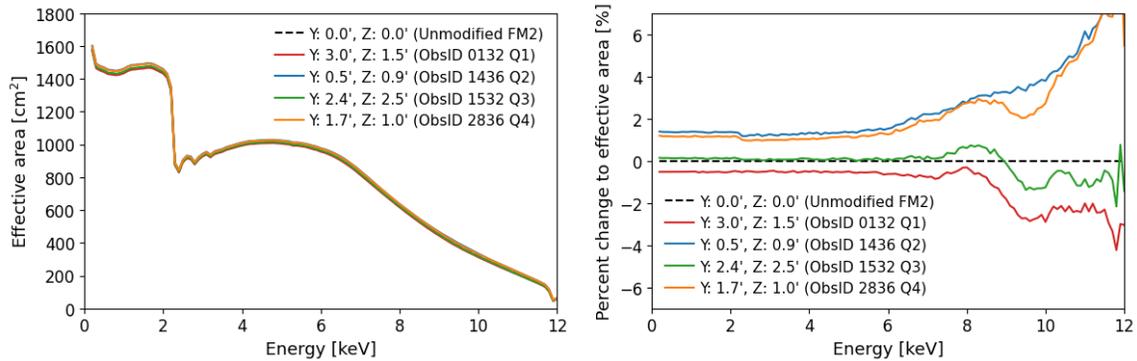


Figure 9: (Left) Simulated nominal on-axis effective area, and (Right) percent change in effective area compared to the unmodified FM2 model for the mirror assembly tilts matching the innermost single-reflection arc in four Sco X-1 pointings. The largest decrease in effective area ($\sim 1\text{--}2\%$) was found when matching the Q1 ObsID 0132 observation, while the Q2 ObsID 1436 and Q4 ObsID 2836 observations show small increases ($\sim 2\text{--}4\%$).

4.3 Matching single-reflection arcs by tilting individual mirror shells

Since tilting the entire mirror assembly did not give a satisfactory result, the next step was to test tilting individual mirror shells around the Y - and Z -axis. The analysis here focused on matching the first 14 mirror shells (shells 2 to 15), since these produced the most distinct and circular single-reflection arcs. When including more mirror shells beyond this, the automatic fitting algorithm struggled due to the arcs appearing slightly elliptical, overlapping, and noisy, and the results for the first 14 shells were already not particularly promising.

The analysis was limited to matching the 14 single-reflection arcs in the Q1 quadrant, using the 3rd, 4th and 5th highest off-axis observations with ObsID 0132, 0133, and 0134. In the simulations, the origin of each single-reflection arc was known from `ssimfilter`, but this was not the case in the observations. When doing the matching, the observed arcs were therefore assumed to follow the expected ordering, with the innermost arc corresponding to shell 2, the next to shell 3, and so on. In practice, this assumption is likely not strictly correct, as shell distortions and tilts can change the ordering of the arcs significantly. Nevertheless, it was adopted here for simplicity and in the absence of a better alternative.

Figure 10 shows the results for the second, third, and fourth highest off-axis pointings in each direction. As expected, the observed and simulated single-reflection arcs align almost perfectly across all three off-axis angles in the Q1 quadrant, which was deliberately matched. Several shells, such as shells 2–7 and 9–13, furthermore appear in two Q1 pointings simultaneously and still match consistently in both. However, when the same parameters are applied to the other



quadrants, the single-reflection arcs no longer align. As with tilting the entire mirror assembly, this result indicates that simple tilting is not sufficient to simultaneously match the single-reflection arcs across all observations and quadrants. Instead, other parameters in the FM2 model may also need to be updated, such as the exact shapes or distortions of the individual mirror shells. Another possible explanation is that no global tilt of each mirror shell exists, and that different parts of each shell may instead somehow be deformed and tilted by different amounts.

