

The Rapid Optical Variability of the Nearby Radio-loud AGN Pictor A: Introducing the Quaver Pipeline for AGN Science with TESS

Journal Article**Author(s):**

Smith, Krista Lynne; Sartori, Lia F.

Publication date:

2023-12-01

Permanent link:

<https://doi.org/10.3929/ethz-b-000648548>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

The Astrophysical Journal 958(2), <https://doi.org/10.3847/1538-4357/acff5c>



The Rapid Optical Variability of the Nearby Radio-loud AGN Pictor A: Introducing the Quaver Pipeline for AGN Science with TESS

Krista Lynne Smith^{1,2} and Lia F. Sartori³

¹ Department of Physics and Astronomy, Texas A&M University, College Station, TX 77845, USA

² Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, TX 75205, USA

³ Institute for Particle Physics and Astrophysics, ETH Zürich, Wolfgang-Pauli-Strasse 27, CH-8093 Zürich, Switzerland

Received 2023 February 7; revised 2023 September 28; accepted 2023 September 29; published 2023 November 27

Abstract

The sampling strategy of the Transiting Exoplanet Survey Satellite (TESS) makes TESS light curves extremely valuable for investigating the high cadence optical variability of active galactic nuclei (AGN). However, because the TESS instrument was primarily designed for exoplanet science, the use of the satellite for other applications requires careful treatment of the data. In this paper, we introduce *Quaver*, a new software tool designed specifically to extract TESS light curves of extended and faint sources presenting stochastic variability. We then use this new tool to extract light curves of the nearby radio-loud AGN Pictor A, and perform a temporal and power spectral analysis of its high-cadence optical variability. The obtained light curves are well fit with a damped random walk (DRW) model, exhibiting both stochastic AGN variations and flaring behavior. The DRW characteristic timescales are $\tau_{\text{DRW}} \sim 3\text{--}6$ days during more quiet periods, and $\tau_{\text{DRW}} \sim 0.8$ day for periods with strong flares, even when the flares themselves are masked from the DRW fit. The observed timescales are consistent with the dynamical, orbital, and thermal timescales expected for the low black hole mass of Pictor A.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Radio active galactic nuclei (2134); Astronomy software (1855)

1. Introduction

Optical variability of active galactic nuclei (AGN) provides one of the very few direct observational probes of physics within the accretion disk. The high-amplitude, coherent variations frequently seen in ground-based AGN optical monitoring projects are ascribed to a number of phenomena intrinsic to the accretion process, such as magnetic phenomena due to turbulence and reconnection (Balbus & Hawley 1998), the damped random walk (DRW) of flux within the disk (e.g., Kelly et al. 2009), reprocessing of rapid X-ray variations in the corona, and overall variations in the mass accretion rate.

Approximately 10% of AGN exhibit powerful jets observed primarily in the radio, often extending far beyond the extent of the host galaxy and are considered an important driver of galaxy evolution via AGN feedback. Why some AGN are able to launch such jets while others do not, remains an active area of research. The ability of an AGN to produce a jet may depend on the spin of the black hole (e.g., Blandford & Znajek 1977; Wilson & Colbert 1995) or upon the accretion state of the system, analogous to X-ray binaries (e.g., Kording et al. 2006).

With this in mind, it is important to study and compare the properties of the variability, and thus the accretion physics, of radio-loud and radio-quiet AGN. To this end, we have extracted and analyzed a high-cadence, high-precision optical light curve of the nearest broad-line radio galaxy, Pictor A (PKS 0518-548), from the Transiting Exoplanet Survey Satellite (TESS). This galaxy has a pronounced double-lobed radio structure that has been known for decades (Maltby & Moffet 1962), and is the closest such source with a Type 1 optical spectrum at $z = 0.035$. Studied extensively by Simkin et al. (1999), Pictor A has a compact, parsec-scale radio

jet confirmed in very long baseline interferometric imaging, and gas kinematics indicating a strong interaction of the radio jet with the interstellar medium. The galaxy is also a known gamma-ray emitter, with variable gamma-ray flux (Brown & Adams 2012), and X-ray knots co-spatial with radio knots in the jet, which also show strong variability (Marshall et al. 2010). The black hole mass in Pictor A is likely quite low, similar to the mass of Sagittarius A*, $M_{\text{BH}} = 5.9 \times 10^6 M_{\odot}$ (Koss et al. 2022).

The TESS instrument (Ricker et al. 2015), designed primarily to search for transiting exoplanets using long-term monitoring with high photometric precision of stars across nearly the entire sky, provides optical monitoring at rapid cadence: every 30 minutes in the early cycles (2018–2019) and every 10 minutes or 2 minutes in later and current cycles (2020 to the present). Because the spacecraft does not contend with diurnal or seasonal gaps, coverage is nearly continuous and much more complete than ground-based monitoring.

These properties make TESS a potentially transformative instrument for the timing of high-energy processes like accretion onto supermassive black holes (SMBHs), magnetic flares, and jetted phenomena. However, because the TESS instrument was primarily designed for exoplanet science, the use of the satellite for other applications requires careful treatment of the data, as indicated in past studies of AGN using the very similar Kepler satellite (e.g., Kasliwal et al. 2015; Smith et al. 2018; Moreno et al. 2021).

In this paper, we introduce a new software tool designed specifically for sources with stochastic, or even nonstationary, variability and that can accommodate sources in extended host galaxies: the *Quaver* program. *Quaver* is designed to be fully interactive, transparent, and flexible; a *Swiss army knife* for the extraction and correction of TESS light curves. Novel features include clickable interactive interfaces for the selection of the extraction aperture and masking of severely compromised cadences, and a sophisticated matrix regression method



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

for the correction of background light and electronic systematics.

In the following, we first discuss the TESS observations and describe the new software package (Section 3), and explain some caveats of the method (Section 4) before undertaking a temporal and power spectral analysis of the TESS light curve of Pictor A as a demonstration of the software (Section 5). We then discuss the implications of our results in Section 6, before concluding with a summary in Section 7.

2. TESS Data

The baseline for a TESS source depends upon its ecliptic latitude, ranging from 27 days near the ecliptic plane to approximately 1 yr at the ecliptic poles. The main output product of TESS are the Full Frame Images (FFIs), from which the user extracts their own light curves. In the mission’s Cycles 1 and 2 (2018 July to 2020 June), the cadence of these FFIs was 30 minutes. In Cycles 3 and 4 (2020 July to 2022 August, the cadence was raised to 10 minutes), and in the current Cycle 5 and Cycle 6 (2022 September to 2024 September), the cadence is 200 s. The frequency of data downlinks has also increased from the earliest cycles, from 13.7 day intervals to 7 day intervals. TESS monitored the section of sky covering Pictor A as part of its routine Cycle 1 survey operations between 2018 October 18 and 2019 January 7 with a 30 minute cadence. This span of time consists of three separate sectors (Sectors 4–6) of 27 days each, for a total baseline of 81 days. Each TESS sector consists of two orbits, in between which the satellite downlinks to Earth to transmit data; such a transmission also occurs between sectors. There is therefore an ~ 1 day gap approximately every 2 weeks during the monitoring. Additionally, during Sector 4 there was an instrument anomaly that paused data collection for 3 days, and the first 3 days of Sector 6 were used for calibration with no science data collected; see the TESS Data Release Notes⁴ for detailed information.

Pictor A was further observed during Cycle 3 Sectors 31 and 33, from 2020 October 21 to 2020 November 19, and again from 2020 December 17 to 2021 January 13, with a 27 day gap in between. These observations had the new, faster Cycle 3 cadence of 10 minutes. A star-tracker anomaly caused Sector 31 to terminate science data collection 2 days early.

The TESS bandpass is very wide, with no capability to collect light curves in different bands, similar to that of Kepler (Borucki et al. 2010). In contrast to Kepler’s *white-light* bandpass, TESS instead spans a *pink* wavelength range of 600–1000 nm to best encapsulate the spectral energy distribution of M dwarfs.

3. Data Reduction with Quaver

3.1. Motivation

Data collected from satellites primarily concerned with detecting the periodic signal of exoplanet transits are not optimized for the study of stochastic variability in sources like AGN. In addition, the majority of detrending techniques used to reduce the data to search for periodic signals in planet hosts or asteroseismic targets are inappropriate for treating AGN variability, frequently resulting in overfitting.

Problematically, many of the systematic trends in space-based optical light curves, in both the Kepler and TESS data,

can mimic true AGN variability and are difficult to remove using simple background-subtraction techniques. Furthermore, the fact that many AGN host galaxies are extended sources and not point sources like stars requires special attention to be paid to the extraction aperture, rather than relying on simple point-spread function (PSF) modeling.

In answer to these challenges, we have developed the software package *Quaver*, which we have made publicly available.⁵ *Quaver* delivers an interactive, transparent, and customizable TESS extraction procedure for quasar-like variability, but can be used for any type of source that may benefit from a tailored extraction approach. The software page linked below has a full user guide available, with walk-throughs of several source cases. While not yet practical for use on large numbers of sources that can be extracted using automated optimization procedures, *Quaver* is intended for sources in which a careful and deliberate approach is required, allowing the user to inspect the removed components and explore the effect of different correction techniques and apertures on the robustness of the extracted light curve.

3.2. Aperture Selection

The extraction process begins with an interactive interface for choosing a custom extraction aperture by click-selecting desired pixels. This is an especially useful feature for AGN, which typically reside in extended host galaxies that may overlap in unpredictable ways with preset aperture shapes or complicate PSF-fitting techniques.

There are several effects that can cause the selected aperture to encompass more or less of the wings of the flux distribution as a function of time, introducing spurious variability in the light curves. These include:

1. Pointing jitter, which causes the distribution of the source PSF across the pixels and their different response functions to change;
2. Differential velocity aberration, an effect of the spacecraft’s motion around the solar system barycenter that causes the source position to drift over a sector; and
3. Thermal breathing, in which the temperature of the optics changes during the orbit and causes the PSF to become larger or smaller.

It is important to realize that although the TESS pixels are quite large, the PSF of even a point source typically extends into nearby pixels. Indeed, the PSF varies across the field of view of the detector, from ~ 0.9 pixels near the center to as much as 2.76 pixels near the edges (Oelkers & Stassun 2018). It is therefore recommended that the user attempt to encompass the *entire host galaxy* in the aperture, if it is feasible. In fields that are not crowded, it can be useful to include a buffer pixel around one’s source; however, for faint sources, this can introduce significant background noise that may substantially dilute the light curve. It is also important to realize that increasing the aperture size will dilute the variable signal from the nucleus relative to the baseline flux of the light curve, changing the overall percent of variability.

In general, it is best to make the aperture sufficiently large to encompass the host galaxy flux and its likely spillover into nearby pixels, while not encompassing so much background as to raise the noise level of the light curves beyond what your desired variability threshold can handle. Experimentation with different apertures

⁴ https://archive.stsci.edu/tess/tess_dm.html

⁵ <https://github.com/kristalynnesmith/quaver>

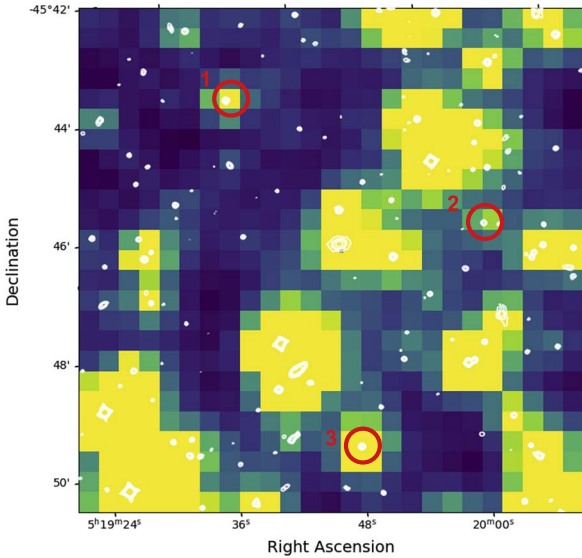


Figure 1. TESS target pixel file image of the Pictor A field of view, with contours from the DSS overlaid. The three reference stars extracted by the pipeline are labeled in red circles; their light curves are shown in Figure 3. The scale is $21''$ per pixel. Pictor A is the galaxy at the center of the frame.

around extended sources is recommended: generally, a well-chosen aperture will result in an output light curve that retains its shape, and only grows more noisy, when additional pixels are added (assuming all other extraction parameters remain the same). An aperture that is too small will result in false variations when drift or thermal effects cause portions of the source flux distribution to enter or leave the aperture. An aperture that is too large will be excessively noisy, and of course, an aperture should not be so large as to include other sources.

The custom aperture also allows the user to avoid nearby contaminants or determine the variability of nearby stars or background regions. The user selects pixels for inclusion in the aperture from a clickable map, generated by TESSCut (Brasseur et al. 2019). The process is assisted by an overlay of contours from the Digitized Sky Survey (DSS) atop the very large $21''$ TESS pixels, so that the user can identify any potential contaminants and better assess the extent of the host galaxy. An image of the type of map shown during this selection is given in Figure 1.

3.3. Systematics Correction

Once the aperture is selected, the Quaver program utilizes both principal component analysis (PCA) and matrix regression to account for systematics. It makes frequent use of algorithms developed for Lightkurve, originally developed for the Kepler mission and now adapted for TESS analysis (Lightkurve Collaboration et al. 2018).

3.3.1. Classification of Pixels

Once the aperture is chosen, the pixel map is divided into three regions: the source aperture plus a 1 pixel border buffer zone (to prevent light from the source from leaking into pixels used to build correction vectors), *faint* background pixels, and *bright* pixels that are likely to contain other sources. This cut is made at 1.5σ above the median flux of the whole pixel map by

default, but can be changed by the user if desired. In a crowded field, the number of background pixels will be fewer than in a field that is mostly empty. If the field is very crowded, you may wish to increase the 1.5σ threshold, otherwise, several actual faint background sources may be included in the *empty* pixels.

3.3.2. Cadence Masking

PCA is performed on the faint pixels. This next step allows the user to see what the dominant effects of this background look like. Quaver will plot the eigenvalues of the principal components of the faint pixels. The tunable parameter `sys_threshold` sets the value that determines whether this plot appears; it is set at the arbitrary value of 0.2 by default, but can be altered. If any of the components exceeds this value at any point in the light curve, the user will be asked whether they wish to interactively mask out periods of monitoring. This is entirely optional. The user can then click-select regions to exclude from the light curve (and fitting) moving forward. This is frequently useful, as there are many TESS sectors where brief periods of spacecraft attitude corrections, erroneous pointing solutions, or necessary thruster firings cause rapid, erratic behavior that is very hard or impossible to remove. The beginnings of the two orbits in each sector are prone to thermal effects, as the spacecraft recovers equilibrium after data downlinks at orbit perigees; the beginnings and ends of orbits are especially prone to increased scattered light from Earth. Quaver removes these gradual effects successfully in most cases; however, masking these edge cadences can be useful to prevent spurious behavior resembling partial flares or rising/falling flux near sector edges from remaining in the light curve.

The user should think carefully about the effects of removing portions of the light curve with masking. For example,

1. Removing significant portions of the light curve edges reduces the overall baseline, shortening the frequency range to which the analysis is sensitive. For AGN, this can mean the inability to detect low-frequency turnovers.
2. Removing various stretches of the light curve throughout renders the TESS sampling much more uneven, meaning that methods such as the Lomb–Scargle periodogram are more appropriate than, say, interpolation over the orbit gap followed by simple Fourier power spectral analysis.

A conservative approach to masking is to decline the option to mask at first, and allow the program to try and correct the systematics. If the final output light curve remains visibly affected by a known issue from the Data Release Notes or serious thermal edge effects, then the program can be rerun with the masking option utilized.

3.3.3. Matrix Regression Correction

Once the aperture is chosen and any desired cadences are masked out, Quaver moves on to the systematics correction itself.

The Quaver pipeline has three options for systematics correction. In order of complexity and aggressiveness of fitting, they are a simple PCA method, a simple hybrid method, and a fully hybrid method. The hybrid methods are so-named because they handle the additive background effects and the multiplicative instrumental effects separately, in a hybrid approach.

