

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering & Physical Sciences
School of Physics & Astronomy

**Investigating the Variability of Low-Mass
Active Galactic Nuclei in the X-ray, UV, and
Optical Regimes**

by

Max William John Beard

MPhys

*A thesis for the degree of
Doctor of Philosophy*

December 2022

University of Southampton

ABSTRACT

Faculty of Engineering & Physical Sciences
School of Physics & Astronomy

Doctor of Philosophy

INVESTIGATING THE VARIABILITY OF LOW-MASS ACTIVE GALACTIC NUCLEI IN THE X-RAY, UV, AND OPTICAL REGIMES

by Max William John Beard

Active galactic nuclei (AGN) display variability in their emission of radiation that can be used to probe the geometry and physical processes that drive them. In this thesis I examine the variability found in the emissions of X-ray, UV, and Optical radiation from AGN to constrain the properties of these extreme systems.

To begin with I present optical interband time lags of the low-mass low-accretion rate AGN NGC 4395 using multiwaveband Liverpool Telescope and HiPERCAM Gran Telescopio Canarias data. The lags found here are consistent with X-ray reprocessing and suggest the first ever detection of the edge of the accretion disc in an AGN.

I then present a study of X-ray/UV observations taken from *XMM-Newton* with simultaneous ground-based optical observations to measure the X-ray/UV/Optical lags of NGC 4395. These support the previous observations and gather further evidence for the existence of a truncated accretion disc in this object.

Next I examine the long-term optical power spectrum of NGC 4395 using multiple light curves observed across different timescales. In addition I use simulation techniques to extend the previous X-ray observations and use reverberation modelling to create synthetic optical light curves from these. Comparison of the model fits of the data and the simulations suggest the existence of a secondary long-term source of variability in addition to X-ray reprocessing.

Finally I investigate the X-ray to UV time lags of a different object, NGC 4593, using simultaneous *XMM-Newton*, *Swift*, and *AstroSat* observations. Here I find evidence of timescale-dependent lags. When filtered for short timescales the lags are consistent with X-ray reprocessing, but this suggests other long-term processes also exist within the object. I also note the presence of previously described contributions from the Broad Line Region's Balmer and Paschen continua.

Contents

| | |
|--|--------------|
| List of Figures | ix |
| List of Tables | xiii |
| Declaration of Authorship | xv |
| Acknowledgements | xvii |
| Abbreviations | xxi |
| Physical Constants | xxiii |
| 1 Introduction | 1 |
| 1.1 Active Galactic Nuclei | 2 |
| 1.1.1 The Unification of AGN classes | 3 |
| 1.2 Black Hole Accretion | 5 |
| 1.3 The Emission Mechanisms of AGN | 11 |
| 1.3.1 X-ray Emission | 11 |
| 1.3.2 UV/Optical Emission | 14 |
| 1.4 The Variability of AGN | 14 |
| 1.4.1 Intrinsic Disc Variations | 16 |
| 1.4.2 X-ray Reprocessing | 16 |
| 1.5 The Power Spectra of XRBs & AGN | 19 |
| 2 First detection of an AGN Accretion Disc's Outer Edge in NGC 4395 | 21 |
| 2.1 Abstract | 21 |
| 2.2 Introduction | 21 |
| 2.3 Liverpool Telescope Observations | 24 |
| 2.3.1 The Observations | 24 |
| 2.3.2 Interband Correlations and Lags | 26 |
| 2.4 GTC HiPERCAM Observations | 31 |
| 2.4.1 Observations | 31 |
| 2.4.1.1 Seeing Effects | 35 |
| 2.4.2 Fractional Variability | 36 |
| 2.4.3 Interband Lags | 38 |
| 2.4.4 Referencing the HiPERCAM lags to the X-ray band | 39 |
| 2.5 Discussion | 41 |

| | | |
|----------|---|-----------|
| 2.5.1 | Preliminary Conclusions: An Edge to the Emission Region and little BLR contribution | 41 |
| 2.5.2 | Full Response Modelling of the Lags: The Disc Edge and The Colour Correction Factor | 44 |
| 2.5.3 | BLR contribution to the UV/optical variability | 47 |
| 2.5.4 | Physical Origin of the Truncation Radius | 49 |
| 2.5.5 | Alternative explanations of the shorter than expected long wavelength lags | 50 |
| 2.5.6 | Implications for other AGN | 51 |
| 2.5.6.1 | The X-ray / UV disconnect | 51 |
| 2.5.6.2 | Overall lag modelling and mass estimation | 52 |
| 2.6 | Conclusions | 52 |
| 3 | X-ray, UV, and Optical Reverberation Mapping of NGC 4395 | 55 |
| 3.1 | Abstract | 55 |
| 3.2 | Introduction | 55 |
| 3.2.1 | Theory | 55 |
| 3.2.2 | Previous Observation Campaigns | 57 |
| 3.3 | Observations | 58 |
| 3.3.1 | Logistics | 58 |
| 3.3.2 | Challenges | 59 |
| 3.4 | Data Reduction | 60 |
| 3.4.1 | X-ray/UV | 60 |
| 3.4.2 | Optical | 61 |
| 3.4.2.1 | Comparison Star Selection | 61 |
| 3.4.2.2 | Reduction Method | 65 |
| 3.4.2.3 | Misalignment and CALI | 67 |
| 3.5 | Results | 69 |
| 3.5.1 | Lag Calculation Methods | 69 |
| 3.5.2 | Javelin Results | 69 |
| 3.5.2.1 | First Observation Set | 69 |
| 3.5.2.2 | Second Observation Set | 71 |
| 3.5.2.3 | Third Observation Set | 72 |
| 3.5.2.4 | Fourth Observation Set | 73 |
| 3.5.3 | Cross-Correlation Function Results | 74 |
| 3.6 | Analysis | 77 |
| 3.6.1 | Lag Analysis | 77 |
| 3.6.1.1 | Optical Lags | 77 |
| 3.6.1.2 | UV Lags | 78 |
| 3.6.2 | Modelling | 79 |
| 3.6.2.1 | Shakura & Sunyaev | 79 |
| 3.6.2.2 | KYNreverb Response Lags | 81 |
| 3.6.2.3 | KYNreverb Response Function Convolution | 84 |
| 3.7 | Conclusions | 88 |
| 4 | Multi-waveband Power Spectral Densities from observations of NGC 4395 | 91 |
| 4.1 | Abstract | 91 |

| | | |
|----------|---|------------|
| 4.2 | Introduction | 91 |
| 4.3 | Observations | 92 |
| 4.4 | Methodology | 95 |
| 4.4.1 | PSD Fitting | 95 |
| 4.4.2 | Synthetic Optical Light Curves | 97 |
| 4.4.2.1 | Simulating Long-Term X-ray Light Curves | 98 |
| 4.4.2.2 | Model Optical Light Curves from X-ray Simulations | 99 |
| 4.5 | Results & Analysis | 99 |
| 4.5.1 | Optical Data Power Spectra | 99 |
| 4.5.2 | X-ray and Synthetic Optical Power Spectra | 101 |
| 4.5.3 | Testing for double-bend compatibility | 103 |
| 4.6 | Conclusions | 106 |
| 5 | Timescale-dependent X-ray to UV time lags of NGC 4593 | 109 |
| 5.1 | Abstract | 109 |
| 5.2 | Introduction | 109 |
| 5.3 | Observations | 113 |
| 5.3.1 | <i>XMM-Newton</i> Observations | 113 |
| 5.3.2 | <i>Swift</i> Observations | 114 |
| 5.3.3 | <i>AstroSat</i> Observations | 114 |
| 5.4 | Lag Determination | 116 |
| 5.5 | Results and Analysis | 117 |
| 5.5.1 | <i>XMM-Newton</i> Lag Results | 117 |
| 5.5.2 | <i>Swift</i> Lag Results | 117 |
| 5.5.2.1 | Reanalysis of CCF lags | 117 |
| 5.5.2.2 | Detrending with LOWESS | 119 |
| 5.5.2.3 | Testing alternate explanations | 121 |
| 5.5.3 | <i>AstroSat</i> Lag Results | 124 |
| 5.5.4 | Model Lag Spectra | 130 |
| 6 | Conclusions | 133 |
| 6.1 | Main Results | 133 |
| 6.1.1 | Optical Interband Lags of NGC 4395 | 133 |
| 6.1.2 | X-ray to UV/Optical Lags of NGC 4395 | 134 |
| 6.1.3 | Multiwaveband Power Spectra of NGC 4395 | 134 |
| 6.1.4 | Timescale-dependent Lags of NGC 4593 | 135 |
| 6.2 | Future Work | 136 |
| 6.3 | Final Remarks | 137 |
| | Appendix A Comparison of KYNxltr lags with our in-house lags | 139 |
| | Appendix B Cross-Correlation and Javelin Methodology | 141 |
| | Appendix B.1 Javelin | 141 |
| | Appendix B.2 Cross-Correlation Function | 142 |
| | Appendix C NGC 4593 Residual Model Plot | 143 |
| | References | 145 |

List of Figures

| | | |
|------|---|----|
| 1.1 | A diagram of the Unified Model of AGN. | 4 |
| 1.2 | A Typical X-ray Spectrum of an unabsorbed AGN. | 12 |
| 1.3 | A Typical AGN SED. | 14 |
| 1.4 | A simplified diagram of X-ray reprocessing. | 17 |
| | | |
| 2.1 | Optical Spectrum of NGC 4395. | 22 |
| 2.2 | g -band image of NGC 4395 from the Liverpool Telescope. | 26 |
| 2.3 | Ratio of star 2 and star 0's g -band count rate. | 27 |
| 2.4 | Night 1: Lightcurves of NGC 4395 from the Liverpool Telescope. | 27 |
| 2.5 | Night 2: Lightcurves of NGC 4395 from the Liverpool Telescope. | 28 |
| 2.6 | Night 3: Lightcurves of NGC 4395 from the Liverpool Telescope. | 28 |
| 2.7 | Night 4: Lightcurves of NGC 4395 from the Liverpool Telescope. | 29 |
| 2.8 | Night 5: Lightcurves of NGC 4395 from the Liverpool Telescope. | 29 |
| 2.9 | DCF of the B-band lag by the g band for LT night 2 (Fig. 2.5). | 30 |
| 2.10 | DCF of the r band lag by the i band for LT night 3 (Fig. 2.6). | 30 |
| 2.11 | JAVELIN lag distribution of the B and g LT night 2 lightcurves (Fig. 2.5). | 31 |
| 2.12 | JAVELIN lag distribution of the g and i LT night 3 lightcurves (Fig. 2.6). | 31 |
| 2.13 | HiPERCAM u_s -band image. | 32 |
| 2.14 | Tracking and seeing plots of the HiPERCAM data. | 33 |
| 2.15 | Normalised ratio of the g_s -band count rates of stars 1 and B. | 34 |
| 2.16 | Observed u_s, g_s, r_s, i_s , and z_s HiPERCAM lightcurves for NGC 4395. | 35 |
| 2.17 | Polynomial fit to the g_s -band HiPERCAM lightcurve | 36 |
| 2.18 | Residuals from Fig. 2.17 vs seeing. | 37 |
| 2.19 | The first 8000s of the g_s -band lightcurve and the corrected lightcurve. | 37 |
| 2.20 | HiPERCAM lag probability distributions from JAVELIN. | 40 |
| 2.21 | UV and Optical lags vs formulated Kammoun et al. (2021a) model. | 43 |
| 2.22 | UV and Optical lags vs KYNreverb model. | 47 |
| 2.23 | UV and Optical lags vs varying KYNreverb f_{col} | 47 |
| 2.24 | UV and Optical lags vs low-mass KYNreverb model. | 48 |
| | | |
| 3.1 | Finding chart of the stars around NGC 4395. | 63 |
| 3.2 | Star 2 compared to the other three comparison stars. | 64 |
| 3.3 | Star 1 compared to the other comparison stars. | 65 |
| 3.4 | UVM2 time lag distribution for the first observation period. | 69 |
| 3.5 | Optical time lag distribution for the first observation period. | 70 |
| 3.6 | UVM2 CCF time lag distribution for the first observation period. | 70 |
| 3.7 | The full light curves from the night of December 13th in all 3 wavebands. | 70 |
| 3.8 | UVM2 time lag distribution for the second observation period. | 71 |

| | | |
|------|---|-----|
| 3.9 | Optical time lag distribution for the second observation period. | 71 |
| 3.10 | UVM2 CCF time lag distribution for the second observation period. . . . | 71 |
| 3.11 | The full light curves from the night of December 19th in all 3 wavebands. | 72 |
| 3.12 | Optical time lag distribution for the third observation period. | 72 |
| 3.13 | The full light curves from the night of December 31st in X-ray and optical. | 73 |
| 3.14 | U-band time lag distribution for the fourth observation period. | 73 |
| 3.15 | Optical time lag distribution for the fourth observation period. | 73 |
| 3.16 | U-band CCF time lag distribution for the fourth observation period. . . | 74 |
| 3.17 | The full light curves from the night of January 2nd in all 3 wavebands. . | 74 |
| 3.18 | X-ray to UVM2 Confidence Contours for the first observation period. . . | 76 |
| 3.19 | X-ray to UVM2 Confidence Contours for the second observation period. | 76 |
| 3.20 | X-ray to U-band Confidence Contours for the fourth observation period. | 76 |
| 3.21 | JAVELIN lags against Shakura & Sunyaev model, k-error. | 79 |
| 3.22 | JAVELIN lags against Shakura & Sunyaev model, a-error. | 80 |
| 3.23 | HiPERCAM SDSS transmission curves. | 80 |
| 3.24 | JAVELIN lags against Shakura & Sunyaev model, forced zero intercept. | 81 |
| 3.25 | Simulated KYNxiltr lag spectra, showing various R_{out} values. | 83 |
| 3.26 | X-ray to Convolved Model lags, model comparison. | 84 |
| 3.27 | $1900R_g$ KYNxiltr half-light point lag spectra, f_{col} variation. | 85 |
| 3.28 | X-ray to Convolved Model lags using the $1900R_g$ response functions. . . | 86 |
| 3.29 | Each night's UV/Optical data plotted with the model light curves. . . . | 87 |
| 3.30 | JAVELIN + HiPERCAM lags against Shakura & Sunyaev model. | 90 |
| 3.31 | JAVELIN + HiPERCAM lags against KYNxiltr model. | 90 |
| 4.1 | <i>XMM-Newton</i> light curve for NGC 4395 from January 2nd 2019 | 92 |
| 4.2 | <i>Swift</i> UVOT light curve for NGC 4395 from Cameron et al. (2012) | 93 |
| 4.3 | TESS light curve for NGC 4395 using data from Burke et al. (2020) | 94 |
| 4.4 | SDSS g' band light curve for NGC 4395 from LT, LCOGT, ZTF, and Zowada. | 94 |
| 4.5 | GTC HiPERCAM SDSS g' band light curve. | 95 |
| 4.6 | A 140ks sample of the simulated X-ray light curve. | 98 |
| 4.7 | A 140ks sample of the synthetic g' -band light curve. | 99 |
| 4.8 | The optical data PSD with the best-fit for an unbending power law. . . . | 100 |
| 4.9 | The optical data PSD with the best-fit for a single-bend power law. . . . | 101 |
| 4.10 | The X-ray data PSDs with the best-fit for a single-bend power law. . . . | 102 |
| 4.11 | The X-ray data PSDs with the X-ray simulation PSDs. | 102 |
| 4.12 | The synthetic g' -band PSD with the best-fit for a single-bend power law. | 103 |
| 4.13 | The optical data PSD with the best-fit for a double-bend power law. . . . | 104 |
| 4.14 | Simulated 10 day light curve using the Timmer and König (1995) method. | 105 |
| 4.15 | Synthetic g' plus TK Sim PSD with best-fit for a double-bend power law. | 105 |
| 5.1 | Optical Spectrum of NGC 4593. | 112 |
| 5.2 | <i>XMM-Newton</i> EPIC pn X-ray and OM UV light curves for NGC 4593. . . | 114 |
| 5.3 | <i>Swift</i> XRT X-ray and UVOT UV light curves for NGC 4593. | 115 |
| 5.4 | <i>AstroSat</i> SXT X-ray and UVIT UV light curves for NGC 4593. | 116 |
| 5.5 | <i>XMM-Newton</i> EPIC pn to OM UVW1 FR CCF Lag Distribution. | 117 |
| 5.6 | <i>XMM-Newton</i> EPIC pn and shifted OM UVW1 light curves. | 118 |
| 5.7 | <i>Swift</i> XRT 0.5-10keV to UVOT UVW1 raw lightcurve CCF Lag. | 119 |

| | | |
|--------------|--|-----|
| 5.8 | <i>Swift</i> XRT to UVOT 2.5-day and 7.5-day Lowess Detrended CCF Lags. . . | 120 |
| 5.9 | Lags vs Lowess Smoothing Fraction for the <i>Swift</i> UVOT lightcurves. . . | 120 |
| 5.10 | <i>Swift</i> XRT to UVOT 5-day Lowess Detrended and Smoothed CCF Lags. . | 121 |
| 5.11 | <i>Swift</i> UVOT UVW1 5-day Lowess Detrended and Smoothed Light Curves. | 122 |
| 5.12 | <i>Swift</i> XRT to UVOT High Intensity Half & Low Intensity Half CCF Lags. | 123 |
| 5.13 | <i>Swift</i> XRT and UVOT plotted along with the <i>XMM-Newton</i> light curves. | 123 |
| 5.14 | <i>Swift</i> XRT to <i>XMM-Newton</i> OM UVW1 CCF Lag Distribution. | 124 |
| 5.15 | <i>AstroSat</i> SXT and UVIT plotted along with the <i>XMM-Newton</i> light curves. | 125 |
| 5.16 | <i>AstroSat</i> SXT to FUV & <i>XMM-Newton</i> EPIC pn to FUV CCF Lags. | 126 |
| 5.17 | <i>AstroSat</i> SXT vs FUV & SXT vs shifted FUV light curves. | 126 |
| 5.18 | <i>AstroSat</i> FUV 4-day Lowess Subtracted and Smoothed Light Curves. . . | 127 |
| 5.19 | <i>AstroSat</i> SXT to Lowess-Subtracted FUV & SXT to NUV CCF Lags. . . . | 128 |
| 5.20 | <i>XMM-Newton</i> EPIC pn to Subtracted FUV & EPIC pn to NUV CCF Lags. | 129 |
| 5.21 | <i>AstroSat</i> SXT to <i>XMM-Newton</i> OM UVW1 CCF Lag. | 129 |
| 5.22 | Our UVW1 and UVIT CCF lags compared to the model lag spectra. . . . | 131 |
| 5.23 | Our, McHardy et al. (2018) 's, & Cackett et al. (2018) 's lags vs model. . . . | 131 |
| Appendix A.1 | As for Fig. 2.22 but the model lags are from our original in-house code with zero spin. | 140 |
| Appendix A.2 | As for Fig. 2.22 but the model lags are from our original in-house code with maximum spin. | 140 |
| Appendix C.1 | NGC 4593 Residual data lags from the Max Spin Model, and the BLR lag shape from Korista & Goad 2019. | 143 |

List of Tables

| | | |
|-----|--|----|
| 2.1 | Log of LT observations of NGC 4395. | 25 |
| 2.2 | Fractional variability of the HiPERCAM observations | 38 |
| 2.3 | Lags for the first 8000s of the HiPERCAM observations. | 41 |
| 3.1 | Space-based Observations from <i>XMM-Newton</i> | 61 |
| 3.2 | Optical magnitudes for the objects, taken from SDSS. | 61 |
| 3.3 | Ground-based Observations from multiple telescopes | 62 |

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as: [McHardy et al. \(2022\)](#) and [Beard et al. \(2022\)](#)

Signed:.....

Date:.....

Acknowledgements

My first and biggest thanks goes my supervisor Ian McHardy, whose breadth of knowledge and enthusiasm for the subject is unending. His support and advice over the past four years have taught me lessons I will carry with me for the rest of my life, and it has been a honour to be involved in research work alongside him.

Thank you to my consistently supportive family and incredible friends for always being there for me throughout this process, keeping my spirits up during tougher times, and always encouraging me when I needed it.

I'd like to extend specific thanks to Robert Rosca for being my eternally helpful software guru, and to Daniel Andrews, Michael Overy, and James Baker for their always valued insight and willingness to help out with proof-reading.

Finally I would like to thank the Science and Technology Funding Council for the financial support that allowed me to engage in this research work.

*Dedicated to Henry Haslam, the best grandad I could have ever
hoped for.*

Abbreviations

| | |
|-------------------|--|
| AGN | Active Galactic Nuclei |
| ADAF | Advection Dominated Accretion Flow |
| ARIES | Aryabhata Research Institute of Observational ScienceES |
| BLR | Broad Line Region |
| CCF | Cross Correlation Function |
| DCF | Discrete Correlation Function |
| EPIC | European X-ray Imaging Camera |
| FR | Flux-Randomisation |
| FWHM | Full Width Half Maximum |
| GTC | Gran Telescopio Canarias |
| HST | Hubble Space Telescope |
| ICCF | Interpolated Cross-Correlation Function |
| ISCO | Innermost Stable Circular Orbit |
| JAVELIN | Just Another Vehicle for Estimating Lags In Nuclei |
| LCOGT | Las Cumbres Observatory Global Telescope |
| LINER | Low-Ionisation Nuclear Emission-Line Region |
| MCMC | Markov Chain Monte Carlo |
| NLR | Narrow Line Region |
| OM | Optical Monitor |
| PDF | Probability Density Function |
| PSD | Power Spectral Density |
| QSO | Quasi-Stellar Object |
| RIAF | Radiatively Inefficient Accretion Flow |
| RSS | Random Subset Selection |
| SAS | Scientific Analysis System |
| SDSS | Sloan Digital Sky Survey |
| SNR | Signal-to-Noise Ratio |
| SXT | Soft X-ray Telescope |
| UVIT | Ultra Violet Imaging Telescope |
| UVOT | Ultra Violet and Optical Telescope |
| XMM-Newton | X-ray Multi-Mirror Mission Newton |
| XRT | X-Ray Telescope |

Physical Constants

| | | | |
|--------------------------------------|---------------|-----|---|
| Boltzmann Constant | k_B | $=$ | $1.380\,648\,52 \times 10^{-23} \text{ m}^2 \text{ kg s}^{-2} \text{ K}^{-1}$ |
| Electron volt | eV | $=$ | $1.602\,176\,57 \times 10^{-19} \text{ J}$ |
| Euler's number | e | $=$ | $2.718\,281\,828$ |
| Gravitational Constant | G | $=$ | $6.674\,08 \times 10^{-18} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$ |
| Parsec | pc | $=$ | $3.085\,677\,6 \times 10^{16} \text{ m}$ |
| pi | π | $=$ | $3.141\,592\,65$ |
| Planck Constant | h | $=$ | $6.626\,070\,04 \times 10^{-34} \text{ m}^2 \text{ kg s}^{-1}$ |
| Proton mass | m_p | $=$ | $1.672\,621\,58 \times 10^{-27} \text{ kg}$ |
| Solar Mass | M_\odot | $=$ | $1.988\,55 \times 10^{30} \text{ kg}$ |
| Stefan Boltzmann Constant | σ_{SB} | $=$ | $5.670\,367 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ |
| Thomson Cross-Section of an Electron | σ_T | $=$ | $6.652\,458\,734 \times 10^{-29} \text{ m}^2$ |

*"The most important step a man can take. It's not the first one, is it?
It's the next one. Always the next step."*

Brandon Sanderson, *Oathbringer*

Chapter 1

Introduction

Overview

The aim of this thesis is to provide further insight and understanding into the physical processes driving X-ray, UV, and Optical emission from Active Galactic Nuclei (AGN) by studying the variability properties of such emissions and linking the observations to theory by use of modelling. A large part of this investigation will be focused on a particularly low-mass AGN, designated NGC 4395, whose low-mass allows us to study processes that in larger objects would only reveal themselves over decades or centuries. Towards the end, we also shift focus onto another object, designated NGC 4593, and investigate the apparent presence of timescale-dependent variability properties.

In this chapter, I discuss the current theoretical understanding of AGN derived from both observation and modelling to give a framework with which we can then describe the new observations presented in this work. This overview will cover radiation emission and absorption mechanisms, the physical processes of accretion, the unified model of AGN, and the variability properties of AGN, including the use of the power spectrum as a method of characterising this variability. The power spectrum will also highlight the similarities between X-ray binaries (XRBs) and AGN, and display how the physics of accretion appears to drive both types of objects. I will then give an overview of this thesis and address its main aims.

1.1 Active Galactic Nuclei

When astronomers measure the flux emitted from galaxies, they sometimes find that there is a strong source of emission in the galaxies core. Upon further investigation it has been found that this emission is both very powerful ($> 10^{41} \text{ erg s}^{-1}$) and varies on timescales from seconds to months. This implies an extremely physically compact emission region (due to the speed of light limiting the variability timescale) that is simultaneously generated huge amounts of energy. Therefore it has become widely accepted among the astronomical community that at the centre of most galaxies is a supermassive black hole, ranging in mass from thousands to billions of solar masses, and mostly accreting matter from their surrounding environments. A majority of these black holes appear to be accreting at low rates, and are faint relative to the rest of their host galaxy, these black holes are described as being quiescent. However there is a smaller, but still significant (as much as $\sim 43\%$, [Ho et al. \(1997\)](#)), fraction of these galaxies that accrete at higher rates and appear to be strong emitters of X-ray and UV light, sometimes enough to outshine the entire rest of their galaxy. These high accreting black holes are called Active Galactic Nuclei.

The range of AGN bolometric luminosities is huge, at low-end are objects like the focus of this investigation, NGC 4395, while at the high-end you have objects like the recently discovered SMSS J114447.77-430859.3 ([Onken et al., 2022](#)) which has a luminosity $> 10^{47} \text{ erg s}^{-1}$. AGN classification has traditionally be done via observational characteristics, though we will see later that these do not necessarily distinguish galaxies by actual physical characteristics.

Low redshift AGN have historically been distinguished by their optical features, due to the limitations of human observations before the advent of satellite observatories, as AGN also emit large amounts of X-rays and UV light. A large percentage of local AGN are Low-Ionisation Nuclear Emission-line Regions (LINERs) which tend to be dominated by low-ionisation optical emission lines, have low accretion rates, and for the most part are hard to directly observe.

Around a quarter ([Ho et al., 1997](#)) of local AGN are referred to as "Seyfert galaxies" (Named after Carl Seyfert who discovered the object prototypical of its class, NGC 1068) and possess very bright nuclei, strong optical emission lines, and typically have higher accretion rates than LINERs. Of these Seyfert galaxies, they are designated between Seyfert 1 and Seyfert 2 galaxies depending on the type of emission lines they display. Seyfert 2 galaxies display relatively narrow emission lines, whereas Seyfert 1 galaxies display extremely broad emission lines with narrow lines observed within them. [Osterbrock \(1977\)](#) showed it is also possible to define intermediate Seyferts between 1 and 2 based on the ratio of $\text{H}\alpha$ and $\text{H}\beta$ in the galaxy's optical spectrum. There are also AGN that appear very similar to Seyferts in their optical spectrum but

also have large amounts of radio emission (defined as being "radio-loud"). Those analogous to radio-loud Seyfert 1's are classed Broad-Line Radio Galaxies (BLRGs) and those analogous to radio-loud Seyfert 2's are classed Narrow-Line Radio Galaxies (NLRGs).

At higher redshifts, we observe much more luminous galaxies than even Seyferts, with nuclei brighter than the entire host galaxy, hence why we are able to observe them at such large distances. These objects are referred to as Quasi-Stellar Objects (QSOs) or simply as Quasars. In addition to their huge X-ray, UV, and Optical emission, some Quasars appear to be radio-loud as well, with some displaying enormously bright radio emission being separately classified as "Blazars". These luminosity distinctions between Quasars and Seyfert galaxies appear to be mostly arbitrary and the result of the limitations of observational equipment, rather than a fundamental physical difference, as it is believed that all of these processes are driven by the accretion of material onto a supermassive black hole.

1.1.1 The Unification of AGN classes

The similarities between these many classes of AGN seem to outweigh their differences, and given that we suspect they are all driven by the same physical process of accretion onto the central engine, there is a mechanism by which we can explain the differences between these many otherwise similar classes of objects. Modern observations have shown that the radio emission of AGN originates in powerful jets of material ejected, most commonly, perpendicular to the accretion disc. While the majority of the higher frequency emission from Optical up to X-rays is produced much closer to the black hole, either by the accretion disc itself or very close to it.

Many differences in the radiation emitted from AGN can be ascribed to the accretion rate and mass of the black hole. Increasing the accretion rate will generally yield a hotter disc, while increasing the mass of the central engine will yield a larger and cooler disc. This will therefore affect the spectrum (due to disc temperature changes), the variability timescales (due to the size of the disc), and the overall luminosity (due to the combination of the two). Therefore we can classify that quasars can be explained by high-mass black holes with high accretion rates, and LINERs and Seyferts can be classified by low-mass black holes with low and intermediate accretion rates, respectively. These two factors however do not explain the differences between the Seyfert subclasses, or the radio-loud/radio-quiet dichotomy, and so the question of whether these objects can be unified under a single model naturally arises.

A proposal for such a unifying model for AGN based on the angle with which the AGN is viewed was first proposed by [Antonucci and Miller \(1985\)](#). The principle

behind this model is to make the minimum number of assumptions required to explain each class of AGN, as such the model proposed is diagrammed in Fig. 1.1.




FIGURE 1.1: A diagram of the Unified Model of AGN.

This model proposes that, as expected, at the centre of the AGN is the black hole accreting matter onto itself in the form of an accretion disc. Further out from the disc is gas that is ionised by the X-ray emission of the central engine that orbits at relativistic velocities and leads to emission with broad lines, due to Doppler Broadening. This region is labelled the Broad Line Region (BLR). Much further out are lower-density slow-moving gas clouds which are responsible for narrow emission lines and is therefore labelled as the Narrow Line Region (NLR). A very important part of this model is then a cooler dusty "torus" that surrounds the BLR and absorbs a large amount of emission from the AGN when viewing the object through it. In addition, the AGN may or may not have a jet that is responsible for radio emission.

This is then designed to explain the differences in the classes of AGN. When the viewing angle (relative to the plane of the disc) of the AGN is large, the central region is unobscured and broad lines would be observed, which in radio-quiet galaxies would appear as a Seyfert 1 galaxy. Alternately, when the viewing angle is small, the "torus" obscures the BLR and only narrow emission lines are observed, therefore leading to the classification as a Seyfert 2 galaxy.

In radio-loud galaxies, the same viewing angle differences would classify large viewing angle galaxies as BLRGs and small viewing angle galaxies as NLRGs, though

there is a further distinction. In [Fanaroff and Riley \(1974\)](#), the authors broke down radio-loud AGN into two classes, today referred to as Fanaroff-Riley-I (FR-I) and Fanaroff-Riley-II (FR-II) galaxies based on the strength of their jet. FR-I galaxies are radio-brighter towards the centre of the galaxy and get fainter towards the edge of the jet's lobes ("edge-darkened"), whereas FR-II galaxies are radio-brighter in the lobes ("edge-brightened"). Quasars aligned so that their jets are parallel to our line of sight towards the object are extremely bright objects called Blazars, with FR-I blazars being classified as "BL Lacterae Objects" (named after the prototypical object of its class) and FR-II blazars classified as Flat-Spectrum Radio Quasars (FSRQs).