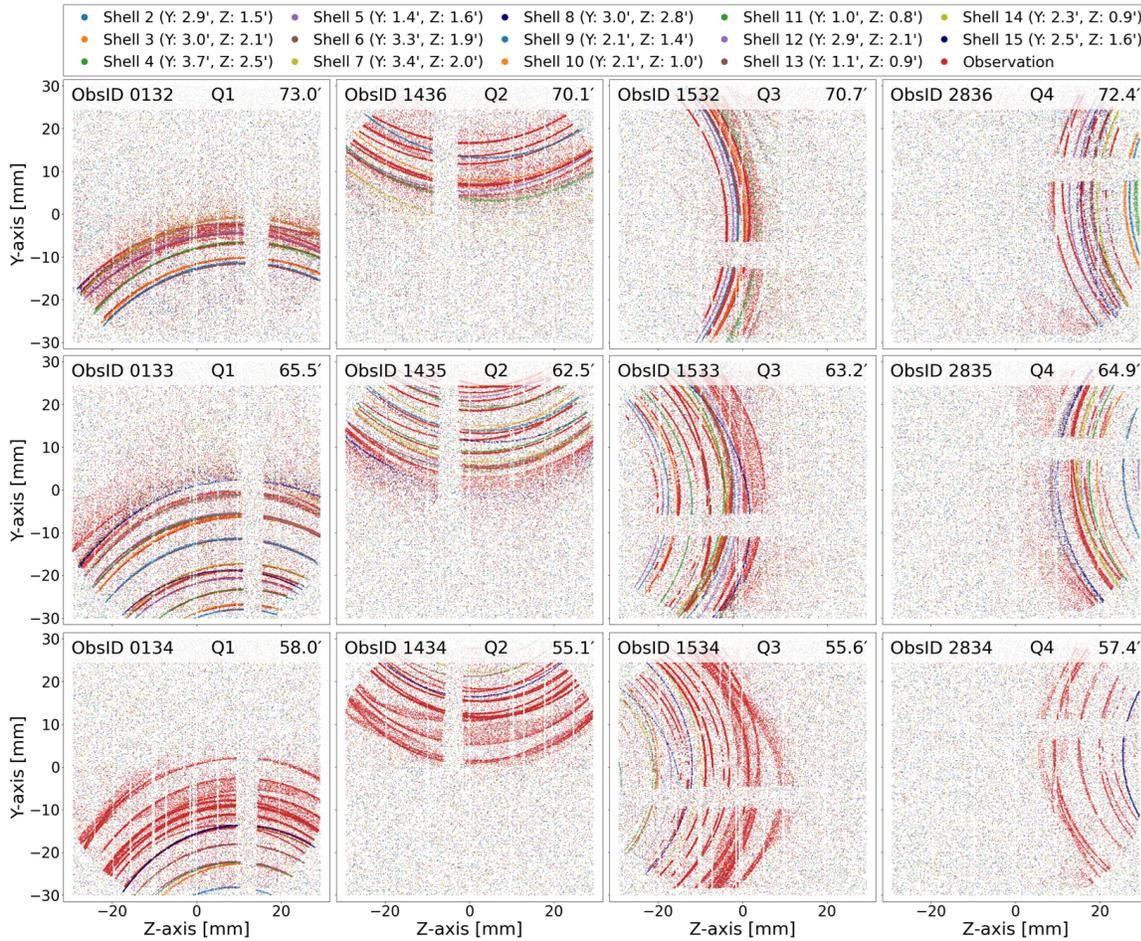


Figure 10: Comparison between observations (red) and simulations (various colors) for three pointings in each of the four quadrants. The tilts applied to individual mirror shells 2 to 15 are shown in the legend, and were found by matching the single-reflection arcs in the Q1 quadrant. While the simulated arcs match the observations well in the Q1 quadrant, the agreement is much worse in the other quadrants, potentially suggesting that the mirror distortions implemented in the FM2 model are inaccurate or that individual shells may somehow be misaligned differently on different sides of the telescope.



4.4 Replicating asymmetry by shifting X-ray baffles

Due to the challenges of matching the single-reflection arcs by tilting the mirrors, attention was finally turned to the alignment of the X-ray baffles, which act as collimators by blocking the majority of incoming stray light. By visually inspecting the observations in Figure 2, a clear asymmetry can be seen between Q4 and the other three quadrants, with the Q4 direction showing significantly weaker and less distinct single-reflection arcs, which could potentially be explained by X-ray baffle misalignments.

In order to investigate this, both X-ray baffles were manually shifted together in both the Y- and Z-directions from -0.5 mm to +0.5 mm using a step size of 0.1 mm. Visual inspections showed that a shift of approximately +0.3 mm in the Z-direction could replicate a big part of the observed asymmetry, with certain mirror shells effectively being blocked in the Q4 direction while remaining unchanged in the other directions. Figure 11 shows detector images of the seven Q4 pointings from the observations, unmodified FM2 simulations, and modified +0.3 mm Z-axis baffle shift simulations, where the simulations with the baffle shift visually correlate a lot better with the observations, particularly at lower off-axis angles.