The Quaver User Guide available with the code offers a detailed description of each reduction method and a chart explaining their workflows. The best choice of method depends upon the user’s science case. In general:

1. The simple PCA version is the least rigorous but most intuitive and similar to existing methods in the literature. This method is prone to both overfitting long-term trends as well as missing subtle, typically electronic multiplicative effects. In this method, all pixels (regardless of brightness or faintness) outside the source aperture and its buffer zone are fit by a number of principal components set by the user. These components are stored in a design matrix, which is then used by `Lightcurve's RegressionCorrector` class to correct the light curve extracted from the source aperture.
2. The simple hybrid method handles the background effects and instrumental systematics in the *bright* source-dominated pixels separately, but does not apply a fit to remove the background light from the source itself, instead using simple background subtraction. The light curve of the *faint* pixels (as defined by the user-adjustable 1.5σ threshold) is extracted directly with those pixels as the aperture, scaled multiplicatively to reflect the difference in the number of pixels between the background and source apertures, and then subtracted from the source light curve. Since no fitting is performed, this is less likely to remove true long-term variability, which can coincidentally mimic the source's true behavior. Once the background is subtracted, the principal components of the faint background pixels are calculated and used to correct the bright pixels (excluding the source and buffer zone). Then, the principal components of the background-corrected bright pixels are calculated, and stored in a design matrix which is used by `RegressionCorrector` upon the background-subtracted source light curve, accounting for the multiplicative systematics. Light curves extracted using this method show excellent similarity to simultaneous (but more sparsely sampled) ground-based light curves.
3. The fully hybrid method is the most aggressive. Like the simple hybrid method, it first creates a design matrix housing the principal components of the faint background pixels, uses this matrix to correct the bright pixels, and then derives the principal components of the background-corrected bright pixels and stores them in a design matrix for the multiplicative systematics. However, instead of performing a simple background subtraction, it combines the additive background and multiplicative systematics design matrices into a hybrid `DesignMatrixCollection` object, which `RegressionCorrector` then uses to correct the light curve extracted from the source aperture. This method does a very thorough job of removing all systematics, but can overfit long-term behavior in some cases (which is generically true of correction at a higher order than background subtraction). This method is most appropriate for users who wish to study detailed high-frequency behavior, such as flare modeling or searching for rapid quasi-periods, since the subtle effects of electronic crosstalk noise and other high-frequency systematic signals are removed in a sophisticated fashion. This method (along with the PCA method) is risky for users interested in low-frequency power, especially in cases where light curves of multiple sectors will be stitched together.

Figure 2 shows the results of each of the three correction methods compared to ground-based light curves of blazars with good simultaneous sampling by the Zwicky Transient Facility (ZTF; Masci et al. 2019). In order to plot the ZTF and TESS light

curves together, some rescaling is required. We convert the ZTF light curves from magnitudes to a generic flux via $F = 10^{m/-2.5}$. We then locate data points from the TESS light curve that correspond to the ZTF sampling times, and minimize the function $\sum_{t=t_{\text{TESS},i}}^{t_{\text{TESS},f}} (F_{\text{TESS}} - A * F_{\text{ZTF}})^2$ to find the best-fitting scale factor A (where $t_{\text{TESS},i}$ and $t_{\text{TESS},f}$ are the first and final timestamps of the TESS monitoring). This multiplicative factor is applied to the generic ZTF flux. One can see that the simple hybrid method best preserves long-term trends while the fully hybrid and PCA methods can result in a flatter light curve. Although in many cases these methods do not overfit (as in PKS 0048-09), other objects experience various degrees of overfitting with methods in which the background is treated using principal component regression (as in 1ES 1011+496, significantly). This overfitting occurs when the principal components have too many degrees of freedom, allowing them to frequently match the object's true variability, which is then removed. Each of the light curves shown in Figure 2 was fit using three principal components; the user of `Quaver` is free to use fewer or more depending on their preferences.

As an indication that the pipeline is capable of removing instrumental and background systematics from targets, we have also extracted the light curves of three stars in the same field of view as Pictor A, labeled in Figure 1. Their light curves for Cycle 1 and Cycle 3 are shown in Figure 3. Each star was extracted with the same `Quaver` parameters as Pictor A. The stellar light curves are flat, as expected for inactive stars, with very few residual effects remaining. Occasional single-cadence spikes remain in all pipeline methods; the simple hybrid method is prone to leaving slight ($<1\%$) ramps across the light curve.

The most dominant term removed by the PCA components is always the scattered light background (see Appendix A). Additional terms are caused by all of the effects listed in Section 3.2. Occasionally, a bright nearby source with strong variability may contaminate the target light curve due to overlapping pixel response functions. In this case, `Quaver` is capable of removing the contaminant trend (Appendix B).

3.4. Multisector Light Curves

If the light curve spans multiple sectors, then each sector is displayed to the user in turn for custom aperture selection, cadence masking, and extraction. At the end of this process, the sectors are additively stitched together into one light curve, simply by matching the mean of the 10 fluxes at either side of the sector gap. Although this method is not ideal, in that it ignores true variability that occurred within the sector gaps, it is the least likely to introduce spurious variability. Tests of more complex methods to model the long-term behavior using, for example, multinoded splines, resulted in considerable added variability not innate to the source. So, this simplistic stitching remains our current method of choice. Indeed, many authors working on a number of other science applications have also contended with this problem (e.g., Anthony et al. 2022; Avallone et al. 2022; Claytor et al. 2022). In any case, each sector's light curve is also output separately from the full stitched light curve, so the user is free to analyze them separately or stitch them using different methods, as they prefer.

3.4.1. Assessing the Correction

The main output panel for a single-sector `Quaver` run is shown in Figure 4. It illustrates the corrected light curve, the additive and multiplicative components identified in the PCA fits, and the chosen aperture.

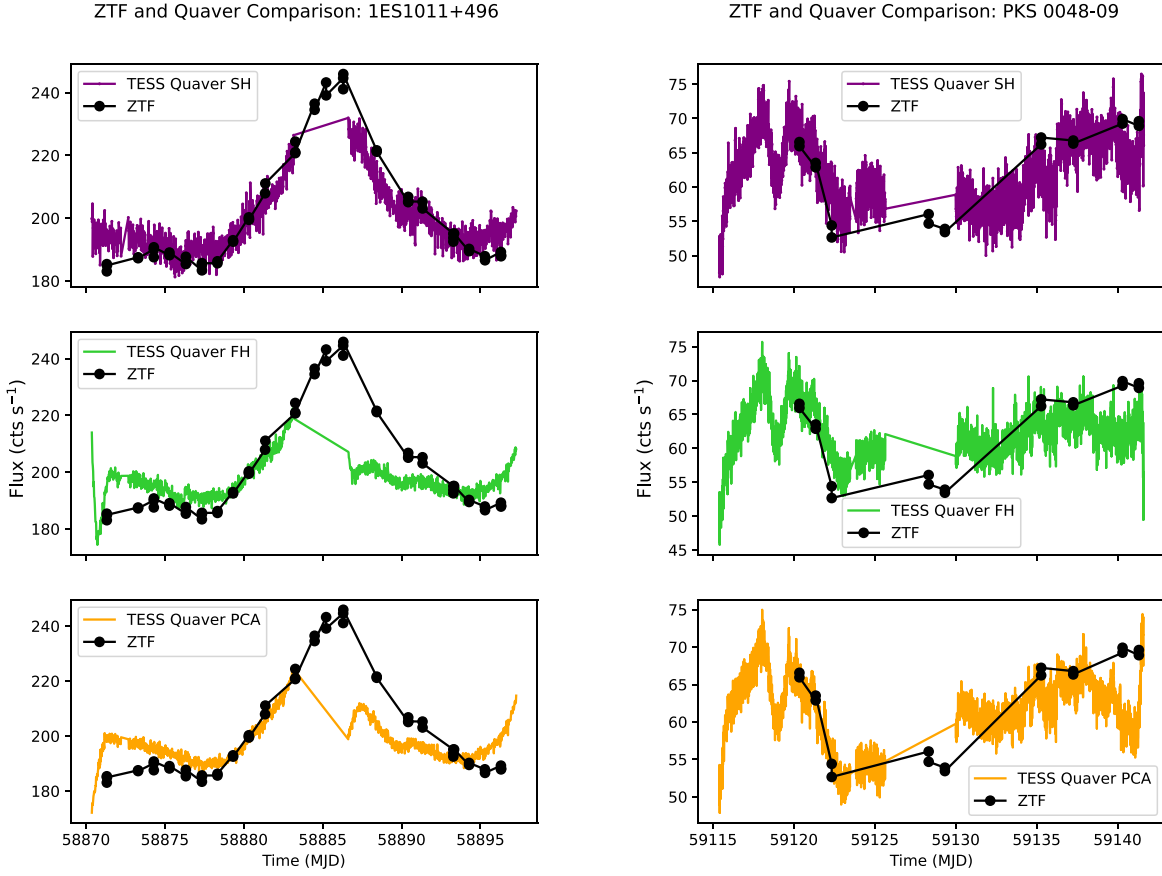


Figure 2. Comparison of the three *Quaver* reduction methods (colored light curves as indicated in legend) and ground-based light curves from ZTF (black). Light curves have been rescaled to have similar normalizations for comparison of long-term behavior. Conversion of TESS counts to standard magnitude systems is not straightforward due to its unusual bandpass; instead, ZTF magnitudes have been converted to generic flux counts via $f = 10^{m/-2.5}$ and rescaled to a similar normalization as the TESS light curves; see the text for details.

We include two further outputs from the program to assist users in deciding which method is best for their target, and the overall quality of the correction.

First, in addition to inspecting the correction vectors used in the fit in the standard output panel, it is useful to see the final result of the combined fit that was ultimately removed from the light curve. *Quaver* produces three output plots, showing the uncorrected light curve, the corrected light curve, and the final form of the fit that was removed (i.e., the linear combination of the constituent vectors with the coefficients that resulted from the fit). Examples of all three methods are shown in Figure 5 for Sector 5. Corrections for the other sectors can be found in Appendix C. We also show these results for two of the stars and sectors from Figure 3.

All three pipeline versions remove similar long-term trends; the background light curve subtracted by the simple hybrid method is the dominant term in the total fit in the other methods, as expected. One can also see that the variable star that dominates the second multiplicative correction vector (Figure 4) has not received any significant weighting in the final fit (it cannot be seen in the multiplicative model); therefore, this signal was not injected into the source light curve.

The second output useful for assessing the fit is the overfitting metric defined in *Lightkurve*,⁶ which compares the broadband power in a Lomb–Scargle periodogram before and after correction. This metric value is output to the terminal after each

sector is completed, for each reduction method. Higher values are better; a metric of 0.5 indicates that power from injected noise due to fitting is equal to power from light curve uncertainties. The documentation for *Lightkurve* indicates that values above 0.8 mean that no significant overfitting has occurred. The user can compare the overfitting metric of the three *Quaver* methods in determining which to use, or simply to assess the amount of overfitting present in their chosen output. We present the results of these metric tests in Table 1.

4. Cautions and Tips

4.1. Spacecraft Systematics

We caution that *Quaver* is not able to completely remove and identify all spacecraft systematics. To the extent that pointing jitter, velocity aberration, and thermal breathing affect all sources in the field in a similar way, *Quaver* can identify and remove these trends. However, because AGN hosts are typically not point sources (while most of the field sources used to determine the correction vectors are likely to be stars), and because individual pixels are not identical in their responses across the detector, traces of these effects can remain, especially in sources where a small extraction aperture was required.

4.2. Underfitting

The reduction of any light curve involves striking a balance between overfitting and underfitting, an especially challenging

⁶ <https://docs.lightkurve.org/tutorials/2-creating-light-curves/2-3-how-to-use-cbvcorrector.html>

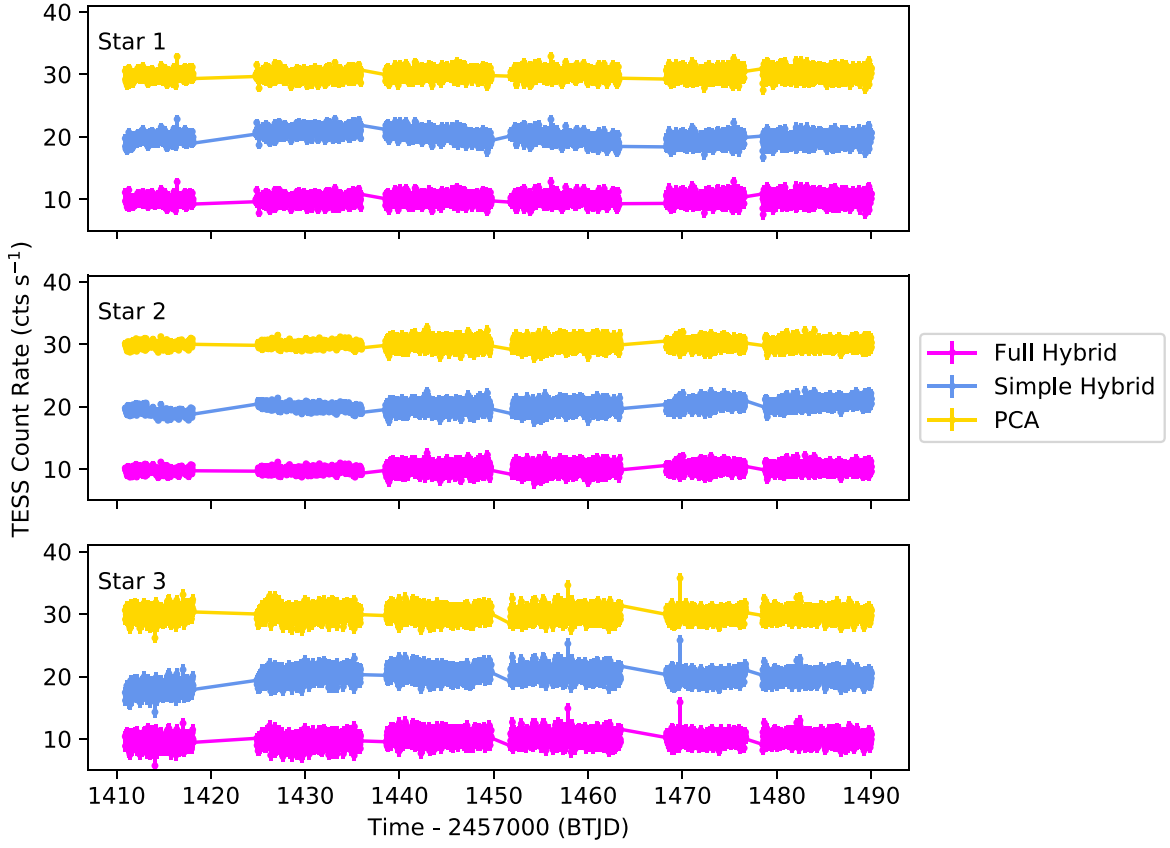


Figure 3. Cycle 1 light curves of three non-variable reference stars extracted by the *Quaver* pipeline. Each panel shows the light curves of stars 1–3, as labeled in Figure 1, from top to bottom. In each panel, all three *Quaver* methods are used, with colors given in the legend. Light curves have been arbitrarily offset for visible distinction. Differences in the noise spread of the light curves is due to the different apertures used to extract the stars and variable level of background light from sector to sector.