Though this model has been very successful explaining much of the differences between these objects, there are still many things that remain unclear. Details of the accretion disc such as its exact shape or size, the location and features of the BLR, NLR, and "torus" are similarly not well constrained. In addition, [Elvis \(2000\)](#) added a wind blown up from the accretion disc to the model for AGN displaying absorption as well as emission lines both with large radial velocities. As such, observations that can constrain these values will help to determine the true nature of these objects.

1.2 Black Hole Accretion

In order to determine how energy is extracted from accreting material, we must look at the physical processes that occur as the accretion process takes place. To begin with we can represent material as a mass m orbiting a black hole with mass M at a distance of R . The energy, E of the mass can be given by the simple Newtonian expression of:

$$E = \frac{1}{2}mv^2 + \left(-\frac{GMm}{R} \right) \quad (1.1)$$

where the first term is the kinetic energy K and the second is the gravitational potential energy V . As we are dealing with a particle orbiting at a constant distance, we can define the velocity using the assumption of circular motion such that

$$v^2 = \frac{GM}{R} \quad (1.2)$$

which then allows Eqn. 1.1 to be reformulated such that the kinetic energy term can be represented as

$$K = \frac{GMm}{2R} \quad (1.3)$$

As the mass falls towards the black hole we can calculate the change in K and V by differentiating them with respect to R such that

$$\frac{dK}{dR} = -\frac{GMm}{2R^2} \quad (1.4)$$

and

$$\frac{dV}{dR} = \frac{GMm}{R^2} \quad (1.5)$$

Combining Eqs. 1.4 and 1.5 we therefore find that

$$\frac{dK}{dR} = -\frac{1}{2} \frac{dV}{dR} \quad (1.6)$$

Therefore to maintain an orbit, only half of the gravitational potential energy of the particle can be converted into kinetic energy, meaning that the other half must be lost through other avenues, namely radiation emission. However unlike with accretion onto an object with a physical surface such as a neutron star, it is not possible to extract any more energy once the particle crosses the event horizon, and therefore once the particle reaches the innermost stable circular orbit (ISCO) any remaining energy of the particle is lost. We can therefore represent the energy ΔV that can be extracted as

$$\Delta V = \int_{R_{ISCO}}^{\infty} \frac{GMm}{R^2} dR = \frac{GMm}{R_{ISCO}} \quad (1.7)$$

where R_{ISCO} is the spin-dependent radius of the ISCO. For a black hole with zero spin, also called a Schwarzschild black hole, R_{ISCO} can be written as

$$R_{ISCO} = \frac{6GM}{c^2} \quad (1.8)$$

which can be substituted into Eqn. 1.7 to give

$$\Delta V = \frac{1}{6}mc^2 \quad (1.9)$$

We can therefore see that one sixth of the rest mass of the particle is converted from potential energy to other forms. As we know half of this goes into the kinetic energy of the particle, we can therefore determine that 1/12th of the rest mass-energy is released, which gives an efficiency, η , of $\eta = 0.0833$. The liberated energy for any given spin then can be given by

$$E = \eta mc^2 \quad (1.10)$$

For a spinning black hole, also called a Kerr black hole, the ISCO is smaller, which means that maximally spinning black holes can have η up to ~ 0.42 , therefore making accretion one of the most efficient mechanisms of energy production in the Universe, allowing the production of the huge luminosities that AGN are observed to have.

Now while this process works perfectly well in theory on a single test mass, in reality an accretion disc has a lot of material which must all radiate in order to fall towards the black hole. This means that some of the radiation from the innermost material must interact with the material further out in the accretion disc in order to escape out into the Universe. This radiation pressure could potentially prevent more material from falling into the disc, much as main sequence stars are prevented from collapsing by their internal radiation pressure. Unlike a star however, if the process of accretion were halted in this way there would be no more luminosity produced, which we do not observe, and therefore the infalling matter must be radiating in a way that does not interfere with the process of accretion in AGN.

We can examine this by looking at the momentum carried by an emitted photon, p_γ , which can be represented as

$$p_\gamma = \frac{E_\gamma}{c} \quad (1.11)$$

where E_γ is the energy of the photon.

We can then define the total force exerted by the radiation pressure, F_{tot} , of x photons as

$$F_{tot} = \frac{d}{dt} p_{\gamma,tot} = \frac{d}{dt} \sum_{n=1}^x p_{\gamma,x} = \frac{1}{c} \frac{d}{dt} \sum_{n=1}^x E_{\gamma,x} \quad (1.12)$$

We also can define the Bolometric Luminosity, L_{bol} , as the total power of the radiation, written as

$$L_{bol} = \frac{d}{dt} E_{\gamma,tot} = \frac{d}{dt} \sum_{n=1}^x E_{\gamma,x} \quad (1.13)$$

We can then substitute this into Eqn. 1.12 to find

$$F_{tot} = \frac{1}{c} L_{bol} \quad (1.14)$$

Assuming isotropic radiation emission we can then find the total radiation pressure at distance d using the relation $F_{tot} = P_{rad} A_d$ where P_{rad} is the pressure and A_d is the surface area at distance d and therefore

$$P_{rad} = \frac{L_{bol}}{A_d c} = \frac{L_{bol}}{4\pi d^2 c} \quad (1.15)$$

Which can then give us the force that the radiation pressure exerts on an individual particle, F_{rad} , as

$$F_{rad} = \sigma_T P_{rad} = \frac{\sigma_T L_{bol}}{4\pi d^2 c} \quad (1.16)$$

where σ_T is the Thomson cross-section of a given particle.

In order for accretion to be prevented, this force needs to be as large as the gravitational force exerted on a particle of mass m . The gravitational force is given by

$$F_g = \frac{GM_{BH}m}{d^2} \quad (1.17)$$

where M_{BH} is the mass of the black hole.

The Thomson cross-section and mass both change depending on the particle that is affected. If we assume that the accreting material is molecular hydrogen then we have a simple atom with one proton and one electron. Each particle contributes far more than the other to each variable, with the mass of the proton, m_p , being far larger than the electron's mass, and conversely the Thomson cross-section of the electron is far larger than the proton's. As such we use the most contributing values for each variable, and by equating Eqns. 1.16 and 1.17 to find this maximum luminosity of accretion, which is known as the **Eddington Luminosity**, L_{Edd} , we find it is given by

$$L_{Edd} = \frac{4\pi GM_{BH}m_p c}{\sigma_T} \quad (1.18)$$

and therefore we can define the maximum rate of accretion, known as the Eddington accretion rate, as

$$\dot{M}_{Edd} = \frac{L_{Edd}}{\eta c^2} = \frac{4\pi GM_{BH}m_p}{\sigma_T \eta c} \quad (1.19)$$

Although this is a simplified model, and spherically symmetric (or "Bondi" [Bondi \(1952\)](#)) accretion and pure molecular hydrogen are not upheld assumptions in real AGN systems, it is still a very useful approximation, and the ratio of the Bolometric and Eddington luminosities (creatively called the Eddington ratio) can still be used to approximate the accretion rate of an object.

One of the problems with Bondi accretion is that it would require infalling material to have zero initial angular momentum (which would not be true in a real gas cloud) and would also not be able to produce the large L_{bol} we see in AGN as it would have low radiative efficiency and therefore accretion rate. For a more realistic model, we can look at [Shakura and Sunyaev \(1973\)](#) which, while initially developed for XRBs, we will see can be applied to AGN.

The argument presented in this paper posited that in an XRB, as it is accreted from a companion star, the material would possess a large amount of angular momentum which needs to be lost somehow as the material falls inwards. Therefore, due to outward transport of this angular momentum, their model forms a steady-state disc with the material in Keplerian orbits.

This disc has viscosity v which can be formulated as

$$v = \alpha c_s H \quad (1.20)$$

where H is the height of the disc, c_s is the speed of sound in the disc, and α is a collective parameter which accounts for other unknown aspects of physics that interact with this viscosity.

This model then tends to give very small heights to the disc, at least relative to the radius of the disc elements, and is therefore referred to as the "thin-disc" or " α -disc" model. This disc model is then much more radiatively efficient and therefore compatible with the L_{bol} we observe for real AGN.

We can retrieve a temperature profile from this model by making the assumption that it is optically thick and therefore approximates a blackbody. This means we can then equate the total energy emitted by the blackbody disc, given by Planck's law ($\sigma_{SB}T^4$), with the energy dissipated at radius R . The local dissipation rate was found by [King et al. \(2002\)](#) to be independent of α and given by

$$D(R) = \frac{3GM_{BH}\dot{m}}{8\pi R^3} \left[1 - \sqrt{\frac{R_0}{R}} \right] \quad (1.21)$$

where \dot{m} is the accretion rate and R_0 is the inner radius of the accretion disc. Therefore this equality gives us a temperature profile of

$$T(R) = \left(\frac{3}{8\pi\sigma_{SB}} \frac{GM_{BH}}{R^3} \dot{m} \left[1 - \sqrt{\frac{R_0}{R}} \right] \right)^{1/4} \quad (1.22)$$

In addition, [Novikov and Thorne \(1973\)](#) described a relativistic generalisation of this model (the "Novikov-Thorne model") that uses a Kerr black hole in place of a Schwarzschild black hole which allows it to introduce spin to the model. It also takes into account material that has crossed the ISCO and is therefore free-falling towards the black hole in the so-called "plunging zone". Though the difference in this model compared to [Shakura and Sunyaev \(1973\)](#) are not particularly relevant to the work presented here compared to other factors we will present later.

However, for low accretion rates we would expect the material to be optically thinner and therefore radiatively inefficient, which would make the assumption of an α -disc model not apply. Instead, this radiatively inefficient accretion flow is thought to consist of a geometrically-thick disc that consists of ionised plasma. The ions in this plasma are poor at radiating and therefore gain higher temperatures, causing the height of the disc to increase, while the electrons, which radiate much more easily, are much cooler.

Narayan and Yi (1994) presented one of the best-known models of this kind of accretion, where some of the energy that is released from accretion is advected into the black hole, and causing energy that otherwise went into heating the disc material to be lost instead of radiated. This model is called advection-dominated accretion flow, commonly referred to as ADAF.

1.3 The Emission Mechanisms of AGN

1.3.1 X-ray Emission

Almost all AGN emit strongly in the X-ray, and these emissions are also the most rapidly variable. This implies that the source of emission is physically compact, and therefore likely very close to the black hole. This X-ray source is now believed to be a "corona" of very high-energy electrons which lies above the inner region of the accretion disc (and is likely an extension of it in some way) and produces X-rays through Inverse Compton scattering ("Compton up-scattering") of UV seed photons. The heating mechanism of this corona is not known, but similar to the Solar corona, these high temperatures could result from magnetic reconnection driven by turbulence in the accretion disc, as proposed in di Matteo (1998).

This Comptonisation process leads to a spectrum with a power-law form (and a high-energy cut-off), such that it gives a photon number, N_γ , distribution of

$$N_\gamma(E) = nE^{-\Gamma} \quad (1.23)$$

where E is the photon energy, n is the normalisation, and Γ is called the "photon index". This then leads to a photon energy, E_γ , distribution of

$$E_\gamma(E) = nE^{-\alpha} \quad (1.24)$$

where α is the spectral index such that $\alpha = \Gamma - 1$. Lightman and Zdziarski (1987) modelled this Comptonisation of thermal UV photons to find an expression for Γ (which is dependant on the energy distribution of the electrons) given by

$$\Gamma = \left[\frac{9}{4} + \frac{1}{\tau(k_B T / m_e c^2)(1 + \tau/3)} \right]^{1/2} - \frac{5}{2} \quad (1.25)$$

where T is the electron cloud temperature and τ is the scatter depth. Therefore, when the corona is hotter, Γ is lower, and the spectrum is flatter with higher-energy emission.

Another method for X-ray production in AGN is Synchrotron radiation, however this emission also produces a power-law spectrum and can therefore be difficult to distinguish from the Comptonised X-rays. In principle it is possible due to polarisation of Synchrotron emission, but few X-ray polarimeters exist to test this.




FIGURE 1.2: A Typical X-ray Spectrum of an unabsorbed AGN and its disentangled components. From Ricci (2011).

When we observe real X-ray spectra however, though we see these power-law components, we also see other features, as can be seen in Fig. 1.2. A particularly striking feature is a strong line at ~ 6.4 keV which is a fluorescent K_{α} line of Iron. As the energy of an emission line depends on the ionisation state of the iron atom, we can determine the state of the emitting atoms by the energy of the line. This particularly energy implies that the iron is largely *not* ionised, which would therefore suggest that it is relatively cold ($< 10^6$ K).

The fact that this line comes from a heavy element and also tends to be broadened implies that this line originates from the X-ray illumination of the accretion disc. We also observe significant *gravitational* redshift in these iron lines in addition to Doppler broadening, which suggests that it originates in the inner accretion disc.

In addition to this, we see a high energy hump in the spectrum. This can be explained by the way that the X-ray emission interacts with the accretion disc. For a corona that

is emitting isotropically, around half of its flux will interact with the (relatively) cold and non-ionised accretion disc, at which point it will interact with the matter in different ways depending on its energy.

At low energies ($E < 10$ keV) the photons will be photoelectrically absorbed by the metals in the disc, so we should not see much extra X-ray flux in this regime. At intermediate energies ($10 < E < 40$ keV) the photons will instead be Thomson scattered, a fraction of which will be scattered up and away from the disc and therefore be observable. At high energies ($E > 40$ keV) a regime is reached where the Thomson cross-section, which is a classical limit, is no longer valid, and one must instead use the more general quantum mechanical Klein-Nishina cross-section to describe the scattering. Unlike Thomson scattering however, Klein-Nishina scattering is preferentially forward directed, leading to scattered photons penetrating further into the disc and losing energy. This loss of energy means that photoelectric absorption once again becomes important and the initially high energy photons are absorbed by the disc.

The result of these scattering and absorption processes of the disc is that some of the incoming X-ray radiation is scattered back out into the Universe with a modified spectral shape, peaked at ~ 30 keV where Klein-Nishina effects and photoelectric absorption are not as important. This leads to the feature that is typically called the "Compton hump" which is observed in AGN X-ray spectra.

The final feature is a so-called "Soft excess" which appears below 1 keV. Unlike the other two, the origin of this feature remains unsolved, though there are two popular explanations of this excess. The first is that there is a second non-thermal emission region, a so-called "warm corona" (as opposed to the hot corona) further out from the black hole, perhaps this warm Comptonisation region represents the very inner edge of the accretion disc, a transition point from thermal disc to non-thermal corona. The second is that the excess instead originates from a collection of highly blurred X-ray fluorescent line emission from very close to the black hole, where gravitational redshift and Doppler broadening are significant enough that the lines become indistinguishable from one another.

Recently [Ursini et al. \(2020\)](#) used the Nuclear Spectroscopic Telescope Array (NuSTAR) and *XMM-Newton* to observe the AGN HE 1143-1810. With NuSTAR's high-energy sensitivity to constrain the power-law, *XMM-Newton* observed the softer emission where the excess is seen, and using the combination broadband X-ray coverage they examined how the spectrum changed with time and found that their observations favoured the warm Comptoniser origin over the line emission model. However, [García et al. \(2019\)](#) analysed Markarian 509 with NuSTAR and Suzaku found that while both models were acceptable statistical fits to their data, they bring up an issue of physics which is that a warm Comptoniser should produce strong

absorption features that are not observed, and therefore favoured the blurred line model. Even this though does not entirely settle the debate as it has been shown more recently that absorption lines can be avoided if the warm Comptoniser has an internal source of heating. Therefore as to the exact nature of the Soft excess there is still much debate and ongoing research.

1.3.2 UV/Optical Emission

As we saw in Section 1.2, the infalling matter of an accretion disc will thermally emit blackbody photons with a temperature profile given by Eqn. 1.22 and will therefore create a total spectrum given by the sum of the blackbody component at each radius with the luminosity of each component scaled by the area of each of its emitting region. The peak for this total spectrum in AGN is in the UV which means the disc emits strongly in the UV and Optical regimes, this component can be seen in the SED of the AGN as shown in Fig. 1.3. There is also UV and Optical emission from the BLR and NLR which can be seen from the emission lines.




FIGURE 1.3: A Typical Spectral Energy Distribution (SED) of an unabsorbed AGN and its disentangled components. Courtesy of F. Shankar.

1.4 The Variability of AGN

The emitted flux from an AGN can be varied in many ways, the strength/absence of the jet will be the main source of radio flux, however for this work will be focusing on

the X-ray, UV, and Optical variability of AGN and so will not be discussing radio emission much. One way in which the flux of a particular AGN can change is through a change in accretion rate. As we saw in Eqn.1.22, by increasing accretion rate, the material in the accretion disc will increase in temperature, which increases both the flux and the peak emission wavelength, as it is emitting as an approximate blackbody. However in addition to this, this should also increase the amount of seed photons entering the corona and therefore cool it down, resulting in a softer power-law emission.

Another avenue of variability is the geometry of the corona, increasing the size of the corona would increase the X-ray flux while also hardening the X-ray spectrum as more seed photon paths would interact with the corona, increasing the number being Compton up-scattered. Disc geometry is also important, as we know from [Bardeen \(1970\)](#) that the efficiency of accretion, η , is dependent on the inner radius such that

$$\eta = 1 - \sqrt{1 - \frac{2}{3r_{inner}}} \quad (1.26)$$

where $r_{inner} = R_{inner}/R_{grav}$, R_{inner} is the inner radius of the accretion disc, and R_{grav} is the gravitational radius such that $R_{grav} = GM/c^2$. Therefore by increasing the inner radius of the accretion disc, the flux of the disc decreases as the efficiency of the accretion drops. This would mostly affect the high frequency emission which would therefore increase the wavelength of the disc spectrum's peak.

As we have discussed so far, there seem to be cases in AGN where the sources of flux at different wavelengths are somewhat dependent on one another, and therefore you would expect some form of causal link between them. Investigation of the causal links between variability at different wavelengths can therefore tell us a lot about the physical processes going on in an AGN. The physical size of an AGN is measurable, even for the smallest AGN, on the scale of light-seconds or more. We would therefore expect any links between different waveband components to exhibit a time delay or "lag", from the component that drives a variation to the component that is varied as a result. Methods of detecting and measuring these lags will be discussed in Section 3.5.1.

There are two main mechanisms of UV/Optical variability that have gained widespread support. The first is that intrinsic disc perturbations, caused by inwardly propagating accretion rate variations, drive UV seed photon variations, which in turn drive X-ray variations ([Arévalo et al., 2008](#); [Lira et al., 2011](#)). The second is the reprocessing of coronal X-ray emission by the accretion disc, with the X-ray variations instead driving the UV/Optical variations ([Sergeev et al., 2005](#); [Breedt et al., 2009](#)). It

is likely that both of these processes occur in reality, but testing which one is dominant can tell us about the structure of the system.

1.4.1 Intrinsic Disc Variations

As we discussed in Section 1.3.1, UV seed photons from the disc are Compton up-scattered to produce a non-thermal X-ray spectrum. One would therefore expect that the X-ray flux would be delayed relative to the disc photons caused by the additional travel time between the disc and the corona, as well as the timescale of the scattering process itself. UV emission from the very central portion of the disc is the dominant source of seed photons for the corona, so we would expect fairly short lags with the UV emission leading the X-rays.

However, this process assumes that the inner disc is cool enough to be emitting in the observed UV/Optical bands. If this is not true however, the disc could still drive X-ray variations by inwardly propagating accretion rate variability. As they travel inwards, eventually these fluctuations will reach the central disc and begin to produce seed photons. The timescale of this kind of lag is much longer than the simple travel time however, and is driven primarily by the viscous propagation timescale. This timescale is thought to be months or even years for standard accretion discs for black holes $\sim 10^6 M_\odot$ and is given in Lyubarskii (1997) as

$$t_{visc} = \frac{1}{\alpha} \left(\frac{R}{H} \right)^2 t_{dyn} \quad (1.27)$$

where α is the α -disc viscosity parameter, R and H are the radius and height of the disc, and t_{dyn} is the dynamical timescale. This would then cause X-ray variability to significantly lag behind the UV/Optical by months, even for a very small AGN like NGC 4395 this timescale would be on the scale of weeks. As in reality the disc is a continuous structure, this would likely be blurred and given the timescales involved, very hard to detect in observations.

1.4.2 X-ray Reprocessing

The existence of the Compton reflection hump in the X-ray spectrum implies that the corona is above the accretion disc, or at the very least has a significant component that extends above the accretion disc (i.e. has a larger height), so that the accretion disc is illuminated by this X-ray flux. Once again looking at the structure of the Compton hump, we know that the X-ray energies that do not appear in the reflected spectrum

are photoelectrically absorbed, either directly in the case of soft X-rays or by further energy loss caused by Klein-Nishina scattering in the case of very hard X-rays.




FIGURE 1.4: A simplified diagram of X-ray reprocessing, demonstrating the origin of the time delay in the reprocessed emission.

This absorption of X-rays will then heat the disc by an amount proportional to the amount of energy absorbed (i.e. the intensity of the incident radiation). As we are dealing with an approximate blackbody, the result of this is simple to see, an increase in both the luminosity and peak wavelength of the emitting material. As we observe variability in the Comptonised X-ray emission, we know that the disc will experience variability in its irradiation by X-rays, which will in turn cause variability in the disc's luminosity. Therefore we should observe UV/Optical variability which is correlated to X-ray variability, time-lagged by the travel time of the X-ray photons from the corona to the point on the disc where the energy is reprocessed, this is visualised in Fig. 1.4. Other factors that can affect the reprocessed variability lag are the inclination of the disc relative to the observer and the timescale of the reprocessing process, though in the case of the latter it is likely small enough to be negligible, even in very small AGN such as NGC 4395.

Using the temperature profile of the disc and the lags at each point on the disc, we can then produce a transfer function which describes the flux produced by the disc from a driving X-ray impulse. The light travel time (and therefore lags) are dependent on mass and can vary from being on the order of days for the larger AGN, down to a few minutes for the smallest, though the relation is also affected by the accretion rate of the object (Breedt et al., 2009).

As we have seen, the delay, τ , we expect at any given radius is a function of the radius, R , of that point, and from Eqn. 1.22 we know that $T \propto R^{-3/4}$. From Wien's Displacement Law we know that for a blackbody the peak emission wavelength is

itself inversely proportional to the temperature, and so from this we can then find the relation

$$\tau(\lambda) = A\lambda^{4/3} \quad (1.28)$$

where A is some normalisation factor.

To quantify how this reprocessed flux changes the temperature profile of our disc, we look at the energy injection the X-rays are providing. The energy of each unit area of the disc from X-ray absorption depends on the X-ray luminosity and the angle of the incoming flux. This incoming X-ray flux at disc radius R can be given as

$$F_X(R) = \frac{L_X}{A_s(R)} \quad (1.29)$$

where $A_s(R) = 4\pi(H^2 + R^2)$ is the area of an imaginary sphere that is centred on the source of the X-rays (which is raised above the disc to height H) and extending out to R . We know that the area of a disc annulus can be given by $2\pi R dR$ where dR is the width of the annulus, however this area is reduced in accordance to the X-ray source height, which leads to a modification of the area by a factor of $\cos\theta$ where θ is the angle between the polar direction and the incoming flux direction. Using simple trigonometry we can see in this situation that $\cos\theta = H/\sqrt{H^2 + R^2}$ which therefore leads to a projected area of

$$A_{proj}(R) = \frac{2\pi R dR H}{\sqrt{H^2 + R^2}} \quad (1.30)$$

Therefore, the total energy which is absorbed by an annulus at a radius R is

$$E(R) = \frac{(1 - A)L_X H}{4\pi(H^2 + R^2)^{3/2}} \quad (1.31)$$

where A is not the area but the albedo of the disc. This extra energy modifies the temperature profile of the disc so that it has an extra term to include the X-ray irradiation. This transforms Eqn. 1.22 so that it becomes

$$T(R) = \left(\frac{1}{\sigma_{SB}} \right)^{1/4} \left(\frac{3}{8\pi} \frac{GM_{BH}}{R^3} \dot{m} \left[1 - \sqrt{\frac{R_0}{R}} \right] + \frac{(1-A)L_X H}{4\pi(H^2 + R^2)^{3/2}} \right)^{1/4} \quad (1.32)$$

1.5 The Power Spectra of XRBs & AGN

On the surface the power spectrum is a simple technique, it is just the squared modulus of a lightcurve's Fourier Transform. It is, however, a very powerful tool in characterising the variability of an radiating object. These power spectra allow us to see the relative contributions to the total variability from variations that take place over different timescales, displayed in a power spectrum in the form of a frequency. An example of a simple power spectrum would be from a single periodic signal such as sine wave, this would generate a power spectrum with a single sharp peak at the frequency of the signal. A slightly less defined signal (such as if there is a slight change in period or the shape of the signal pulse is broad) would lead to a broader peak and is known as a 'Quasi-Periodic Oscillation' or QPO.

Similarly if one had a light curve with uncorrelated random variability (such as from Poisson noise) it would lead to a completely flat power spectrum and is known as a 'white-noise' process. On the other hand if the variability is correlated, which is observed in AGN, we find what is known as a 'red-noise' power spectrum, which has the characteristic shape of a power-law such that

$$P(\nu) \propto \nu^{-\alpha} \quad (1.33)$$

where α is the slope of the power spectrum.

Power spectral techniques were first used on X-ray Binaries (XRBs) which exhibit features that are indicative of a combination of QPO, white-noise, and red-noise in their power spectra on the timescales of their observations. High-frequency QPOs (kHz) are features of neutron star XRBs and are associated with the spin frequency of the star, while black hole binaries exhibit lower frequency QPOs whose physical origin is unresolved but believed to be the result of a relativistic effect called Lense-Thirring Precession ([Ingram and Motta, 2019](#)).

It was observed that the variability that contributes most to the variance of XRB lightcurves is red-noise, displaying power-law power spectra at high frequencies, though this flattens to white-noise on long timescales. The frequency at which this flattening happens scales with the frequency of the low-frequency QPOs which

suggests they are correlated, likely to the accretion rate if the origin of the low-frequency QPOs is tied to it. Above this break frequency, α is usually greater than 1, though QPOs can make measuring the true slope difficult. Some XRBs also display a higher-frequency break, above which α increases to ~ 2 , the frequency of this break can move and is thought to be associated with accretion rate as well, perhaps due to the inner edge of the disc moving inwards at higher accretion rates.

It is now accepted that XRBs and AGN are very similar physical systems just on different mass scales. We should therefore expect to see similar variability properties but with the timescales of the processes scaled up by the increase in black hole mass. As we have seen the gravitational radius of a black hole system, which scales linearly with mass, is a very important feature of black hole systems, which means that the geometry of AGN and XRB X-ray coronae should be similar. The environment should only have a minimal effect since the corona is heavily dominated by the gravitation of the black hole and therefore should not be affected by the presence or lack of a companion object, such as a star in an XRB's case, which might otherwise affect outer regions of an accretion disc.

The main challenge with AGN variability studies are the longer timescales involved even for X-ray emission, and the fact that they are much fainter than XRBs and therefore have poorer signal-to-noise. The earliest AGN power spectra shown in [McHardy and Czerny \(1987\)](#) and [Lawrence et al. \(1987\)](#) showed a red-noise power-law slope, but no breaks or QPOs. QPOs are actually seemingly lacking from AGN power spectra, which is a problem if the process of accretion is supposed to be scale-invariant. Though recent observations such as [Ashton and Middleton \(2021\)](#) have claimed detection of QPOs in AGN, they are still very elusive and an ongoing area of study.

As for the red-noise slope of AGN, it has been shown that they are quite steep with $\alpha \sim 2$ ([Green et al., 1993](#); [Papadakis and Lawrence, 1993](#)) which implied that these slopes lie above the high-frequency break, though at the time, the length of the light curves available was on the scale of AGN variability timescales and so it was difficult to measure low enough frequencies to see a flattening of the power spectrum. We know that theoretically these red-noise spectra *must* flatten at some point to an α less than 1 or else the total variance (and therefore the integrated power) of the light curve would diverge to infinity. The methodology of fitting a power spectrum, the issues with measuring it, and the shape we should expect will be further elaborated on in Chapter 4.

Chapter 2

First detection of an AGN Accretion Disc's Outer Edge in NGC 4395

2.1 Abstract

In this chapter, I will present fast sampled optical photometry of NGC 4395 with the Liverpool Telescope, followed up by very fast multiwaveband simultaneous optical monitoring with HiPERCAM on the 10m Gran Telescopio Canarias. These observations provide the fastest ever multiband photometry of an AGN and will allow very high precision optical interband lag measurements.

2.2 Introduction

The origin of UV/optical variations in AGN, their relationship to X-ray variations and what those variations together can tell us about the inner structure of AGN, has been a matter of major observational activity for over two decades. Initially, combined X-ray monitoring from RXTE and optical monitoring from the ground (e.g. Uttley et al., 2003; Suganuma et al., 2006; Arévalo et al., 2008, 2009; Breedt et al., 2009, 2010; Breedt, 2010; Lira et al., 2011) typically showed optical variations lagging the X-rays by about a day ($\pm 0.5d$).

These observations are consistent with reprocessing of X-rays from around the central black hole by a surrounding accretion disc, but usually without sufficient multiband detail to map out the temperature structure of the disc (Blandford and McKee, 1982). Some multiband optical monitoring (e.g. Sergeev et al., 2005; Cackett et al., 2007) showed wavelength-dependent lags consistent with the expectations of reprocessing from an accretion disc but without simultaneous X-ray monitoring.

More recently, X-ray and UV/optical multiwaveband monitoring, based mainly around Swift observations (e.g. Cameron et al., 2012; Cameron, 2014; Shappee et al., 2014; McHardy et al., 2014; Edelson et al., 2015; Troyer et al., 2016; Fausnaugh et al., 2016; McHardy et al., 2018; Cackett et al., 2018; Edelson et al., 2019; Cackett et al., 2020; Hernández Santisteban et al., 2020; Vincentelli et al., 2021, 2022), but also with XMM-Newton (e.g. McHardy et al., 2016), have confirmed that the UV/optical variations are mostly well explained by reprocessing of high energy radiation from an accretion disc, but have also noted a reprocessed contribution from the gas in the broad line region (BLR).

The lags expected from the BLR were predicted by Korista and Goad (2001, 2019) and are particularly large in the u band, where combined Balmer emission lines form the Balmer continuum, and also in the i band from the Paschen continuum (see particularly Cackett et al., 2018). The combined signatures of reprocessed disc and BLR emission are clearly seen in NGC 4593 (McHardy et al., 2018). Here the response functions (Horne et al., 2004) required to explain the UV and optical lightcurves by reprocessing of X-rays, show both a short-timescale (disc) peak and long-timescale (BLR) tail. In Mrk110, Vincentelli et al. (2021) also show a combination of short-timescale (disc) and long-timescale (BLR) lags.

The BLR and the Narrow Line Region (NLR) are structures whose effects can be seen in the optical spectrum of an AGN by viewing the emission lines produced from each region. The BLR tends to be out at 10^3 to $10^5 R_g$ with gas velocities from 3,000-10,000 km/s, reflected in the Doppler Broadened emission lines it produces. The NLR on the other hand is much further out closer to $10^9 R_g$ with gas velocities from only 300-1,000 km/s, resulting in its namesake much narrower emission lines. Looking at the spectrum from NGC 4395 in Fig. 2.1 we can see only very narrow bands, which implies a limited contribution from the BLR.