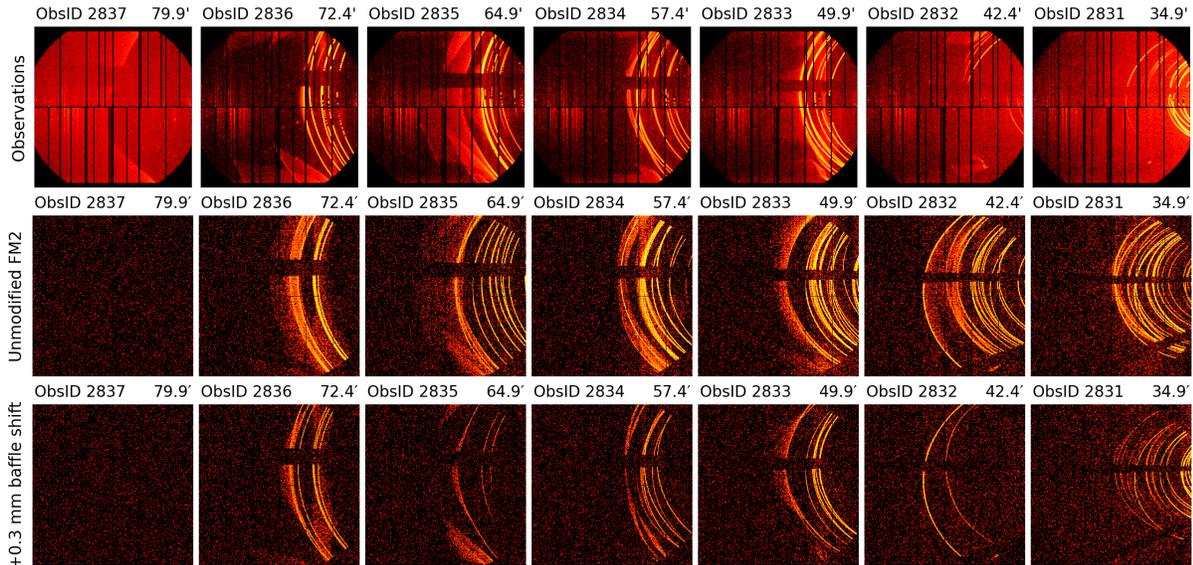


Figure 11: Detector images of the seven Q4 pointings from the observations (top), the unmodified FM2 simulations (middle), and simulations with both X-ray baffles shifted by +0.3 mm along the Z-axis (bottom). While still not perfect, the simulations with the baffle shift are found to match the observations significantly better than the unmodified FM2 simulations, especially for the lower off-axis pointings in ObsID 2831, 2832, and 2833. This suggests that a small baffle misalignment could help explain part of the observed asymmetry between Q4 and the other three quadrants.



With a reasonable shift to the X-ray baffles in the FM2 model file now determined, the effect on the nominal on-axis effective area was also computed for Z-axis shifts in the range +0.0 mm to +0.5 mm, which are shown in Figure 12. It was here found that using the assumed baffle shift of +0.3 mm, the effective area was reduced by just 2-3% depending on the energy. Like with the mirror assembly tilt, the baffle misalignment is not enough to explain the desired ~20% effective area loss, however it could potentially explain a small part of the discrepancy between XMM-Newton and NuSTAR.