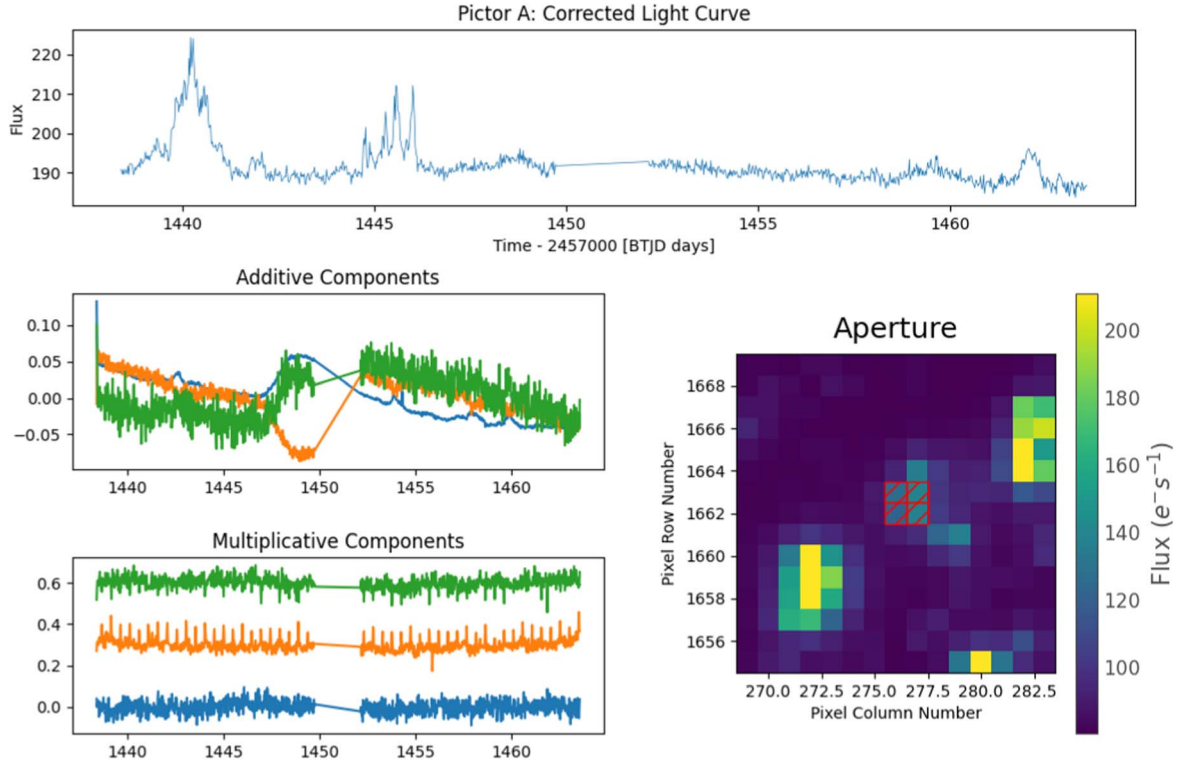


Figure 4. Example *Quaver* output panel for Pictor A in Sector 5. Top: corrected light curve created with the simple hybrid method. Bottom left: additive and multiplicative components removed by the hybrid method. Different colors correspond to different PCA components. Bottom right: Postage stamp of the TESS Target Pixel File. The source is well detected in the center of the image, and the extraction region is cross-hatched in red. The pixel scale is $21'' \text{ pixel}^{-1}$.

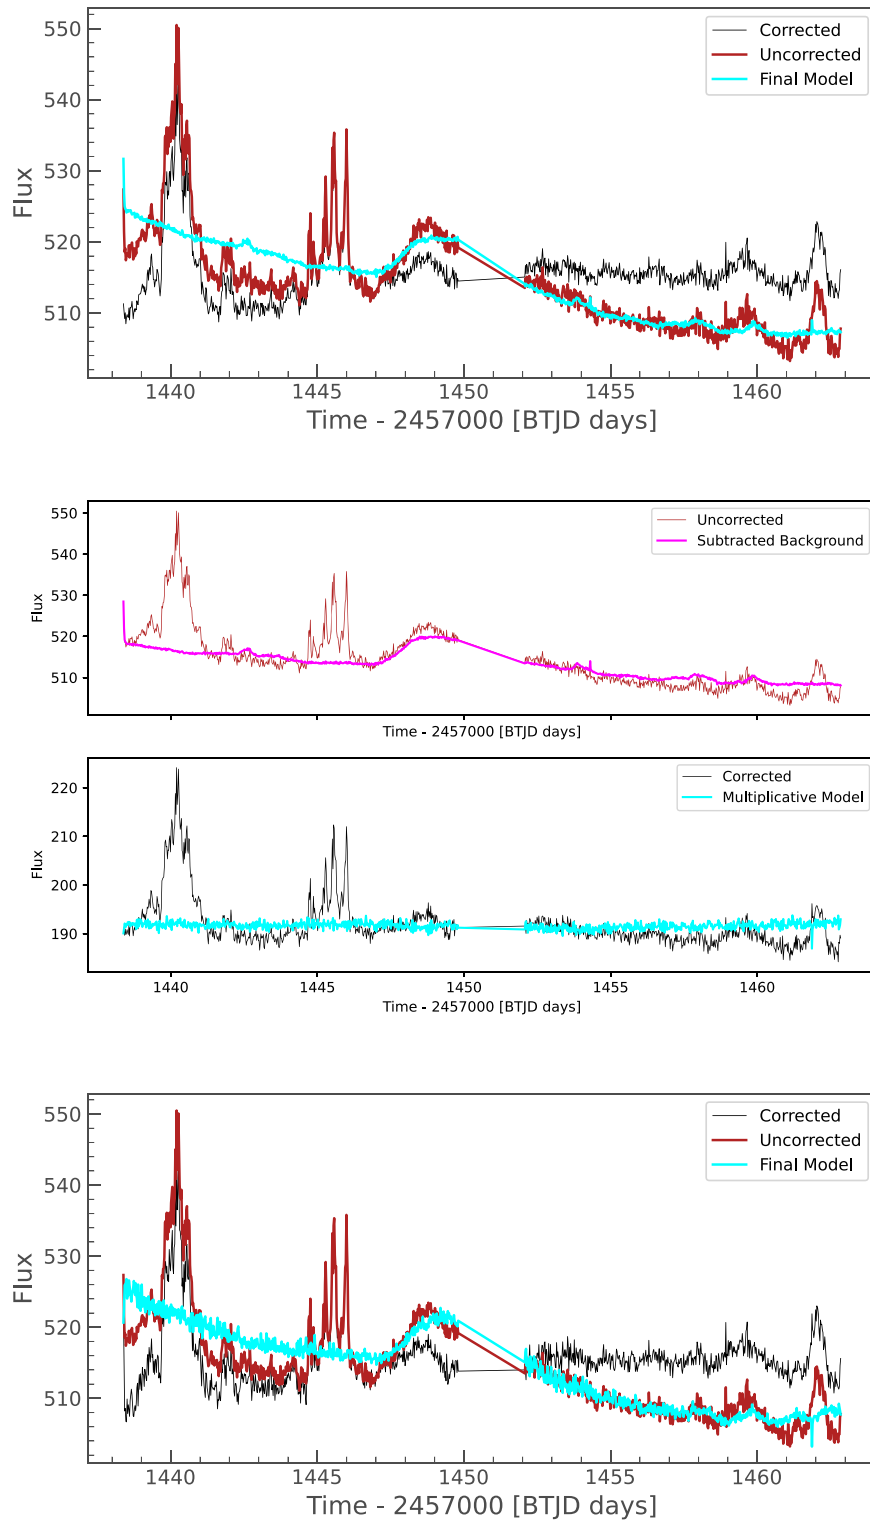


Figure 5. Fit output diagrams for Sector 5 for each of the three methods. Top: PCA fitting only. Middle: simple hybrid model, with subtracted background light curve (scaled to match the median flux of the raw light curve) shown with the original in the top panel, and the background-subtracted light curve shown with the corrected multiplicative model in the lower panel. The amplitude of the subtracted background can be estimated by comparing the raw and background-subtracted curves. Bottom: fully hybrid model.

problem to diagnose when your target varies stochastically and without quiet periods or predictive behavior. While using fewer PCA vectors can reduce the chances of overfitting, using too few can result in underfitting, in which systematics remain in the light curve of *Quaver*'s methods, the simple hybrid approach is the most prone to underfitting long-term trends, as a simple

background subtraction is more likely to leave residual trends in the light curve than a fitted approach.

There is an underfitting metric provided by *Lightkurve*, alongside the overfitting metric mentioned in the previous section. However, it compares the output light curve to data reduced by the TESS SPOC, and warns that it is inappropriate

Table 1

Overfitting Metrics Obtained Using the Comparison of Broadband Power from Lomb–Scargle Periodograms

Sector	Simple PCA	Simple Hybrid	Fully Hybrid
4	0.73	0.86	0.74
5	0.65	0.88	0.52
6	0.82	0.91	0.79
31	0.84	0.81	0.83
33	0.80	0.81	0.70

Note. Higher numbers indicate better fits with a lower likelihood of overfitting; numbers over 0.5 are acceptable, while numbers over 0.8 are good.

to compare these to light curves reduced in other ways. These targets are also almost certainly much brighter than the majority of AGN, so we do not include this metric in *Quaver* at present. Future versions may include a custom underfitting metric designed for *Quaver* output.

4.3. Signal Injection Using PCA

In some rare cases, if a strongly variable source in the field is the most dominant component in the regression matrix, such a signal could theoretically be *injected* into the light curve during the correction. It is important for the user to inspect the correction vectors in the output image (e.g., Figure 4), as well as the form of the final fitted model that will be removed from the light curve. It is occasionally the case that a variable target will manifest in the multiplicative components, as in the orange vector in the bottom panel of Figure 4. However, when the set of vectors is regressed against the light curve, if these signals are not present, those vectors are severely downweighted. This can be seen by inspecting the final models shown in Figure 5, which do not include this periodic component.

4.4. Effects of Scientific Analyses of AGN

The purpose of this pipeline is to create a user-friendly tool that allows for flexibility and transparency in dealing with TESS data. It is not intended to automatically create trustworthy light curves using the default parameters.

The process, therefore, includes a number of subjective steps (aperture selection, cadence masking, number of principal components, type of reduction method, etc), which are often efficiently decided automatically using predefined criteria in other TESS pipelines (for example, *eleanor*; Feinstein et al. 2019). When making use of the additional flexibility of *Quaver*, the user must think carefully about the effects of each reduction choice on the desired scientific result.

Using too many principal components is likely to result in overfitting, especially of low-frequency behavior from long-term trends. This has a major effect on the shape of the power spectrum at low frequencies, and can even induce a break timescale where one did not exist before. Because these break timescales are often the goal of AGN timing, as they potentially offer an independent measure of the black hole mass or accretion behavior, it is extremely important to avoid overfitting gradual trends as much as possible. This is the motivation behind our choice of the simple hybrid method for this work, as it is the least likely to overfit at low frequencies.

Underfitting also affects the shape of the power spectrum, by leaving power at many frequencies that should have been removed. *Quaver* is effective at removing high-frequency

behavior even in the simple hybrid method, so underfitting at high frequencies is unlikely, but very short-term spikes or jitter can contribute to high-frequency power and artificially shallow the slope of the PSD.

Finally, we stress that TESS is an instrument of relative photometry. In cases like AGN, when the a priori behavior is difficult or impossible to model and the ground truth may not be known, and especially when detrending methods and subjective aperture sizes can affect the baseline flux, it is important to avoid using the data for science cases requiring *absolute* photometry, such as estimating intrinsic AGN luminosities. In these cases, it is important to use ground-based data like ZTF to normalize the light curve accurately before proceeding.

4.5. Light Curve Errors

The photometric errors on the raw extracted light curve are propagated for each cadence from the errors of the flux in each pixel. The errors given in the final, corrected light curve can either be propagated through the reduction and correction process, or can simply reproduce the original photometric errors from the raw light curve. The tunable parameter `prop_error_flag` allows the user to decide whether to allow *Lightcurve* to propagate the errors through the matrix regression process, or to retain the original errors. Propagation increases the program runtime.

For the purposes of this paper, the original photometric errors are maintained, as this is the typical *Lightcurve* use case. The magnitude of the errors does not change by more than 0.2 dex, in the case of Pictor A, when the errors are propagated.

5. Variability Measurements

The light curves of Pictor A as extracted with *Quaver* with all three methods are shown in Figure 6. We extracted these light curves using *Quaver*’s default fitting parameters and our interactively selected 2×2 pixel extraction aperture (Figure 4). Our present science application is a low-mass Seyfert galaxy, with expected variations on timescales of days to weeks. Based on the reasoning in the preceding section, we have selected the simple hybrid method. Table 1 indicates that the simple hybrid method is least likely to overfit this target, as it is the only method with metrics above 0.8 in all sectors. We are interested in not only the high-frequency slope of the power spectrum, which is well preserved in all methods, but also whether the power spectrum turns over toward low frequencies at some characteristic timescale. The risk of overfitting low-frequency trends that is evident in the fully hybrid method and removing intrinsic low-frequency power is not appropriate for such an investigation. Although the simple hybrid method can be prone to underfitting these trends, inspection of the three different methods in Figure 5 indicates that the simple hybrid method does a good job of removing the same low-frequency behavior as the simple PCA method.

In the following sections, we describe different analyses we performed in order to characterize the variability we observe in the TESS light curves of Pictor A. Because of the uncertainties arising when trying to stitch together the light curves obtained for different sectors, we consider each sector separately, therefore concentrating on timescales up to 26 days (in the galaxy’s rest frame).

5.1. Flux Distributions and the Rms–Flux Relation

Figure 7 shows the flux distributions for the individual sectors, which we attempted to fit with Gaussian, bimodal, and log-normal

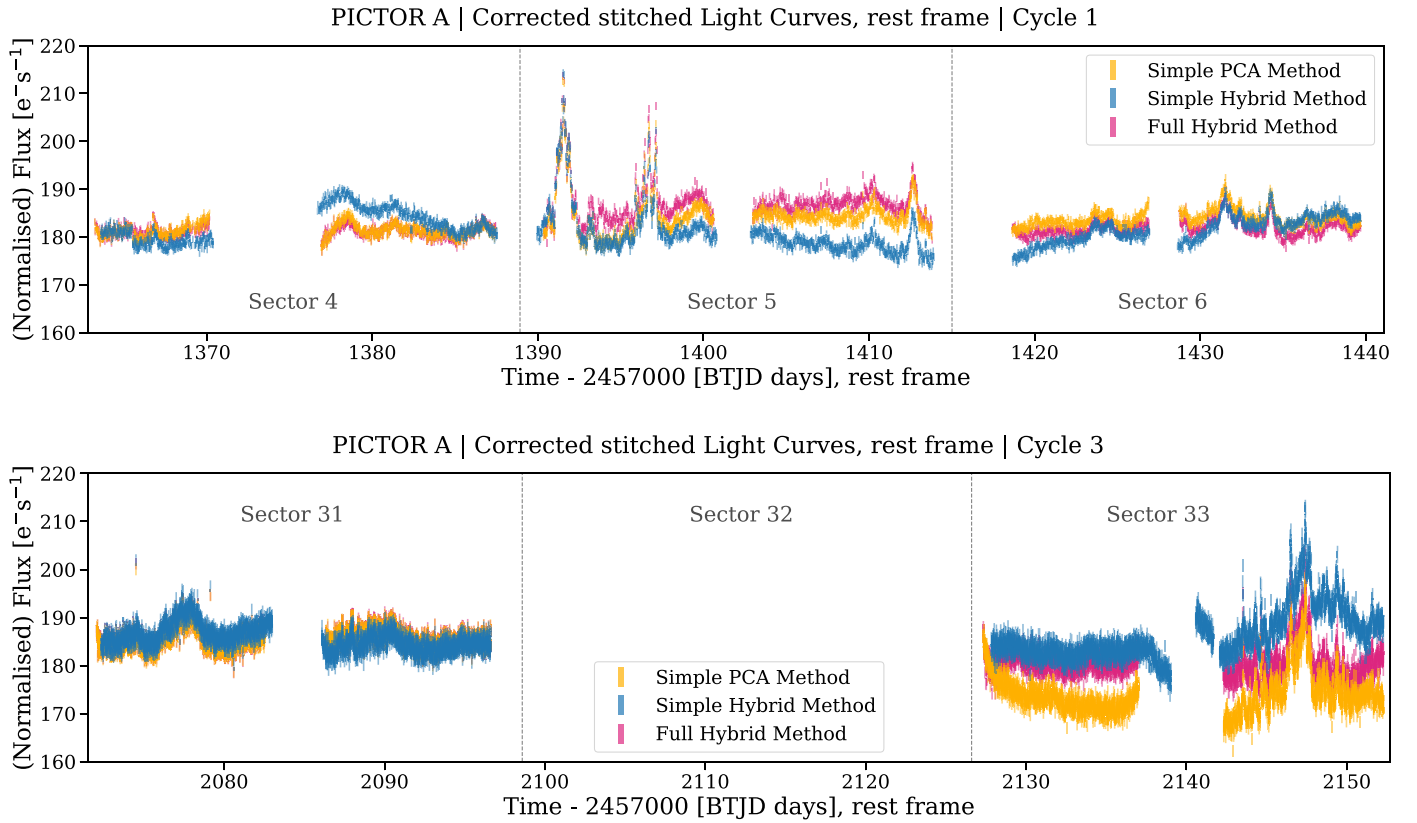


Figure 6. Stitched light curves of Pictor A in Cycle 1 (Sectors 4–6; top panel) and Cycle 3 (Sectors 31 and 33; bottom panel). The time axis has been shifted to the galaxy’s rest frame. The colors correspond to the different extraction methods: simple PCA (orange), simple hybrid (blue), and fully hybrid (purple). In addition to the normalization applied for the stitching (see the text for details), the simple PCA and fully hybrid light curves are artificially normalized by subtracting a factor of ~ 285 to match the simple hybrid light curve. We notice that while the main features (e.g., flares) are similar in each light curve, the simple hybrid one shows more variability at timescales comparable to the sector’s length. For the following analysis, we will use the light curve extracted with the simple hybrid method.

distributions, as these have all been observed for AGN light curves obtained at different wavelengths (e.g., Uttley et al. 2005; MacLeod et al. 2010; Smith et al. 2018; Bhatta & Dhital 2020). Log-normal distributions are indeed predicted by multiplicative fluctuation models, but are *not* themselves necessarily indicative of such underlying processes (Scargle 2020). Gaussian and bimodal distributions may be observed in the case that multiple components are contributing to the observed variability. While Sectors 5, 31, and 33 are well fit with a bimodal distribution, the flux distribution in Sectors 4 and 6 is less well defined. This may arise from the fact that in these sectors more cadences had to be excluded because of systematic effects (see Section 3), and the number of observations is therefore too low to confidently probe the underlying flux distribution.