FIGURE 2.1: The Optical Spectrum of NGC 4395. From Ho et al. (1995).

Given the many observations and many suggested explanations, there is increasing interest in modelling the lag spectra. Initially most modelling was based on reprocessing of X-rays by a disc with the temperature profile described either by Shakura and Sunyaev (1973) or (with GR corrections) by Novikov and Thorne (1973).

Although the broad reprocessing picture is now established, this modelling highlighted a variety of problems including:

(1) The observed optical lightcurves are much smoother than predicted (e.g. Berkley et al., 2000; Arévalo et al., 2008) unless the X-ray emitting corona is much larger (~ 100 gravitational radii, R_g) than estimated from X-ray reverberation ($\sim 4R_g$, e.g. Emmanoulopoulos et al., 2014; Cackett et al., 2014) or microlensing observations ($\lesssim 10R_g$ Dai et al., 2009; Chartas et al., 2012; Mosquera et al., 2013)

(2) The disc sizes estimated from almost all Swift monitoring are larger than predicted, with analytic formulae (e.g. Fausnaugh et al., 2016) giving larger discrepancies ($> 3\times$ larger) than numerical disc modelling ($\lesssim 2\times$, McHardy et al., 2018). However numerical modelling which better includes relativistic effects (Kammoun et al., 2021b) gives no discrepancy, albeit with an X-ray source height ($\sim 20 - 70 R_g$) much larger than that of most previous modelling.

(3) The lag between the X-ray and far UV band is much larger than expected, compared with the lags between the various UV and optical bands (e.g. summary in McHardy et al., 2018). Possible solutions include scattering through an inflated inner disc (Gardner and Done, 2017) or including a more distant reprocessor such as the BLR McHardy et al. (2018). Indeed Netzer (2021) claims that the BLR dominates the lags. Alternatively, if all lags are measured relative to a reference UV band, then the X-ray to UV lag may be increased if the X-ray auto-correlation function is broad (Kammoun et al., 2021a).

Reprocessing, of course, cannot explain all aspects of UV/optical variability. There are long term (months-years) trends in the UV/optical which are not mirrored in the X-rays and which may result from inwardly propagating accretion rate fluctuations in the disc (e.g. Arévalo et al., 2008; Breedts et al., 2009, 2010). If these trends are filtered out then the shorter-timescale reprocessing signature can be revealed (McHardy et al., 2014, 2018; Pahari et al., 2020). However in this investigation we are concentrating on the short-timescale variations.

Most AGN which have been monitored so far have been of similar black hole masses, mostly between $\sim 10^7$ and $\sim 10^8 M_\odot$. Although accretion rates are notoriously hard to measure accurately, the majority of AGN monitored so far have had accretion rates of few per cent of Eddington, although a small number have higher rates, eg Mrk110, $\dot{m}_E \sim 40\%$ (Vincentelli et al., 2021), NGC7469, $\dot{m}_E \sim 50\%$ (Pahari et al., 2020) and Mrk142 $\dot{m}_E \sim 100\%$ (Cackett et al., 2020). There have been very few observations of very low mass or very low accretion rate AGN, largely because most of them are too faint.

However to test models, it is important to extend the range and hence objects at the extreme edges of the distribution, where disc and BLR properties are likely to be most different, are particularly useful. For example, at very low accretion rates, the disc scale height is likely to be lower (Treves et al., 1988) thus affording a less obscured view to the outer disc.

Here we present lag measurements of NGC 4395. Although its exact mass is currently a matter of some debate, with values between $3.6 \times 10^5 M_\odot$ (Peterson et al., 2005) and $1 \times 10^4 M_\odot$ (Woo et al., 2019) having been claimed, there is complete agreement that its mass is definitely very low. For the mass noted by Peterson et al. (2005), the accretion rate is also very low, 0.12% Eddington so its luminosity is low. However it is close enough to us (3.85 Mpc, Tully et al., 2009) that it is bright enough for accurate lag measurements to be made.

Early Swift observations (Cameron et al., 2012) showed that the B-band lagged the X-rays in NGC 4395 by less than < 45 min, with a tentative estimate of the UVW2 band lagging the X-rays by 400s. XMM-Newton observations, using the optical monitor in fast readout mode, showed that the UVW1 band lagged the X-rays by 473^{+47}_{-98} s and parallel ground based observations gave a lag of the X-rays by the g band of 788^{+44}_{-54} s. There have, however, been no simultaneous multiband lag measurements and so serious comparison with reprocessing models has not yet been possible.

In Section 2.2 we discuss *ugriz* monitoring of NGC 4395 with the robotic 2m Liverpool Telescope (LT) which showed variability on few hour timescales, consistent with the few hundred second lags noted above, but were not able to measure the lags accurately. However based on those observations we obtained very high sensitivity simultaneous *ugriz* photometry over one night with HiPERCAM on the 10.4m Gran Telescopio Canarias (GTC) with time resolution down to 3s (Section 2.3), thus allowing very accurate measurement of interband lags. In 2.4 we compare the resultant lag measurements with reprocessing models. Conclusions are presented in Section 2.5.

2.3 Liverpool Telescope Observations

2.3.1 The Observations

NGC 4395 was observed with the IO:O instrument on the robotic Liverpool Telescope (LT, Steele et al., 2004) for approximately 6 hours on each of 5 separate nights in 2017 March. The IO:O¹ is a CCD imaging camera with a 10×10 arcmin field of view and an unbinned pixel size of 0.15 arcsec, here used in 2×2 binning mode.

The observations were made by cycling through successive exposures in groups of filters. Given the previously measured approximate lag of 330s between the UVW1 (290 nm) and g band (McHardy et al., 2016), the observations were arranged in two groups so that cycle times between the same filter did not exceed 200s, thus allowing reasonable sampling of likely lag times. To allow measurement of lags in all bands relative to the g-band, one group consisted of the g, r, i and z bands and the other

¹<https://telescope.livjm.ac.uk/TelInst/Inst/IOO>

| Night | Date (UTC) | Start Time | End Time | Filters and Exposure Times | Lightcurve Figure |
|-------|------------|------------|----------|---|-------------------|
| 1 | 2017-03-23 | 00:49:05 | 03:23:10 | <i>g, r, i, z</i> all 10s | Fig. 2.4 |
| 2 | 2017-03-24 | 01:03:46 | 07:01:45 | <i>u</i> -30s, <i>B</i> -20s, <i>g</i> -20s | Fig. 2.5 |
| 3 | 2017-03-26 | 00:40:01 | 06:39:45 | <i>g, r, i, z</i> all 20s | Fig. 2.6 |
| 4 | 2017-03-28 | 00:34:54 | 06:31:33 | <i>u</i> -30s, <i>B</i> -20s, <i>g</i> -20s | Fig. 2.7 |
| 5 | 2017-03-29 | 00:33:37 | 06:32:06 | <i>g, r, i, z</i> all 20s | Fig. 2.8 |

TABLE 2.1: Log of LT observations of NGC 4395. Start Time is the start time of the first observation in the series and End Time is the start time of the last observation.

consisted of the *g*, *B* and *u* bands. The *g*, *r*, *i*, *z* group was observed on nights 1, 3 and 5. The *g*, *B*, *u* group was observed on nights 2 and 4. Due to bad weather, these nights were not consecutive. The log of the observations is given in Table 2.1.

To remove any non-astrophysical variations, such as those due to changes in atmospheric transparency which would affect all objects on a CCD image, lightcurves were produced using the standard method of differential photometry relative to one or more comparison stars.

An LT IO:O *g*-band image of NGC 4395 is shown in Fig. 2.2. In this figure a number of comparison stars are labelled, together with the AGN nucleus. Fig. 2.2 is centred midway between comparison stars 0 (RA 12 25 54.63 Dec +33 26 51) and 2 (RA 12 25 55.87 Dec +33 33 58). However the observations reported in this investigation were centred on the AGN. Star 2 was always on the field of view, well away from the edge of the CCD, but star 0 was outside of the field of view. Star 2 was therefore used as the comparison star. Also labelled in Fig. 2.2 are stars 1 and B which are within the smaller HiPERCAM field of view.

For each CCD image, an aperture of radius 4.25 arcseconds, which is 14 LT (binned) pixels, was placed around the AGN and a similar aperture was placed around star 2. A background annulus of inner radius $3\times$ the source radius and outer radius $4\times$ was also placed around each object. Lightcurves were derived as the ratio of the background subtracted counts in the AGN compared to those of the comparison star.

In later observations which are reported in Chapter 3, both stars 0 and 2 were on the IO:O field of view. In Fig. 2.3 we show one of the IO:O lightcurves of star 2 compared to star 0. The middle half of this lightcurve is constant (reduced $\chi^2 = 0.94$) but including the ends there is a very slight fractional decrease of 0.006. Stars 0 and 2 have different colours. From SDSS magnitudes $g-r = 1.33$ for star 2 and 0.98 for star 0. We have not corrected the lightcurves for colour related changes in atmospheric transmission during these observations but it is quite likely that this very small change might be explained in that way. We therefore conclude that both stars are non-variable over the timescales of the observations reported here and can be used safely as comparison

stars. Note that the zero point in Fig. 2.3 is zero in order to demonstrate the relative constancy of the ratio.

We performed differential photometry using two different pipelines. The first pipeline was an in-house Southampton University pipeline and the second is the HiPERCAM pipeline.² Both pipelines use very similar methods, centroiding on the objects of interest and then performing aperture photometry using a background annulus surrounding the object. The results were very similar. Here we present the results using the HiPERCAM pipeline, in large part because this pipeline is well established and well documented and the code is already publicly available thus allowing future researchers to check our results more easily.




FIGURE 2.2: *g*-band image of NGC 4395, 40s exposure, with the IO:O CCD imager on the Liverpool Telescope. The comparison stars, 0 and 2, used in the production of the LT lightcurves are labelled, as are the comparison stars, 1 and B, used in the production of the HiPERCAM lightcurves. The AGN nucleus is also labelled. The field of view is 10×10 arcmin.

2.3.2 Interband Correlations and Lags

The lightcurves for nights 1, 2, 3, 4 and 5 are shown in Figs. 2.4, 2.5, 2.6, 2.7 and 2.8 respectively. Here, unlike in Fig. 2.3, the ordinate does not start at zero. Note that the fluxes typically vary by $\sim 10\%$, which is much larger than the 0.6% trend in the comparison stars. In all cases the individual lightcurves on each night had similar

²<http://deneb.astro.warwick.ac.uk/phsaap/hipercam/docs/html/>

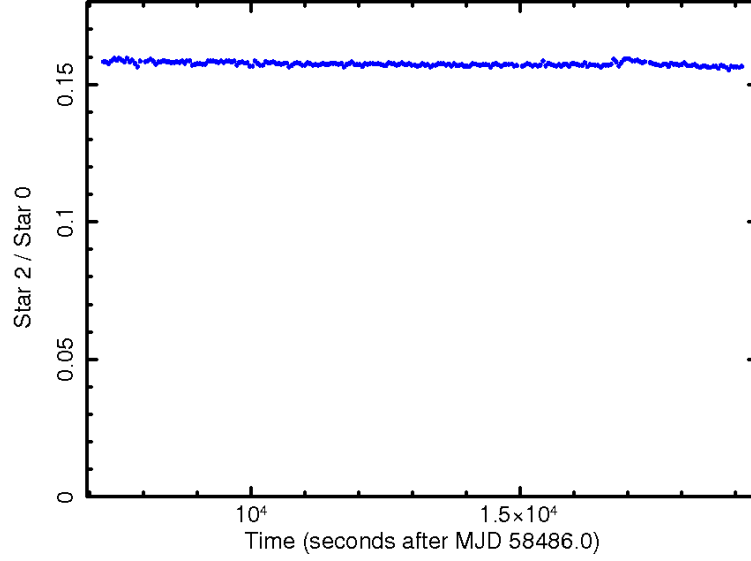


FIGURE 2.3: Ratio of the background subtracted count rate of star 2 compared to that of star 0 in the g-band.




FIGURE 2.4: Night 1: Lightcurves of NGC 4395 in, from the top, the *g*, *r*, *i* and *z* bands, from the Liverpool Robotic Telescope. Fluxes are normalised to the mean. UT times are given in Table 2.1.

shapes and are well correlated. However on nights 1 and 5, the *griz* lightcurves were dominated by slow trends, without enough prominent features to measure a lag between the bands. Similarly, the *ubg* lightcurves on night 4 were also dominated by slow trends. However the *griz* lightcurves on night 3 and the *ubg* lightcurves on night 2 did show more features.

In Figs. 2.9 and 2.10 we show, as representative, the simple discrete correlation functions (DCFs: [Edelson and Krolik, 1988](#)) between the B and g bands on night 2 and between the *r* and *i* bands on night 3. We also show the 68, 90 and 95% confidence contours. The observed correlations easily exceed 95% confidence and so are unlikely




FIGURE 2.5: Night 2: Lightcurves of NGC 4395 in, from the top, the u , B and g bands from the Liverpool Robotic Telescope. Fluxes are normalised to the mean. UT times are given in Table 2.1.




FIGURE 2.6: Night 3: As for Fig. 2.4.

to have occurred by chance. The confidence contours are generated as described in [Breedt et al. \(2009\)](#) and [McHardy et al. \(2018\)](#). This process involves simulating 1000 realisations of the shorter wavelength lightcurve based on its power spectrum. Specifically here we note that as the power spectrum derived from the observed lightcurve is not of sufficient quality to measure accurately its low and high frequency slopes, these values are fixed at -1 and -3. Varying these parameters does change the confidence contours, but not by a great deal.

The DCF is useful for establishing the significance of a correlation but not for measuring the lag. For that purpose we use the Flux Randomisation / Random Subset




FIGURE 2.7: Night 4: As for Fig. 2.5.




FIGURE 2.8: Night 5: As for Fig. 2.4.

Selection (FR/RSS) method of [Peterson et al. \(1998\)](#) and JAVELIN ([Zu et al., 2011, 2013](#)). In Figs 2.11 and 2.12 we show the lag distributions from JAVELIN between the B and g lightcurves from night 2 and the g and i lightcurves from night 3 respectively. Formally the g band lags the B-band by 268^{+192}_{-155} s and the i band lags the g band by 267^{+309}_{-280} s. The FR/RSS method gives similar median values, though with larger uncertainties. JAVELIN gives a lag of the g band by the r band of -74^{+294}_{-226} s whilst FR/RSS gives $+28 \pm 436$ s. For the lag of the r band by the i band JAVELIN gives a lag of 273^{+223}_{-365} s and FR/RSS gives 300 ± 500 s. Due to the low signal to noise (S/N) of both the u and z-band lightcurves, both FR/RSS and JAVELIN, although showing correlations, give very broad lag distributions with uncertainties > 1000 s for any correlations involving those bands.




FIGURE 2.9: DCF showing the lag of the B-band by the g band in seconds for LT night 2 (Fig. 2.5). From bottom to top, the 68%, 90% and 95% confidence contours are shown.




FIGURE 2.10: DCF showing the lag of the r band by the i band in seconds for LT night 3 (Fig. 2.6). From bottom to top, the 68%, 90% and 95% confidence contours are shown.

We conclude that the LT lightcurves show a strong correlation between all wavebands. The lag of the g band by the B band is statistically robust (Fig. 2.11) and there are suggestions of wavelength-dependent lags between some of the other wavebands. However the lag uncertainties for the other bands are too large to allow accurate measurement of the lags.




FIGURE 2.11: Lag distribution from JAVELIN between the B and g lightcurves from LT night 2 observations (Fig. 2.5).




FIGURE 2.12: Lag distribution from JAVELIN between the g and i lightcurves from LT night 3 observations (Fig. 2.6).

2.4 GTC HiPERCAM Observations

2.4.1 Observations

Following the proof of concept LT observations, NGC 4395 was observed for just over 6 hours from 2018-04-16 20:58:14.665 UTC (MJD 58224.8737808) to 2018-04-17 03:06:07.266 UTC (MJD 58225.1292507) with HiPERCAM on the 10.4m Gran Telescopio Canarias (GTC)³.

³<http://www.gtc.iac.es/>




FIGURE 2.13: HiPERCAM u_s -band image. The light blue lines designate the four quadrants of the CCD which are read out individually. The AGN together with comparison stars 1 and B, also shown in Fig. 2.2, are labelled. The inner circles, green around the AGN and black around the comparison stars, of radii 52 pixels (4.21 arcsec), define the object apertures. The sky background lies between the two surrounding solid red circles (radii 84 and 140 pixels, i.e. 6.80 and 11.34 arcsec respectively). The dashed red circles indicate stars in the background regions which have been masked out. The dashed straight lines simply connect the masked areas to the central object that they are associated with. For the relatively bright star to the east of the AGN, a large masking radius is chosen so as to encompass also a nearby fainter star. However only the part of the image which is also within the background annulus is masked. The object aperture is unaffected. Note that, opposite to the conventional projection used in Fig. 2.2, east is to the right of north. The HiPERCAM major axis was tilted very slightly east of north.

HiPERCAM (Dhillon and Marsh) is a quintuple-beam high-speed astronomical imager. Incoming white light is split into 5 colours, each sent to a separate camera, with each camera containing a 2048×1024 pixel frame-transfer CCD which can read out at up to 1000 frames per second. The 5 filters have bandpasses which very closely resemble the standard Sloan $ugriz$ bands but have improved throughput, particularly in the u_s band where throughput is 41% higher. These Super SDSS filters⁴ are referred to as u_s , g_s , r_s , i_s , and z_s . The plate scale is 0.081 arc seconds per pixel giving a field of view of 2.8×1.4 arcminutes. The integration time was 3s for the g_s , r_s , i_s , and z_s bands and 15s for u_s .

We arranged the field of view to encompass the two comparison stars, 1 and B, which are shown in Fig. 2.2. The data were reduced using the HiPERCAM software suite⁵. A HiPERCAM u_s image is shown in Fig. 2.13. Object (green or black) and sky apertures (solid red) are shown around all objects. Masked regions, to remove stars from background areas (dashed red lines) are also shown. Details are given in the figure

⁴<http://www.vikdhillon.staff.shef.ac.uk/ultracam/filters/>

⁵<http://deneb.astro.warwick.ac.uk/phsaap/hipercam/docs/html/>

caption. Note that this image is mirrored around the north-south axis so east is to the right of north, not the usual left. The software then produces count rates in the various apertures, X and Y coordinates and FWHM of targets for all of the separate images thus allowing the production of lightcurves.




FIGURE 2.14: Tracking and seeing plots of the HiPERCAM data.

During the observations the sky was dark and photometric, but the seeing was quite variable, ranging between ~ 1.5 and 4 arcseconds. The seeing, defined as the FWHM of the unresolved core of the AGN, together with the X and Y CCD pixel coordinates of the AGN, are shown in Fig. 2.14. We use the FWHM of the AGN rather than that of star 1 simply for computational convenience but the values are almost exactly the same.

In the g-band discussed here the ratio of the FWHM of the AGN to that of the star is constant at $1.02 \times$ with a variance of 8.4×10^{-5} so any variations are tiny compared to the factors of a few in the seeing itself. The field of view remained as shown in Fig. 2.13 until the source approached transit, which occurred close to the middle of the observation. At this time the field passed very close to the zenith where tracking is not so precise and, at the time of these observations, there was no autoguider at the HiPERCAM focus of the GTC. Hereafter we refer to times after the start of the observation. After 8000s, the field of view drifted slowly to the west, as shown by the large change in the Y-coordinate of the AGN (Fig. 2.14).

After 12000s the field of view was manually adjusted back to the original position. During these movements, star B crossed into the CCD quadrant containing the AGN and was then moved back. However star 1 remained within its original CCD quadrant at all times. Additionally, as the background surrounding star 1 is empty, unlike that surrounding star B, star 1 is preferred as the comparison star. There are a number of small bumps in the lightcurves between 10000 and 14000s, most of which can be associated with rapid changes in field location and so are not considered to be

physically real. Data between these two times is therefore not considered to be reliable and so is not considered in this analysis.




FIGURE 2.15: Normalised ratio of the g_s -band count rates of stars 1 and B.

At 14650s an unplanned instrument rotator unwrap occurred. As it is not possible to follow the motion of objects through that rotation, the lightcurve was produced in two parts, before and after that time. There is a simultaneous small ripple in all lightcurves at that time. Between 16780 and 16840s there was a small change in the Y-positions but this change occurred sufficiently slowly that we were able to accurately follow the movement so the photometry was not affected.

In Fig. 2.15 we show the normalised ratio of the count rates between star 1 and star B. This figure only covers the time when star B was in the CCD quadrant in which it started. There is a gap in the middle when star B was in a different CCD quadrant, i.e. the one in which the AGN is always located. During the first part of the observation when the AGN and both stars remained on their original CCD quadrants, there is no change in the ratio of star 1 to star B. Formally the error-subtracted excess variance of this part of the lightcurve is -6.68×10^{-8} and the average standard deviation is 1.37×10^{-3} . There is no evidence for variability in either comparison star on the timescales of the observation and so, in principle, either star may be used to measure changes in AGN intensity. We do, however, note a small change in the stellar ratio after the gap shown in Fig. 2.15 and around 16800s, corresponding to a sudden change in position of the objects. As star 1 remains on its original CCD quadrant throughout the observation and also has a lower surrounding background and so is less likely to suffer from spurious position-induced changes in intensity, we use star 1 to determine changes in the AGN. The resultant lightcurves are shown in Fig. 2.16.




FIGURE 2.16: Observed u_s , g_s , r_s , i_s , and z_s lightcurves, normalised to the mean, for NGC 4395 from HiPERCAM on the GTC binned to 15s exposures. Statistical errorbars are too small to be seen and are, from u_s to z_s respectively, 0.0047, 0.00165, 0.00175, 0.00243 and 0.00356. The lightcurves between 10,000 and 14,000s, as indicated by the dashed grey lines, are affected by loss of tracking and motion of the target and comparison stars over the CCDs, as described in the text. For each of the g_s to z_s bands, the u_s lightcurve, with variability amplitude scaled to be the same as for the appropriate g_s to z_s band, is underlaid in grey, for comparison.

2.4.1.1 Seeing Effects

Although the lightcurves are generally quite smooth, there are occasional changes of flux on timescales of a \sim few hundred seconds. As the seeing can vary on timescales at least as short we investigated whether seeing changes might cause flux changes which might affect lag measurement. Assuming that the slower changes represent real physical changes in the luminosity of the AGN, we fitted 6th order polynomials to the lightcurves and measured the residuals. An example fit to the g_s lightcurve, with residuals, and seeing, is shown in Fig. 2.17. Whilst recognising that this fit does make assumptions about what a real physical variation in the AGN is, a lower-order polynomial does not fit the long-timescale variations as well and a higher order does not improve the fit. We therefore take it as a first attempt at determining the effect of seeing on flux.

In Fig. 2.18 we plot the residuals against seeing. Although there is considerable scatter, there is a definite trend for the residuals to become larger as the seeing improves, i.e. the residuals and seeing are anti-correlated. Changes with seeing in the fraction of light in the object apertures should be the same for the AGN core and for the comparison star. Therefore this anti-correlation may be more related to changes in the light in the background apertures and may not necessarily be repeatable in imaging of other objects in different environments. We calculate the best fit to that relationship and apply the resultant correction to the observed lightcurve, producing a corrected lightcurve, Fig. 2.19. The amplitudes of small-timescale variations are lower in the corrected lightcurves. In the next section we consider the effect of this correction on lag measurements.




FIGURE 2.17: *Top panel:* g_s -band normalised flux (red errorbars) with 6th order polynomial fit (black) for the first 8000s of the GTC HiPERCAM observation. *Middle panel:* Flux residuals relative to the polynomial fit. *Bottom panel:* Seeing, i.e. FWHM of the AGN point spread function (PSF). An anti-correlation with the residuals is apparent.

2.4.2 Fractional Variability

We list, in Table 2.2, the percentage fractional variability (F_{var}) in the different bands as we are aware that these numbers may well be of interest to other observers in other contexts. The numbers in Table 2.2 are not corrected for seeing. Correction for seeing reduces these numbers by only 4%. We list the values as derived for the first 8000s and for the complete lightcurves with the section from 10,000 to 14,000s removed. The u_s band is approximately twice as variable as the g_s band but the decrease in F_{var} with increasing wavelength is not so pronounced at longer wavelengths. The i_s band is, in fact, slightly more variable than both the r_s and z_s bands. These observations are consistent with the standard paradigm of the longer wavelengths coming from a larger area of the disc, more distant from the black hole. As F_{var} is larger on longer





FIGURE 2.18: Residuals from Fig. 2.17 vs seeing.

FIGURE 2.19: *Top Panel* The first 8000s of the g_s -band lightcurve. *Bottom Panel*. Corrected g_s -band lightcurve, using the correction shown in Fig. 2.18

timescales, we conclude that these observations are sampling the high frequency part of the optical variability power spectrum where the power spectrum is steep. The power spectrum of these observations will be discussed elsewhere (Chapter 4).

We also list, in Table 2.2 the magnitudes of Star 1 in the various bands as taken from SDSS. The SDSS coordinates are 12 25 51.28 +33 31 26.76 (J2000). We also list the count rate ratios of the AGN to Star 1 from the first 8000s of the observations and hence the derived magnitudes of the AGN. These numbers should be treated with caution and treated as indicative rather than precise as we did not do detailed multiband photometric calibration of the HiPERCAM observations.

Although we do not show plots here, we have carried out flux/flux analysis similar to that shown in Figs.10 and 11 of [McHardy et al. \(2018\)](#). We again see linear variability of the observed fluxes in each band against a model driving lightcurve and so can decompose the lightcurves into contributions from a red, constant, component (host galaxy) and a blue and variable component (disc) which can explain some observed colour variations. When the overall flux is highest at the end of the observation, the variation in the u_s -band is greater than that in the z_s -band. However in the lower amplitude variation towards the beginning of the observation, the amplitudes of variability were similar. This behaviour of bluer when brighter is simply explained as the combination of a red constant component and a blue variable component.

| Band | F_{var} (%) First 8000s | F_{var} (%) All bar 10-14ks | Ratio AGN/Star1 | Star 1 Mag | AGN Mag |
|-------|--|--|--------------------|---------------|------------|
| u_s | 2.92 | 7.71 | 1.1482 | 17.41 | 17.26 |
| g_s | 1.44 | 4.01 | 0.73566 | 16.23 | 16.56 |
| r_s | 1.09 | 2.74 | 0.6153 | 15.71 | 16.24 |
| i_s | 1.17 | 3.05 | 0.4666 | 15.56 | 16.39 |
| z_s | 1.04 | 1.71 | 0.5108 | 15.45 | 16.18 |

TABLE 2.2: Fractional variability of the HiPERCAM observations as a function of both waveband and timescale together with the SDSS magnitude of Star 1, the count rate ratio of the AGN compared to Star 1 as derived from the first 8000s of the observations and the derived magnitude of the AGN.

2.4.3 Interband Lags

Given the concerns regarding possible spurious count-rate variations induced by rapid positional changes, we initially restrict our lag measurement to the first 8000s of the observation which are largely unaffected. In Table 2.3 we give the lags, relative to the u_s band, for the other HiPERCAM bands for the first 8000s of the observations. We use lightcurves with 15s sampling as that is the minimum available for the u_s band.

We list the lags as measured using JAVELIN ([Zu et al., 2011, 2013](#)) and using the centroid lags as calculated using the FR/RSS correlation function method ([Peterson et al., 1998](#)). Here we apply both flux randomisation and random subset selection. Thus the errors here are larger than would be obtained if we chose only one of the methods.

We list the lags as measured both on the uncorrected data and on data corrected for seeing effects. There is no significant difference between any of the methods except that the uncertainties for the FR/RSS method are larger as, in the RSS method, only 70% of the available data are used and gaps are introduced into the lightcurves. Thus the associated uncertainties are very conservative. However, compared to all previous

AGN lag measurements, as presented in the many RXTE and Swift-based papers listed in the first paragraph of the Introduction here, the errors are still very small.

We have also calculated the lags for the full lightcurve, excepting the middle region from 10,000 to 14,000s where the tracking was particularly poor. Although these data are therefore not of quite as good quality as that for the first 8000s, they provide a preliminary check on whether the lags might depend on which small section of the lightcurves one chooses, which might be affected by some small non-astronomical variation which we have not properly accounted for. We have not attempted to correct this full lightcurve for seeing variations. The only possible difference between the first 8000s and the full lightcurve is that the average luminosity level is lower in the first 8000s than in the full lightcurve. The measured lags, using JAVELIN, relative to u_s are 415^{+13}_{-4} , 622^{+13}_{-14} , 689^{+30}_{-45} , and 669^{+9}_{-5} s respectively for the g_s , r_s , i_s and z_s bands.

Although the lags of the g_s band and to a lesser extent also the r_s band are a little longer than in the first 8000s of the observations (Table 2.3), the lags of i_s and z_s are quite similar. These small differences are in agreement with what we expect from a higher illuminating luminosity.

There are some short-timescale bumps in the lightcurves which remain even after correction for seeing variations. For example there is a bump at around 4800s which is quasi-simultaneous in all bands. It is possible that the seeing correction model has not been sophisticated enough as the bump does correspond to a decrease in seeing. We have therefore measured lags excluding the section from 4500 to 8000s but it does not make a noticeable difference to the lags. Overall we conclude that the lags listed in Table 2.3 are a fair reflection of the true lags.

In Table 2.3 we also give the effective wavelength of each band for these particular filters as used on HiPERCAM on the GTC and taking account of atmospheric transmission. We derived these values, which are the ‘Bandpass Average Wavelengths’, using the PYSYNPHOT package⁶ which is specifically designed for HiPERCAM. We note that the values derived here are actually very similar to the central wavelengths listed on the HiPERCAM filter web site⁷, i.e. 3526, 4732, 6199, 7711 and 9156Å respectively for u_s , g_s , r_s , i_s , and z_s for which slightly simpler assumptions are made. Here we use the effective, or Bandpass Average Wavelengths, given in Table 2.3.