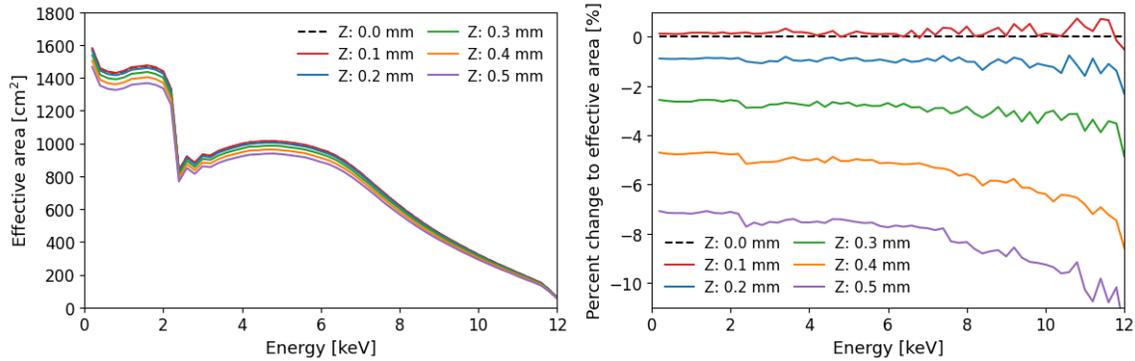


Figure 12: Simulated nominal on-axis effective area (left) and percent change in effective area compared to the unmodified FM2 model (right) for different Z-axis shifts of the X-ray baffles. Using the assumed shift of +0.3 mm, a decrease in effective area of ~2–3% depending on energy is found. While shifting the baffles affects the nominal on-axis effective area, the impact is once again small and not enough to explain the full ~20% flux discrepancy between XMM-Newton and NuSTAR.

5 Summary

The Sco X-1 analysis shows that, despite being able to produce clear single-reflection arcs, the simulations differ noticeably from the observations in both arc position, shape and intensity. Generally the arcs and structures in the simulations appeared simpler and less complex than the observations, indicating that the FM2 model is not completely accurate. This could either be due to the component alignments in the model being outdated, or due to mirror shell shapes and distortions not being modeled properly, or a combination of both. Since any physical changes to the telescope optics will also change its response and effective area, some of the limitations in the FM2 model likely contribute to the cross-calibration discrepancies with NuSTAR.

For the mirror assembly tilts, it was clear that tilting all mirror shells by the same amount was not sufficient to match the simulated single-reflection arcs with the observed ones, strongly indicating that the shells are misaligned with respect to each other. Despite this, all observations still consistently pointed towards a real systematic positive tilt of the entire mirror assembly



around both the Y - and Z -axis. However, determining a single optimal set of values was not feasible due to the variations in the tilts found for the different quadrants.

When tilting the individual mirror shells, the observed and simulated single-reflection arcs aligned almost perfectly in the Q1 quadrant being matched, but showed significant discrepancies in the other three quadrants. The simplest explanation for this is that the exact mirror shell each single-reflection arc originates from is not known for the observations. It was therefore assumed that the arcs followed the expected ordering, with the innermost arc corresponding to mirror shell 2, the next arc to mirror shell 3, and so on. However, if the mirror shell distortions or tilt misalignments are sufficiently large, which appears to be the case in the simulations, this assumption fails, and some simulated arcs could therefore unintentionally have been matched to the wrong observed arcs. This would result in the single-reflection arcs matching almost perfectly in the quadrant of interest, but failing in the others, exactly like the analysis revealed. However, if this was the sole reason for the discrepancies, it would have been highly likely that at least one of the 14 simulated single-reflection arcs would have been matched to its correct observed arc and therefore worked in all observations, which did not appear to be the case. Another potential explanation could therefore also be that the mirror shell distortions in the FM2 model are oversimplified and no longer representative of the true in-orbit distortions. Depending on the severity, this could also cause different parts of the mirrors to appear as though they are tilted by different amounts, explaining why the arcs only match in the Q1.

When investigating the X-ray baffle alignment, a simple Z -axis shift of just +0.3 mm reproduced much of the observed asymmetry between Q4 and the other three quadrants. While appearing to yield a reasonable result, this estimated shift was based purely on subjective visual inspections of the images. A more quantitative approach for the future would instead be to measure the flux ratios between each of the four quadrants, and attempt to reproduce it in the simulations using baffle shifts. Furthermore, during the analysis it was assumed that the two X-ray baffles had no tilt and were shifted together by the same amount. However, it is entirely possible that the true misalignments are due to a combination of tilts and shifts, and that the baffles are not perfectly aligned with each other. While incorporating these effects could in theory lead to better results and larger effective area decreases, it would also significantly increase the amount of free parameters for the simulations, exponentially increasing the required computation time. Furthermore, trying to distinguish between baffle shifts and tilts might also prove challenging, as these were found to be degenerate and had almost the same effect on the resulting simulations.