Similarly to what was observed in Kepler AGN light curves (e.g., Smith et al. 2018), our data do not show any evidence of the rms–flux relation exhibited by the X-ray light curves of AGN (e.g., Vaughan et al. 2003; McHardy et al. 2004) and X-ray binaries (e.g., Uttley & McHardy 2001).

5.2. Periodograms

AGN periodograms (or power spectra) are usually well fit by either a single or broken power law, where the break timescale may be associated with the characteristic timescales of the underlying variability process. Since the extracted TESS light curves are not evenly spaced (because of the observing gaps and excluded cadences), we computed the periodogram of each

sector following the generalized Lomb–Scargle definition (Cumming et al. 2008 with formalism from Zechmeister & Kürster 2009), which is less sensitive to gaps in the light curve compared to standard Fourier techniques. In Figure 8, the measured Lomb–Scargle periodogram of each sector is shown in gray, while the black lines show the corresponding binned periodograms. As we will better quantify below, each periodogram shows a bent power-law behavior at low frequencies and flattens at higher frequencies due to Poisson noise.

Modeling the power spectrum of short light curves in the frequency domain is complicated by spectral distortions such as red noise leak (Deeter & Boynton 1982; Deeter 1984) and aliasing (Kirchner 2005). We, therefore, prefer to model the light curve directly in the time domain assuming a DRW⁷ (Kelly et al. 2014) model using the `EzTao` Python package (Yu & Richards 2022), and to compute the corresponding power spectrum, which we then compare to the Lomb–Scargle periodogram as a sanity check, as done in, for example, Burke et al. (2020). An example of DRW fit results, including residual and posterior distributions, is shown in Appendix D. We further tested that the DRW is a good model for our data by applying a Durbin–Watson test (Durbin & Watson 1950, 1951) to verify that the residuals are not autocorrelated. The Durbin–Watson test returns values between $d = 2$ and $d = 2.5$ for four

⁷ We note that the DRW model corresponds to the first-order continuous autoregressive model, CARMA(1,0), for irregularly spaced observations.

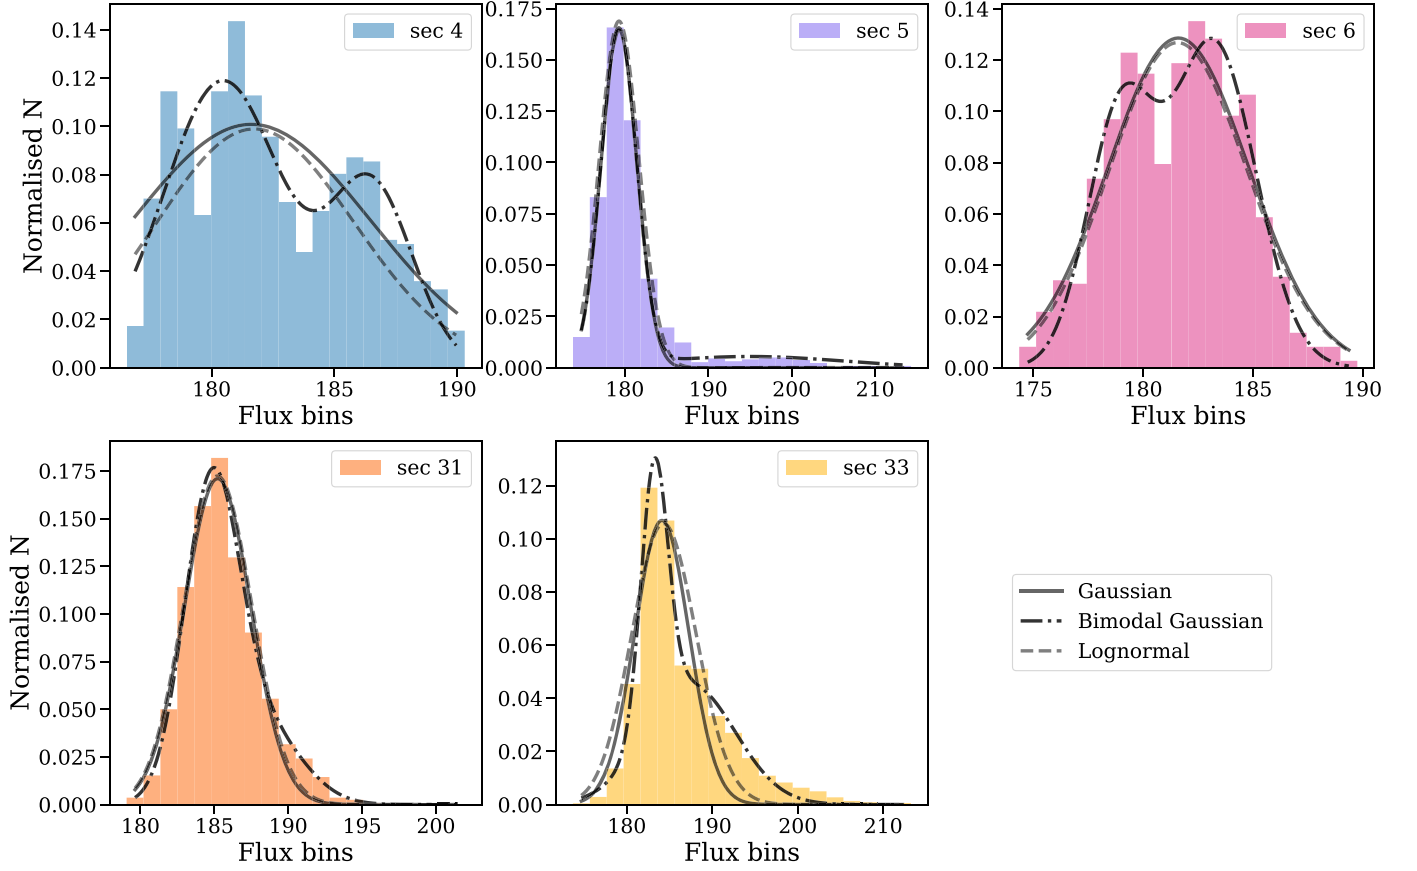


Figure 7. Flux distribution histograms in the individual sectors with Gaussian (solid line), bimodal Gaussian (dashed–dotted line), and log-normal (dashed line) fits.

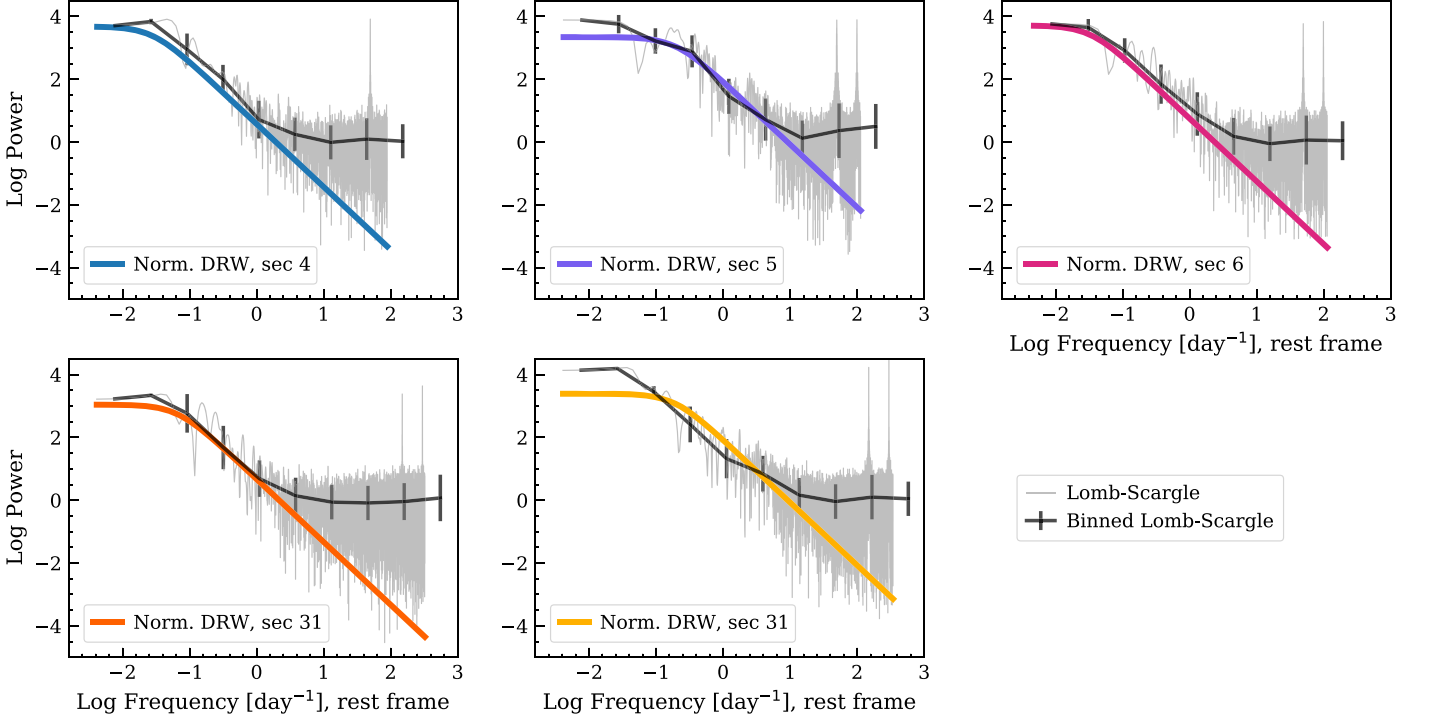


Figure 8. Lomb–Scargle periodograms computed for the sectors separately. Gray shadow: total computed periodograms. The peaks at the highest frequencies are aliases due to the time resolution of the TESS observations (30 minutes in Cycle 1 and 10 minutes in Cycle 3). Black line: binned Lomb–Scargle periodograms. Here, the error bars represent the scatter within each bin. Colored lines (same color coding as Figure 7): normalized periodogram for the underlying DRW model. These models are more reliable than the Lomb–Scargle estimates (see Section 5.2 for more details).

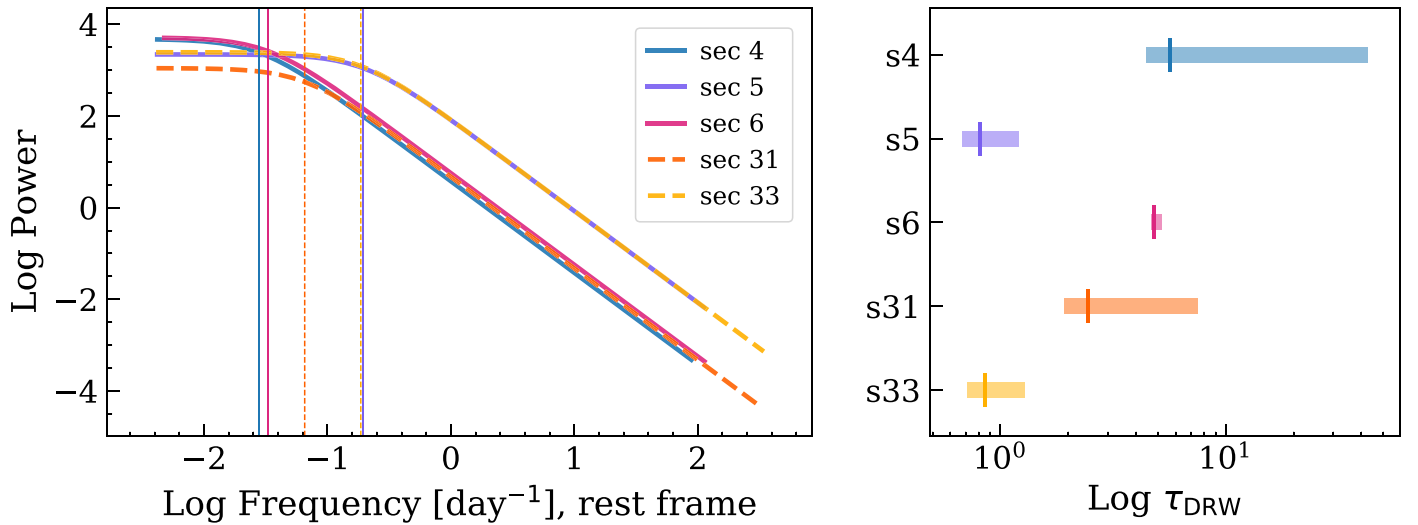


Figure 9. Left: periodograms from the DRW fit of the different sectors. Right: corresponding τ_{DRW} . Depending on the sector, τ_{DRW} ranges between 0.8 and 5.5 days.

out of five of our TESS Sectors, and $d = 2.8$ for the remaining Sector 5. As a rule of thumb, $d = 2$ (or close to this value) indicates no autocorrelation, while $d > 3$ ($d < 1$) may be a cause of alarm for negative (positive) autocorrelation. Since the values we obtain for our data are in the range $2 < d < 3$, we are confident that the residuals are not (strongly) autocorrelated. This, together with the fact that the residuals are normally distributed, supports our hypothesis that the DRW is an acceptable model for our data. The obtained periodograms are shown in Figures 8 and 9, along with the corresponding characteristic timescales τ_{DRW} . The DRW slope in the regime which is not dominated by Poisson noise and/or spectral distortions is consistent with the slopes of the Lomb–Scargle diagram in all the sectors. In addition, we observe that the τ_{DRW} values are different for Sectors 5 and 33, in which some flares are prominent, and Sectors 4, 6, and 31, which are mostly devoid of flares: $\tau_{\text{DRW}} \sim 0.8$ day for Sectors 5 and 33, a τ_{DRW} of ~ 3 –6 days for the others. This may signify that the AGN is switching between two different variability modes, as we will discuss in Section 6. It is important to notice that τ_{DRW} can usually be fully trusted only if it is shorter than 10% of the considered light curve length L_{lc} , in the case of our observations 2.6 days (rest frame). Moreover, if the underlying DRW process has a τ_{DRW} longer than L_{lc} , then the fit will return a value that is 20%–30% of L_{lc} (e.g., Kozłowski et al. 2010; Emmanoulopoulos et al. 2013; Burke et al. 2020). The characteristic timescales we measure lie in this critical regime, and should therefore be treated carefully. To check the robustness of our results, for each sector, we first simulated 5000 light curves with the same time steps and DRW parameters as the real observations, and then fit them in the same way as done for our data. For each sector, the characteristic timescale τ_{DRW} measured from the real data and from the simulated ones correspond well within 1σ . As a complementary test, we run simulations to determine the probability of obtaining the measured τ_{DRW} in the case that the real τ_{DRW} is actually longer than the observed time baseline. Specifically, for each sector, we simulated three sets of 1000 DRW time series each with the same time baseline as our observations and $\tau_{\text{DRW}} = 10, 50$, and 100 days, respectively, and fit them following the same procedure as for the real data. For each simulation set, we then determined the fraction

of simulations for which the fit returns τ_{DRW} values within a given percent of the τ_{DRW} measured from the real observations. In summary, we see that for $\tau_{\text{DRW}} = 10$ days, less than 10% (20%) of the simulated τ_{DRW} are within $\pm 10\%$ ($\pm 20\%$) of the τ_{DRW} measured from real data for the considered period. These percentages strongly decrease for $\tau_{\text{DRW}} = 50$ –100 days.⁸ Although the performed tests do not allow us to fully rule out the possibility that the measured τ_{DRW} is due to biases in the modeling, our results are consistent with the idea that the measured τ_{DRW} is real. This is further supported by the agreement between the break frequency observed in the Lomb–Scargle and in the DRW periodograms (Figure 8). In addition, as we will elaborate in Section 6, the measured τ_{DRW} are consistent with the orbital and thermal timescales computed for Pictor A, and are therefore physically meaningful. Finally, we tested that the characteristic timescales measured for the sectors that exhibit strong flares (Sectors 5 and 33) are not driven by the flares themselves. For this purpose, we repeated the DRW fit of Sectors 5 and 33 after masking the regions with the most prominent flares. The masking was achieved by fitting the five highest-amplitude flares in these sectors with a typical flare profile of

$$F(t) = F_c + \frac{2F_{\text{max}}}{e^{(t_0-t)/T_r} + e^{(t-t_0)/T_d}},$$

where F_c is a constant baseline, F_{max} is the peak flux of the flare, t_0 is the peak time of the flare, and T_r and T_d are the rise and decay times. This asymmetric exponential form was originally used to model gamma-ray flares of blazars in Fermi data (Abdo et al. 2010), and was used to model the flares of the only monitored blazar in the Kepler field, W2R 1926+42, by Li et al. (2018). Once the flares were modeled, we simply removed all data points within the span of the flare’s modeled baseline and repeated the DRW fit.