2.4.4 Referencing the HiPERCAM lags to the X-ray band

The HiPERCAM lags are all measured relative to the u_s band but to obtain full lag-wavelength coverage it is necessary to reference the HiPERCAM lags to the X-ray

⁶https://github.com/StuartLittlefair/ucam_thruput

⁷<http://www.vikdhillon.staff.shef.ac.uk/ultracam/filters/>

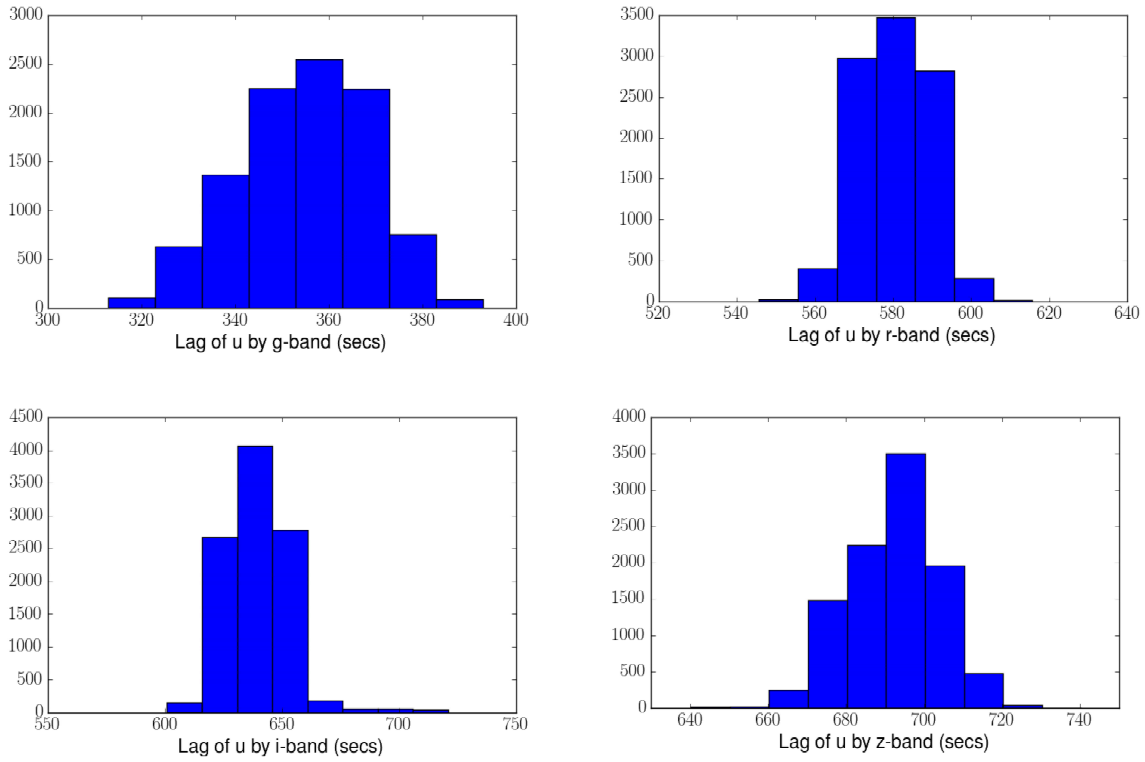


FIGURE 2.20: HiPERCAM lag probability distributions from JAVELIN for the g_s (top left), r_s (top right), i_s (bottom left), and z_s (bottom right) bands relative to the u_s band. The distributions are for the first 8000s of the observation, with lightcurves corrected for seeing changes.

band. There has been a previous measurement with XMM-Newton of the lag of the UVW1 band (2910 Å), observed using the Optical Monitor in fast mode, relative to the EPIC X-ray 0.5-10 keV band, of 473^{+47}_{-98} s [McHardy et al. \(2016\)](#). The X-ray and UVW1 lightcurves overlap almost completely and are both of good quality, resulting in a correlation exceeding 99% significance. The X-ray and UVW1 lightcurves are still available so we have repeated the lag analysis. We confirm the original result, measuring a lag of 430^{+50}_{-80} s. Although it is of lower statistical significance, this result is quite consistent with the measurement of a ~ 400 s lag of the UVW2 band (1920 Å) relative to the X-rays ([Cameron et al., 2012](#)).

Ground-based g -band observations accompany most of the X-ray and UVW1 lightcurves from which a g -band to X-ray lag of 788^{+44}_{-54} s was derived. However the g band lightcurve had to be made by adding together lightcurves from 6 separate ground-based observatories, all with different S/N, and mostly of lower S/N than either the X-ray or UVW1 lightcurves so the confidence level of the correlation barely reached 90%. We also note that although the JAVELIN lag distribution for the X-ray to UVW1 is symmetric, with no long tails, the X-ray to g -band lag has a long tail up to longer lags (~ 900 s). Unfortunately the g -band lightcurve can no longer be found so we cannot repeat this analysis. We therefore consider the X-ray to g -band lag to be less

reliable than the X-ray to UVW1 lag and therefore we reference the HiPERCAM lags to the UVW1 lag.

The effective wavelength of the UVW1 filter on the XMM-Newton Optical Monitor, including the entire in-flight telescope response, is 2910 Å. This effective wavelength is not very different from that of the HiPERCAM u_s band (3694 Å), but is different enough that we should estimate the small offset expected in the lags. In any reasonable scenario, i.e. whether the UV/optical variability is dominated by reprocessing from a disc or from the BLR, we expect this lag to be positive. Scaling the lag to the 4/3 power of wavelength (Cackett et al., 2007), as expected from reprocessing by a disc with the temperature profile defined by Shakura and Sunyaev (1973), would give a lag of UVW1 by u_s of ~ 150 s, i.e. u_s lagging the X-rays by 580s. However there is sometimes an additional lag between the X-rays and far-UV which would reduce the predicted UVW1 to u_s lag. On the other hand, the contribution from the BLR may increase the lag slightly. Thus we take a round figure of 600s for the X-ray to u_s -band lag. Although this figure is almost certainly not absolutely correct, it probably is not very wrong. We discuss the implications of incorrect referencing in Section 2.5. Correct referencing of the HiPERCAM lags relative to the X-ray band is important and is the subject of further observations in Chapter 3.

| Band | Central Wavelength (Angstroms) | JAVELIN Lag Uncorrected | JAVELIN Lag Corrected | Centroid ICCF Uncorrected | Centroid ICCF Corrected |
|-------|--------------------------------|-------------------------|-----------------------|---------------------------|-------------------------|
| u_s | 3694 | 0 | 0 | 0 | 0 |
| g_s | 4778 | 325^{+9}_{-7} | 360^{+10}_{-15} | 377 ± 20 | 390 ± 18 |
| r_s | 6201 | 578^{+12}_{-5} | 578^{+11}_{-5} | 565 ± 20 | 585 ± 19 |
| i_s | 7640 | 606^{+13}_{-3} | 633^{+15}_{-14} | 571 ± 23 | 615 ± 23 |
| z_s | 9066 | 676^{+10}_{-16} | 693^{+10}_{-8} | 650 ± 30 | 698 ± 33 |

TABLE 2.3: Lags, relative to the u_s band, in seconds for the first 8000s of the HiPERCAM observations. The first column gives the band. The second column gives the effective, or Bandpass Average wavelength of the band. The third and fourth columns are the lags as measured by JAVELIN using the original lightcurves and using the seeing-corrected lightcurves respectively. The fifth and sixth columns are the centroid interpolation cross correlation function (ICCF) lags as measured using the FR/RSS method (Peterson et al., 1998) on the original lightcurves and on the seeing-corrected lightcurves respectively.

2.5 Discussion

2.5.1 Preliminary Conclusions: An Edge to the Emission Region and little BLR contribution

Our best estimates of the multiwavelength lags in NGC 4395 are given in Fig. 2.21. Without any comparison with models, two important conclusions can be drawn

immediately. Firstly the g_s -band lags considerably behind the u_s -band, unlike in almost all other AGN. This observation implies that the BLR is not a major contributor to the UV/optical variability of NGC 4395. Secondly, the lag spectrum flattens at long wavelengths, implying an edge to the emission region. These two observations will be considered in more detail later.

As a preliminary comparison with lag models we plot the model lags expected from disc reprocessing, using the formula given in equation 8 from Kammoun et al. (2021a) based on mass, accretion rate and illuminating X-ray luminosity.

We plot, as solid red lines, the lags expected for both zero and maximum black hole spin, using the mass and accretion rate from Peterson et al. (2005), i.e. mass $3.6 \times 10^5 M_\odot$ and accretion rate 0.12% Eddington. The formula from Kammoun et al. (2021a) requires the 2-10 keV luminosity as a percentage of the Eddington luminosity. Taking an observed 2-10 keV flux of 5.6×10^{-12} ergs cm $^{-2}$ s $^{-1}$ from Vaughan et al. (2005) which is very similar to the value from McHardy et al. (2016) and Vincentelli et al (in preparation), and a distance of 3.85 Mpc, we calculate the 2-10 keV luminosity, divide it by the Eddington luminosity based on a mass of $3.6 \times 10^5 M_\odot$ and multiply by 100 to obtain the required number, which is 0.025. We adopt an X-ray source size of $6R_g$, consistent with the values derived from X-ray reverberation studies (Emmanoulopoulos et al., 2014).

This value is smaller than most of the values, of $> 10R_g$, preferred by Kammoun et al. (2021b) when fitting the lag spectra of higher mass AGN, but changing from 6 to $10R_g$ makes little difference to the predicted lags as the light travel time across $1 R_g$ for this mass is only 1.78s. This simple analytic model fit would give a zero lag at very short wavelengths, ie in the X-ray band, but the minimum lag for any reprocessed radiation, even at very short wavelengths, is limited by the height of the X-ray source above the disc. Even though the minimum lag here, for a height of $10R_g$ is only < 35 s, we do not plot the lags below 1000\AA to avoid implying that the model lags should go through the zero point.

Model lags can be increased by increasing the illuminating X-ray luminosity and accretion rate. We noted previously (McHardy et al., 2018) that, for most AGN, changing the luminosity by a factor 3 does not typically change the lags by more than a few per cent and, using the formula in Kammoun et al. (2021a), increasing that luminosity by a factor 10 only leads to an 8% reduction in the lags. The observed lags are therefore shorter than can be produced by this mass, using the formula in Kammoun et al. (2021a), whatever the spin or illuminating luminosity. When we move on to detailed response modelling, we will consider the illuminating luminosity and accretion rate in more detail.

We therefore also plot, in dashed green lines, the lags expected from an alternative, lower, mass. Here we use $4.9 \times 10^4 M_\odot$ (Edri et al., 2012). We refer readers to Edri et al.

(2012) for a discussion of the relative merits of the two masses, including how a different interpretation of the data presented by Peterson et al. (2005) would result in the smaller mass. We scale the accretion rate and X-ray luminosity parameter according to fixed bolometric and 2-10 keV luminosities. The maximum-spin model lags are then in reasonable agreement with the observations out to about the r_s -band. We again note, as stated in McHardy et al. (2018) and emphasised also by Kammoun et al. (2021b), that spin has a major effect (factor 2) on lag predictions so if other disc parameters are known, and the modelling is well understood, lag measurements can provide a good method of measuring black hole spin.

We note, however, that the lag formula in Kammoun et al. (2021a) is derived from more detailed response modelling of AGN with larger masses and without including any effects of an edge to the emission region. It also has a stated lower limit of applicability of $10^6 M_\odot$ so we are below that limit. Next, we therefore carry out more detailed modelling of the lag spectra using full impulse response calculations.




FIGURE 2.21: The blue squares are X-ray and XMM-Newton OM UVW1 lags and ground based g-band lags from McHardy et al. (2016). The red circles are the HiPERCAM JAVELIN *ugriz* lags, corrected for seeing changes and the dark grey star is the B-band lag from the LT observations. All lags are referenced to the X-ray band, with the HiPERCAM u_s -band lag placed at 600s, as described in the text. The red lines are the predicted model lines using the formulation of Kammoun et al. (2021a) for a mass of $3.6 \times 10^5 M_\odot$ both for zero spin ($a=0$) and maximum spin ($a=0.998$). The green dashed lines are the predictions for mass $4.9 \times 10^4 M_\odot$.

2.5.2 Full Response Modelling of the Lags: The Disc Edge and The Colour Correction Factor

We used the KYNxilrev code⁸, previously used by Kammoun et al. (2021a), to produce response functions at a number of different wavelengths over our observed lag wavelength range. We make the standard assumption that the lag is defined as the mean of the response function. We do have an in-house code (Veresvarska, see Appendix A) which provides broadly similar results but, unlike the KYNxilrev code, it does not contain General Relativistic modifications. Hence, as it is also not yet publicly available, we base our results on the KYNxilrev code.

The determination of the relevant accretion rate and illuminating luminosity is not entirely clear in the case of NGC 4395. Its ratio of X-ray to optical luminosity is larger than that of the higher mass AGN which have been monitored recently with Swift. For example, using the BAT 105 month survey, the ratio of the X-ray flux of NGC 4593 to that of NGC 4395 is 3.2. However using our own observations here and previous observations by McHardy et al. (2018) which are all small aperture nuclear fluxes, the ratio of u-band fluxes is 26. McHardy et al. (2018) do not provide g-band but do give B (7mJy) and V (13mJy) fluxes of NGC4593. For a simple linear interpretation to g of 9.5 mJy we obtain a ratio to NGC 4395 of 9, again a good deal more than the X-ray ratio. NGC 4395 has a harder X-ray spectrum than most other AGN, except for LINERS, so may well have a jet contribution to the X-ray luminosity (e.g. King et al., 2013).

Regarding the bolometric luminosity, there are a number of estimates derived from optical measurements, via far from certain bolometric correction factors, e.g. Peterson et al. (2005) gives 5.4×10^{40} ergs/s, Moran et al. (1999) gives $5.2 - 8.8 \times 10^{40}$ ergs/s and Brum et al. (2019) gives 9.9×10^{40} ergs/s. The most recent estimate, including properly high energy (NuSTAR) observations together with optical gives 1.9×10^{41} ergs/s Kammoun et al. (2019). These authors also separately give a 3-10 keV luminosity of $0.9 - 2.4 \times 10^{40}$ ergs/s and a photon index of 1.6. Taking a mean value, extrapolation to 14-195 almost exactly agrees with the BAT value of 8×10^{40} ergs/s. This value also agrees with extrapolation of the XMM-Newton observations of Vaughan et al. (2005). Thus the illuminating luminosity is reasonably well agreed upon. Extrapolating the BAT luminosity to the whole illuminating X-ray band gives 1.6×10^{41} ergs/s⁹.

From the optical SED provided by our observations we have made a preliminary estimate of the disc luminosity. The details of the modelling will be presented elsewhere when completed but our preliminary estimate is $L_{disc} = 3.3 \times 10^{40}$ ergs/s. Adding that luminosity to the X-ray luminosity we arrive at the same overall bolometric luminosity of 1.9×10^{41} ergs/s as do Kammoun et al. (2019). Thus we

⁸<https://projects.asu.cas.cz/stronggravity/kynreverb>

⁹<https://swift.gsfc.nasa.gov/results/bs105mon/>

believe that we have a reasonably robust estimate of the overall bolometric luminosity. However the effective accretion rate through the disc which acts to determine the disc temperature structure and hence the lags, is less clear.

In a very simple scenario for material accreting through a disc, roughly half of the gravitational potential energy is radiated and half is turned into rotational motion which will either just be lost into the black hole or may power the corona. Multiplying the model disc luminosity by two then gives a rough estimate of the accretion rate as 0.146% of Eddington. In fact our more detailed SED modelling gives a more accurate measure as 0.1% of Eddington. For historical reasons we began our modelling using the accretion rate estimate from [Peterson et al. \(2005\)](#) of 0.12% of Eddington. As this value is close to the estimates from the SED we remain with the Peterson value in this report but note that it is uncertain. How we account for the remaining part of the bolometric luminosity, it is not clear. Perhaps some of the accretion occurs in a corona over the surface of the disc which does not participate in heating the disc.

The flattening of the lag spectrum at long wavelengths requires an edge to the emission region which we model here as a truncation in the outer disc radius; see also Fig.22 of [Kammoun et al. \(2021a\)](#). However the shape and amplitude of the lag spectrum also depends greatly on the assumed disc colour correction factor. This factor takes account of the fact that at high temperatures, in the inner disc, hydrogen and helium will be ionised so that the deeper, and hotter, parts of the disc are visible. So the disc will appear hotter than expected on the basis of the standard [Shakura and Sunyaev \(1973\)](#) and [Novikov and Thorne \(1973\)](#) models. The colour correction factor is discussed at length by [Done et al. \(2012\)](#) who provide prescriptions for the correction factor in different temperature regimes. We refer to their standard prescription as the DONE colour correction factor. This factor is approximated as 2.4 for $T > 10^5\text{K}$, 1 for $T < 10^4\text{K}$ and decreases between those two temperatures.

Finding the most appropriate combination of mass, accretion rate, spin, truncation radius, colour correction factor and even inclination is not trivial or unambiguous. Thus our aim here is not to find precise best-fit values but to find a range of acceptable values. We begin by testing the values of mass ($3.6 \pm 1.1 \times 10^5 M_\odot$ with 1σ uncertainty) and accretion rate (0.12% of Eddington) given by [Peterson et al. \(2005\)](#). This mass is the same as that ($4^{+8}_{-3} \times 10^5 M_\odot$ with quoted 3σ uncertainties) derived by [Brok et al. \(2015\)](#) from detailed gas dynamical modelling. We take an intermediate inclination of 45 degrees, which is in good agreement with the value of 37 degrees estimated by [Brok et al. \(2015\)](#) and gives almost identical lags. We assume a lamp-post X-ray emission geometry with a source height of $10R_g$. As noted earlier, changing this value even by a factor of 2 makes little difference to the predicted lags. We also choose zero spin. Thus the remaining variables are truncation radius and colour correction factor.

Following some experimentation, it became clear that a relatively small truncation radius, $\lesssim 2000R_g$, is required to produce enough flattening of the lags at long wavelengths, almost independent of the other variables. Assuming the DONE colour correction factor, which is the best physically motivated factor we show, in Fig. 2.22, models for a range of disc outer radii which broadly bracket the observed lag data and have roughly the right shape. A truncation radius of 1600-1700 R_g agrees reasonably with the data.

To investigate the effect of the colour correction factor, we then fix the outer radius at $1600R_g$ and vary the colour correction factor (Fig. 2.23) with the other parameters the same as for Fig. 2.22. Although the DONE prescription varies the colour correction factor as a function of disc temperature, the alternative implementations simply apply the same colour correction factor over the whole disc. A factor of 2.4 is used, following [Ross et al. \(1992\)](#), by [Kammoun et al. \(2021a\)](#). A broadly similar correction is derived, though with a different physical model, by [Petrucci et al. \(2018\)](#).

The model using the DONE colour correction factor, which is very close to almost no colour correction factor (i.e. $f_{col}=1$ or 1.2), is closest to the data for this mass, accretion rate and spin. The lag models with higher values of f_{col} differ greatly from the data, having quite different shapes. [Davis and El-Abd \(2019\)](#) give a prescription for f_{col} as a function of mass, accretion rate and stress parameter. For both the large and small masses, and for values of α between 0.001 and 1, f_{col} remains close to unity, ie close to the DONE value.

We have not shown model fits for maximum spin as the lags for maximum spin are $\sim 20\%$ less than those for zero spin and so, for this mass and accretion rate, lie well below the observed lags.

As a smaller mass, $\sim 4 \times 10^4 M_\odot$ has been proposed ([Edri et al., 2012](#)) we also show, in Fig. 2.24, some model lags for that mass with appropriate accretion rate of 1.08% of Eddington. Even for zero spin and a very large truncation radius, the observed lags are almost twice those for the DONE correction factor. We therefore experimented with the large colour correction factor, 2.4, used by [Kammoun et al. \(2021a\)](#). We show a range of cutoff radii. Although not lying as close to the data as the low colour correction factor models for the larger mass, the colour correction 2.4 model with cutoff radius of $14000R_g$, is in approximate agreement with the data.

We conclude from this preliminary investigation that disc reprocessing models based on the most commonly accepted mass of $3.6 \times 10^5 M_\odot$ for zero spin and for the most physically motivated disc colour correction factor (DONE factor), are in reasonable agreement with the observed lags but the disc must be truncated with an outer disc truncation radius of $\sim 1700R_g$.

Models based on a smaller mass, $4 \times 10^4 M_\odot$, can also be brought into moderate agreement with the observed lags but do require a very high colour correction factor of 2.4 over the whole disc, not just the inner hot part, and an outer truncation radius of $\sim 14000 R_g$.

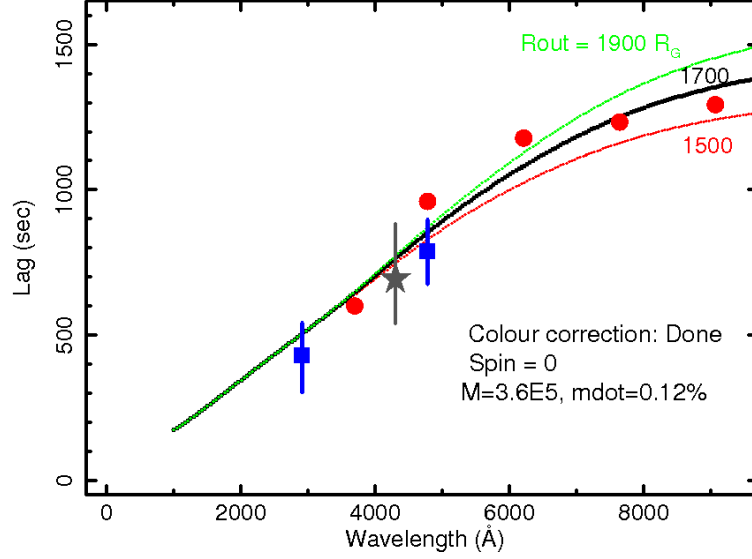


FIGURE 2.22: The lag observations are as described in Fig. 2.21. Here we assume a mass of $3.6 \times 10^5 M_\odot$, $\dot{m}_E=0.12\%$ Eddington, X-ray source height $10R_g$, zero spin and the DONE colour correction factor. The lines are lag predictions for different disc outer radii of 1500, 1700 and $1900R_g$.




FIGURE 2.23: As for Fig. 2.22 but with the disc outer radius fixed at $1600R_g$ and showing a range of colour correction factors, as described in the text.

2.5.3 BLR contribution to the UV/optical variability

In almost all other well studied AGN, with the possible exception of Mrk817 (Kara et al., 2021), the u_s -band lag is longer than, or at the very least about the same as, that




FIGURE 2.24: The lag observations are as for Fig. 2.21. The model lines are for a mass of $4.0 \times 10^4 M_{\odot}$, $\dot{m}_E=1.08\%$ Eddington, zero spin. The solid lines are for a colour correction factor of 2.4 and the dashed line is for the DONE colour correction factor. A range of truncation radii are noted.

in the surrounding wavebands. This u_s -band excess is usually attributed to a lag contribution from reprocessing in the more distant BLR (Korista and Goad, 2001, 2019). Here, however, the u_s -band lag is considerably shorter than that in the g_s -band. This result depends only on HiPERCAM observations so is completely robust. Although a weak contribution from the BLR is still consistent with the data, and the lag from the BLR can be reduced by a suitable choice of gas density and ionisation parameter (Lawther et al., 2018), the conclusion here is that the disc rather than the BLR is the dominant contributor to the UV/optical variability, at least on the few hour timescales observed here, in NGC 4395.

An HST UV spectrum of NGC 4395 is presented by Filippenko et al. (1993) and optical spectra by Lira et al. (1999). Most of the emission lines are dominated by an unresolved core but some of the permitted lines have weak broad bases. Line fluxes are given in Tables 2 and 3 from Lira et al. (1999). For example, the flux from the broad component of H_{β} was observed between 14 and 18×10^{-15} ergs cm $^{-2}$ s $^{-1}$, compared with between 22 and 28×10^{-15} ergs cm $^{-2}$ s $^{-1}$ for the narrow component, both of which variations are within the quoted calibration uncertainties of 30%. Typical continuum levels below 4000 Å have varied between 0.5 and 1×10^{-15} ergs cm $^{-2}$ s $^{-1}$ Å $^{-1}$. Thus when considered over a typical $ugriz$ bandpass of ~ 1000 Å, emission line variability is not a large contributor to the total flux variations. This limited contribution from the BLR may explain why we do not see the large excess u -band lag in NGC 4395.

2.5.4 Physical Origin of the Truncation Radius

There are two natural radii which might be considered: the self-gravity radius and the dust sublimation radius. Considering first the self-gravity radius, the physics here is not straightforward and simple. As one moves out in the disc the first change is when the gravity of disc material to itself exceeds the gravity to the central object. Disc structure may become more fragmented but there will still be a disc, available for reprocessing. Still further out cooling may become efficient enough that the disc will collapse into small clumps which will not be effective reprocessors.

There is considerable modelling of this behaviour which depends on a number of assumptions, e.g. regarding metallicity, illumination and viscosity, which we shall not attempt to repeat. We refer the interested reader to [Gammie \(2001\)](#) and to (e.g. [Shore and White, 1982](#); [Laor and Netzer, 1989](#); [Collin-Souffrin and Dumont, 1990](#); [Collins and Huré, 1999](#); [Lobban and King, 2022](#)). Those papers provide formulae enabling the reader to estimate the self-gravity radius. However, of relevance to NGC 4395 is that [Shore and White \(1982\)](#) note that the ‘cliff’ which appears at the self-gravity radius, beyond which the disc scale height is much reduced and reprocessing from the disc will stop does not appear for masses $< 3 \times 10^7 M_{\odot}$. For lower masses, only a smoothly flared disc results. We conclude that, although further work is needed to be certain, the self-gravity radius may not be so important in defining an effective outer radius to the disc for reprocessing in low mass AGN like NGC 4395 and we do not consider it further.

Considering next the sublimation radius, [Baskin and Laor \(2017\)](#) state that, at the radius at which the disc temperature falls to below the dust sublimation temperature, the atmosphere of the disc will become dusty, its opacity will greatly increase and the disc can again become radiation-pressure supported. The resultant inflated dusty disc obscures the region behind it, leading to an edge to the part of the disc which can reprocess radiation. A dusty disc wind, driven by radiation pressure, will also form, which would constitute the inner radius of the BLR and give rise to the wind component seen in some emission lines, such as CIV (e.g. [Coatman et al., 2016](#)). The dust sublimation radius is therefore a natural edge to reprocessing either from the disc or the BLR.

The dust sublimation radius, R_{sub} , depends on the dust composition and on the temperature at which it will sublime. Taking a sublimation temperature of 2000K [Baskin and Laor \(2017\)](#) with an X-ray source height of $10 R_g$ an illumination luminosity of $1.6 \times 10^{41} \text{ ergs cm}^{-2} \text{ s}^{-1}$, an albedo of 0.2 and zero spin we calculate sublimation radii for dust on the surface of the illuminated disc of $4300 R_g$ and $1900 R_g$ respectively for the large and small masses. For the smaller mass this radius is not too much larger than the radius determined from the lags, and may be consistent within the many uncertainties of the parameters, but for the large mass the sublimation

radius is larger. For the small mass a dusty wind might therefore obscure the outer part of the disc and provide a suitable edge.

For the larger mass the sublimation radius is a factor of almost 3 further out than the truncation radius derived from the lags. We might try to drop the sublimation radius by modelling the lags with a maximum spin disc. Although the sublimation radius then drops by about 16% as the disc is cooler due to a lower accretion rate in kg/sec for the same accretion rate in Eddington units, the lags drop by about 10% too, so that does not greatly help. The sublimation radii of cooler, unilluminated discs are lower, closer to the truncation radii measured from the lags, but we require illumination to produce the variable UV/optical emission. Also, for an almost unilluminated disc, the lags drop by almost a factor of 2, at least at short wavelengths. If the disc continues out to the sublimation radius and maybe beyond, we therefore require material to shield the outer portion of the disc from the central illumination. If this material is in the form of a disc wind we would need it, for the large mass, to leave the disc at a local temperature of $\sim 4 - 5000\text{K}$, i.e. above the dust sublimation temperature. Probably this material would have to be gaseous rather than dusty as, once raised from the disc and exposed fully to the illuminating radiation, the dust would be rapidly vapourised. A line-driven wind might be applicable in this case (e.g. [Proga and Kallman, 2004](#); [Higginbottom et al., 2014](#)). We comment further in Section 2.5.5.

2.5.5 Alternative explanations of the shorter than expected long wavelength lags

The i_s and z_s lags here are shorter than expected from disc reprocessing. Shorter than expected lags do not necessarily imply reprocessing material closer to the illuminating source than expected, we simply require that the path from the illuminating source, via the reprocessor to the observer is shorter than expected. A wind, close to the line of sight (e.g. [Hönig, 2019](#)) might, depending on the orientation, add very little extra distance for the light path from the central source to the observer and might produce very short near-IR lags. The base of the wind, if optically thick, might provide an effective wall which might mimic a truncated disc, at least for the longer wavelengths.

The model AGN lag spectra from the BLR by [Netzer \(2021\)](#) contain a peak due to the Paschen continuum at 8200\AA and then become shorter at longer wavelengths. Although this decrease, compared to the expectation from a standard untruncated disc, is in the same sense as our observed decrease, the change which we observe is even more extreme. However a difficulty with interpreting the present observed lag spectra as being mainly due to reprocessing in the BLR is that here the g_s -band lags a long way after the u_s -band, whereas for BLR reprocessing we expect those two lags to be the other way around.

However using the relationship between the H_β BLR radius and the 5100Å luminosity for low luminosity AGN derived by [Bentz et al. \(2013\)](#), and taking a luminosity of $5.9 \times 10^{39} \text{ ergs s}^{-1}$ from [Peterson et al. \(2005\)](#), we derive a radius, in light-travel time, of 1600s. Given the many uncertainties involved, this figure is quite close to the lag at long wavelengths which our measured lags seem to be asymptotically approaching. Thus the observations are consistent with the inner edge of the BLR being cospatial with the outer edge of the emission region and thus providing an edge to the emission region.

A possible scenario then is that a wind arises from the disc which both shields the outer regions of the disc from central illumination and acts, itself, as a reprocessor. If the disc temperature is above the dust sublimation temperature this may not be a dusty wind but winds can be driven by other mechanisms, eg line driving, leaving the disc at higher temperatures than the sublimation temperature. Such winds can still shield the outer disc from X-ray irradiation ([Proga and Kallman, 2004](#); [Higginbottom et al., 2014](#)). The wind may also act as a reprocessor and contribute to the variable UV/optical emission and there may even be some feedback of radiation from the wind onto the disc. None of these effects are considered here and clearly, further modelling is needed to test the wind scenario.

2.5.6 Implications for other AGN

The observations, and possible interpretation, presented here have implications for our understanding of the previously puzzling lag spectra of some other AGN.

2.5.6.1 The X-ray / UV disconnect

As noted in the Introduction, one of the major problems in interpreting AGN lag spectra has been the observed excess lag between the X-rays and the shortest wavelength UV band, compared to lags between the UV and optical bands. Whilst it may not be the only cause of this excess lag, fitting a disc model without truncation to observations from a disc that actually is truncated will automatically produce an apparent excess lag. A truncated disc model is then consistent with a direct view of the disc from the central X-ray source rather than requiring additional scattering through, e.g., an inflated inner edge of the disc ([Gardner and Done, 2017](#)), to explain the excess X-ray to UV lag.

2.5.6.2 Overall lag modelling and mass estimation

As we can see from Figs. 2.22 and 2.24, if the UV/optical variability really does originate from reprocessing of X-ray emission by a truncated accretion disc, fitting a model from an un-truncated disc will produce nonsense and completely incorrect mass estimates. The very short wavelength bands, being the least affected by disc truncation, are the most sensitive to the mass and accretion rate.

Here we note that response functions for longer wavelength bands than B show a cliff at long timescales but the shorter wavelength responses are unaffected. The red bands are most sensitive to the truncation radius. A flattening of the lag spectrum with increasing wavelength is an indication of disc truncation so it is important to define the lag spectral slope well at short wavelengths. [Cackett et al. \(2018\)](#), using HST to extend lag measurements down to $\sim 1100\text{\AA}$, finds a steepening at wavelengths below 2000\AA in NGC 4593 which is consistent with disc truncation. With only a small number of bands, there will be degeneracy in the fits but with a full range of wavelengths quite sensitive mass estimates, and disc truncation radii, may be made.