When quantifying the impact of the adjustments to the FM2 model, it was found that mirror assembly tilts alone reduced the nominal on-axis effective area by roughly 1–2% at most, while the baffle shift caused an additional 2–3% decrease. This leads to a combined loss of at most 3–5%, which falls significantly short of the ~20% required to explain the entire cross-calibration discrepancy between XMM-Newton and NuSTAR. Since the changes implemented in this analysis were not able to recreate all single-reflection arcs in all observations simultaneously, it is possible that additional adjustments improving the match further could reduce the effective area even more. For example, if it is found that the two X-ray baffles are shifted independently and



misaligned with each other, or if additional and more severe distortions are discovered and implemented into the FM2 model, both would likely reduce the effective area further. Despite this, it still seems unlikely that these further adjustments could account for an additional 15–17% reduction, and misalignments of the optical components therefore appear to only explain part of the $\sim 20\%$ discrepancy between XMM-Newton and NuSTAR.

6 Outlook

Looking ahead, there are several ways to improve and expand on the analysis and results in the future. The first of these is to attempt to model and update the mirror shell distortions currently implemented in the FM2 model, since these appear to be inaccurate or oversimplified, which makes matching the positions of individual single-reflection arcs very difficult. Unfortunately, these distortions are described using hundreds of parameters for each mirror shell, and updating them might therefore prove to be unfeasible. Here abandoning the shell matching entirely and instead focusing solely on the observed asymmetry and X-ray baffle alignment could also be a reliable strategy, since this appeared to be simpler and have a larger impact on the nominal on-axis effective area. Furthermore, incorporating additional off-axis observations at more intermediate off-axis and azimuthal angles of Sco X-1 would also be useful in the future. This could significantly help in determining the originating mirror shell for each single-reflection arc, as it would make it easier to follow the individual arcs around the telescope. Because of this, it is the suggestion of this TN that future Sco X-1 observations be performed at several additional azimuthal angles, so that the intermediate directions between the four currently observed quadrants are also covered. For example, this could be every $\sim 15^\circ$ to 45° rather than every $\sim 90^\circ$. Ideally, these observations would primarily focus on the single-reflection arcs appearing at the highest off-axis angles around $\sim 71.5'$, as these were found to be the simplest and most straightforward to investigate. If additional future observations prove to be unfeasible, old observations of other bright X-ray sources showing single-reflection arcs could also be used, with one example being archival pointings of the X-ray binary GX 5-1, which show clear single-reflection arcs from parts of the mirror module not covered by the Sco X-1 dataset.

Once a satisfactory result has been obtained for the EPIC-pn mirror model, it is possible to extend the same analysis using the same dataset to the mirror modules in front of the two EPIC-MOS cameras on XMM-Newton, improving their model and effective area as well. While this does come with the additional challenge of the RGS instruments diverting some photons in the optical path, it should be possible to simulate these effects in SciSim, however the exact accuracy and complexity of this remains to be seen. After the three XMM-Newton mirror models are fully understood and modeled, the same analysis and methodology could also be applied to other current in-orbit X-ray missions, such as NuSTAR and Chandra, to further improve the knowledge about the alignment of their optical components and effective area. Likewise, applying the same approach to future missions like NewAthena shortly after launch would also be especially valuable, since potential misalignments could be identified and corrected early on.



References

- Lumb, D.H., Schartel, N., & Jansen, F.A. 2012, *Opt. Eng.* 51, 011009
- Navarro, V. et al. 2020, *ASP Conf. Ser.* 527, 287
- Gulbahar, E. et al. 2025, *Astron. Comput.* 53, 100969
- Gabriel, C., Ibarra, A., & Hoar, J. 2005, *Proc. SPIE* 5898, 58981O
- Buitrago-Casas, J. C. et al. 2017, *Proc. SPIE* 10399, 103990J
- Spiga, D. 2015, *Proc. SPIE* 9603, 96030H
- Fürst, F. 2022, XMM-SOC-CAL-TN-0230 (Issue 1.1), *Empirical correction of the EPIC effective area based on NuSTAR observations*
- Bakker, J. & Jalota, L. 2005, *EPIC/RFC Camera Simulator (ESIM/RFCSIM) 4.0*
- Beijersbergen, M.W. 2005a, *Mirror Module Simulator (MSIM) 4.0*
- Beijersbergen, M.W. 2005b, *SciSim Tools 4.0*
- Beijersbergen, M.W. & Siddiqui, H. 2005, *Ray Generator for SciSim (GSIM) 4.0*
- SciSim Team 2005, *SciSim User Guide 4.0*
- Siddiqui, H. & Bakker, J. 2005, *Spacecraft Simulator (SPSIM) 4.0*