The obtained timescales, $\tau_{\text{DRW}} \sim 1$ –1.2 days, are consistent within 1σ with the values measured for these sectors before flare masking ($\tau_{\text{DRW}} \sim 0.8$ days) and well below the values measured

⁸ Example: if the measured τ_{DRW} for the considered sector is ~ 5 days, then less $< 10\%$ of the simulations will return a τ_{DRW} within 5 ± 0.5 days, i.e., 4.5–5.5 days.

for the sectors without strong flares ($\tau_{\text{DRW}} \sim 3\text{--}6$ days). We also applied the Durbin–Watson tests to the residuals in these cases and obtained $d = 2.35$ for both. These results show that the DRW fit is not driven by the presence of flares, and are consistent with the possibility that the flares may indeed occur during a different variability regime.

6. Discussion

As presented in Section 5, the flux distributions obtained for the single sectors can mostly be fit by bimodal distributions, instead of simple Gaussian or log-normal functions. In addition, we do not find any evidence of the rms–flux relation observed in X-ray light curves (e.g., Uttley & McHardy 2001; Vaughan et al. 2003; McHardy et al. 2004). This may indicate that the optical/UV variability observed by TESS on day-month timescales is due to multiple physical processes. In addition, as found in Smith et al. (2018), the lack of correlation between flux and rms in the Kepler and TESS light curves may indicate that the optical/UV variability is not a simple reprocessing of the more rapid variability observed in the X-ray.

As shown in Figure 9, the DRW periodograms obtained by fitting the sectors with prominent flares have shorter characteristic timescales ($\tau_{\text{DRW}} \sim 0.8$ days) than the sectors without obvious flares ($\tau_{\text{DRW}} \sim 3\text{--}6$ days). This result persists even when the flares are removed from the light curve, indicating that the flares themselves are not responsible for the shorter timescales. While the TESS sector boundaries are of course artificial, as we are examining the sectors separately, they provide a simple, consistent baseline for comparison between flaring and non-flaring periods, which is otherwise arbitrary. This may signify that the AGN is switching between two different variability modes, as it is often observed in X-ray binaries (e.g., Coriat et al. 2011; Eckersall et al. 2015), and is expected to happen also in AGN, although at different timescales (e.g., Alexander & Hickox 2012; Sartori et al. 2016).

6.1. Characteristic Timescales

To put the observed characteristic timescales into context, it is interesting to investigate the different variability timescales expected for Pictor A given its black hole mass of $M_{\text{BH}} = 5.9 \times 10^6 M_{\odot}$ (Koss et al. 2022). The *light-crossing timescale* t_{lc} is defined as the time needed for light to cross a region of radius r , and is, therefore, the relevant timescale for irradiated disks (e.g., Gaskell & Klimek 2003). For a standard accretion disk (Shakura & Sunyaev 1973), this is parameterized as

$$t_{\text{lc}} \equiv \frac{r}{c} \simeq 0.875 \left(\frac{r}{150 r_g} \right) M_8 \text{ [d]},$$

where c is the speed of light, M_8 the black hole mass in units of $10^8 M_{\odot}$, and r_g is the gravitational radius, which for the SMBH in Pictor A corresponds to $r_g = 8.7 \times 10^9$ m. For a near-UV-to-optical continuum emitting region with $r = 150 r_g$ we obtain $t_{\text{lc}} = 0.05$ days, which is much shorter than the observed timescales and within the Poisson noise for these light curves in any case.

Another important timescale is the *dynamical timescale* t_{dyn} , which is the time needed for the accretion disk to achieve

hydrostatic equilibrium:

$$t_{\text{dyn}} \equiv \frac{r}{v} \simeq \frac{1}{\Omega} \simeq 10 \left(\frac{r}{150 r_g} \right)^{3/2} M_8 \text{ [d]},$$

where $\Omega = (GM_{\text{BH}}/r^3)^{1/2}$ is the angular velocity of the material within the disk (Netzer 2013) and G the gravitational constant. This is also considered to be the shortest timescale on which it is possible to observe large-scale physical changes in the disk. For the near-UV-to-optical continuum emitting region with $r = 150 r_g$ we obtain $t_{\text{dyn}} = 0.6$ day.

Closely related to the dynamical timescale is the *orbital timescale*, defined as

$$t_{\text{orb}} = 2\pi t_{\text{dyn}}.$$

For the same near-UV-to-optical continuum emitting region as above we obtain $t_{\text{orb}} = 3.7$ days.

Finally, the *thermal timescale* t_{therm} is the time needed for energy to redistribute due to dissipative and cooling processes within the disk, and it is related to t_{dyn} and the viscosity α as

$$t_{\text{therm}} = \frac{t_{\text{dyn}}}{\alpha}.$$

Assuming $\alpha = 0.1\text{--}0.4$ (King & Pringle 2007) and the dynamical timescale computed above we obtain $t_{\text{therm}} = 1.5\text{--}5.9$ days.

It is important to recognize that the characteristic timescale τ_{DRW} simply parameterizes the phenomenological damped random walk model to fit the data, in a statistical sense. While this may generate a range for a physical *characteristic timescale* of variability, it is *not* predicted by any specific physical models. We can therefore only determine whether the τ_{DRW} is plausibly consistent with physically relevant timescales; it cannot be used to eliminate or support different physical origin models of the variability. The purpose of our analysis is to demonstrate the potential and limitations of TESS data in studying these phenomena. The characteristic timescales resulting from our DRW analysis of Pictor A are consistent with the dynamical, orbital, and thermal timescales for the best-available black hole mass estimate, as found in numerous previous studies of AGN variability. Indeed, recent studies pointed out that the variability observed in the optical AGN spectrum, both as regards stochastic variability and changing-look (of changing-state) AGN, is mostly consistent with the thermal timescale (e.g., Kelly et al. 2009; Stern et al. 2018; Parker et al. 2019). This would imply that variability may be mostly due to rapid temperature changes in large parts of the accretion disk (e.g., Dexter & Agol 2011; Ricci & Trakhtenbrot 2023).

7. Summary and Conclusions

In this paper, we present *Quaver*, a new software tool designed specifically to extract TESS light curves of extended and faint sources exhibiting stochastic variability, as is the case for AGN. The code is designed to be fully interactive and flexible, allowing the user to choose the best extraction parameters for their data, and is publicly available on Github.⁹

As a first example, we used *Quaver* to extract the TESS light curves available for the radio-loud AGN galaxy Pictor A, and perform a variability analysis. The obtained light curves are well fit with a DRW model. We find that the source exhibits

⁹ <https://github.com/kristalynnesmith/quaver>

characteristic timescales of $\tau_{\text{DRW}} \sim 3\text{--}6$ days during periods when flares are not dominant, and $\tau_{\text{DRW}} \sim 1$ day during periods with numerous flares, even when the flares are removed from the DRW analysis. The observed timescales are consistent with the dynamical, orbital, and thermal timescales expected for the black hole mass of Pictor A, as observed also in other sources by recent optical/UV variability studies.

This work demonstrates the power of *Quaver* for extracting TESS light curves and therefore enabling accurate high-cadence optical variability analysis of AGN.

Acknowledgments

This work is supported by NASA grant number 80NSSC22K0741. K.L.S. gratefully acknowledges the staff of the K2 Guest Observer office at NASA Ames, especially C. Hedges, for their assistance and advice in adapting the matrix regression methods. L.F.S. acknowledges the financial support of the Swiss National Science Foundation (SNSF). L.F.S. acknowledges E.O. Garvin for support in the choice and application of statistical tests. All of the data presented in this paper were obtained from the Mikulski Archive for Space Telescopes (MAST) at the Space Telescope Science Institute. The specific observations analyzed can be accessed via DOI:[10.17909/k6mn-kp63](https://doi.org/10.17909/k6mn-kp63). Finally, we thank the anonymous referee for comments that substantially improved the clarity of the manuscript, as well as led to modifications of *Quaver*, which will improve the user’s ability to assess the validity of the fitting process.

Facilities: TESS (Ricker et al. 2015), PO:1.2m.

Software: *Quaver* (Smith 2023), *TESSCut* (Brasseur et al. 2019), *Astropy* (Astropy Collaboration et al. 2013, 2018), *Astroquery* (Ginsburg et al. 2019), *Matplotlib*

(Hunter 2007), *Lightkurve* (Lightkurve Collaboration et al. 2018), *EzTao* (Yu & Richards 2022).

Author Contributions

K.L.S. developed and tested the *Quaver* tool, performed the analysis of the flares, and wrote part of the manuscript. L.F.S. extracted the TESS light curves of Pictor A, performed the variability analysis, and wrote part of the manuscript. The authors collaborated throughout the project by discussing the possible analyses and the obtained results.

Appendix A

Handling of Scattered Light Background

By far the dominant trend in uncorrected TESS light curves is the contribution from the scattered light from the moon and Earth, which can leak into the field of view of the detector and varies significantly over the course of a given sector.

In the simple hybrid method, this light is removed using a simple subtraction of the light curve of the faint/empty pixels.

In the PCA and fully hybrid methods, the scattered light curve is always the dominant (first) principal component, and is removed in the matrix regression as that component. We demonstrate this in Figures 10 and 11, where we show two example sectors: the extracted light curve of the inactive galaxy 2dFGRS TGS145Z214 (see Appendix B) using the PCA method with only one PCA chosen, and the extracted light curve of Pictor A using the fully hybrid method with only one additive PCA allowed. If this is compared to Figure 12, it is clear that the primary contribution to both light curves is the same as the scattered light trend for the appropriate camera and sector.

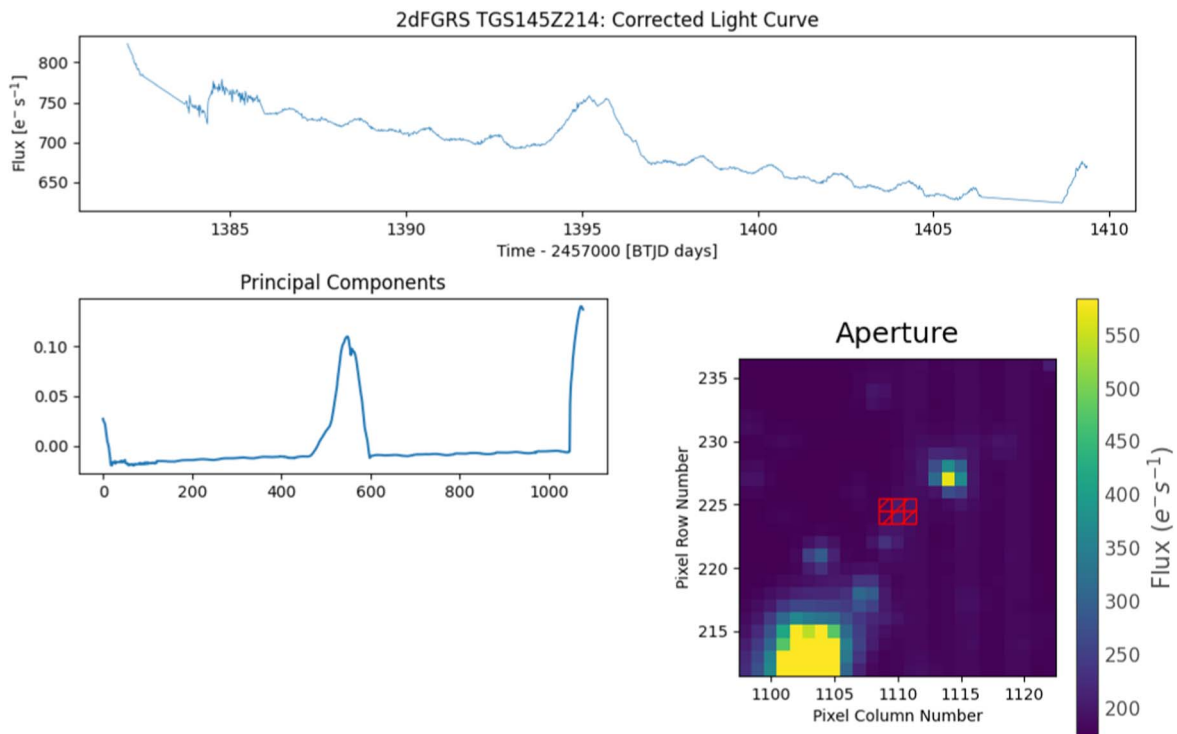


Figure 10. TESS Sector 3 light curve of the non-AGN galaxy 2dFGRS TGS145Z214, restricted to a one-component PCA analysis with insufficient parameters for background and contaminant correction.

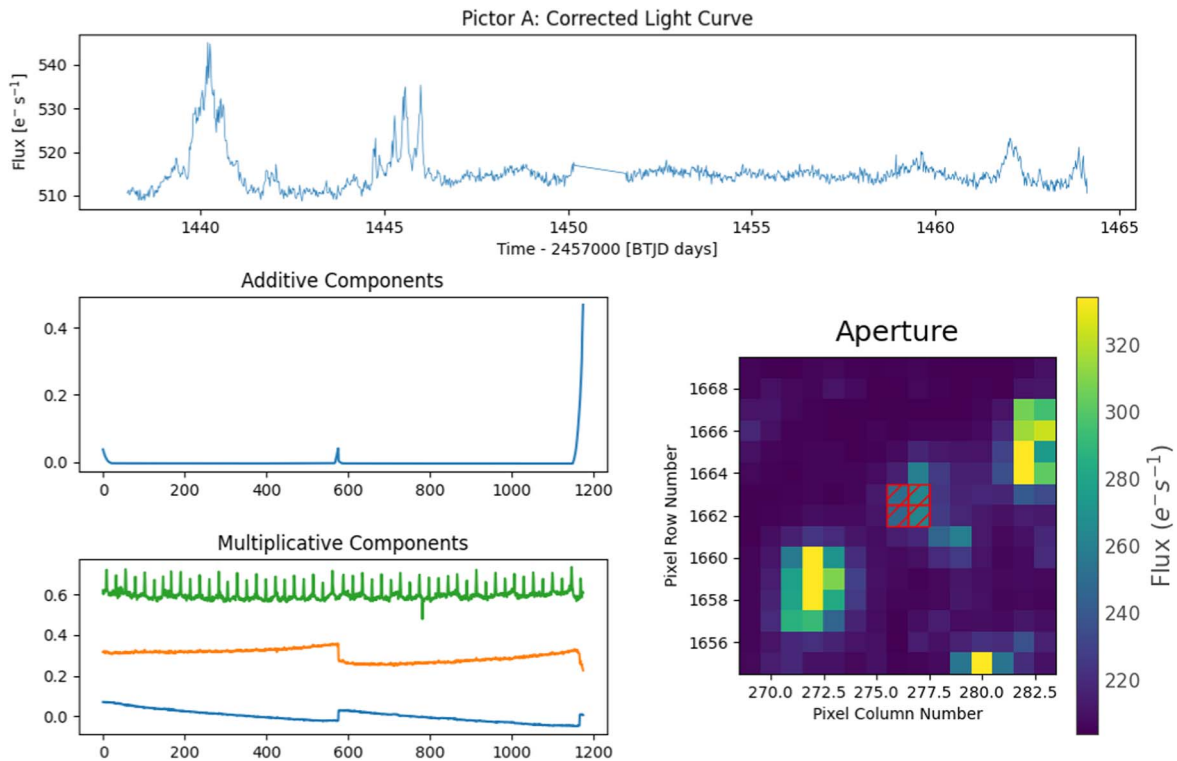


Figure 11. TESS Sector 5 light curve of Pictor A, using the fully hybrid method and restricted to only one additive PCA component.