If there is only a mild degree of truncation, affecting only the longest wavelengths, then lag fits of the form $\text{lag} \propto \lambda^\beta$, which do not include any truncation, could produce values of β nearer to unity than $4/3$, e.g. as seen in NGC 5548 [Fausnaugh et al. \(2016\)](#).

2.6 Conclusions

We present three nights ($\sim 15\text{hr}$) of u, g, r, z and two nights ($\sim 12\text{hr}$) of u, B, g imaging of the low mass AGN NGC 4395 on the robotic LT. These data provide lightcurves which show clear correlated variations on few hour timescales but with interband lags too short to be measured accurately given the $\sim 200\text{s}$ sampling. A follow-up 6 hr observation with HiPERCAM on the 10.4m GTC provided superb simultaneous 5-band (u_s, g_s, r_s, i_s and z_s) imaging, with 15s sampling in u_s and 3s in all other bands, the fastest yet reported for an AGN. The observation is dominated by two smooth, almost sinusoidal, variations, similar to those seen in the LT observations.

Very clear correlated variability is seen in all HiPERCAM bands and wavelength dependent lags between bands are measured relative to the u_s -band with much higher precision than in any other AGN. Two main points are immediately obvious from the lags measured from the HiPERCAM observations. Firstly, and unlike in almost all other AGN, the g_s -band lags considerably after the u_s -band. Thus the u_s -band is not dominated by Balmer continuum emission from a distant BLR and hence reprocessing in the BLR is not a large contributor to the UV/optical continuum variability in this AGN. Secondly, the lags between the r_s, i_s and z_s -bands are very small, so that the lag

spectrum flattens off at long wavelengths. These observations imply an edge to the emission region which limits the outward movement of the centroid of the emission regions at long wavelengths. A likely possibility is a truncated accretion disc.

Although we do not discuss it at length here, we observe that when the source is brighter it is bluer and also the lags, at least at short wavelengths, are longer. These results are in fairly good agreement with the expectation of disc reprocessing with a higher illuminating luminosity, around a factor 2 or 3. Variations in the X-ray luminosity by even larger factors on comparable timescales are commonly seen in X-ray observations (Vaughan et al., 2005) and could easily explain such variations. We hope to repeat this analysis with higher quality HiPERCAM data to test whether these results are repeatable and will then look into the colour and lag variations in more detail.

We have combined the HiPERCAM lags with lags measured previously between the X-ray, UV and g band (McHardy et al., 2016) and have fitted the combined lag spectra with truncated accretion disc models using lags derived from the KYNxilrev code and also from our own in-house models. Similar results are obtained with all codes. We have not tried to perform a full statistical fit over all parameters but have compared the observed lag spectrum with models based on the most commonly assumed mass of $3.6 \times 10^5 M_{\odot}$ (Peterson et al., 2005; Brok et al., 2015) and also on a mass a factor of 10 lower (Edri et al., 2012).

We demonstrated that the disc colour correction factor has a large effect on the model lags. The DONE temperature-dependent factor is probably the best physically motivated correction factor and is also similar to the colour correction factor near unity which can be derived for NGC 4395 using the formulation from Davis and El-Abd (2019). Using that factor we find that model lags based on a truncation radius of $\sim 1700 R_g$ agree with the observed lags reasonably well for the mass of $3.6 \times 10^5 M_{\odot}$ and zero spin. For a mass of $4 \times 10^4 M_{\odot}$ the model lags based on the DONE colour correction factor are too short but lags based on a factor of 2.4 applied over the whole disc are similar to the observed lags tho, visually, the agreement between the observations and the models is not quite as close as for the DONE factor with the larger mass.

Regarding possible reasons for the truncated disc, Shore and White (1982) note that for low masses such as those considered here the disc does not suffer catastrophic collapse at large radii due to self-gravity effects and so self-gravity is probably not the cause of the truncated disc. For the small mass the disc truncation radius measured from the observed lags is about at the same radius as the dust sublimation radius. However for the large mass the sublimation radius is almost a factor 2-3 larger than the truncation radius measured from the lags. If our calculations are correct, a shield of some sort is required to cut off illumination to the outer disc. A dusty disc wind might work for

the smaller mass but for the larger mass, where the dust sublimation radius lies beyond the measured disc truncation radius, a different type of wind, eg line-driven, might be more applicable (Proga and Kallman, 2004; Higginbottom et al., 2014).

We note that, if the relationship between optical luminosity and BLR radius for low luminosity AGN (Bentz et al., 2013) is applicable to NGC 4395, the inner edge of the BLR should be at a light radius of 1600s, which is close to the maximum lag measured from our data. Thus the inner edge of the BLR may be close to the edge of the UV/optical emission region and may be at least partially responsible for the apparently truncated disc.

There are a number of caveats. In particular there are a large number of variables involved in the lag modelling and we have not carried out an extensive parameter search. It may very well be that a different combination of variables would produce model lags which are closer to the observed lags. However we do not believe that such modelling is yet merited by the data. Although the lags between the HiPERCAM bands are determined with very high precision, we only have one set of measurements and the observations should be repeated to test for long term stability. We also caution that referencing the HiPERCAM lags to the X-ray band requires an estimate of the lag between the XMM-Newton UVW1 band (2910Å) and the HiPERCAM u_s -band (3695Å). Due to the closeness of these wavebands the lag is unlikely to be large ($\lesssim 150$ s), but it is not zero. We also only have one measurement of a lag between the X-rays and any UV band, here UVW1. However any change in the registration between the X-ray frame and the HiPERCAM frame would mainly result only in slightly changing the best estimate of the truncation radius and cannot affect the main conclusions of this report. To address the registration concern, further X-ray and UV observations from four full XMM-Newton orbits (> 400 ks), together with ground based g-band observations, have been made and are reported in Chapter 3.

We note that, immediately prior to submission, we have been informed that another group, using different methods, are working on a similar conclusion regarding the outer edge of the disc (Starkey et al, private communication).

We finally note that the HiPERCAM observations define the optical high frequency power spectrum very well. A description of the overall optical power spectrum, including long-timescale (months-year) monitoring from the LT and intermediate scale (hours-weeks) observations from TESS (Burke et al., 2020) is found in Chapter 4.

Chapter 3

X-ray, UV, and Optical Reverberation Mapping of NGC 4395

3.1 Abstract

In this chapter, I will expand on the results from Chapter 2 with the addition of multiple simultaneous high cadence X-ray and UV observations from *XMM-Newton* and ground-based optical telescopes to gain a more complete picture of the X-ray to UV/Optical time lags in NGC 4395.

3.2 Introduction

3.2.1 Theory

NGC 4395 is of particular interest for Reverberation Mapping because it is an extremely low mass AGN. Measurements taken by [Peterson et al. \(2005\)](#) give a mass value of $(3.6 \pm 1.1) \times 10^5 M_{\odot}$ which would make it one of the smallest AGN ever measured. Therefore by probing its structure, we will be answering questions on whether the disc geometry of an accreting supermassive black hole system changes more than would be expected in the models we will test by simply a change in mass. In other words, by collecting data on a fringe case like 4395, we are testing the limits of existing models and, where the data diverges from them, proposing mechanisms by which these differences might be explained.

One of the most famous models for an accreting SMBH is the one proposed by [Shakura and Sunyaev \(1973\)](#) which simply models a geometrically-thin, optically-thick radiating disc. In this relatively simplistic model the source of the temperature gradient across the disc is explained by the release of gravitational potential energy as the accretion material moves inwards. This generates a temperature profile of

$$T \propto (M\dot{m}/R^3)^{1/4}$$

where M is the mass of the black hole, \dot{m} is the accretion rate, and R is the radius from the black hole. From this we can see that in this model the X-ray emission boosts the thermal emission already present in the disc, so we should be able to find a lag between the X-rays and the re-radiated lower energy photons, dependent on the wavelength, such that

$$\tau = R/c \propto (M^2\dot{m}_E)^{1/3}\lambda^\beta$$

where τ is the lag, λ is the wavelength, \dot{m}_E is the accretion rate in Eddington units, and $\beta=4/3$. [Cackett et al. \(2007\)](#)

To test that our lags are consistent with the Shakura & Sunyaev model, they should be consistent with $\beta=4/3$, which has been observed in [Cackett et al. \(2007\)](#) for the UV/Optical component, but was not measured to the X-rays. As well as this simple time lag, we should expect the UV emission to be much smoother with lower amplitudes of the variations than the X-ray emission, and similarly for the Optical relative to the UV emission as the photons for each come from progressively larger emission areas as the radius increases.

There are complications to this simple model however. It's possible that the Broad Line Region, radially further out than the accretion disc, is also producing reprocessed UV/Optical emission. This BLR emission would equally contribute to the observed light curves, due to it being unresolved by our telescopes, and so perhaps interfere with lag measurements obtained from them. In fact previous campaigns, for example [Edelson et al. \(2015\)](#), have observed a larger lag for the u-band than for the surrounding wavebands. This then could be evidence of contribution to the variation in the light curve by the BLR.

In addition, many campaigns such as [McHardy et al. \(2014\)](#) have noted that if the Shakura & Sunyaev model is assumed then the disc should be larger than we expect. This is consistent with several gravitational microlensing campaigns such as [Morgan et al. \(2010\)](#) which independently measure accretion discs to be larger than expected.

Other concerns have been raised about the exact structure of the corona. A basic model is one called the 'lamp post model' where the X-ray corona is a small emitting region above the black hole, illuminating symmetrically over the disc. However observations such as [Arévalo and Uttley \(2006\)](#) have noted that this coronal geometry does not lead to enough of the X-ray emission interacting with the disc to produce the observed optical variability. There are two proposed solutions to this. The first being that the X-ray corona is an extended object with a structure that produces significant effects on its emission variability.

The other is one proposed by [Gardner and Done \(2017\)](#) which is that a majority of the X-ray emission does not shine onto the outer parts of disc, and instead heats up the inner edge of the disc. This inner edge then inflates in height, once again moving away from a simple Shakura & Sunyaev disc, and re-radiates in hard UV. This hard UV is then what illuminates the rest of the disc, rather than the coronal X-ray emission. This should be measurable by an excess lag between the X-ray and UVW1 bands, representing the thermal time-scale it takes the X-ray heating energy to pass through the inner disc before being re-radiated out the other side. This excess lag has been observed by both [Gardner and Done \(2017\)](#) and [Edelson et al. \(2017\)](#) for NGC 5548 and NGC 4151 respectively. In NGC 4151 however, this excess is dependent even on the energy of the incident X-rays, increasing with higher energy X-rays. This may be due to greater scattering of the higher energy X-rays in the inflated inner disc before the re-radiation occurs.

3.2.2 Previous Observation Campaigns

Previous X-ray observations have used RXTE for lag measurements, however due to the low mass of NGC 4395, and therefore the smaller time lags, its time resolution is not sufficient for accurate lag measurements. However, measurements of the necessary level of accuracy are achievable using *Swift* and *XMM-Newton*. Both observatory's X-ray telescopes, *Swift*'s X-Ray Telescope (XRT) and *XMM-Newton*'s EPIC, operate within the 0.5-10 keV range. Both also allow simultaneous UV/Optical coverage from the UV-Optical Telescope (UVOT) and the Optical Monitor, respectively. Both these cameras can observe in three UV filters (UVW1, UVW2, and UVM2) and three Optical filters (U, B, and V).

One such set of *Swift* observations is from [Cameron et al. \(2012\)](#). Using *Swift* for both long-term and short-term X-ray/UV/Optical observations of NGC 4395, they managed to detect short-term lags, although the errors on these values were quite large and therefore their conclusion uncertain. To determine a value for their lag they perform a Discrete Correlation Function and found a ~ 400 s lag of the UVW2 from the X-ray, although it is a single DCF and has rather large error bars, so this value is rather low confidence.

Another set of relevant observations of NGC 4395 is by [McHardy et al. \(2016\)](#). This observation campaign, similar to our campaign, used a combination of *XMM-Newton* and ground-based observations. This consisted of two ~ 53 ks observations from *XMM-Newton* on the 28th and 30th of December 2014, with complimentary ground-based observations from six ground-based telescopes (LCOGT McDonald Observatory, Texas; Whipple Observatory, Arizona; LCOGT Haleakala Maui; Kanata Observatory, Japan; ARIES observatory, India and the Wise Observatory, Israel).

To get more accurate lags than with a DCF, the lags were calculated using the Javelin program. The determined lags were 473^{+47}_{-98} s and 788^{+44}_{-54} s, respectively, which when plotted against wavelength can be fitted with a β value. By forcing the fit through zero a straight line (representing $\beta=1$) fits the best. However due to the error bars a fit of $\beta=4/3$ is also consistent with the data, and therefore is also consistent with the Shakura & Sunyaev model. These observations are therefore consistent with X-ray reprocessor driven UV/Optical emission.

3.3 Observations

3.3.1 Logistics

Our observation campaign was designed as a successor to [McHardy et al. \(2016\)](#) and as such contains similar elements but expanded in both time and repetition. The X-ray and UV components of this campaign is a set of 4 *XMM-Newton* observations, with each observation being ~ 117 ks each, on the 13th, 19th, and 31st of December 2018, and one more on the 2nd of January 2019. The X-ray telescope used was the EPIC instrument, operating within the 0.5-10keV range, similar to the 2016 campaign.

The simultaneous UV observations were taken with the Optical Monitor (OM) telescope, and as we had 4 separate observation periods it allowed us to observe the UV component in different filters for different nights. The filters chosen for each night were UVM2 for December 13th and 19th, UVW1 (as was used in [McHardy et al. \(2016\)](#)) for December 31st, and for the final night, the OM's U-band Optical filter. This allows us to probe different UV/Optical regions on each night, in principle giving us a more complete picture of the lag spectrum for NGC 4395.

OM has two operating modes, Imaging Mode and Fast Mode. Imaging Mode produces accumulated images with no photon timing data. These have a minimum exposure time of 800s, which when combined with the readout time of ~ 300 s, means the *minimum* time resolution you can achieve is ~ 1100 s which is far too large to observe UV/Optical variation in an AGN as low-mass as NGC 4395 which is expected to be on the time-scale of <1000 s based on previous observations.

Fast Mode on the other hand operates much like an X-ray detector; instead of generating a two-dimensional accumulated image it measures individual photon's time of arrival at each pixel location. This mode has an integration time of 500ms (smaller times are possible but are not necessary for our purposes, based on light travel times of a AGN of 4395's size) and so can effectively be binned up to any time resolution we require, thus giving us the ability to probe sub-kilosecond lags. The downside of this is the spatial resolution is reduced to 23x23 pixels, but given our measurements only require measurement of photon count changes, this is not an issue as long as the object is on the Fast mode section of the CCD.

For our ground-based Optical observations we unfortunately cannot perform a continuous 32.5 hour observation from a single telescope and so we must utilise multiple observatories spread longitudinally across the planet to maintain simultaneous optical data with the long continuous X-ray/UV observations. These observations were all taken using the SDSS g-band filter.

They were taken using the 0.7m Centurion 28 (C28) Telescope at the Wise Observatory of Tel Aviv University in Israel, the IO:O instrument on the 2m Liverpool Telescope at Roque de los Muchachos Observatory in La Palma, the Keplercam CCD on the 1.2m telescope at the Fred Lawrence Whipple Observatory (FLWO) in Arizona, the 1.3m Telescope at Devasthal at the Aryabhata Research Institute of Observational Sciences (ARIES) in India, and two telescopes that are part of the Las Cumbres Observatory (LCOGT) network of telescopes: a 2m Spectral telescope at Haleakala Observatory in Hawaii and a 1m Sinestro telescope at McDonald Observatory in Texas.

The difference in telescope reflector size was another factor we had to take into account. The larger 2m telescopes could collect a good quality image with only a 30-40s exposure, which is ideal on the time-scales we are working with. The smaller telescopes however, such as the C28 and the Sinestro required triple the exposure time to retrieve a similar quality of image for the g-band exposures, which based on the previous observations we expected an x-ray lag of ~ 800 s they are still short enough exposures to give us useful information. Details of the observations can be seen on Tables 3.1 and 3.3.

3.3.2 Challenges

Despite extensive planning, observations using multiple sets of equipment simultaneously for large time intervals are inevitably going to invite practical challenges. However, given the large number of independent observers as well as some thoughtful data selection we minimised the impact of these challenges on the final results.

One of the few things that could not be avoided was issues with the *XMM-Newton* observations themselves; as our only source of X-ray/UV observations, any issues with the data from the satellite would have inevitably caused issues. And unfortunately, this was the case for us on the observation that took place on December 31st. The reduced spatial resolution of the OM's Fast Mode means accurate object tracking is absolutely necessary to capture the source, and for this night's observations a fault with the OM's centroiding used for tracking the object caused it to not appear in the frame for the entire duration of the observation, leading to the loss of the entire UV observation for this night. This carries the additional consequence of being the night we were observing using the UVW1 filter, used in [McHardy et al. \(2016\)](#), which therefore prevents us from making a direct comparison of our UV lags with those from that campaign. In addition to this we had some other more limited gaps in our *XMM-Newton* data. On the night of December 13th a ground station outage interrupted the observation and, although it was successfully restarted, caused a 1.2hr gap in all coverage from the satellite. Then on the night of December 19th an instrument anomaly caused the loss of two OM exposures, equating to a gap of 8800s in the UV data.

For our ground-based data, most of it was unaffected with the exception of a few tracking losses which still kept the objects in frame and a few technical issues like file ordering which could be relatively easily overcome. The two exceptions to this are, first, our LCOGT Sinestro observations on the 20th of December suffered from excessive lunar sky brightness which made the first 59 frames not suitable for photometry. The final 27, however, were perfectly usable as they were taken after the moon had set below the horizon. The second is the Sinestro observation on the 31st of December which has no objects whatsoever in frame. This was due to the shutter on the fa05 Sinestro camera becoming stuck closed for a time that night that included our observation period, which unfortunately means that no usable data was recorded.

3.4 Data Reduction

3.4.1 X-ray/UV

XMM-Newton data reduction is a fairly simple ordeal, as luckily the ESA provides a software package called Science Analysis System (SAS), which automates a lot of the more complicated aspects of the process. Using the Observation Data Format (ODF) files provided, SAS can generate photon event files. From there it is relatively trivial to obtain light curves for the X-ray and OM components of the observations, with any

| Observation Period | Observation Start Time (UTC) | Observation End Time (UTC) | OM Filter | Observation Notes |
|--------------------|------------------------------|----------------------------|-----------|--|
| XMM-Newton | | | | |
| 1 | 2018-12-13: 06:16:19 | 2018-12-14: 14:46:19 | UVM2 | 1.2hr gap in observation for all instruments due to 'ground station outage' |
| 2 | 2018-12-19: 05:51:22 | 2018-12-20: 14:18:02 | UVM2 | ~3hr gap in OM data (two 4400s exposures lost) because of 'Instrument Anomaly' |
| 3 | 2018-12-31: 05:11:45 | 2019-01-01: 13:45:04 | UVW1 | OM instrument not pointed at object due to 'Bad centroiding tables load' |
| 4 | 2019-01-02: 05:03:24 | 2019-01-03: 13:16:44 | U | No additional notes |

TABLE 3.1: Space-based Observations from *XMM-Newton*

| Object Name | SDSS Filter | | | | | u-r colour |
|-------------|-------------|-------|-------|-------|-------|------------|
| | u | g | r | i | z | |
| AGN | 16.42 | 14.55 | 14.01 | 13.72 | 14.08 | 2.41 |
| Star 1 | 17.41 | 16.23 | 15.71 | 15.56 | 15.45 | 1.70 |
| Star 2 | 17.77 | 15.29 | 13.96 | 15.77 | 12.68 | 3.81 |
| Star 3 | 16.85 | 15.60 | 15.21 | 15.08 | 15.04 | 1.64 |
| Star 4 | 14.92 | 14.99 | 12.88 | 12.72 | 13.25 | 2.04 |

TABLE 3.2: Optical magnitudes for the four comparison stars and the AGN taken from SDSS

desired cadence, and after that if the light curves requiring further binning this can be done afterwards. We chose to create light curves with 10s bins, which given our variation time-scale of hundreds of seconds is a sufficient level of time resolution.

3.4.2 Optical

3.4.2.1 Comparison Star Selection

One inevitable challenge that comes with using ground-based data is that due to atmospheric turbulence and weather we cannot simply use the raw photon counts captured from the AGN to measure the variability. In order to counteract these effects we can find a source of otherwise constant flux in the frame, namely a comparison star, to obtain a light curve that correlates heavily to the random variations due to the atmospheric effects, and as such we can counteract the effect these variations have on the AGN's light curve. If multiple stars are present in the frame it then becomes important to determine which star will give us the cleanest removal of the atmospheric variations. Ideally, we want a star that is not variable over our observation timescales,

| Observation Period | Observation Start Time (UTC) | Observation End Time (UTC) | Filter | Observation Notes |
|---|------------------------------|----------------------------|---------|--|
| Liverpool Telescope IO:O | | | | |
| 1 | 2018-12-14: 03:15:59 | 2018-12-14: 06:09:51 | SDSS g' | No additional notes |
| 2 | 2018-12-20: 02:49:48 | 2018-12-20: 06:23:24 | SDSS g' | Tracking failure for last ~2000s, final 17 frames unusable |
| 3 | 2018-12-31: 05:08:45 | 2018-12-31: 06:29:30 | SDSS g' | No additional notes |
| 3 | 2019-01-01: 02:09:38 | 2019-01-01: 05:27:36 | SDSS g' | Tracking failure for last ~4000s but all frames are usable |
| 4 | 2019-01-02: 05:02:50 | 2019-01-02: 06:33:21 | SDSS g' | No additional notes |
| 4 | 2019-01-03: 02:00:41 | 2019-01-03: 05:18:59 | SDSS g' | Tracking failure for last ~4000s but all frames are usable |
| Wise Centurion 28 | | | | |
| 1 | 2018-12-13: 23:58:44 | 2018-12-14: 03:31:30 | SDSS g' | No additional notes |
| 3 | 2018-12-31: 22:41:30 | 2019-01-01: 03:40:20 | SDSS g' | No additional notes |
| 4 | 2019-01-03: 00:08:36 | 2019-01-03: 03:32:22 | SDSS g' | No additional notes |
| Las Cumbres Observatory Haleakala Spectral | | | | |
| 1 | 2018-12-13: 12:51:04 | 2018-12-13: 15:18:30 | SDSS g' | No additional notes |
| 1 | 2018-12-14: 12:47:12 | 2018-12-14: 15:18:01 | SDSS g' | No additional notes |
| 3 | 2018-12-31: 11:33:07 | 2018-12-31: 15:36:17 | SDSS g' | No additional notes |
| 4 | 2019-01-03: 11:26:10 | 2019-01-03: 15:39:14 | SDSS g' | No additional notes |
| Las Cumbres Observatory McDonald Sinestro | | | | |
| 1 | 2018-12-14: 12:11:38 | 2018-12-14: 12:46:08 | SDSS g' | No additional notes |
| 2 | 2018-12-20: 08:52:02 | 2018-12-20: 12:39:53 | SDSS g' | First 59 frames unusable due to excessive lunar sky brightness but final 27 frames taken post-moonset are good |
| 3 | 2018-12-31: 08:08:52 | 2018-12-31: 09:57:49 | SDSS g' | No images as shutter on Sinestro camera was stuck closed |
| 3 | 2018-12-31: 11:00:33 | 2018-12-31: 12:52:18 | SDSS g' | No additional notes |
| 3 | 2019-01-01: 11:02:10 | 2019-01-01: 12:56:30 | SDSS g' | First 66 frames unusable |
| Fred Lawrence Whipple Observatory Keplercam | | | | |
| 3 | 2018-12-31: 09:02:02 | 2018-12-31: 13:27:10 | SDSS g' | Not flat fielded but still usable |
| ARIES Devasthal Telescope | | | | |
| 4 | 2019-01-02: 20:54:00 | 2019-01-03: 00:44:09 | SDSS g' | No additional notes |

TABLE 3.3: Ground-based Observations from multiple telescopes

is isolated from other objects that would cause interference with photometry, and is bright enough to give us as high a signal-to-noise ratio as possible.




FIGURE 3.1: Finding chart of the stars around NGC 4395, generated using an LCOGT Spectral exposure.

Above we show a finding chart generated using an LCOGT Spectral exposure in Fig. 3.1. It displays the AGN as well as the four useful stars in the frame, labelled 1 to 4 from their right to left positions on the image, respectively. There is a fifth star in the frame (labelled Star B), however, this star is both not particularly bright. On top of this, it is close enough to the centre of the galaxy that there is significant background contamination of the aperture used to capture photometry data, and this background flux varies rapidly enough in its surroundings that it introduces undesirable levels of uncertainty into its light curve. Therefore, it was excluded from selection immediately.

Of the four remaining stars then, looking at our requirements for a comparison star, they are all far enough away from the galaxy centre that they have a minimal background that is approximately constant over the scale of an aperture, with any anisotropies that would interfere with background subtraction small enough to be masked out. Comparing all of their light curves with one another they also all appear to be non-variable stars. All four independent stars could be variable stars that vary at exactly the same rate consistently, but this is so unlikely as to be functionally impossible.

With the first two criteria fulfilled it would seem then that the brightest star of the four would be the most suitable. Star 4 is the brightest by an order of magnitude, with Star 2 then being $\sim 1.5\times$ brighter than Star 3, and finally Star 3 is $\sim 2\times$ brighter than Star 1. However there are several complications to this basic arrangement, the first and most prominent being that the chip on the LT images is not positioned to include star 3 and 4. This effectively means we cannot use either of these stars at all for comparison as

the Liverpool Telescope data makes up a significant portion of our ground-based observations for every single night. The brightest star then that appears on all exposures is Star 2, but this also presents a problem.

Star 2 appears constant as all the others and otherwise produces light curves similarly to the other stars, with the one exception of data taken from the LCOGT Sinestro which appears to show a decrease in flux compared with all other instruments, as can be seen in Fig. 3.2. Care was taken to ensure the aperture sizes corresponded with the published 0.389 arcsec/pixel resolution of the Sinestro CCD, ensuring any binning was accounted for as well, and all of this was checked multiple times with multiple reruns of the reduction pipeline to ensure no mistake was being made. However, despite this extra attention being given, the disparity remains.




FIGURE 3.2: Star 2 compared to the other three comparison stars, demonstrating the disconnect between the Sinestro light curves and the other telescopes.

Looking at the magnitudes of each star in SDSS data (shown in Table 3.2) shows Star 2 to be far redder than the other stars, with a colour $u - r = 3.81$ compared to the values of 1.70, 1.64, and 2.04 for Stars 1, 3, and 4 respectively. However, this should only be a factor if the filter used on Sinestro was different from the rest, but consulting both the observation logs and the headers of the .fits files shows that the filter used was the standard LCOGT SDSS g' filter. It could then be argued that the LCOGT SDSS g' filter is somehow different from those used by the other telescopes, however this argument also does not hold up as the LCOGT Spectral uses identical standard filters and yet does not demonstrate this disparity. It appears then that we cannot explain this disparity, which means we cannot rigorously correct for it without unintentionally introducing other sources of error.

In this case then in order to use Star 2 we would have to exclude the Sinestro data which makes up a significant portion of our ground observing time on two nights and will therefore be vital in making our lag calculations more accurate. Alternately we could use Star 2 on the two nights (13th of December and 2nd of January) with little or

no Sinestro observing time, however using different comparison stars for half of the nights introduces an unnecessary source of error to our measurements.




FIGURE 3.3: In contrast to Fig. 3.2 it can be seen here that Star 1 is consistent with the brighter comparison stars, therefore making it a suitable comparison star.

Additionally, the source of the disparity is unknown and we have no way to know what other effects this could have on the data that may be less obvious to the naked eye. For these reasons we chose to use Star 1 as the comparison star, as although it is less bright than Star 2, our exposure times ensure the S/N is sufficient to get a high-quality light curve. Despite Stars 3 and 4 not being usable as comparison stars, their brightness allows us to use them, as in Fig. 3.3, to demonstrate that Star 1 has sufficiently constant emission and thus is the best choice of comparison star.

3.4.2.2 Reduction Method

Although coming from several different telescopes, all the optical data arrives simply as a series of `.fits` format images. In order to extract the count rates for the objects in the images and therefore construct a light curve, we must perform some form of photometry on the images. The most suitable way to do this for our data is aperture photometry, however this process can be handled in multiple ways. For consistency between the different telescopes we should use the same photometry method for each.

Previous to this campaign, we performed reduction on data from the Gran Telescopio Canarias in La Palma taken using the HiPERCAM imaging system (Chapter 2), which has its own reduction pipeline written by Vik Dhillon and Tom Marsh ([Dhillon and Marsh](#)). This system proved itself to be very useful and intuitive to use, able to generate consistent and well tracked apertures, and given its high quality it seemed ideal to use this system for other data. Additionally, this pipeline contains a script that allows for the conversion of non-HiPERCAM data into the `.hcm` format the pipeline uses (which is simply a `.fits` file with a modified header). By default, there is only a single mode which converts Liverpool Telescope data into this `.hcm` format, but it is trivial to modify the script in order to add additional modes to convert any telescope's `.fits` data into this `.hcm` format. With these modifications, this single system allowed

fast, easy reduction of all the data for every telescope, removing the potential introduction of discrepancies from using different pipelines for different telescopes.

As we need to combine the light curves from different telescopes into one large light curve, we must ensure during the reduction process that the apertures we extract are taking light from exactly the same areas of sky. This is because to our detector stars are technically infinite point spread functions, and so the point at which the aperture should ideally be placed is subject to multiple factors, particularly how good your S/N is and how polluted the background around the star is.

After some experimentation with reductions of the Liverpool Telescope and Wise C28 data (our largest and smallest telescopes respectively) we determined a suitable aperture radius of 14 binned pixels from the Liverpool Telescope data. Using the telescope's arcsecond per pixel value of 0.3037, this translates into a 4.2518 arcsecond radius for each aperture. We similarly then used every other telescope's published arcsecond per pixel value in order to calculate the pixel radii of 4.2518 arcsecond radius apertures.

To actually perform a reduction in the HiPERCAM system requires three steps, first is the generation of the aperture file. In order to do this we first generate an average frame from a small number of frames at the beginning of the observation using the `averrun` command. This is not strictly necessary but simply makes the image clearer and makes it easier when placing the apertures. The exception to this is when using the longer exposure time data such as Wise C28, as an average over too much time can cause smearing, and the exposures were bright enough on their own.

Next, we load the target frame using the `setaper` command, choosing target and sky aperture radii, which allows placing of apertures with simple mouse pointing and single key commands. A centroiding function ensures these 'arbitrarily' placed apertures are in fact centred on the object. Then using simple mouse placement, we can mask out anisotropies in the sky aperture. These masks are locked to the centre of the target aperture ensuring they always cover the same element of the background. To ensure identical masks for each telescope, the aperture files were copied between them to preserve the mask positions. This copying still allows for loading in a new target frame, changing the aperture sizes, and re-centering of the apertures, meaning the only thing preserved is the relative mask positions.

The second step is to use the `genred` command to generate a reduction file that contains information not found in the apertures themselves such as whether the apertures are fixed or movable, and parameters around object tracking for movable apertures. Generally, for a project such as ours where trying to be consistent across all telescopes is necessary, this process only needs to be performed once, after which using the same reduction file with new apertures specified each time will suffice.