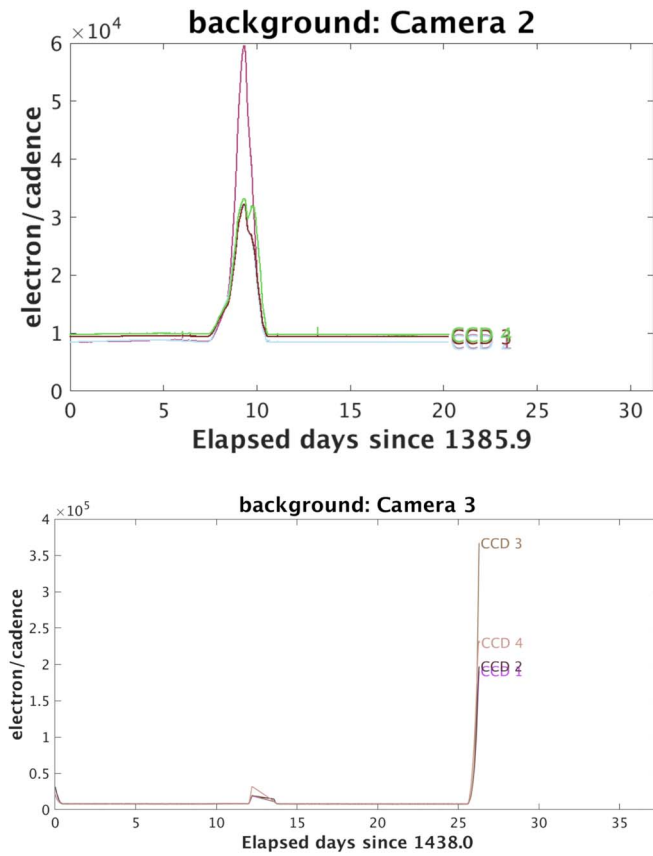


Figure 12. Plots of the scattered background light over time for Sector 3, Camera 2 (left) and Sector 5, Camera 3 (right), from the TESS Data Release Notes (https://archive.stsci.edu/teess/teess_drn.html). Differences in plot formatting are due to differences as presented in the Sector 3 and 5 release note documents.

Appendix B Contamination from Variable Sources

The large pixel response function of TESS means that even at apparently large distances from the target source, variable stars of sufficient magnitude can contribute significantly to the flux in the target aperture. In some cases, it is sufficient to cull one's sample of all targets in close proximity to potential contaminants, as in the recent work selecting AGN using TESS variability by Treiber et al. (2023). However, the number of AGN below TESS's magnitude limit is already only a few thousand, and those with long light curves number in the hundreds. It is therefore desirable to be able to correct for nearby source contamination.

The custom aperture selection is the first line of defense, allowing the user to manually exclude sources from nearby bright sources and use the DSS overlay to ascertain the proximity of even faint nearby stars. However, it is not obvious by eye which stars may still be contributing to flux inside the target aperture.

The hybrid methods mentioned above provide an important safeguard against strongly variable signals: all pixels containing a source (i.e., containing a source above the user-defined threshold) are background corrected, using either subtraction in the case of the simple hybrid method or using a separate regression matrix in the fully hybrid method. Then, these background-corrected sources are *themselves* fitted by principal components, to create the second correction matrix. Any strongly variable sources will contribute to these components, and this contaminant variability will be corrected for.

To demonstrate the effectiveness with which Quaver removes these trends, we extract the light curve of the bright ($V=8.4$) α^2 CVn variable star AP Sculptor (Balona 2022), and the light curve of the nearby non-AGN (and thus non-variable) galaxy 2dFGRS TGS145Z214 ($V=15.98$). The objects are separated by $4''.8$, and the star is very bright, with a flux distribution across many pixels.

In Figure 13, we show a map of the variable star and galaxy. In Figure 14, we show the result of a Quaver extraction of AP Sculptor, to demonstrate its strong variability. In Figure 15 we show the Quaver-corrected light curve of 2dFGRS TGS145Z214 (using the fully hybrid method, with three principle components for the additive and multiplicative components).

The second component in the multiplicative components panel is contaminant variability AP Sculptor. This has been removed from the galaxy light curve, which is flat, as expected.

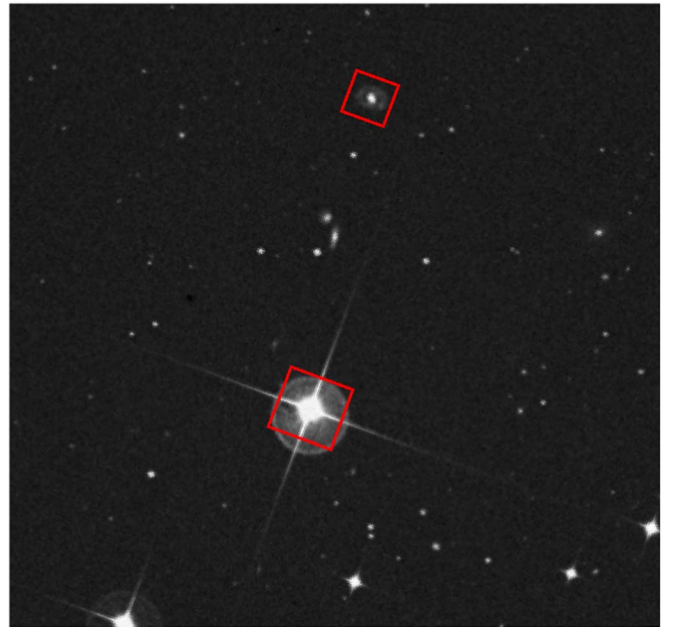


Figure 13. Map of the variable star AP Sculptor and the nearby galaxy 2dFGRS TGS145Z214, each shown in a red box. The sources are separated by $4''.79$.

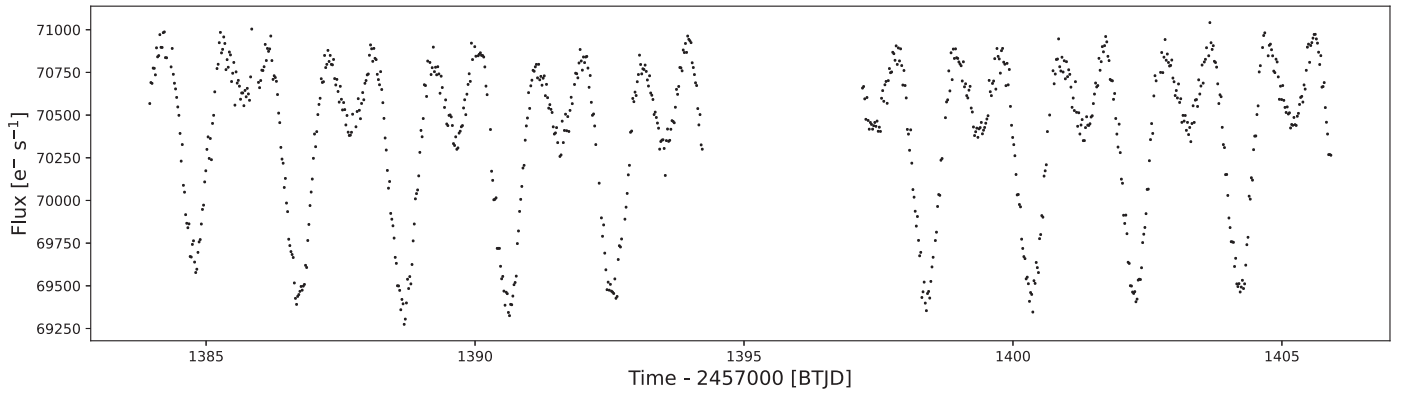


Figure 14. TESS Sector 3 light curve of the variable star AP Sculptor, extracted with *Quaver*.

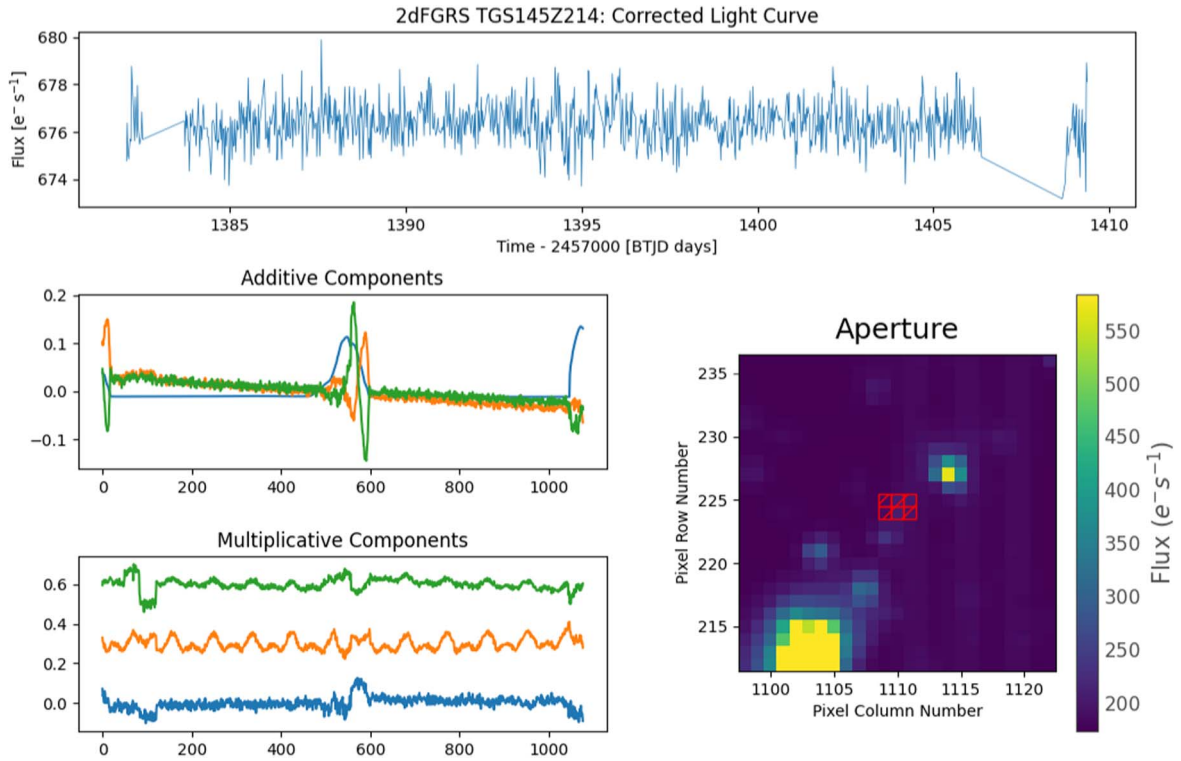


Figure 15. TESS Sector 3 light curve of the non-AGN galaxy 2dFGRS TGS145Z214, extracted with *Quaver*'s full-hybrid method.

Appendix C

Light Curve Correction Figures

Here we provide the output from *Quaver*, which shows the light curves before and after correction, along with the final models derived by the matrix regression. This Appendix includes these plots for Star 1, Sector 4 (Figure 16) and Star 3,

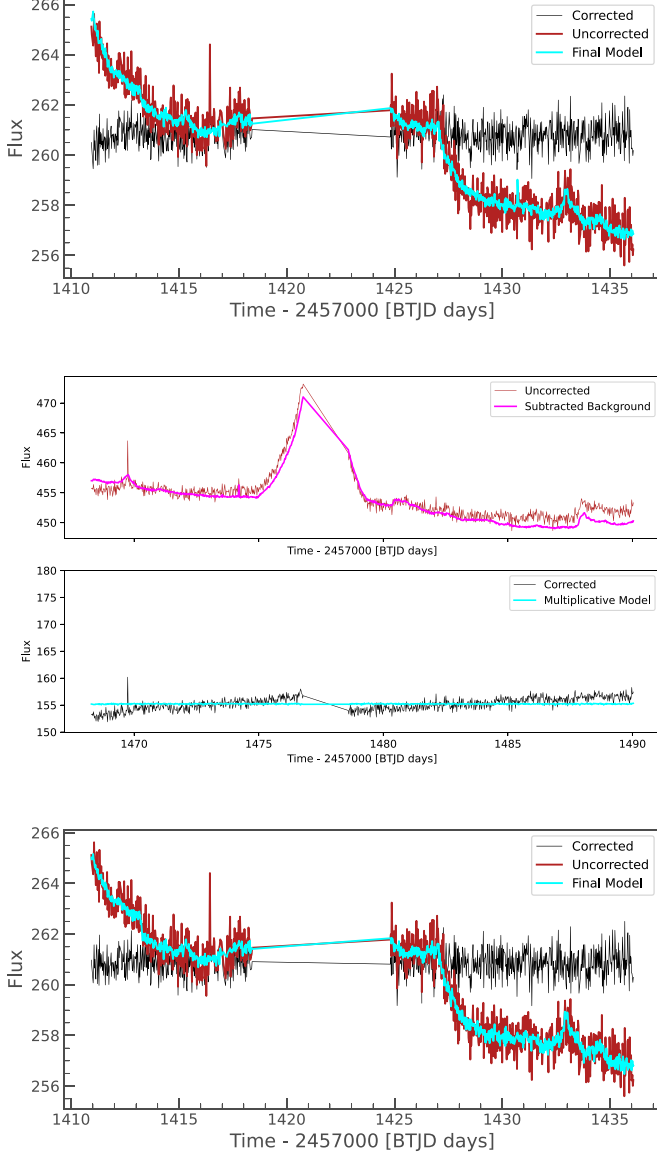


Figure 16. Star 1, Sector 4.

Sector 6 (Figure 17) from Figure 3, as well as for Pictor A (except for Pictor A Sector 5, which was shown in the main part of this paper) in Figures 18, 19, 20, and 21. In each figure, the top panel shows the results of the simple PCA fitting; the middle pane shows the results of The simple hybrid method by showing the uncorrected light curve and the subtracted background in the top panel, and the background-subtracted and corrected light curve with the removed multiplicative model in the bottom panel; and the bottom panel shows the results of the fully hybrid method.

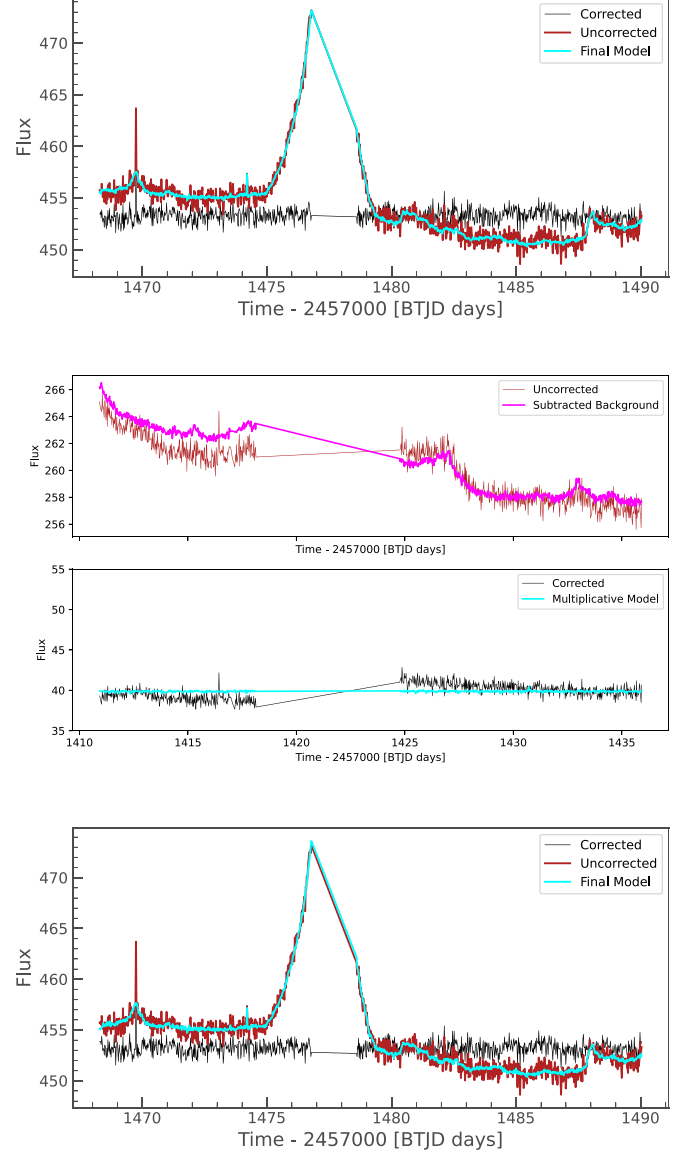


Figure 17. Star 3, Sector 6.