The third step is the actual reduction. The HiPERCAM pipeline has a very nice live reduction interface that generates a live light curve, x and y target aperture positions, and FWHM curves, as well as a live display of the images and positions of the apertures. These features allow us to manually perform quality checks for problems while the reduction is happening, and although the live light curve generation feature can be slightly temperamental for non-HiPERCAM data, the rest of the information is still extremely useful.

Once the reduction is complete it is saved to a log file that contains the settings used to generate it and all the information taken from each frame, including aperture counts, times, FWHMs, and aperture positions. From here a few HiPERCAM commands in Python can extract the relevant data as a 'tseries' object from the log. Initially it is practical to keep the data in this form as these tseries objects automatically propagate the light curve errors when divided by the comparison stars' tseries. Once this is complete the data can be transferred into whatever format is needed and can be used to plot the light curve.

The full light curves generated for each night in the X-ray, UV, and Optical can be found in Section 3.5.

3.4.2.3 Misalignment and CALI

Despite ensuring consistent reduction methods and ensuring identical apertures during the reduction down to 5 decimal places, some of the overlapping data points in the final two observation sets are slightly misaligned. The reason for this is difficult to determine as so many of the factors have been calibrated to be completely identical, and so it may simply be subtle atmospheric differences due to the difference in locations of the telescopes.

For the third observation set, the misalignment comes during an overlapping section of KeplerCam, LCOGT Sinestro, and LCOGT Spectral data. As the LCOGT Spectral and Sinestro data are perfectly aligned, and that there is more Sinestro data later in the observation period, it seems clear that it must be KeplerCam that is the one that is misaligned. Since this misalignment could interfere with the lag calculations, we need to find a way to try and fix it while still maintaining the integrity of our data.

A way to do this is a program called CALI [Li et al. \(2014\)](#) which inter-calibrates the separate telescopes light curves with a Monte-Carlo Markov Chain method and generates a reconstructed light curve with the overlaps correctly calibrated. This system works very well for Night 3's light curves, including filtering out some high frequency noise in the WISE data, and can be seen in Fig. 3.13.

The fourth observation set was slightly trickier however. The misalignment occurs in the centre section during an overlap with Aries, Wise C28, and Liverpool Telescope data. Here, unlike on Night 3, it appears that all three observations are out of line with one another, which makes prescribing an adjustment to ammend this issue more difficult. The saving grace on this night though is that the first section of data is also Liverpool Telescope data, and so it makes sense that the relative positions of the two datasets would have a correct relationship, which gives us an anchor for adjustments.

To support this, there appears to be a gradually increasing long term trend in both the UV and the Optical. This is supported by the difference in the positions of the first and last datasets in the Optical, and by using this it can be seen visually (Fig. 3.17) that the central Liverpool Telescope data appears to fall in line with this trend, whereas the Wise and Aries data do not.

Unfortunately, it is precisely this trend that makes the adjustment with CALI difficult, given the disparate gaps between the sections of data, CALI assumes the trend is actually caused by misalignment and 'flattens' out the whole light curve, which is visually obvious to be an incorrect adjustment. The best solution to this we found was to simply adjust the light curve beforehand, ensuring the overlapping Aries/Wise and Wise/Liverpool Telescope sections were aligned. However due to the Wise noise on these sections and noise overall, running it through CALI creates a reconstructed light curve where those 'joint' sections are one consistent light curve. This gave us the best results and can be seen in Fig. 3.17.

3.5 Results

3.5.1 Lag Calculation Methods

In order to determine the time lags between the light curves in this investigation, we have used two different methods. The first is the program Javelin written by [Zu et al. \(2013\)](#) which fits a Damped Random Walk generated continuum light curve model to each lagged light curve. The second is a FORTRAN-based Cross-Correlation Function method. These methods are described in more detail in Appendix B.

3.5.2 Javelin Results

Javelin produces a distribution of values for the input and output parameters of the process. For the most part the non-lag distributions are unimportant for us, other than as a consistency check. The most important distributions are the UV and Optical time lags. From these distributions, another pattern that appears is repeating peaks every ~ 100 seconds in the UV time lags, this is a side effect of the 100 second binning of the light curves meaning that Javelin also finds correlation a couple of data points either side of the true value. For this section then we will be plotting the time lag peaks over a time range that makes the peaks clearer to the eye, consistently binned to 4 seconds.

3.5.2.1 First Observation Set

For the first night the Javelin lag distributions can be seen in Figs. 3.4 and 3.5. As can be seen from these distributions they give lag values of $511.7^{+8.3}_{-9.7}$ and $696.4^{+3.8}_{-54.3}$ seconds for the UVM2 and Optical lags, respectively. It can be seen in the optical peak that the existence of a second smaller peak in the optical throws off the distribution error slightly, but the main peak is very strong at ~ 700 seconds.




FIGURE 3.4: UVM2 time lag distribution for the first observation period which gives a time lag of $511.7^{+8.3}_{-9.7}$ seconds.




FIGURE 3.5: Optical time lag distribution for the first observation period which gives a time lag of $696.4^{+3.8}_{-54.3}$ seconds.




FIGURE 3.6: UVM2 CCF-generated time lag distribution for the first observation period which gives a time lag of 637.9 ± 430.0 seconds.




FIGURE 3.7: The full light curves from the night of December 13th in all 3 wavebands. UV light curve has been binned up to 200 seconds to make them clearer to visually inspect.

3.5.2.2 Second Observation Set

For the second night the lag distributions can be seen in Figs. 3.8 and 3.9. These distributions give lag values of $414.8^{+57.4}_{-81.7}$ and $499.6^{+25.6}_{-36.6}$ seconds for the UVM2 and Optical lags, respectively. These distributions have peaks that are far less defined than Night 1, likely because of a lack of optical data for this night, particularly as most of it occurs during the gap in the UV, limiting the time where all 3 light curves overlap.




FIGURE 3.8: UVM2 time lag distribution for the second observation period which gives a time lag of $414.8^{+57.4}_{-81.7}$ seconds.




FIGURE 3.9: Optical time lag distribution for the second observation period which gives a time lag of $499.6^{+25.6}_{-36.6}$ seconds.




FIGURE 3.10: UVM2 CCF-generated time lag distribution for the second observation period which gives a time lag of 402.8 ± 458.1 seconds.




FIGURE 3.11: The full light curves from the night of December 19th in all 3 wavebands. UV light curve has been binned up to 200 seconds to make them clearer to visually inspect.

3.5.2.3 Third Observation Set

For the third night the g-band lag distribution only can be seen in Fig. 3.12, as there is no UV data for this night. This distribution gives an optical lag value of $828.6^{+9.7}_{-18.5}$ seconds. Due to the large amount of optical coverage on this night, the peak here is very clear despite the lack of an Ultraviolet light curve as a third comparison.




FIGURE 3.12: Optical time lag distribution for the third observation period which gives a time lag of $828.6^{+9.7}_{-18.5}$ seconds.




FIGURE 3.13: The full light curves from the night of December 31st in X-ray, optical, and CALI-applied optical.

3.5.2.4 Fourth Observation Set

For the fourth night the lag distributions can be seen in Figs. 3.14 and 3.15. These give lags of $511.0^{+7.2}_{-6.6}$ and $747.4^{+14.3}_{-9.5}$ s for the U-band and Optical lags, respectively.

Although this is our only U-band coverage, it appears out of place, given the relative positions of the UVM2 and g coverage. Possible causes will be discussed in Section 5.5.




FIGURE 3.14: U-band Javelin time lag distribution for the fourth observation period which gives a time lag of $511.0^{+7.2}_{-6.6}$ seconds.




FIGURE 3.15: Optical time lag distribution for the fourth observation period which gives a time lag of $747.4^{+14.3}_{-9.5}$ seconds.




FIGURE 3.16: U-band CCF-generated time lag distribution for the fourth observation period which gives a time lag of 596.6 ± 208.8 seconds.




FIGURE 3.17: The full light curves from the night of January 2nd in all 3 wavebands and adjustment and CALI applied to the optical. UV light curve has been binned up to 200 seconds to make them clearer to visually inspect.

3.5.3 Cross-Correlation Function Results

The Cross-Correlation Function (CCF) method was not as simple as the Javelin method in the consistency of its outputs. It has been observed before that the CCF Lag determination tends to give larger errors than Javelin does, however having an independent measurement will reassure our confidence in the Javelin lags. All of these CCF lag distributions are consistently binned to 100s.

Unfortunately, we can only have this extra confidence for a small number of our values due to several issues. The first, and primary, issue is that we were unable to determine any significant lags for any of the ground-based SDSS g-band data.

The produced CCFs, for all nights, are many thousands of seconds wide and so cannot produce a centroid with the accuracy required with our lag timescales. We suspect this is due to the sparse sampling of our optical data on most nights, and on the night where we have the most optical data, Night 3, we have no UV data. A direct X-ray to g-band correlation would be weaker than UV because of increased light curve

smoothing further out in the disc where the g-band emission originates, and there are still significant gaps in our data across all nights, and so we cannot determine a lag for even the third night where we have the most coverage.

Fortunately, however, we can determine them between the X-ray and UV bands. In addition, to evaluate the validity of the CCF distributions, we can also generate DCF confidence contours.

For the first two observation periods we can determine an X-ray to UVM2 lag distribution that can be seen in the Figs. 3.6 and 3.10.

While a less defined peak, Night 1 nevertheless remains consistent with the Javelin determined lags, with the CCF lag of 637.9 ± 430.0 seconds being consistent in its error bars with the Javelin value of $511.7^{+8.3}_{-9.7}$ seconds. While the peak is off of this value to a larger margin, and not within the Javelin error bars, this is not to be too unexpected given the width of these CCF centroid distributions and the precision of the Javelin errors.

In addition, we can see from Fig. 3.18 that the centroid peak is consistent with a Discrete Correlation Function (DCF) peak of over 99% confidence. Although this and all the DCF peaks are much wider than the CCF results themselves, these contours show that this lag seems to be indicative of a feature of the light curves outside of the background noise, likely in our case to be X-ray reverberation.

The second night's distribution, in Fig. 3.10, has the cleanest distribution and indeed the confidence of this peak is backed up by the fact that the central value of the determined peak, 402.803 seconds, is almost perfectly aligned with the Javelin determined value of $414.8^{+57.4}_{-81.7}$ seconds, despite the fact that the peaks width gives the value an error range of ± 458.1 . This therefore increases our confidence that the UVM2 lag measurement on Night 2 is correct.

In this case the Javelin is far more informative than the confidence contours, which can be seen in Fig. 3.19, as for Night 2 the contours simply show a consistent, though far wider, peak in the DCF in the same range as our detection peak. While this does not gain us further precision, it is also consistent with our other measurements, further ensuring confidence in our Javelin and CCF values.

As previously mentioned, the third observation period has no UV data and only the lag distribution for Night 4 is shown now, which is an X-ray to U-band lag, shown in Fig. 3.16. The value determined from Night 4's distribution is 596.6 ± 208.8 seconds which is consistent with the Javelin measurement though with an obviously much wider peak. However this value required the CCF to be performed on only the first two thirds of the light curve, as the X-ray flare in the last third was causing the CCF to detect a lag value ~ 1000 seconds. The fact that the Javelin lag is consistent with the




FIGURE 3.18: X-ray to UVM2 DCF plotted with Confidence Contours for the first observation period.




FIGURE 3.19: X-ray to UVM2 DCF plotted with Confidence Contours for the second observation period.




FIGURE 3.20: X-ray to U-band DCF plotted with Confidence Contours for the fourth observation period.

CCF when this flare section is removed suggests then that the flare was interfering with the CCF lag detection. The DCF does not, however, have as strong a peak as it does for the other two nights, with the peak barely making the 90% contour, as can be seen in Fig. 3.20.

3.6 Analysis

3.6.1 Lag Analysis

3.6.1.1 Optical Lags

For all but one night, we have a very well determined g-band time lag, with the values from nights 1, 3, and 4 of $696.4^{+3.8}_{-54.3}$, $828.6^{+9.7}_{-18.5}$, and $747.4^{+84.3}_{-97.5}$ seconds, respectively, all overlapping with at least one of the other's error bars. It seems that what is preventing the two values that are farthest apart from overlapping with one another is simply that Javelin's errors, determined directly from the distribution peaks, determine errors in precision only, and do not seem to factor in errors in accuracy. Therefore, it overvalues a narrow distribution and assumes when there is a narrow distribution that the value is very accurate when instead it might simply be very precise.

The outlier of the g-band lags is the value determined on Night 2 of $499.6^{+25.6}_{-36.6}$ seconds. This value does not come close to being consistent with the other 3, and for reasons we will explore now, it is very unlikely to be physically real. We can compare the value to Night 1, taken 6 days earlier. Being generous and extending the values to the edge of their respective error ranges, we find 642.115 seconds for the first night and 525.232 seconds for the second night. If we assume both these values are representative of the physical state of the AGN, that would require the thermal band to have shifted 116.883 light-seconds further away from the AGN, which would require extra input energy from the Black Hole, and therefore an increase in X-ray flux overall. However if we look at the mean X-ray flux for each night, we see that in fact Night 2 has decreased by $\sim 45\%$ and so this instead suggests that the X-ray flux has been on a decreasing trend between the two nights.

This is likely an unphysical situation and so it seems one of the results is not physically representative of the real state of the AGN. As Night 1's value is far more consistent with the rest of the data, this leads us to conclude that it is Night 2's g-band lag that is at fault. We can see there is not much g-band data on that night, and for a large portion of the g-band exposure there was a gap in the OM, further reducing its reliability.

Comparing these to the previous observations taken in [McHardy et al. \(2016\)](#) which determined a g-band lag of 788^{+44}_{-54} seconds, we can see that this value is consistent

with two of our three determined values, with Night 1's value having a very small positive error bar, possibly due to the precision/accuracy argument made before.

3.6.1.2 UV Lags

Unlike for the optical, for the OM monitoring we do not have four sets of data taken at the same wavelength, therefore we shall analyse the data separately.

To begin with we will look at our two observations in the UVM2 filter, these found lags of $511.6^{+8.3}_{-9.7}$ and $414.8^{+57.4}_{-81.7}$ seconds for Nights 1 and 2, respectively. Now initially it appears that as these values do not overlap in their error bars, that they are not consistent.

However, upon looking at the original peaks that Javelin generated it can be seen that there are secondary peaks generated, seemingly related to the binning of the light curve, as they are spaced ~ 100 seconds apart and the light curves are binned to 100 seconds. We believe whatever feature caused these extra peaks has "stolen" the outer values of the real peak, i.e. a run that would have produced an outlier on the main peak is instead close enough to be detected as one of the secondary peaks. This would have the effect of narrowing the main peak, thus reducing the error on the lag and meaning the error bars don't overlap despite the values both likely being physically consistent.

The Cross-Correlation lags for Nights 1 and 2 are also consistent with both of the Javelin lags, measuring a lag of 637.9 ± 430.0 seconds for the Night 1 and 402.8 ± 458.1 seconds for the Night 2. The large error bars on these values don't make them particularly useful for analysis by themselves, but instead they work as an independent consistency check for the Javelin lags.

Without any data for our UVW1 filter, we only have the Night 4 data in SDSS U to look at. Upon initial inspection, the value obtained by Javelin of $511.0^{+7.2}_{-6.6}$ seconds, it is hard to immediately determine whether this is valid due to a lack of other observations; though there are things we can glean from this value.

The first is that it is remarkably close to the UVM2 lags, which should not be the case as the UVM2 filter has a central wavelength of $\sim 2260\text{\AA}$, whereas U-band has a central wavelength of $\sim 3470\text{\AA}$. It should therefore have a lag somewhere in the middle of SDSS g and UVM2, rather than a smaller lag.

The second is that the Cross-Correlation lag that was determined for this observation set is 596.6 ± 208.8 seconds. This is mostly consistent with the Javelin lag, though the low levels of confidence that were determined from the synthetic light curve DCF Contours imply this value is not produced from a strong correlation. Even so, a future

campaign observing more in this wavelength range would be useful to discover the presence, or lack thereof a U-band excess.

3.6.2 Modelling

3.6.2.1 Shakura & Sunyaev

A final aspect we looked at is how these lags fit against lag distributions that can be generated from modelling. The first model we looked at is a simple Shakura & Sunyaev model of $\tau \propto \lambda^{4/3}$, and so in Figs. 3.21 and 3.22 we have plotted the retrieved Javelin lags in seconds against the central wavelength of the filters, taken from Fig. 3.23. Overlaid we, using `scipy.optimize's curve_fit`, fit a function of the form:

$$y = k\lambda^{4/3} + a$$

where k and a are simply input variables, to see if the values we have produced can be described by such a model. Along with this are plotted lines showing the error variation from the fit in each variable, in Fig. 3.21 we show the k error, and in Fig. 3.22 the a error.




FIGURE 3.21: Javelin-produced Time Lags plotted against the filter central wavelength, fit with a function of the form $\tau = k\lambda^{4/3} + a$, showing variation in k around the error. Central wavelengths of the 3 filters are 226.0nm, 347.0nm, and 477.8nm, respectively.

Plotted on these graphs in yellow is the unreliable U-band value, and the g-band value from Night 2 that appears to be unphysical. It becomes visually apparent from




FIGURE 3.22: Javelin-produced Time Lags plotted against the filter central wavelength, fit with a function of the form $\tau = k\lambda^{4/3} + a$, showing variation in a around the error. Central wavelengths of the 3 filters are 226.0nm, 347.0nm, and 477.8nm, respectively.

these plots that the Night 2 g-band is an outlier as it does not fit well inside the function that appears to fit the rest of the points nicely, taking into account the narrow errors on Night 1's UVM2 observations.




FIGURE 3.23: The filter transmission curves for the SDSS filters made using the Ultracam python module `ucam_thruput`, from which the central wavelengths were calculated.

Interestingly the U-band point seems to fit well into this function, which suggested it is not as much of an outlier as it initially appeared. Further observations at this wavelength could give more insight into this u-band value.

These fits show that the reliable results we have collected in this campaign are in fact consistent with a function that follows the form $\tau \propto \lambda^{4/3}$.

However, it can be seen that these model fits rely on a non-zero y-intercept, which has been observed in higher accretion rate AGN as in [Gardner and Done \(2017\)](#) and is considered to be a feature representative of an inflated inner disc that absorbs the corona's X-ray radiation and re-emits it in the Far UV. This Far UV is then what the rest of the disc is reprocessing, with the y-intercept, i.e. the value of a , representing the time it takes for the X-rays to scatter inside the inflated disc region before being re-emitted as FUV radiation.

The aspect of this that is surprising is that previous observations of NGC 4395 in [McHardy et al. \(2016\)](#) found X-ray to UVW1 and X-ray to g-band lags that were consistent with a $\tau \propto \lambda^{4/3}$ function with a zero intercept. However, as can be seen from our data in Fig. 3.24, the results we found are not consistent with this, and instead the fit suggests an offset from the X-ray in the range of ~ 200 to ~ 400 seconds, implying at least some inflation of the disc.




FIGURE 3.24: Javelin-produced Time Lags plotted against the filter central wavelength, fit with a function of the form $\tau = k\lambda^{4/3} + a$, with an intercept forced through zero. Central wavelengths of the 3 filters are 226.0nm, 347.0nm, and 477.8nm, respectively.

3.6.2.2 KYNreverb Response Lags

As a Shakura-Sunyaev model seemed to not include all the nuance required to explain the lags we observe, we instead turned to a modelling code called KYNreverb, developed by Michal Dovčiak ([Dovčiak](#)).

The version we used here is the KYNxilt code. This simulates a Kerr black hole with a geometrically-thin, optically-thick accretion disc with a radial density profile and a lamp-post model X-ray corona above the black hole, emitting isotropic radiation in the typical power-law distribution observed for the X-ray components of AGN. From here it generates an X-ray pulse and then uses fully general-relativistic ray-tracing to map out the photon paths from the corona to the observer, including those photons which are reprocessed by the disc. This generates response functions from the object for any waveband, and by choosing several over a range of wavebands, it is possible to build up responses for that range.

As explained in [Kammoun et al. \(2021a\)](#), a commonly accepted method of measuring the time lag from a response function is to use the centroid of the response function $\Psi(t, \lambda_c)$, as this is most representative of the peak of a CCF-determined lag. This centroid is defined as:

$$\tau_{cen}(\lambda_c) = \frac{\int t \Psi(t, \lambda_c) dt}{\int \Psi(t, \lambda_c) dt}$$

where $\tau_{cen}(\lambda_c)$ is the time lag for any given central wavelength λ_c . We then plot this lag spectrum and compare it to our data to see how well the model fits.

However this model is very complex and has around 40 parameters, so the exact parameters can very much affect the output. A lot of these parameters are standard or observed values; some do not change the output hugely or have a fairly well constrained value so varying along these values is not always time best spent. Therefore there are few values we used as standard in this model for NGC 4395.

First we set the spin to be minimised and the corona lamp-post height to be 10 gravitational radii, since these values do not vary the responses significantly within our error bars when varied along values constrained by observations. Similarly the flux in the 2-10 keV band at distance of 3.85 Mpc of approximately 3.7×10^{-11} ergs s^{-1} , calculated using XMM PN observations in [Vaughan et al. \(2005\)](#), is relatively well constrained for NGC 4395 and only a very large deviation from this significantly changes the output.

Far more significant are four parameters, those of which end up being at issue: the mass of the central black hole, M_{BH} , the outer radius of the disc, also known as the cutoff radius R_{out} , the accretion rate \dot{m} , and the colour correction factor f_{col} . For the accretion rate we also know that for a constant luminosity:

$$M_{BH} \propto \dot{m}^{-1}$$

Therefore, if we vary the mass we must also transform the accretion rate in the inverse manner.

The initial values for M_{BH} and \dot{m} are the values taken from [Peterson et al. \(2005\)](#) of $M_{BH} = 3.6 \times 10^5 M_\odot$ and $\dot{m} = 1.2 \times 10^{-3} \dot{m}_e$ where \dot{m}_e is the Eddington accretion rate.

We then took our initial guesses at R_{out} and f_{col} from Chapter 2 where the data fit well with $R_{out} = 1,600 R_g$ and a colour correction factor modelled using the temperature (and therefore radius) dependent method from [Done et al. \(2012\)](#), chosen in the simulation code by setting $f_{col} = -1$.

Visible in Fig. 3.25, the initial input values fit our data fairly well, which is encouraging since we are measuring the same object. We see that the U-band lag again appears to be smaller than would otherwise expected, and unsurprisingly the g-band lag from Night 2 is far below expected.




FIGURE 3.25: Simulated lag spectra generated from KYNxilt, showing various R_{out} values using the [Done et al. \(2012\)](#) model f_{col} .

It can be seen from these spectra that unlike with a simple power-law, in this model the outer cutoff of the disc begins to dampen the gradient of the lag spectrum, with a smaller disc lowering the maximum possible lag for any given frequency. We see this effect as increasing or decreasing R_{out} increases or decreases the optical lags while leaving the UV lags relatively unaffected.

The rest of the data appears to be mostly consistent with these lag spectra, although perhaps our g-band data is more consistent with a slightly smaller R_{out} of $\sim 1,400 R_g$. The additional fact that KYNreverber is a much more detailed simulation than a simple power-law implies that the smaller U-band lag cannot be accounted for in a simple

geometric disc model, and therefore likely has its origin in a more complex disc geometry.

3.6.2.3 KYNreverb Response Function Convolution

Another way we can measure lags using these response functions is to create model UV and Optical light curves from the X-ray data. This can be done by convolving the X-ray data with KYNreverb response functions from each waveband. This generates model light curves that are what the UV/Optical light curves would look like if they were purely composed of reprocessed X-ray flux under the conditions of the KYNreverb simulation. If we then retrieve lags from these light curves we have a direct comparison of the X-ray to real UV/Optical light curve lags and the X-ray to model UV/Optical light curve lags.

However, if we do this measurement with the $1400R_g$ response functions we compared before to the data, we find the lags displayed in Fig. 3.26.




FIGURE 3.26: X-ray to Convolved Model lags using the $1400R_g$ response functions plotted with the data lags, alongside centroid and half-light point lag spectra. The error bars have been moved off the data points for clarity.

As can be seen from this plot, with the exception of the Night 2 g-band lag which is dubious anyway, these convolution lags are inconsistent with the data. This initially seems puzzling as they were made from response functions whose lag spectrum appears to align very well with our data. However upon further investigation, we realised that the method we previously used to determine the expected lag from a response function, retrieving the centroid, is not the only method to obtain such a lag.

Another way to measure an expected lag is to measure the half-light point lag, defined as the point at which half of the total response functions light is received and given by:

$$\tau_{hlp}(\lambda_c) = \frac{1}{2} \int_0^\infty \Psi(t, \lambda_c) dt$$

When we create a lag spectrum instead using $\tau_{hlp}(\lambda_c)$ we see that these lags are about 30% shorter than the centroid lags and are consistent with our convolution lags.

By modelling these response functions we are defining them as the functions that, when convolved with our X-ray light curves, give a certain lag. Given then that actually applying that convolution and measuring the lags retrieves values that are consistent with the half-light point lag spectrum and inconsistent with the centroid lag spectrum, it seems that we must use the half-light point lags instead.

This then presents a problem, as the parameters we used in creating these responses no longer give us a lag spectrum consistent with the data, and so we must change them. The UV lags are most affected by the mass, accretion rate, and colour correction factor, while the optical lags are affected by the outer radius as well. The outer radius is something we are mostly fitting to here, and so this is our most free parameter. However the effects that change to the mass, accretion rate, and colour correction factor have on the spectrum are degenerate with one another which makes it very difficult to determine a correct outcome.




FIGURE 3.27: $1900R_g$ KYNxltr half-light point lag spectra, varying fixed f_{col} values and showing the Done Model for comparison.

The mass as determined by [Peterson et al. \(2005\)](#) is $3.6 \times 10^5 M_\odot \pm 1.1 \times 10^5 M_\odot$. Modelling a response function with an identical accretion rate (accounting for the

mass change) and colour correction factor but a free R_{out} , we found that a mass of $\sim 6 \times 10^5 M_\odot$ is needed to fit the data. This is quite far outside the error bars of Peterson's mass and so unlikely. Similarly fixing the mass and colour correction factor but allowing accretion rate to vary, we find that in order to be consistent with the data we would require an accretion rate $\sim 4 \times$ larger, which is a very large increase for such a low accretion rate object. Finally if we fix the mass and accretion rate but change the colour correction factor, we are no longer using the model from [Done et al. \(2012\)](#) and instead using a fixed f_{col} over the entire disc. This best fits our data with an $f_{col} = 1.5$ and an $R_{out} = 1900R_g$ as can be seen in Fig. 3.27.

In reality it is likely that all three of these values are slightly changed to account for this change in the lags, however as the mass and accretion rate are far more well constrained than the colour correction factor, the choice we decided to make here is to keep the mass and accretion rates at their currently accepted values use the fixed colour factor instead. All further convolutions and lag spectra were then be made with these values.

The convolved light curves can be seen in Fig. 3.29.

Looking at the model light curves we can see that they align very well with the UV data, though there are some features, such as the large peak in Night 1 and the peaks at around 30ks in Night 3, which do not seem to be present in the real UV data.




FIGURE 3.28: X-ray to Convolved Model lags using the $1900R_g$ response functions plotted with the data lags, alongside the half-light point lag spectra.

The reasons for this are unclear, though as the disc does not appear to respond to this flare at all, its possible that the X-ray flare we measure is somehow obscured from the disc, perhaps by inflation of the inner edge of the disc or a change in the height of the




FIGURE 3.29: Each night's UV/Optical data plotted with the model light curves made from the X-ray data convolved with the *kynxltr* response functions.

X-ray corona. The SDSS g seems to diverge more from the model, though in the case of Night 2, we once again disregard it because of the lack of data.

Using these model light curves created from our new responses, we can then plot the model lags versus the data again in Fig. 3.28. It can now be seen that the lags retrieved from the convolved model light curves are much more consistent with the data. In fact when comparing the data lags and the model lags, we can see that they are consistent to within 1σ of one another when excluding the outlying Night 2 g-band lag.

We also note that the model U-band lag appears to be smaller than the lag spectrum would suggest, which could hint to a feature of the light curve causing a lag decrease in both the data and the model. However when we look at the other model lags, it is clear there is a distribution in the measured model lags around the half-light point lags and therefore it is inconclusive whether the smaller model light curve is representative of a real feature or simply a random scatter of the measured lag from the half-light point lag spectrum.

3.7 Conclusions

In this report we presented extended multi-wavelength observations of the low-accretion rate and low-mass NGC 4395 using *XMM-Newton*'s EPIC camera for X-ray data, its Optical Monitor system for UV data, as well as six ground-based telescopes for g-band data. We observed over 4 total periods, each one giving us ~ 117 ks of *XMM-Newton* data and some simultaneous ground-based data when available. Using these systems allowed us high enough time resolution data to extract useful time lags even for such a low-mass AGN. For our OM Ultraviolet observations three different filters were chosen to expand on the observations made only in UVW1 in [McHardy et al. \(2016\)](#), with UVM2 being observed for the first two observations, UVW1 being observed for the third observation, and U-band being observed for the final observation. Due to a technical fault however, on the third night the OM was not pointed correctly and so we lost all potential UVW1 data. With the remaining data we used Javelin and a Flux Randomisation/Random Subset Selection Cross-Correlation Function program to extract X-ray to UVM2, X-ray to U-band, and X-ray to SDSS g-band time lags for each night these filters were observed. We unfortunately faced problems retrieving g-band lags using the FR/RSS CCF program, possibly due to the large gaps in the g-band data where constant observation was not possible due to the limits of ground observation.

We were able to determine lags from X-ray to UVM2 of $511.7^{+8.3}_{-9.7}$ and $414.8^{+57.4}_{-81.7}$ seconds from Javelin, with the non-overlap in error bars being attributed to the small error bars on the first value, due to Javelin's overconfidence on the accuracy-error and

due to the small precision-error. To increase our confidence in these we were able to determine FR/RSS CCF lags for those two nights of 637.9 ± 430.0 and 402.8 ± 458.1 seconds, overlapping with the more precise Javelin lags to reinforce their legitimacy. For the U-band lag, we obtained a value of $511.0^{+7.8}_{-91.3}$ seconds from Javelin which both does not line up with the Shakura & Sunyaev disc model, nor does it exhibit the excess lag expected from emission originating in the Broad Line Region. The FR/RSS CCF lag of 596.6 ± 208.8 seconds meanwhile is consistent with the Javelin lag. For the ground-based g-band lags we were able to determine confident lags from the first, third, and fourth observations of $696.4^{+3.8}_{-54.3}$, $828.6^{+9.7}_{-18.5}$, and $747.4^{+84.3}_{-97.5}$ seconds, respectively. The second night's lag value of $499.6^{+25.6}_{-36.6}$ seconds is not consistent with these three values, however this night's lack of g-band data, along with most of the g-band data occurring during a gap in the UV, means this value is likely not physical.

Plotting these values together as seen in Fig. 3.21 shows that the data is consistent with a function of the form $\tau \propto \lambda^{4/3}$ in accordance with the simple disc model proposed in [Shakura and Sunyaev \(1973\)](#). However this consistency demonstrates a likely offset from the origin, implying at least some amount of X-ray scattering occurring in an inflated inner disc component on the scale of several hundred seconds. This has been observed before in higher-mass and higher-accretion rate AGN such as in [Gardner and Done \(2017\)](#) but not in NGC 4395. In fact this campaign's results disagree with the previous observations made of NGC 4395 in [McHardy et al. \(2016\)](#) which showed lags consistent with zero offset and therefore, no inflated disc component.