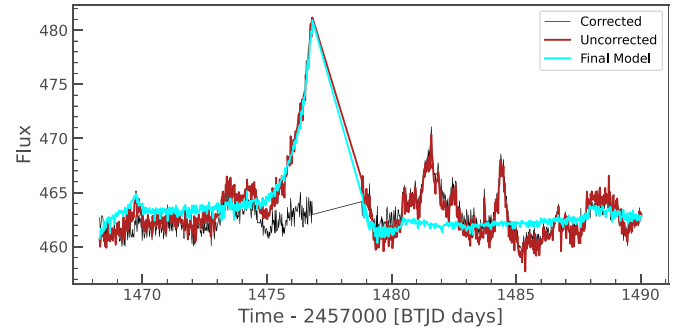
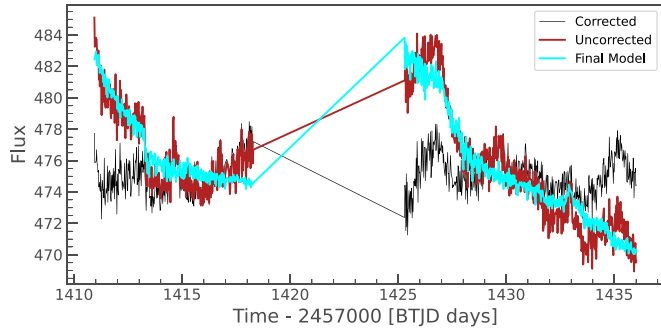
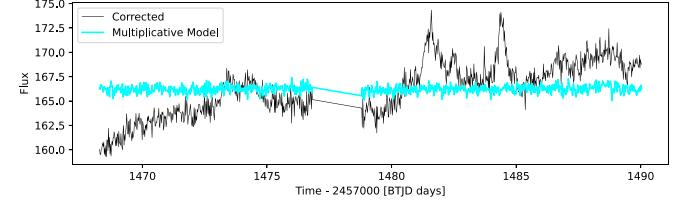
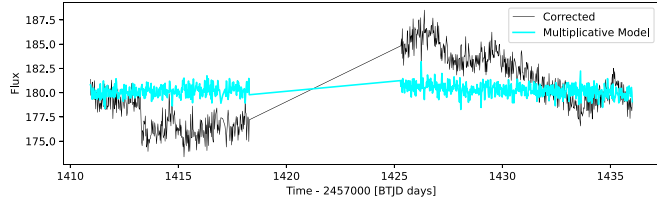
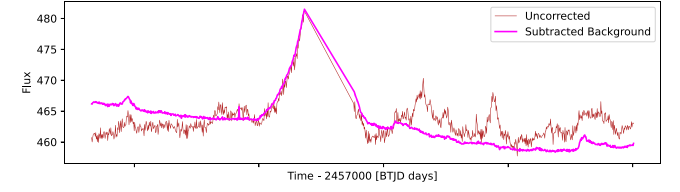
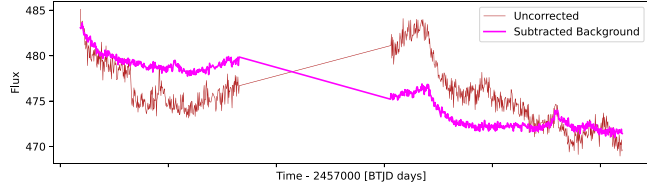
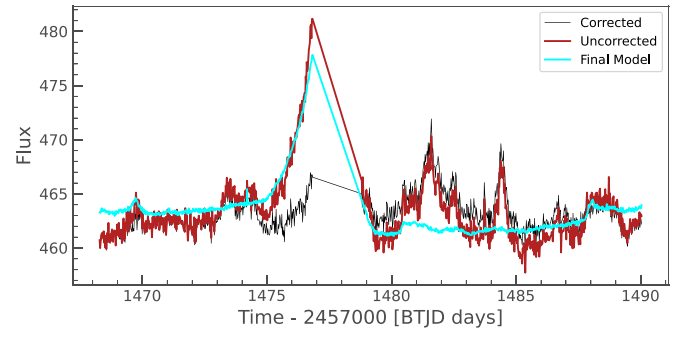
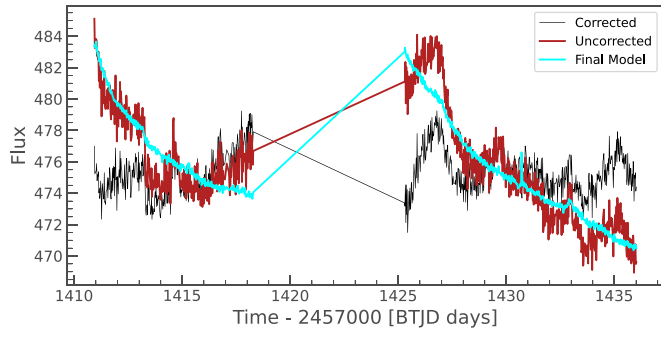


Figure 18. Pictor A, Sector 4.

Figure 19. Pictor A, Sector 6.

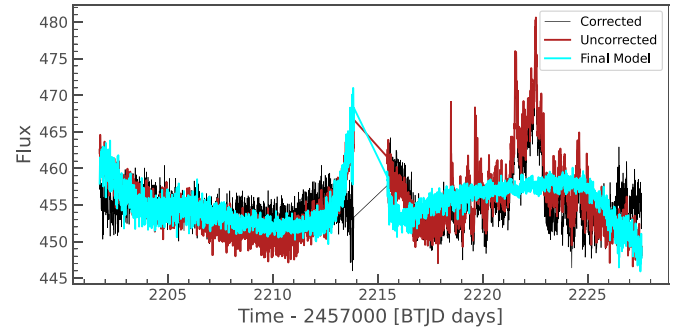
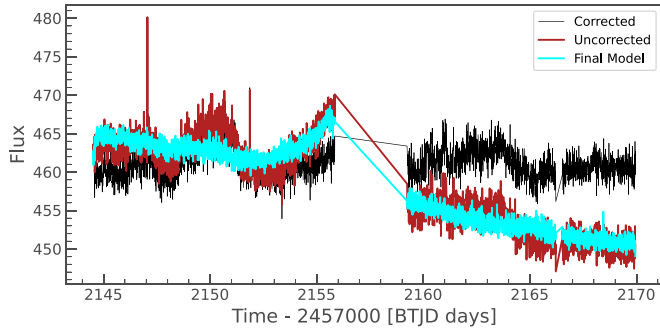
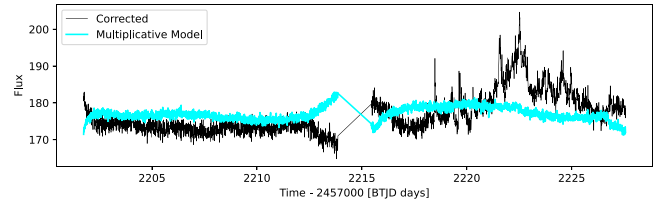
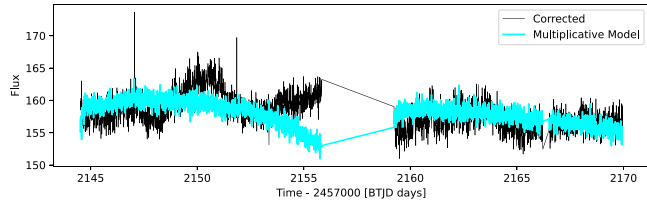
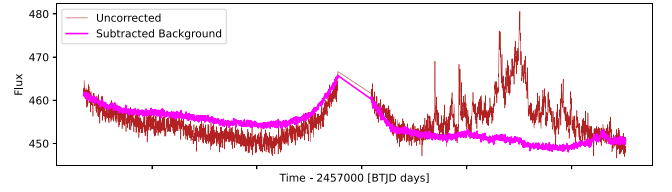
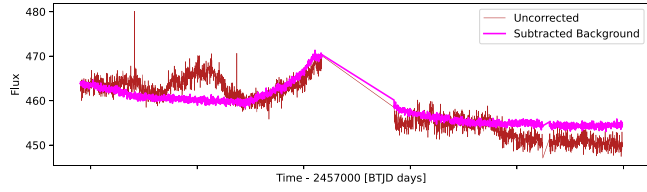
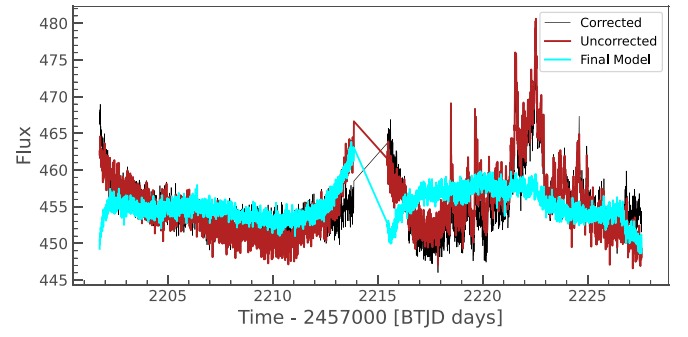
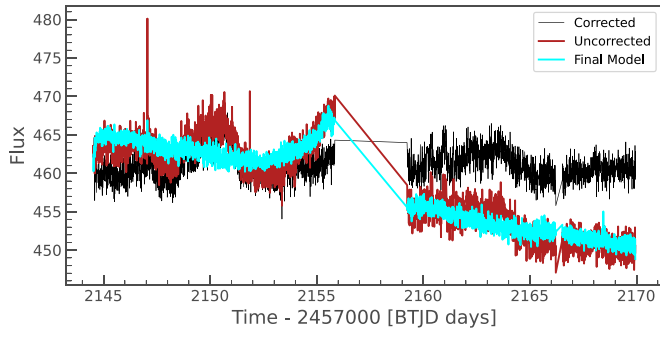


Figure 20. Pictor A, Sector 31.

Figure 21. Pictor A, Sector 33.

Appendix D

DRW fit

Here we provide a few additional figures to support the tests carried out in Section 5.2. In Figure 22, we provide an example of a DRW fit result, with posterior distributions constraining the parameter, and a histogram showing that the residuals of the

fit are normally distributed. This, along with the results of the Durbin–Watson test as described in Section 5.2, supports the statement that the DRW is an acceptable model for the observed variability behavior.

In Figures 23 and 24, we show the DRW fit in the light curve of Sector 5 with and without the masking of the five highest-amplitude flares.

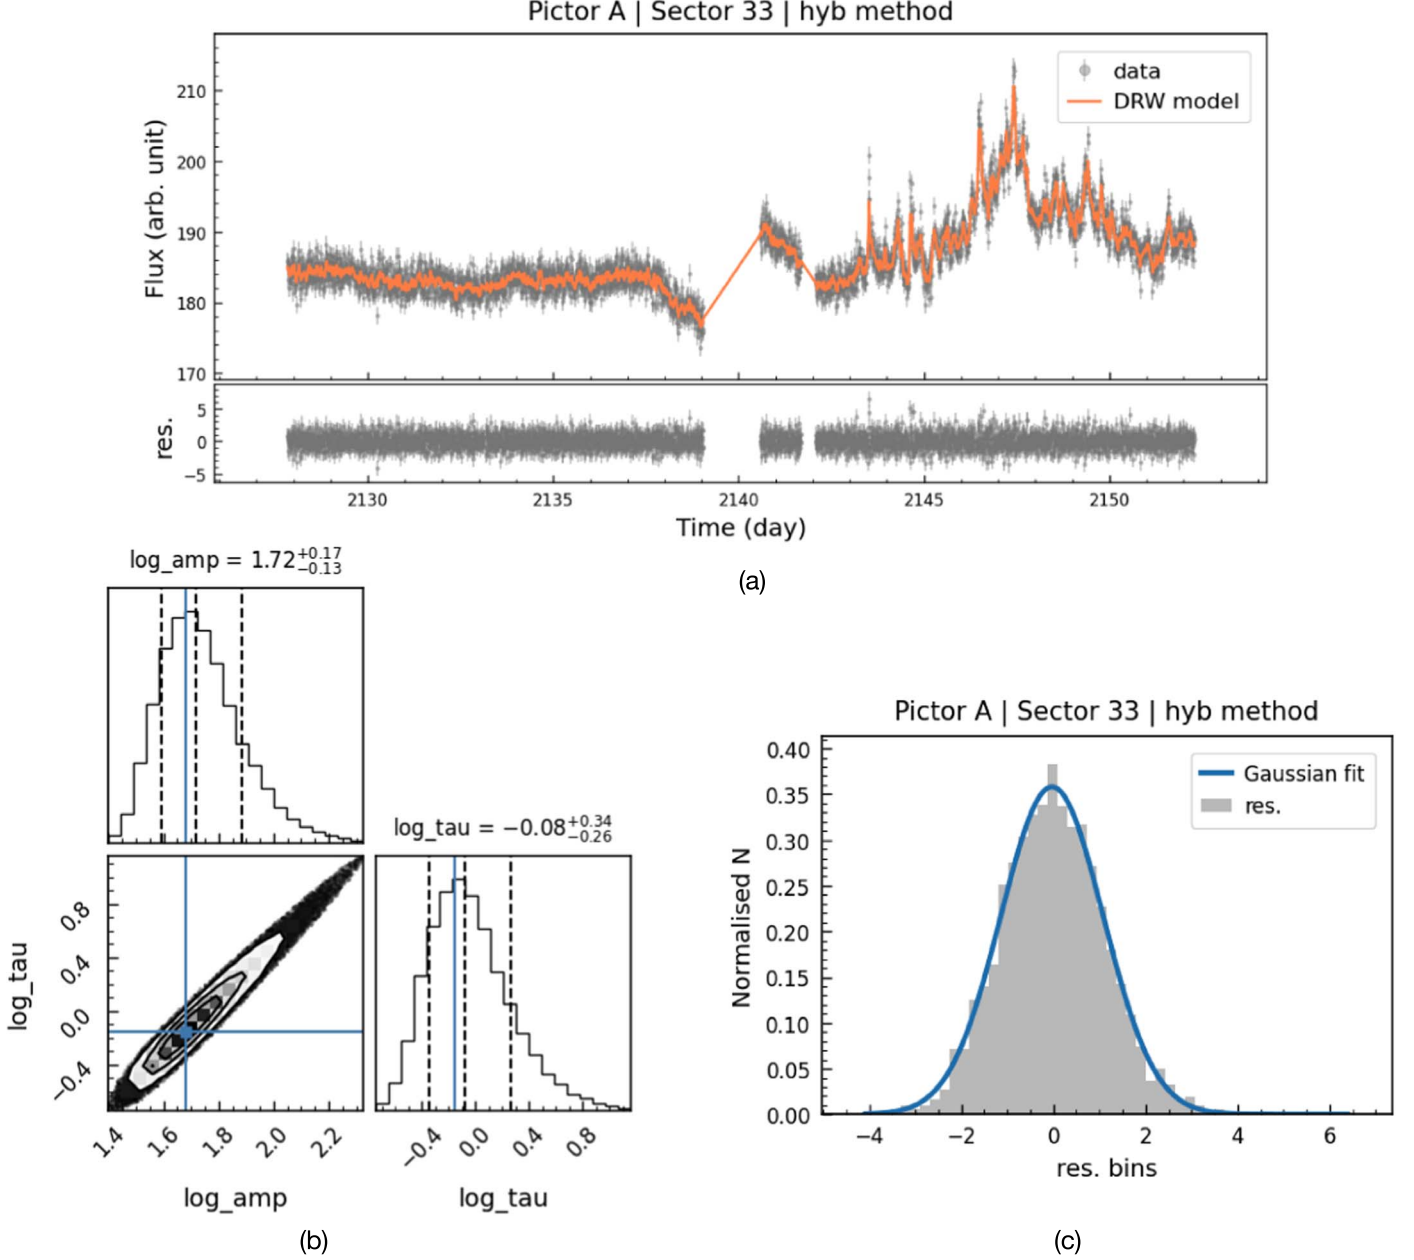


Figure 22. Example of DRW fit performed for the light curve of Sector 33. Panel (a), top: rest-frame TESS light curve for Sector 33 (gray) and best-fit DRW model (orange). Panel (a), bottom: residuals from the DRW fit. As we can see in panel (c), the histogram of the residuals (gray) follows a Gaussian distribution (blue). Panel (b) Distribution of the posteriors from the DRW fit performed with EzTao. Both parameters, i.e., amplitude and τ_{DRW} , are well constrained. See Section 5.2 for details.