The set of our observations in Chapter 2, that used the HiPERCAM system on the Gran Telescopio Canarias, have produced some preliminary lags relative to its SDSS u' light curve for the SDSS bands g', r', i', and z'. While OM's U-band filter is not identical to the SDSS u' band used by Hipercam, they are extremely close in transmission width, shape, and are within 5 nm of each other's central wavelengths. Within the scope of our error bars this difference is negligible.

For this reason these preliminary lags have been plotted alongside this report's lags in Fig. 3.30, with the HiPERCAM lags plotted relative to the X-ray to OM U-band lag. For the most part these can be seen to be consistent with the lags observed here, although the r-band lag is quite high relative to the fit and there overall seems to be a flattening trend in the lags that is not represented in the fit.

We can also use these HiPERCAM lags to explore the longer wavelength domain for the KYNreverb simulations in order to check if the response function lags remain consistent with the data. We do see that the turnover caused by the edge of the disc continues out to the infrared, and that the HiPERCAM data shows that our cutoff radius does indeed need to be $\sim 1,900R_g$ for our set of other parameters.

Of the two models we explored, it appears that the KYNreverb simulation represents more accurate physical features of the data, which is probably to be expected since

Shakura & Sunyaev was a very simplified toy model and this was a full GR ray-tracing simulation. But it is interesting to note that even the KYNreverb curve shows a shape similar to a power-law in the initial far UV, but it appears that the finite size of the disc causes it to turn off of this function as it nears the disc edge.




FIGURE 3.30: Javelin-produced Time Lags plotted against the filter central wavelength, fit with a function of the form $\tau = k\lambda^{4/3} + a$, showing variation in both k and a around the error and plotted with the GTC HiPERCAM lags. Central wavelengths of the 6 filters are 2260Å, 3470Å, 4778Å, 6201Å, 7640Å, and 9065Å.




FIGURE 3.31: Simulated lag spectra generated from KYNxiltr for $R_{out} = 1900R_g$ and $f_{col} = 1.5$, compared to the lags determined from this campaign and those from the GTC HiPERCAM observations.

Chapter 4

Multi-waveband Power Spectral Densities from observations of NGC 4395

4.1 Abstract

In this chapter I will explore another avenue of variability by measuring and fitting the optical power spectra of NGC 4395 to support the previously measured X-ray power spectra. I will also be searching for additional evidence of X-ray reprocessing to complement the lag results from Chapter 3.

4.2 Introduction

Active Galactic Nuclei (AGN) have long been observed to emit X-ray continua that vary rapidly and non-periodically. This has supported the idea that the X-ray emission originates from a compact region (within a few Schwarzschild Radii) around the central black hole. The X-ray emission of these objects has long been observed to be approximately described by what is called a "red-noise" power spectrum which is steep at high frequencies (with a slope ~ 2) and is flatter (a slope < 1) below some bend or break frequency ν_B (Lawrence et al., 1987; McHardy, 1989; Edelson and Nandra, 1999; Nandra and Papadakis, 2001; Uttley et al., 2002).

The location of this bend frequency is believed to scale with Black Hole Mass, as studies performed on X-ray binaries find the bend at proportionally higher frequencies than are found in AGN (McClintock and Remillard, 2003). A bend to a slope of less than 1 is required at some point in theory to prevent the variance of the

light curve diverging to infinity, so a bend is a reasonable feature to be searching for. In addition, NGC 4395 is a very low mass AGN ($M_{BH} \sim 3.6 \times 10^6 M_{\odot}$ Peterson et al. (2005)) which would imply a relatively high bend frequency.

Previous measurements of NGC 4395's X-ray power spectrum in Vaughan et al. (2005) found it consistent with breaking power law with a high frequency slope, low frequency slope, and bend frequency of ~ 2 , ~ 1 , and $\sim 1 \times 10^{-3}$, respectively. This suggests that NGC 4395 is typically of its type, if only unusual in its very small mass.

As our previous observations of NGC 4395 in Chapters 2 & 3 have shown evidence from X-ray to optical lags that the optical emission of NGC 4395 has a component that is consistent with lagged X-ray emission, which implies the existence of X-ray reverberation from the disc in the production of the optical flux.

If then the optical flux is reverberated from the X-ray flux, the features of the X-ray emission should be imprinted on the optical emission. Our previous investigations have tested this connection in the time domain, and so in this study we test the connection in the Fourier domain in order to more completely understand the connection between the coronal X-ray emission and the disc-based optical emission.

4.3 Observations

In this paper we use four separate ~ 117 ks X-ray observations from the *XMM-Newton* EPIC pn taken on the nights of the 13th, 19th, and 31st of December 2018, as well as the fourth that was taken on the 2nd of January 2019. The fourth nights observation is plotted in Fig. 4.1, binned up to 50s for clarity, to show what these light curves look like. The rest of the X-ray light curves are plotted in Chapter 3.




FIGURE 4.1: *XMM-Newton* 0.5-10keV light curve for NGC 4395 using data from the night of January 2nd 2019, binned up to 50s for visual clarity.

For UV/Optical observations, we present two light curves from *Swift* UVOT in the B and V bands, taken from [Cameron et al. \(2012\)](#), a light curve using the Transiting Exoplanet Survey Satellite (TESS) from [Burke et al. \(2020\)](#), and an SDSS g' band light curve from the Zwicky Transient Facility using the ALeRCE ZTF explorer pipeline ([Masci et al., 2018](#)).

We also present two new long term light curves in the SDSS g' band from both the Liverpool Telescope’s IO:O camera ([Steele et al., 2004](#)), the Las Cumbres Observatory (LCOGT) McDonald Observatory’s Sinestro and Spectral cameras ([Brown et al., 2013](#)), and the Zowada Observatory ([Carr et al., 2022](#)). These are supported by very a high cadence 8 ks light curve taken with HiPERCAM ([Dhillon and Marsh](#)) on the Gran Telescopio Canarias (GTC)¹.




FIGURE 4.2: *Swift* UVOT light curve for NGC 4395 using data from [Cameron et al. \(2012\)](#).

For the [Cameron et al. \(2012\)](#) UVOT data, the observations consisted of two halves, the first of which shows an outburst for almost its entire length and the second of which appears to show more standard variability. As we are interested here in linking the long and short timescale variability, using an isolated outburst would not be representative of the usual variability we see at that timescale. As such going forward will be using only the second halves of the UVOT B and V light curves, these light curves can be seen in Fig. 4.2.

As TESS’ light curve is a white filter, we modified its mean to match the Fractional Variation seen in Liverpool Telescope’s data, as well as removing some outlier points. The resultant light curve can be seen in Fig. 4.3.

In Fig. 4.4 the combined g' -band long-term light curves from Liverpool Telescope IO:O, Las Cumbres Observatory’s Sinestro & Spectral, ZTF Palomar P48, and Zowada

¹<http://www.gtc.iac.es/>




FIGURE 4.3: TESS light curve for NGC 4395 using data from [Burke et al. \(2020\)](#) plotted with its original flux and the one adjusted for LT's Fractional Variation.

Observatory are plotted together. Due to seasonal breaks, the light curves are divided into 3 'epochs' which will be referred to as such when it is relevant.




FIGURE 4.4: Long term SDSS g' band light curve for NGC 4395 comprised of data from Liverpool Telescope, Las Cumbres Observatory, Zowada Observatory, and the Zwicky Transient Facility.

In Fig. 4.5 we show the g' -band light curve taken using HiPERCAM from the Gran Telescopio Canarias (GTC), previously shown in Chapter 2. Due to effects described in that paper with the telescope tracking later in the light curve, here we only use the first 8 ks so these effects do not interfere with our PSD determination. We are mostly interested in the very high frequencies with this light curve, so the shortening should not have much of an effect on the PSD we retrieve.




FIGURE 4.5: GTC HiPERCAM SDSS g' band light curve.

4.4 Methodology

4.4.1 PSD Fitting

In order to create the power spectra from the light curves and fit them to a model, we use PSRESP (Power Spectral RESponse) developed by Philip Uttley in [Uttley et al. \(2002\)](#), based on the 'response method' described in [Done et al. \(1992\)](#).

This is necessary, as opposed to simply fitting the raw PSD naively to a power-law model, mainly due to two effects. The first is red-noise leak which is where variability on time scales longer than the length of the observation period leaks power into the PSD, increasing the uncertainty in the shape of the measured power spectrum.

Secondly, as our light curves are sampled non-continuously for limited periods, the data point values are only representative of the flux during the period of observation, but are used as values for an entire 'bin'. As it is unlikely in an object that varies as fast as NGC 4395 for the flux to remain steady over this entire period (as indeed we see in our high time-resolution light curves), this means that the flux values of our observations will most likely not be representative of the true average across the entire bin.

What this means in practice is that for a discretely sampled light curve with time intervals $\Delta\tau$, we can only measure the power spectrum up to the Nyquist frequency $\nu_{\text{Nyquist}} = \frac{1}{2\Delta\tau}$ and any extra power that exists at higher frequencies is redistributed to lower frequencies. This effect is called 'aliasing'.

To confront these issues PSRESP uses a Monte Carlo method to predict the shape and uncertainties in the PSD. To do this it first simulates continuous light curves with a known true PSD shape and then applies the sampling pattern of the data to retrieve the distorted PSDs. It then averages them to retrieve the mean shape of the distorted

model PSD. The spread of the simulated PSDs around the mean can then be used to determine the errors on the observed PSD.

This model average PSD $\bar{P}_{\text{sim}}(\nu)$ and the 'RMS error' on the power at that frequency $\Delta\bar{P}_{\text{sim}}(\nu)$, defined here as the RMS spread of model power around the mean, can be used to determine a statistic χ^2_{dist} such that

$$\chi^2_{\text{dist}} = \sum_{\nu_{\text{min}}}^{\nu_{\text{max}}} \frac{(\bar{P}_{\text{sim}}(\nu) - P_{\text{obs}}(\nu))^2}{\Delta\bar{P}_{\text{sim}}(\nu)^2} \quad (4.1)$$

where ν_{min} and ν_{max} are the minimum and maximum frequencies measured by the observed PSD $P_{\text{obs}}(\nu)$.

With this statistic it can then estimate a goodness of fit. To do this it first re-normalises $\bar{P}_{\text{sim}}(\nu)$ to minimise χ^2_{dist} . With this minimum χ^2_{dist} , it then randomly samples a number of the simulated PSDs (using all of them would be too computationally intensive) to find an accurate estimate of χ^2_{dist} 's distribution and then sorts them into ascending order. The probability that the model can be rejected is then given by the percentile of the simulated χ^2_{dist} distribution above which χ^2_{dist} exceeds that measured for the observed PSD. Going forward a reference to the 'model probability' (i.e. "the best-fit model probability for these parameters is 0.57") will be referring to this value.

Comparing the confidence levels obtained from this method as we vary the input model parameters then can give us the model whose parameters give us the highest confidence, i.e. the best fitting model. The complete description of this method can be found in Uttley et al. (2002).

For the data presented here, we will be using three particular models in our attempts to fit the data.

The first is a simple unbroken power law, given by:

$$P(\nu) = A\nu^{-\alpha} + C \quad (4.2)$$

where A is the normalisation, α is the slope, and C is the Poisson noise level of the PSD.

The second is a bending power law, given by:

$$P(\nu) = A\nu^{-\alpha_L} \left(1 + \left(\frac{\nu}{\nu_B} \right)^{\alpha_H - \alpha_L} \right)^{-1} + C \quad (4.3)$$

where ν_B is the bend frequency, and α_H and α_L are the high-frequency and low-frequency slopes, respectively. As in the regimes of $\nu \ll \nu_B$ and $\nu \gg \nu_B$ the equation simplifies so that it describes a power law with a slope of α_L and α_H , respectively.

The third and final model is a double-bending power law, which following similar logic to the single-bend power law is given by:

$$P(\nu) = A\nu^{-\alpha_L} \gamma^{-1} + C \quad (4.4)$$

with

$$\gamma = \left[\left(1 + \left(\frac{\nu}{\nu_{B1}} \right)^{\alpha_M - \alpha_L} \right) \left(1 + \left(\frac{\nu}{\nu_{B2}} \right)^{\alpha_H - \alpha_M} \right) \right]$$

where α_H , α_M , and α_L are the high-frequency, mid-frequency, and low-frequency slopes. ν_{B1} is the bend frequency between the low and mid-frequency slopes, and ν_{B2} is the bend frequency between the mid and high-frequency slopes. In Section 5.5 we will look at how appropriate each of these models is for our optical power spectra.

4.4.2 Synthetic Optical Light Curves

In addition to fitting the power spectra of the observed X-ray and Optical light curves we are investigating, we are trying to determine whether the optical data can be described as a consequence of the reverberation of the X-ray flux. Therefore it is useful to see if we can reproduce the observed optical power spectrum from our actual X-ray data.

In order to do this we need to do two things. First we need to simulate long-term X-ray light curves from our ~ 140 ks observations in order to probe the low frequency domain, and second we need to convert these simulated X-ray light curves into simulated Optical light curves.

4.4.2.1 Simulating Long-Term X-ray Light Curves

To create long-term X-ray light curves from our data, we use a Python simulation created by Connolly (2015) using the method from Emmanoulopoulos et al. (2013). The reason for the use of this method is that visual inspection of the X-ray light curves of NGC 4395 shows they are non-Gaussian, but previous methods of light curve simulation such as Timmer and König (1995) produce only Gaussian light curves.

This method takes in the observed light curve and initially fits its PSD to a single-bending power law model. It then fits the PDF (Probability Density Function) of the data fluxes to a combined lognormal and gamma distribution, a necessary step as objects like AGN that exhibit ‘burst-like’ behaviour in their X-ray light curves have decidedly non-Gaussian PDFs, typically ones with long tails.

With these two fits the light curves this method produces reproduces both the exact variability properties of the data, as the synthetic light curves follow the input power spectrum, and the exact statistical properties of the data since it uses its PDF. A far more comprehensive description of the algorithm this method uses is available in Emmanoulopoulos et al. (2013).

As these light curves are generated from the PSD and PDF fits, there is no limit on their length, and so once the code fits these models to the X-ray data, we can simulate arbitrarily long light curves. In this case for computational reasons we generate 100 day long, 10 second binned X-ray light curves for each night. An example of a simulated X-ray light curve is found in Fig. 4.6, this shows a 140ks sample of the full light curve to show how it replicates the variability and statistical features of the real data.




FIGURE 4.6: A 140ks sample of the simulated X-ray light curve using data from the night of January 2nd 2019, binned up to 50s for visual clarity.

4.4.2.2 Model Optical Light Curves from X-ray Simulations

Once we have these simulated X-ray light curves we must transform them into simulated Optical light curves. To do this we use a similar method to Chapter 3 where we take a model response function from *KYNreverberb* (Kammoun et al., 2021a) and convolve them with the simulated X-ray light curves. This will then generate a model light curve comprised of X-ray flux reverberated by an accretion disc as modelled by the code.

As our optical data is either g' -band (or calibrated to be), we use a g' -band response function using the same input parameters as found in Chapter 3, which is to say zero spin, $10R_g$ corona height, a colour correction factor, f_{col} , of 1.5, and an outer radius of $1500R_g$. This synthetic g' -band light curve generated from the X-ray simulation is found in Fig. 4.7, showing the same 140ks sample as Fig. 4.6.




FIGURE 4.7: A 140ks sample of the synthetic g' -band light curve created from the simulated X-ray LC in Fig. 4.6, binned up to 50s for visual clarity.

4.5 Results & Analysis

4.5.1 Optical Data Power Spectra

When creating our combined optical PSD, we have our different light curves to cover different frequency ranges. The HiPERCAM data covers the highest frequencies, from a timescale of 15s down to a frequency of about $10^{-3.5}$ Hz. TESS covers the next frequency range from about $10^{-3.5}$ to 10^{-6} Hz. In the next decade, we have the UVOT light curves to cover this range, however we note that our individual epochs for our combined LT/LCOGT/ZTF/Zowada light curves also cover this range, and so we add them as individual epochs as well to cover this range more extensively.

For the longest timescales we of course have our entire 3 epoch light curve, however due to the large seasonal gaps it is not as simple to implement. So to cover the long timescales we set the simulated sampling high to properly sample the data (0.3 days) but have PSRESP then bin it up to 80 days so we get about 2 bins per season and hence the sub-season timescales are above the Nyquist frequency, but all the high frequency variability is still accounted for by the simulation. This then gives us some extra long-term power in the frequency range 10^{-7} Hz down to around 10^{-8} Hz.

With this set up established we begin by inputting our optical light curves into PSRESP and attempting to fit the combined PSD to a simple unbending power law as in Eq. 4.2 to see if it can be described by such a simple model. This fit is plotted in Fig. 4.8.




FIGURE 4.8: The optical power spectra from our data plotted with their best-fit model PSDs as well as the best-fit underlying power law for an unbending power law model.

As can be seen both visually and from the low fit probability, an unbending power law does not appear to be a good fit for our data, with the best fit only have an 0.11 probability with a slope of -1.7 . In fact it can be seen from the plot that some form of bending is required for our underlying power law to fit our data, which gives us confidence that perhaps a bend will improve our fit.

Therefore to try and improve our fit, we apply a model with a bending power law as in Eq. 4.3. The best-fit model for this is plotted in Fig. 4.9. As can be seen, the fit here is far more appropriate than a straight power law, with a best-fit model probability of 0.66, and a fit that visually fits the data much better.

The model parameters suggest a bend around 3.8×10^{-6} Hz with a high frequency slope of -2.1 and a low frequency slope of -1 . This then seems to be a good fit for our optical data, and indeed we should now be able to compare this to our synthetic data




FIGURE 4.9: The optical power spectra from our data plotted with their best-fit model PSDs as well as the best-fit underlying power law for a single-bend power law model.

to search for similar features which should be present if our optical data is created by X-ray reverberation from the accretion disc.

4.5.2 X-ray and Synthetic Optical Power Spectra

Before we look at the synthetic g' -band light curves, we should look at the power spectra made from the raw X-ray data. Fig. 4.10 shows these PSDs plotted with a bending power law as a comparison to the fit found in [Vaughan et al. \(2005\)](#) and find a low frequency slope and bend frequency within the error bars of the measurements found in that paper, though the high frequency slope is slightly outside the errors. The reason for the larger uncertainty in the high frequency slope is likely because the Poisson noise level is high enough that it begins making a non-negligible effect at the bend frequency, which means the error on the high frequency slope is likely compounded with the error on the Poisson noise.

Having looked at the PSDs generated from the data, the next step before we generate synthetic g' -band PSDs is to check that the simulated X-ray light curves generate PSDs that are consistent with the data PSDs. Fig. 4.11 shows that the simulated power spectra are indeed consistent with the data in the regimes where they overlap and are therefore an accurate representation of the X-ray flux being reprocessed by the accretion disc.

As can be seen from the X-ray PSDs, these nights are very consistent with one another (as one would expect from the same object), in order to probe the largest frequency




FIGURE 4.10: The 0.5-10keV x-ray power spectra from the *XMM-Newton* data plotted with their best-fit model PSDs as well as the best-fit underlying power law for a single-bend power law model.




FIGURE 4.11: The 0.5-10keV x-ray power spectra from the *XMM-Newton* data plotted with the PSDs of those same nights' simulated x-ray power spectra.

range possible we want to use the longest and best time sampled light curve, and so fitting to 4 different sets of the same light curve simply increases computational time while giving the same result four times. As such here we simply fit the synthetic g' -band from the night of January 2nd 2019, binned up to 500s and simulated over a length of 10 days, to have the best chance of detecting the synthetic optical bend without using an unnecessary amount of computing time.




FIGURE 4.12: The synthetic g' -band PSD from the *XMM-Newton* night 4 data plotted with its best-fit model PSD as well as the best-fit underlying power law for a bending power law model. This plot uses the same axis limits as the optical data PSD for ease of comparison.

Fig. 4.12 shows that, unlike our optical data, a clear bend is seen $\sim 10^{-4}$ Hz with both a very steep high frequency slope and a very flat low frequency slope. This is problematic as it would seem our data is inconsistent with reprocessing if it does not also exhibit this feature. However one possibility is that this feature does in fact exist in our data, and that the bend we detect at a frequency of $\sim 10^{-6}$ Hz is instead caused by longer-term trends than reprocessing.

4.5.3 Testing for double-bend compatibility

To test this hypothesis we will do two things. First we will try to fit our optical data to a double-bending power law as in Eq. 4.4, with a high frequency second bend around best-fit bend frequency for the simulation, to test if this is a viable model for our data.

In Fig. 4.13 it can be seen that this model does in fact fit our data well with bend frequencies around 10^{-6} Hz and 10^{-4} Hz, and the highest frequency slope very close to the simulated high frequency slope, exactly as we would expect if there was a

higher frequency reprocessing bend. In fact the fit is actually *better* than for a single bend, though admittedly only by a single percentage point. However this does indeed show that this double-bend in the power law is represented in the data, the question is then how does the reprocessed X-ray flux gain its lower frequency bend.




FIGURE 4.13: The optical power spectra from our data plotted with their best-fit model PSDs as well as the best-fit underlying power law for a double-bending power law model.

The second is to test whether the simulated light curves can be compatible with the fit from the data is a longer-term component is added to them. This long term component will be added in the form of a light curve simulated using the [Timmer and König \(1995\)](#) method for a simple power law with a bend at 10^{-6} Hz, a high frequency slope of -2 , and a low frequency slope of -1 . The simulated light curve is displayed in Fig. 4.14. The exact scaling between these light curves is unknown, so we will experiment to see if any combination is possible.

After some experimentation, it seems as if a factor of ~ 10 is needed on the [Timmer and König \(1995\)](#) simulated light curve in order to lift the mid frequency slope high enough to be compatible with our data. It should be noted however that this is very much an estimate as this test is to check if the simulation is at all compatible to our data, rather than making any claims as the nature of this long-term component, if it exists.

We then attempt to fit the PSD of our synthetic g' -band light curve additively combined with the [Timmer and König \(1995\)](#) simulated light curve that has been scaled up by a factor of 10. To do this, we apply the double-bending power law model that we applied to our data, and due to the PSD of our 10 day simulation not going to low enough frequency to accurately measure ν_{B1} or α_L , we fix these values at 10^{-6} Hz




FIGURE 4.14: Simulated 10 day light curve using the [Timmer and König \(1995\)](#) method with an input bending power law with bend frequency 10^{-6} , high frequency slope -2, and low frequency slope -1.

and 0.8, respectively. The reason for the value of ν_{B1} is that it is where we define a low frequency by using the PSD we did to generate the long-term variability simulated light curve, and the reason for the value of α_L is that it is what our data is best-fit to, and this test is to determine compatibility with the data.




FIGURE 4.15: The power spectra from our synthetic g' -band light curve plus $10\times$ the [Timmer and König \(1995\)](#) simulated light curve plotted with its best-fit model PSD as well as the best-fit underlying power law for a double-bending power law model. This plot uses the same axis limits as the optical data PSD for ease of comparison.

Fig. 4.15 shows that in fact this combination of simulated light curves do have a power spectrum that can indeed be described by a double-bending power law. The fit almost exactly reproduces the mid frequency slope, though the high frequency slope is

slightly steeper than the optical data fit would suggest. However as previously mentioned, as this slope occurs entirely within a Poisson noise dominated part of the optical data's power spectrum, the sensitivity required in measuring the noise means that the high frequency slope detected in the best fit could change a large amount for a small change in the measurement of the noise.

It seems then that the optical variability of NGC 4395 can be described by a power spectrum which would seem to originate from both X-ray reverberation *and* a long-timescale varying component. A possibility exists, of course, that there is no reverberation and that the real optical PSD has some completely unrelated physical origin. However many previous measurements of X-ray to optical lags, including our own in Chapter 3, seem to indicate a causal link between X-ray variability and optical variability that is consistent with the process of reverberation. Therefore it is very likely that reverberation does indeed contribute, and that the real PSD is being modified by some additional physical process to produce the power spectrum we see in the data.

4.6 Conclusions

In investigating the variability properties of NGC 4395, we presented four ~ 117 ks X-ray light curves from *XMM-Newton*'s EPIC pn camera previously shown in Chapter 3. In addition I showed multiple optical light curves at different timescales, two *Swift* UVOT light curves from [Cameron et al. \(2012\)](#) in the *B* and *V* bands, a white light curve taken from TESS and used previously in [Burke et al. \(2020\)](#), and a very high cadence g' -band light curve from HiPERCAM ([Dhillon and Marsh](#)), previously shown in Chapter 2.

To support these I also presented very long-timescale g' -band observations using a combination of Liverpool Telescope's IO:O ([Steele et al., 2004](#)), LCOGT's Sinestro and Spectral cameras ([Brown et al., 2013](#)), ZTF's Palomar P48 ([Masci et al., 2018](#)), and data from the Zowada Observatory [Carr et al. \(2022\)](#). From there we made power spectra of all these light curves to create an overall optical power spectrum for NGC 4395 spanning nearly 8 decades. This then allowed us to attempt to fit several power law models to our data using PSRESP ([Uttley et al., 2002](#)). A straight power law, shown in Eq. 4.2 and a power law with a single gentle bend, shown in Eq. 4.3. We discovered that a straight power law gave a very poor model probability of 0.11, while the power law with a single bend at $\sim 10^{-6}$ Hz gave a very good model probability of 0.66. To test how well the data can be described by reverberation of its X-ray flux, we then attempted to create a synthetic g' -band light curve using the convolution of the X-ray light curves and a response function generated from the KYNreverb simulation

(Kammoun et al., 2021a) which had previously been shown to create a model lag spectrum compatible with the lags detected by the data (Chapter 3).

However, our X-ray observations were too short to make a good comparison, so we had to simulate longer X-ray light curves using the method in Emmanoulopoulos et al. (2013). This method allowed us to simulate the non-Gaussian behaviour of the object by fitting both the PSD and the PDF of the data. This allowed us to create 10 day light curves which we then convolved with the response functions to create 10 day synthetic g' -band light curves. Despite ensuring that the X-ray PSDs of the simulations were consistent with the X-ray data, the optical PSDs retrieved from our synthetic g' -band light curves were not compatible with our optical data, detecting a bend two decades above the data bend frequency. Therefore we hypothesised that our data could be described as a combination of reverberation and another long-term process, which would lead to two bends, one at the lower frequency we detected with our single-bend model, and one at the higher frequency where the reverberated simulations detected a bend. To test this we fit a double-bending power law model to the data and found a model probability of 0.67, comparable and even slightly better than our single-bend detection, for bends at 0.8×10^{-6} Hz and 1.7×10^{-4} Hz, suggesting that the data can indeed be described by a double-bending power law. With the establishment that our data could be described by a double-bending power law, we then tested our hypothesis by generating a light curve with long-term variations using the method of Timmer and König (1995) and a power spectrum with a bend at 1×10^{-6} Hz and adding it to our synthetic g' -band light curve. As the ratio of the two fluxes is unknown, we simply tested several ratios and found that the long-term light curves with a scaling factor of $\sim 10 \times$ lifted the mid frequency slope enough to be possibly compatible with our data.

We then fit the PSD of this combination light curve to a double-bending power law, with the low frequency bend fixed at 1×10^{-6} Hz as we had established that by the nature of our simulation, and the low frequency slope fixed at our data's best-fit slope of 0.8, as our simulation did not sample low enough frequencies to accurately measure that slope itself. The model found a good fit probability of 0.68 for the double bending model with mid frequency slope reproducing a similar slope to the one detected by the data, though the high frequency slope was steeper. This could be because our data's high frequency slope measurement is dependent on the precise measurement of the Poisson noise, and even a small deviation could create a larger deviation in the detected slope. These fits then indicate that the optical variability of NGC 4395 can be described by a reverberation of its X-ray flux combined with another source of long-term variability. The physical origin of this long-term variability is unknown. It could be intrinsic disc variability as the reprocessed flux is both a product of the corona and of the disc, however more conclusive answers should be tested in future investigations to determine the origin of this variability component.

Chapter 5

Timescale-dependent X-ray to UV time lags of NGC 4593

5.1 Abstract

In this chapter, I will investigate the X-ray to UV lags from the AGN NGC 4593 using simultaneous observations from *XMM-Newton*, *Swift*, & *AstroSat*, and discuss the possible implications of timescale-dependent lag measurements in X-ray reverberation.

5.2 Introduction

Understanding the inner structures of AGN is one of the main aims of extragalactic research. However in almost all cases, the AGN are too small to resolve by direct imaging and so other techniques must be used. One particularly useful technique is ‘reverberation mapping’ (Blandford and McKee, 1982) where we measure the time lags between the X-ray emission and that in longer wavelength UV and optical wavebands. The assumption is that the X-rays originate from around the central supermassive black hole (SMBH) and illuminate the surrounding material, principally the accretion disc and broad line region (BLR) gas. These structures reprocess the X-rays and re-emit at longer UV/optical wavelengths.

In the absence of external illumination, the accretion disc will have a radial temperature profile which depends on its mass, spin and accretion rate and is given by Shakura and Sunyaev (1973) or, with relativistic corrections, by Novikov and Thorne (1973). X-ray illumination increases the temperature and boosts the emission.

However unless the X-ray illumination on unit area of the disc is comparable to the intrinsic black body emission from the disc, which almost never happens, the centroid of emission at any particular wavelength only moves inwards by a per cent or two. Thus the X-ray to UV/optical lags can be used to map out the temperature structure of the disc. For the BLR both UV/optical line and continuum emission is produced (Korista and Goad, 2001, 2019), with peaks in the u-band due to the Balmer continuum and in the i-band due to the Paschen continuum. The u and i-band lags produced by the BLR are longer than for those produced by the accretion disc as the BLR lies further from the SMBH.

There have now been many studies of correlated X-ray and UV/optical variability. Early studies combined X-ray observations from RXTE with ground based optical observations (e.g. Uttley et al., 2003; Suganuma et al., 2006; Arévalo et al., 2009; Breed et al., 2009, 2010). These studies typically showed that the X-rays led the optical by about one day, consistent with reprocessing, but with large uncertainty (~ 0.5 d) and without sufficient wavelength detail to map out the temperature structure of the reprocessor.

Later studies, based around X-ray/UV/optical monitoring with *Swift* provided much greater wavelength coverage (e.g. Cameron et al., 2012; Shappee et al., 2014; McHardy et al., 2014; Edelson et al., 2015; McHardy et al., 2018; Cackett et al., 2018; Edelson et al., 2019; Vincentelli et al., 2021, 2022), broadly confirming reprocessing from both disc and BLR. However a number of problems with the simple reprocessing scenario emerged. Here we concentrate on just one of these problems, which we shall call ‘the timescale problem’, i.e. the fact that the lag that is measured depends on what the time resolution of the data and whether you have removed long timescale trends or not. This problem is particularly apparent between the X-ray and far-UV bands.

Here it has been noted, since the earliest *Swift* papers (e.g. Shappee et al., 2014), that although a lag spectrum of the form $\text{lag} \propto \lambda^{4/3}$, expected from reprocessing from an accretion disc (e.g. Cackett et al., 2007), was usually a good fit to the lags throughout the UV and optical bands, the extrapolation of that fit back to the X-ray band (i.e. ~ 2 Å) predicted an X-ray to UV lag which was a good deal smaller than the one that was actually observed (see also summary in McHardy et al. (2018)). Thus the lag spectrum is usually steeper between the X-rays and far-UV than between the UV and optical bands.