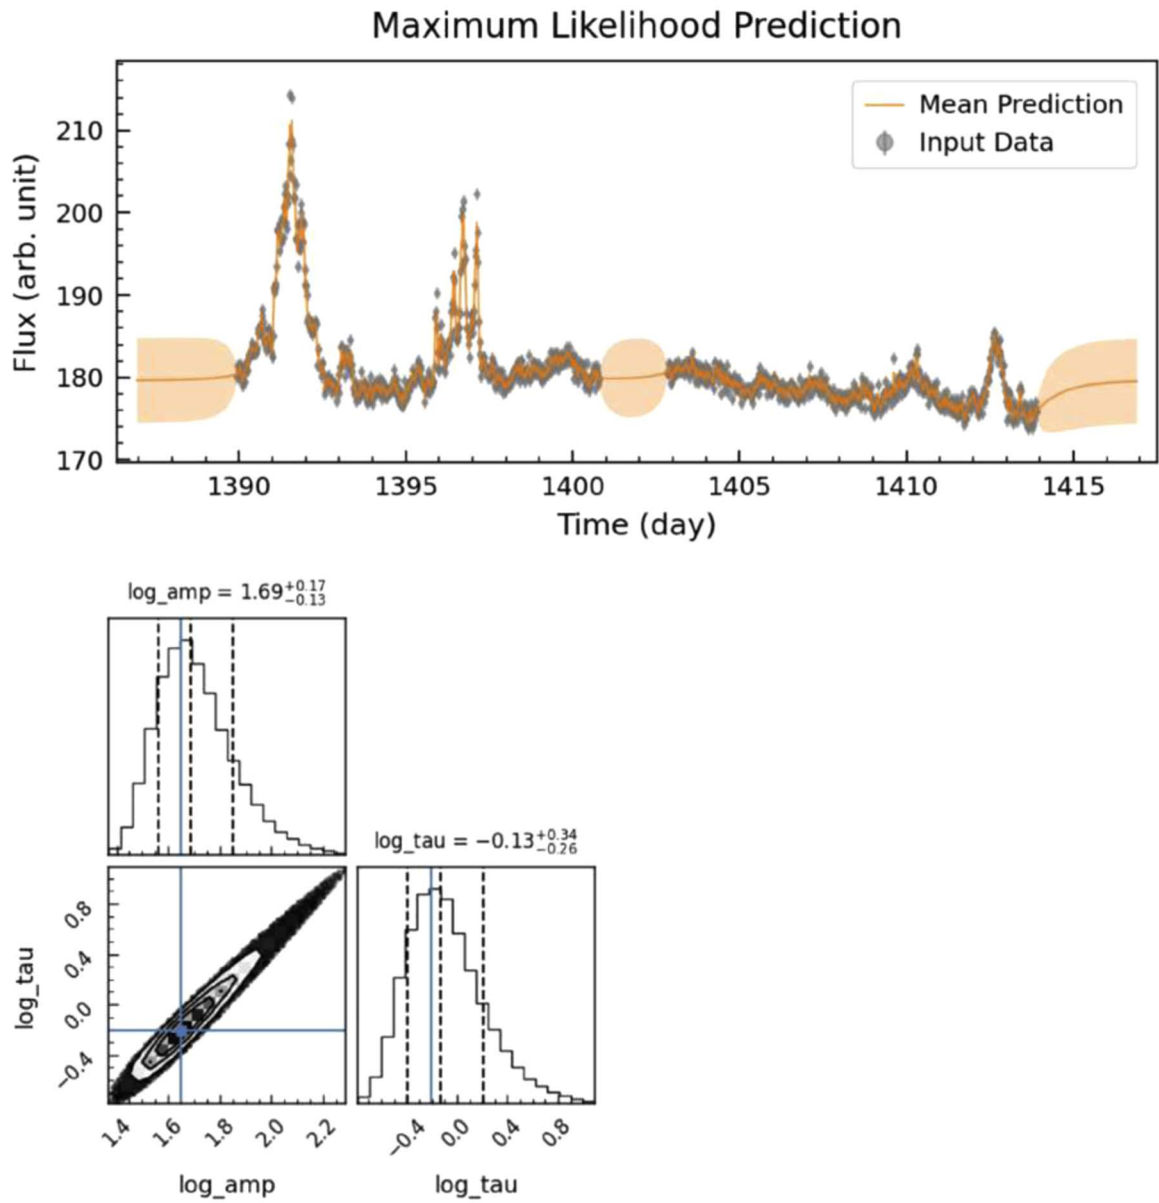


Figure 23. Results of the DRW fit to Sector 5 data as observed.

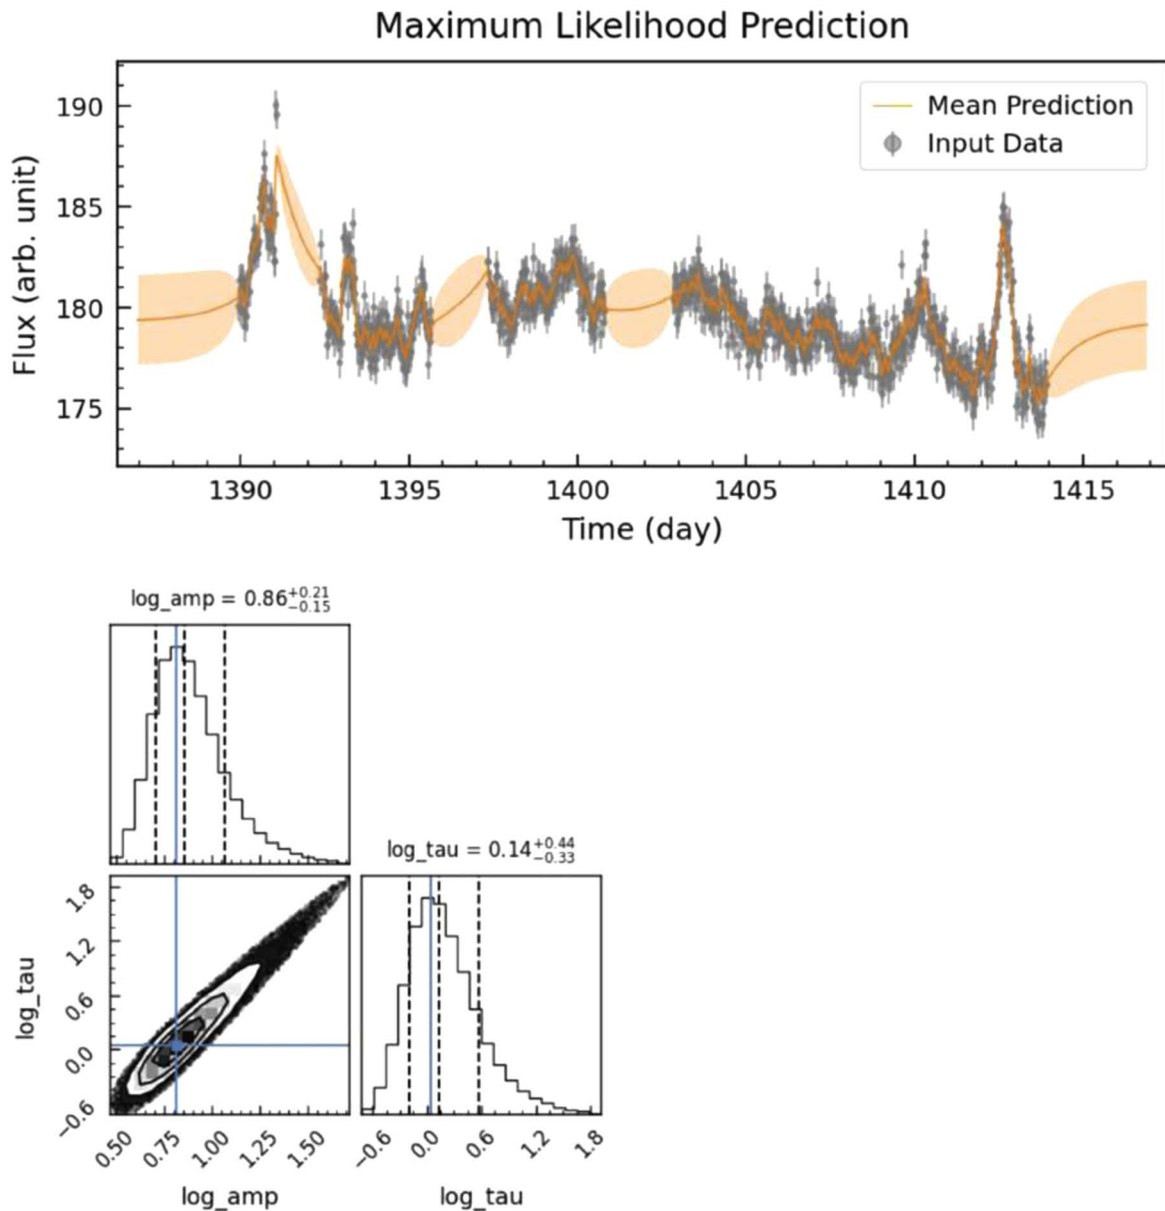


Figure 24. Results of the DRW fit to Sector 5 data after the five highest-amplitude flares have been masked and removed.

ORCID iDs

Krista Lynne Smith <https://orcid.org/0000-0001-5785-7038>
 Lia F. Sartori <https://orcid.org/0000-0001-8020-3884>

References

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, **722**, 520
 Alexander, D. M., & Hickox, R. C. 2012, *NewAR*, **56**, 93
 Anthony, F., Núñez, A., Agüeros, M. A., et al. 2022, *AJ*, **163**, 257
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Avallone, E. A., Tayar, J. N., van Saders, J. L., et al. 2022, *ApJ*, **930**, 7
 Balbus, S. A., & Hawley, J. F. 1998, *RvMP*, **70**, 1
 Balona, L. A. 2022, *MNRAS*, **510**, 5743
 Bhatta, G., & Dhital, N. 2020, *ApJ*, **891**, 120
 Blandford, R. D., & Znajek, R. L. 1977, *MNRAS*, **179**, 433
 Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
 Brasseur, C. E., Phillip, C., Fleming, S. W., Mullally, S. E., & White, R. L., 2019 Astrocut: Tools for creating cutouts of TESS images, Astrophysics Source Code Library, ascl:1905.007
 Brown, A. M., & Adams, J. 2012, *MNRAS*, **421**, 2303
 Burke, C. J., Shen, Y., Chen, Y.-C., et al. 2020, *ApJ*, **899**, 136
 Claytor, Z. R., van Saders, J. L., Llama, J., et al. 2022, *ApJ*, **927**, 219
 Coriat, M., Corbel, S., Prat, L., et al. 2011, *MNRAS*, **414**, 677
 Cumming, A., Butler, R. P., Marcy, G. W., et al. 2008, *PASP*, **120**, 531
 Deeter, J. E. 1984, *ApJ*, **281**, 482
 Deeter, J. E., & Boynton, P. E. 1982, *ApJ*, **261**, 337
 Dexter, J., & Agol, E. 2011, *ApJL*, **727**, L24
 Durbin, J., & Watson, G. S. 1950, *Biometrika*, **37**, 409
 Durbin, J., & Watson, G. S. 1951, *Biometrika*, **38**, 159
 Eckersall, A. J., Vaughan, S., & Wynn, G. A. 2015, *MNRAS*, **450**, 3410
 Emmanoulopoulos, D., McHardy, I. M., & Papadakis, I. E. 2013, *MNRAS*, **433**, 907
 Feinstein, A. D., Montet, B. T., Foreman-Mackey, D., et al. 2019, *PASP*, **131**, 094502
 Gaskell, C. M., & Klimek, E. S. 2003, *A&AT*, **22**, 661
 Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, **157**, 98
 Hunter, J. D. 2007, *CSE*, **9**, 90, <https://ieeexplore.ieee.org/document/4160265>

- Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2015, *MNRAS*, **451**, 4328
- Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, *ApJ*, **698**, 895
- Kelly, B. C., Becker, A. C., Sobolewska, M., Siemiginowska, A., & Uttley, P. 2014, *ApJ*, **788**, 33
- King, A. R., & Pringle, J. E. 2007, *MNRAS*, **377**, L25
- Kirchner, J. W. 2005, *PhRvE*, **71**, 066110
- Körding, E. G., Jester, S., & Fender, R. 2006, *MNRAS*, **372**, 1366
- Koss, M. J., Ricci, C., Trakhtenbrot, B., et al. 2022, *ApJS*, **261**, 2
- Kozłowski, S., Kochanek, C. S., Udalski, A., et al. 2010, *ApJ*, **708**, 927
- Li, Y., Hu, S., Wiita, P. J., & Gupta, A. C. 2018, *MNRAS*, **478**, 172
- Lightkurve Collaboration, Cardoso, J. V. d. M., Hedges, C., et al., 2018
Lightkurve: Kepler and TESS time series analysis in Python, Astrophysics
Source Code Library, ascl:1812.013
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, **721**, 1014
- Maltby, P., & Moffet, A. T. 1962, *ApJS*, **7**, 141
- Marshall, H. L., Hardcastle, M. J., Birkinshaw, M., et al. 2010, *ApJL*, **714**, L213
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003
- McHardy, I. M., Papadakis, I. E., Uttley, P., Page, M. J., & Mason, K. O. 2004, *MNRAS*, **348**, 783
- Moreno, J., Buttry, R., O'Brien, J., et al. 2021, *AJ*, **162**, 232
- Netzer, H. 2013, *The Physics and Evolution of Active Galactic Nuclei* (Cambridge: Cambridge Univ. Press)
- Oelkers, R. J., & Stassun, K. G. 2018, *AJ*, **156**, 132
- Parker, M. L., Scharrel, N., Grupe, D., et al. 2019, *MNRAS*, **483**, L88
- Ricci, C., & Trakhtenbrot, B. 2023, *NatAs*, in press
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, **1**, 014003
- Sartori, L. F., Schawinski, K., Koss, M., et al. 2016, *MNRAS*, **457**, 3629
- Scargle, J. D. 2020, *ApJ*, **895**, 90
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, 337
- Simkin, S. M., Sadler, E. M., Sault, R., Tingay, S. J., & Callcut, J. 1999, *ApJS*, **123**, 447
- Smith, K. L. 2023, Quaver: A Pipeline for AGN Science with TESS, Version 2.0, Zenodo, doi:10.5281/zenodo.8400525
- Smith, K. L., Mushotzky, R. F., Boyd, P. T., et al. 2018, *ApJ*, **857**, 141
- Stern, D., McKernan, B., Graham, M. J., et al. 2018, arXiv:1805.06920
- Treiber, H. P., Hinkle, J. T., Fausnaugh, M. M., et al. 2023, *MNRAS*, **525**, 5795
- Uttley, P., & McHardy, I. M. 2001, *MNRAS*, **323**, L26
- Uttley, P., McHardy, I. M., & Vaughan, S. 2005, *MNRAS*, **359**, 345
- Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, *MNRAS*, **345**, 1271
- Wilson, A. S., & Colbert, E. J. M. 1995, *ApJ*, **438**, 62
- Yu, W., & Richards, G. T., 2022 EzTao: Easier CARMA Modeling,
Astrophysics Source Code Library, ascl:2201.001
- Zechmeister, M., & Kürster, M. 2009, *A&A*, **496**, 577