A variety of solutions have been proposed, including that the X-rays are scattered through an inflated inner edge of the disc, thus introducing an additional delay, before emerging as far-UV photons which then illuminate the outer disc Gardner and Done (2017). Alternatively, including a more distant reprocessor than the disc, i.e. the BLR, can produce the additional lag in some cases (McHardy et al., 2018). It has also

recently been shown (McHardy et al., 2022) that an accretion disc, truncated at the dust sublimation radius, can explain a steeper X-ray to far-UV lag spectrum.

Early on McHardy et al. (2014) it was noted that if the lightcurves were filtered to remove long timescale trends, the X-ray to UV lags derived from the resultant de-trended lightcurves were quite consistent with an extrapolation of the lags within the UV/optical bands back to the X-ray band. Long timescale trends are more apparent in the UV/optical bands than in the X-rays and indicate a second source of slower variations which affects the UV/optical bands but has less effect in the X-ray band. This second source of variability could be accretion rate variations propagating inwards through the disc (Arévalo et al., 2009; Breedt et al., 2009).

Further examples of how the X-ray to far-UV lag is reduced by the removal of long term trends are given in the case of NGC 4593, the source discussed in the present report, by McHardy et al. (2018) and also, for NGC 7469, by Pahari et al. (2020). The latter case is particularly dramatic. From a month or so of almost continuous X-ray observations with RXTE and UV observations by IUE, Nandra et al. (1998) had shown that, in the original, non-detrended lightcurves, the UV led the X-rays by about 4 days, which is entirely contrary to the expectations of X-ray reprocessing. However after removing only the longer timescales (>5 d), Pahari et al. (2020) showed that the far-UV lagged behind the X-rays by 0.37 d and that the longer wavelength UV lagged the short wavelength by similarly short amounts.

Overall, the lag results were entirely consistent with reprocessing from an accretion disc. Thus the reprocessing signal was present but had been hidden by the second source of variations. Welsh (1999) provides a detailed analysis of the errors that can occur when measuring lags from cross-correlation functions if long timescale trends are not removed. Care, of course, must be taken not to filter out timescales of interest. It is noted, entirely predictably, by McHardy et al. (2018) that the measured lags decrease as the timescale of filtering decreases. Thus signals of reprocessing from the BLR, which can be seen in the tails to longer timescales in the MEMecho (Horne et al., 2004) UV and optical response functions of NGC 4593 (McHardy et al., 2018), may be filtered out with over-vigorous filtering.

The analysis described above is important because, if the lags are consistent with simple disc reprocessing, it shows that the accretion disc must be directly visible, without obscuration, to at least part of the X-ray emission region. Thus the inner disc cannot be too highly inflated, and remain optically thick to X-rays.

NGC 4593 is one of the brightest Seyfert galaxies in the sky and has been subject to intense monitoring with *Swift* (McHardy et al., 2018) and HST (Cackett et al., 2018). Whilst the lags between the various UV and optical bands were more or less in agreement with simple disc reprocessing, the observed lag between the X-rays and the

shortest wavelength UV band, i.e. UVW2 (1920 Å) was ~ 66 ks compared to the model lag based on reprocessing from the accretion disc of ~ 8 ks. This difference causes a big problem for simple disc reprocessing models although it was shown that it was possible to reproduce the observed UV/optical lightcurves from the observed X-ray lightcurve as long as the reprocessor had long tail, probably consistent with additional reprocessing from the BLR.




FIGURE 5.1: The Optical Spectrum of NGC 4593. From Kollatschny and Fricke (1985).

If we look at the optical spectrum for NGC 4593, as in Fig. 5.1, we can see that unlike the spectrum of NGC 4395 in Chapter 2, we in fact do see very Doppler Broadened lines. This then supports the idea that the emission of this AGN has a strong BLR component.

The range of lags which is accessible depends on the frequency of sampling. *Swift* observed NGC 4593 almost every orbit (96 min) for 6.4 d from 2016 July 13 to 18 and thereafter every second orbit for a further 16.2 d. Each observation totalled approximately 1 ks although observations were often split into two, or sometimes more, visits. Whilst the sampling rate for the first 6 d is very good by *Swift* standards, the observations were far from continuous, limiting our coverage of short timescale variations.

However, during the 6.4 d of intensive *Swift* observations, we were able to arrange a continuous 140 ks (1.6 d) observation with *XMM-Newton*, providing continuous X-ray coverage and also continuous UV coverage (2910 Å), using the Optical Monitor in Fast Mode. These observations provide superb coverage of short timescale variability. We also observed for ~ 4 of the 6.4 d period with *AstroSat*, providing additional X-ray coverage and also coverage in two further UV bands. In this report we re-examine the

lags between the X-ray and UV emission in NGC 4593 (Section 5.5.1) in the light of this additional data and by re-examining the *Swift* data. We find that the measured lag does indeed depend very significantly on the data and the timescales available for investigation within that data. We conclude by fitting disc reprocessing models to the measured lags (Section 5.5.4), including testing for whether a truncated disc is necessary to explain the lags.

5.3 Observations

5.3.1 *XMM-Newton* Observations

NGC 4593 was observed continuously by *XMM-Newton* for one orbit of 140ks duration. The EPIC PN X-ray camera observed in small window mode and the Optical Monitor (OM) observed in Fast Mode with the UVW1 ($\lambda_{eff} = 2910\text{\AA}$) filter. The advantage of Fast Mode is that readout is continuous with a 500ms integration time which can be binned up as required.

In Imaging Mode the shortest integration time is 800s with 300s gap for readout. A potential disadvantage of Fast Mode is that the target must be fully located within the 10×10 arcsec detector window. However images of the window can be made with longer integrations and, in our case, the target remained fully located within the window.

XMM-Newton data reduction is done using the ESA Science Analysis System (SAS). From the Observation Data Format (ODF) files, SAS can generate photon event files from which we obtain light curves for the PN and OM. Initially we chose 100s bins, which we later increased.

In addition we ran the light curves through a filter to remove outliers more than 3σ from nearby points, particularly a large flare in the OM that was not replicated in the X-rays, and a large flare in the X-rays that was not filtered out by the SAS pipeline's filtering process.

The 0.5-10 keV X-ray and UVW1 light curves, binned to 400s and 1ks respectively, are shown in Fig. 5.2. By eye it can be seen that both light curves have a similar shape but the variations are not simultaneous. For all light curves used in this paper we take the start of the *XMM-Newton* light curves as the time zero point.




FIGURE 5.2: *XMM-Newton* EPIC pn 0.5-10keV X-ray light curve for NGC 4593, binned to 400s, and the *XMM-Newton* OM UVW1 light curve for NGC 4593, binned to 1000s.

5.3.2 *Swift* Observations

We use the *Swift* data that has already been published by [McHardy et al. \(2018\)](#), to whom we refer for observational details. The timing information which is relevant to the present paper is that NGC 4593 was observed for 1 ks nearly every orbit (96 min) from 13 to 18 July 2016 (6.4 days) and every 3 orbits for another 16.2 days. These observations are made using the X-ray Telescope (XRT) in the 0.5-10 keV band and the UV and Optical Telescope (UVOT) in 3 UV (*UVW2*, *UVM2*, and *UVW1*) and 3 optical (*U*, *B*, and *V*) filters. The effective wavelengths for the 6 UVOT bands in the same order as above are 1928Å, 2246Å, 2600Å, 3465Å, 4392Å, and 5468Å. For comparison with the *XMM-Newton* OM data the most relevant UVOT filter here is *UVW1*.

The *Swift* X-ray and *UVW1* light curves are shown in Fig. 5.3.

5.3.3 *AstroSat* Observations

AstroSat observed NGC 4593 from 14 to 18 July 2016 (4.13 days). Observations were made using the Soft X-ray Telescope (SXT [Singh et al., 2016, 2017](#)) in the 0.3-7 keV X-rays and the two cameras of the Ultra-Violet Imaging Telescope (UVIT; [Tandon et al.](#)




FIGURE 5.3: *Swift* XRT 0.5-10keV X-ray and UVOT UVW1 light curves for NGC 4593.

(2017), Tandon et al. (2020)) in the near and far UV. The filters used were the FUV BaF2 filter ($\lambda_{eff} = 1541\text{\AA}$) and the NUV B4 filter ($\lambda_{eff} = 2632\text{\AA}$). Note that though SXT's designed energy range is 0.3-8keV, an instrumental background line at $\sim 7.3\text{keV}$ reduces the useful bandwidth.

AstroSat has a similar low-earth orbit to *Swift*. Although it observed NGC 4593 continuously, earth occultation reduced the on-source time to $\sim 1000 - 1500\text{s}$ per orbit, leaving gaps of $\sim 4000\text{s}$. Here we binned each observing section into one data point giving a light curve with accurately timed data points and $\sim 5600\text{s}$ sampling.

This *AstroSat* data will also be shown by Kumari et al. (2022, in prep). The light curves for the binned X-ray, FUV BaF2 and NUV B4 are shown in Fig. 5.4.




FIGURE 5.4: *AstroSat* SXT 0.3-7keV X-ray, UVIT FUV BaF2, and UVIT NUVB4 light curves for NGC 4593.

5.4 Lag Determination

We determined lags using the Flux Randomisation (FR) Cross-Correlation Function (CCF) method of [Peterson et al. \(1998\)](#). We did not add Random Subset Selection (RSS) as it does not use the full datasets and hence artificially increases the uncertainties. However the central lag measurements do not change if RSS is added.

Another common method of lag determination is the Damped Random Walk modelling method of JAVELIN ([Zu et al., 2013](#)). However with only one large variation visible in the *XMM-Newton* observations, the light curves are too short for JAVELIN to produce a sensible result. Thus although JAVELIN produces more precise lags, here we use only FR CCF method.

5.5 Results and Analysis

5.5.1 *XMM-Newton* Lag Results

The FR CCF lag distribution between the PN X-ray and OM UVW1 lightcurves is shown in Fig. 5.5. This clean distribution has no secondary peaks and gives a centroid lag value of 29.5 ± 1.3 ks.




FIGURE 5.5: *XMM-Newton* EPIC pn to OM UVW1 Lag Distribution determined from our FR CCF trials, with a centroid lag of 29.5 ± 1.3 ks.

As reassurance that this lag is correct and not a result of a strange artefact in one or both of the light curves, in Fig. 5.6 we plot the X-ray light curve (in red) with the UVW1 light curve (in blue), scaled to the same amplitude of variability and shifted back by 29.5 ks, superposed. We can see visually that there is good agreement between the two light curves.

5.5.2 *Swift* Lag Results

5.5.2.1 Reanalysis of CCF lags

Comparing the above newly determined lag to that found in [McHardy et al. \(2018\)](#), we see that there is a large disparity. The value of ~ 30 ks determined from the *XMM-Newton* data is less than half than that determined from the *Swift* data (~ 60 ks).




FIGURE 5.6: XMM-Newton X-ray PN 0.5-10keV light curve with OM UVW1 light curve shifted by the XMM-Newton (blue) and *Swift* (red) determined lag values.

The XMM-Newton X-ray data has a much higher S/N than that from *Swift*, with a mean count rate $> 10\times$ higher. Although *Swift* observed over a longer period (see Sec. 5.3), the observations, unlike with XMM-Newton, were not continuous. Thus the XMM-Newton provides much better temporal sampling of short-timescale variability than *Swift*. However it is hard to see why the improved sampling should have quite such a large effect so we reexamined the *Swift* data to see if we could discover the answer.

Upon performing the same FR CCF lag determination for the *Swift* XRT and UVOT UVW1 we discovered a similar lag to the one in McHardy et al. (2018), i.e. ~ 60 ks. If we also shift the OM data back by this amount (Fig. 5.6) we can see that this lag does not fit our XMM-Newton data at all.

However upon examining the *Swift* lag distribution, it is clear that it's bimodal, showing one peak around 30 ks and another around 70 ks (Fig. 5.7). By fitting a convolution of two asymmetric, decaying gaussian functions to this bimodal distribution using *scipy's* `curve_fit`, we can then disentangle them to retrieve the lags for each peak.

The lags we retrieve from these disentangled distributions are 27.0 ± 8.2 ks and 68.0 ± 10.3 ks.

The mean level of X-ray light curve does not change noticeably over the period of the observations but the mean UV level decreases. To test the idea that there are actually two lags within the *Swift* data, we produce a smoothed UV light curve which we then




FIGURE 5.7: *Swift* XRT 0.5-10keV to UVOT UVW1 raw lightcurve Cross Correlation Lag Distribution with disentangled peaks.

subtract from the original light curve to produce a detrended light curve. We can then measure lags using both the smoothed and detrended UV light curves.

5.5.2.2 Detrending with LOWESS

To do the smoothing we use a method called Locally Weighted Scatterplot Smoothing (LOWESS), which is a non-parametric regression fitting technique and which we used successfully previously (Pahari et al., 2020). We use the Python module `statsmodels.lowess` function. The main input of this function, other than the data itself, is the fraction of the light curve the function will smooth along. Using a very small smoothing fraction will, after detrending, leave only very short-timescale variations. Thus this fraction should be significantly larger than any timescales of interest.

Our total light curve length is ~ 22 days and the lags seen in the original lightcurves are on the order of a day so we investigated the effect of various smoothing timescales longer than a few days.

In the left and right panels of Fig 5.8 we show the lag distributions between the un-detrended Swift X-ray lightcurve and the detrended UVOT lightcurve after detrending with 2.5d and 7.5d timescales respectively. The 7.5d timescale still leaves a large residue of the longer peak. The 2.5d detrending timescale shows a messy

distribution, as the detrending is beginning to remove short-term lag features and therefore reducing the confidence of the lag determination, as we are approaching the timescale of the lags themselves. The resulting lag measurements for different detrending timescales are plotted in Fig. 5.9.

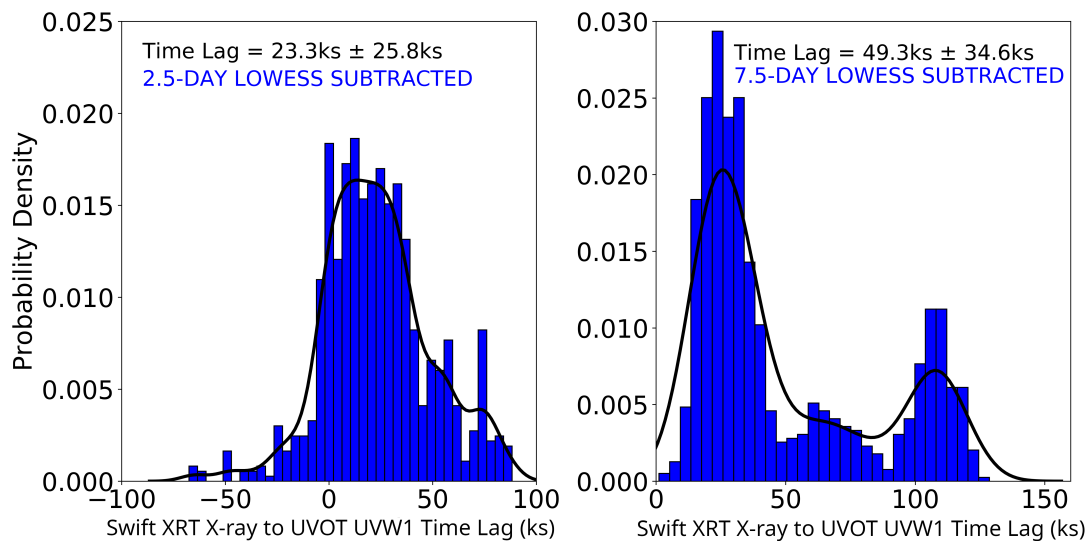


FIGURE 5.8: Cross Correlation Lag distributions from *Swift* un-detrended XRT 0.5-10keV to UVOT UVW1 detrended with 2.5-day (left) and 7.5-day (right) Lowess smoothing timescales.




FIGURE 5.9: Lag Determinations vs Lowess Smoothing Fraction for the *Swift* UVOT lightcurves.

The trend in the UV is larger than in the X-rays and detrending also in the X-rays has usually had little extra effect on the lags. However to check whether detrending the X-rays as well as the UVOT has any effect on the lags here we detrend both bands with a 5d smoothing timescale. The resultant lag distribution is shown in Fig. 5.10, left panel, giving a lag of 23.8 ± 21.2 ks, which is almost identical to that where only the UVOT was detrended.




FIGURE 5.10: Cross Correlation Lag Distributions from *Swift* XRT 0.5-10keV to UVOT UVW1 where both lightcurves have been 5-day Lowess Detrended (left) and Smoothed (right), respectively.

To determine whether the longer (68ks) lag found in the raw, un-detrended UVOT and X-ray lightcurves (Fig. 5.7) might result from the slower variations which we have subtracted from the detrended lightcurves, we performed a lag measurement using the smoothed X-ray and smoothed UVW1 light curves. The detrended and smoothed UVOT lightcurves are shown in Fig. 5.11 and the lag distribution is shown in Fig. 5.10, right panel. The resultant lag is 99.1 ± 17.2 ks, but with no sign of a peak in the 20-30 ks range. This lag probably represents the longer lag seen in the full light curves.

5.5.2.3 Testing alternate explanations

To check whether the different lags were the result of energy dependent lag components, we carried out the above detrending analysis separately on the soft (0.5-2keV) and hard (2-10keV) *Swift* light curves but found the same result.

In addition to check that the two peaks are not related to the sampling pattern, e.g. with the initial high sampling intensity period of 6.4 days producing a short lag and the remaining 16.2d of lower intensity sampling producing a longer lag, we carried out a separate lag analysis for both periods (Fig. 5.12). Both peaks are present in both




FIGURE 5.11: *Swift* UVOT UVW1 5-day Lowess Detrended and Smoothed Light Curves.

sections although the reduction in data leads to messier distributions. This analysis shows that the *Swift* data does contain a short lag consistent with that seen in the XMM data. There is a small caveat that although the *Swift* UVOT and XMM OM are very similar instruments with similar filters, there has been some loss of UV throughput in the OM. This loss has resulted in the central wavelength of the UVW1 filter now being $\sim 2910\text{\AA}$ whereas that of *Swift* is $\sim 2600\text{\AA}$. This small wavelength difference may result in a similarly small lag difference. However the uncertainty in the *Swift* lag is too large for us to detect that difference.

To check whether the bimodal lags from the raw light curves were a result of stray points resulting from minor irregularities in the data reduction process, we compared these light curves, generated with an in-house University of Southampton pipeline for [McHardy et al. \(2018\)](#), to lightcurves from the same observation period generated with the University of Leicester’s online product builder ¹. The lag distributions we determined for these regenerated light curves were identical to those found with our in-house light curves for both the raw, detrended, and smoothed light curves, showing that the distributions we detect are therefore not the result of any issue in the reduction process.

¹https://www.swift.ac.uk/user_objects/index.php




FIGURE 5.12: *Swift* XRT 0.5-10keV to UVOT UVW1 High Intensity Half only (left) and Low Intensity Half (right) CCF Lag Distributions.




FIGURE 5.13: *Swift* XRT 0.5-10keV and UVOT UVW1 plotted along with the XMM-Newton EPIC pn and OM UVW1 light curves.




FIGURE 5.14: *Swift* XRT 0.5-10keV to *XMM-Newton* OM UVW1 Cross Correlation Time Lag.

The *Swift* observation period overlaps entirely with that of XMM. The XMM and *Swift* X-ray and UV light curves are plotted together in Fig. 5.13, showing good agreement, particularly in the X-rays where S/N is greatest. The much better sampling of the XMM data in both bands is apparent.

As a check, we can perform a lag analysis of the *Swift* XRT light curve with the XMM OM. The resulting lag distribution is shown in Fig. 5.14 and the measured lag is 23.3 ± 3.2 ks with no other peaks at larger lags. This lag is very similar to that provided by both the XMM data on its own and by the detrended *Swift* data. This result implies that the detrended *Swift* lag correctly represents the true lag that would be measured in short, well sampled, light curves.

5.5.3 *AstroSat* Lag Results

As a further diagnostic of the short timescale lags we can use another overlapping light curve set from *AstroSat*.

The central wavelengths of the *AstroSat* NUV B4 and FUV BaF2 filters are 2632Å and 1541Å respectively. The central wavelength of the NUV filter is similar to that of the *Swift* UVOT UVW1 filter. In Fig. 5.15 we compare the *AstroSat* SXT, FUV and NUV lightcurves with the X-ray and UVW1 lightcurves from XMM. As with *Swift*, the




FIGURE 5.15: *AstroSat* SXT 0.3-7keV, FUV BaF2, and NUVB4 plotted along with the *XMM-Newton* EPIC pn and OM UVW1 light curves.

AstroSat sampling does not match that of XMM but similar broad features are visible in both the X-ray and UV lightcurves.

We begin to measure the lags using the original, un-detrended data. Here in Fig. 5.16 the lag distribution of SXT to FUV shows a larger lag than would be expected for its short central wavelength. On the right of the same figure, we determine the FUV lag relative to the *XMM-Newton* EPIC pn X-rays and similarly find a larger than expected main lag.

If we look at the FUV data points overlaid onto the SXT data, we can see from the top panel of Fig. 5.17 that the FUV visually appears to be compatible with a small lag. Additionally it can be seen from the bottom panel that applying the 33ks lag detected from the raw lightcurves (Fig. 5.16, left panel) to the FUV light curve is larger and does not visually align the light curves.




FIGURE 5.16: *AstroSat* SXT 0.3-7keV to UVIT FUV BaF2 (left) and XMM EPIC pn 0.5-10keV to UVIT FUV BaF2 (right) Cross Correlation Time Lags.




FIGURE 5.17: *AstroSat* SXT 0.3-7keV plotted with UVIT FUV BaF2 (top) and SXT 0.3-7keV plotted with UVIT FUV BaF2 shifted by -33ks (bottom) light curves.

Therefore with the evidence of long-timescale trends found in the *Swift* light curves, we can apply a LOWESS detrending to see if it alters the FUV lags. In this case, though the light curve is shorter than *Swift*, variability power can still ‘leak’ into a light curve if there is variability on timescales longer than the length of our observation (Deeter and Boynton (1982)), which according to our *Swift* observations there appears to be. In this case then we can perform a LOWESS detrending over the entire length of the light curve. Though the light curve is only ~ 4 days long, it can be seen in Fig. 5.9 that even an averaging window down to 2.5 days does not change the result much as it is still a timescale an order of magnitude larger than our expected lags, therefore a much larger 4 day smoothing should be sufficient for our detrending.




FIGURE 5.18: *AstroSat* UVIT FUV BaF2 4-day Lowess Detrended and Smoothed Light Curves.

In Fig. 5.18 the results of the detrending can be seen. Note that this detrending detects a very similar trend to the *Swift* over the equivalent time period, giving more confidence that the trend being removed is physically real.

We note that we for the NUV B4 lag distributions, the lags we detect appear to be consistent with the rest of the lag spectrum with regards to their wavelength and Lowess detrending does not affect the detected lag. We speculate that a reason for this could be that as the NUV is a longer wavelength band, the damping effects we see at larger responding wavelengths have reduced the magnitude of the long-term contribution. This can be visually seen in the light curves, as the trend on the first

150ks of the FUV light curve results in an overall flux decrease of $\sim 18\%$, whereas in the equivalent period of NUV the decrease appears to only be $\sim 8\%$ of the flux. Though NUV is approximately the same wavelength as *Swift*'s UVW1, the light curve is far shorter than the *Swift* light curve and thus the lag determination may not be sensitive enough to detect the long-term variability's effect in such a short light curve at this waveband, which would be consistent with what we see in the even-shorter *XMM-Newton* UVW1 light curves.

Therefore, the lag distributions between the *AstroSat* SXT 0.3-7keV X-rays & the Lowess-Detrended FUV, and between the SXT 0.3-7keV X-rays & the NUV are shown in Fig. 5.19, finding lags of 23.0 ± 4.9 ks and 34.2 ± 9.0 ks, respectively. This FUV lag is slightly longer than might be expected for its wavelength, but as we will examine in Section 5.5.4, there may be other factors than disc reverberation in NGC 4593. It is also noted that the FUV shows physical similarities to the *Swift* UVW2 light curve, which also shows similar lag taken relative to UVW1's lag, so this result is consistent with the appearance of the light curve. The NUV lag is in agreement with the short timescale lags found with *Swift* and *XMM-Newton*.




FIGURE 5.19: *AstroSat* SXT 0.3-7keV to Lowess-Detrended UVIT FUV BaF2 (left) and SXT 0.3-7keV to UVIT NUVB4 (right) Cross Correlation Time Lags.

We also measure the lags of the Lowess-Detrended FUV and the NUV relative to the high-intensity XMM EPIC pn X-rays. These lags produce the distributions found in Fig. 5.20, reporting lags of 24.1 ± 5.8 ks and 35.5 ± 7.2 ks, for the Lowess-Detrended FUV and the NUV respectively. Both of these lags are consistent with their SXT counterparts within the uncertainties and considering the differences in sampling patterns.

Finally, we can calculate the lag between the *AstroSat* SXT and the XMM OM UVW1 light curve. The lag distribution is shown in Fig. 5.21. The measured lag is 28.5 ± 2.9




FIGURE 5.20: *XMM-Newton* EPIC pn 0.5-10keV to *AstroSat* Lowess-Detrended UVIT FUV BaF2 (left) and EPIC pn 0.5-10keV to UVIT NUVB4 (right) Cross Correlation Time Lags.

ks which agrees with all of the previous lags derived from short datasets, adding further confirmation to the conclusion that, if measured from short lightcurves or with long timescale trends removed, the lag between the X-rays and UVW1 wavelengths is around 20-30 ks.




FIGURE 5.21: *AstroSat* SXT 0.3-7keV to XMM OM UVW1 Cross Correlation Time Lag.

5.5.4 Model Lag Spectra

To test whether the lags measured here are consistent with reprocessing of X-rays directly by a surrounding accretion disc we compare the lag measurements to model predictions. Here we use the analytic model approximation given by Kammoun et al. (2021a), which is derived from the simulation code KYNreverb, developed by Dovčiak². We assume a coronal lamp-post height of 6 gravitational radii, consistent with that found from X-ray reverberation mapping (Emmanoulopoulos et al., 2014; Cackett et al., 2014) and microlensing (Dai et al., 2009; Mosquera et al., 2013), a mass of $7.63 \times 10^6 M_{\odot}$ and an accretion rate in Eddington units of 8.1% as listed in Table 3 of McHardy et al. (2018). From the same source we also derive an illuminating 2-10keV luminosity of $6.31 \times 10^{42} \text{ ergs s}^{-1}$, taken from the 0.1-195keV ionising BAT luminosity of 3.0×10^{43} using a correction factor of 0.288 assuming a spectrum with an index of $\Gamma = 2$. The predicted lags for both minimum ($a=0$) and maximum ($a=0.998$, but approximated in the model by $a=1$) spin, together with the observed lags, are shown in Fig. 5.22.

The lags in Fig. 5.22 are not completely independent as the same UV observations are used in conjunction with different X-ray observations. We indicate the relevant UV observations on the figure. The lags using the *Swift* UVW1 (though with very large error) and the *XMM-Newton* OM observations favour the high spin option but the *AstroSat* NUVB4 observations favour low spin. An intermediate spin would fit these data better, with the exception of FUV BaF2, whose lag appears too large but will be explored in our full lag spectrum. We do not draw strong conclusions from these data except to note that, whatever the spin, the previously accepted mass and accretion rate are mostly consistent with these lags. We also note that a lag fit that goes straight to the origin is quite consistent with the data.

To investigate the lag modelling further we add the UV inter-band lags from McHardy et al. (2018) and Cackett et al. (2018), referenced to our new X-ray to UVW1 lag, taking the weighted mean of all of those lags. It should be noted that although these previous lags were not taken from smoothed light curves, we have not seen evidence of bimodality in inter-UV lags, and so as the previous *Swift* and HST lags used here were all measured relative to the UVW2 band, the lack of smoothing should not interfere. The resultant combination is shown in Fig. 5.23, with all lags plotted relative to UVW2.

By referencing the *Swift* UVOT and HST lags to our new best estimate of the lag of the X-rays by the UVW1, we see more clearly that the whole lag spectrum does continue smoothly to the X-ray origin without need for an additional offset such as has previously been commonly found. We also see that, assuming the previously accepted mass of $7.63 \times 10^6 M_{\odot}$ and an accretion rate in Eddington units of 8.1%, a low spin

²Available from <https://projects.asu.cas.cz/stronggravity/kynreverb>




FIGURE 5.22: Our UVW1 and UVIT CCF lags compared to the model lag spectra. The XMM to *AstroSat* UVIT and SXT to OM data points have been shifted 40\AA to the right for visibility.




FIGURE 5.23: *Swift* lags for NGC 4593 relative to [McHardy et al. \(2018\)](#) UVW2 lags, as well as HST lags from [Cackett et al. \(2018\)](#) and our *AstroSat* lags.

model could fit the lags out to $\sim 3500\text{\AA}$. However such a model overpredicts the lags at longer wavelengths. A high spin model slightly underpredicts the majority of the lags, except around 5000\AA where there are fewer emission lines. There are, however, clear excesses in the $2000\text{--}4000\text{\AA}$ and $7000\text{--}8000\text{\AA}$ region. These excesses have been noted previously as probably being due to Balmer and Paschen continuum emission, respectively (e.g. [Cackett et al., 2018](#)) (see Appendix C). This emission is believed to originate from reprocessing of high energy emission in the Broad Line Region ([Korista and Goad, 2001, 2019](#)). This also helps explain the excess lag in the FUV measurements, as by comparing it to the lags around it from [Cackett et al. \(2018\)](#) it appears this BLR component is contributing to lag excess at this waveband. Thus, overall, a high spin model with additional BLR contribution provides the best explanations of the observed lags.

Chapter 6

Conclusions

6.1 Main Results

Over the course of this thesis I have presented four studies on AGN variability in order to expand on the knowledge of accretion physics and the geometry of the environment close to supermassive black holes. I have looked at the variability of the X-ray, UV, and Optical emission of the low-mass low-accretion rate black hole NGC 4395, as well as the effect of timescale-dependence of the time lags of NGC 4593. Chapters 2 and 5 will be published in refereed journal papers as [McHardy et al. \(2022\)](#) and [Beard et al. \(2022\)](#), respectively. The rest will be submitted for publishing in the near future.

This chapter will be a summary of the results and implications from each study, and then a reflection on the kinds of future work that could be done to expand on the information that has been presented here in order to further the understanding of AGN systems.

6.1.1 Optical Interband Lags of NGC 4395

In Chapter 2, I presented fast (~ 200 s sampling) u, g, r, i, z photometry of the very low mass AGN NGC 4395 with the Liverpool Telescope, followed up by even faster (3s sampling) u_s, g_s, r_s, i_s and z_s simultaneous monitoring with HiPERCAM on the 10 metre Gran Telescopio Canarias. Unlike in all other AGN, g_s lags u_s by a large amount. This observation is contrary to the expectation of reprocessing in the Broad Line Region but consistent with disc reprocessing. I also saw, with high precision, that there is very little increase in lag between the u_s -band and r_s, i_s , and z_s bands, indicating an edge to the reprocessing region. I then fit truncated disc reprocessing models to the combined HiPERCAM lags and earlier lags between the X-ray and

UVW1 (2910 Å) bands. Though I have not tested a full range of masses, for the normally accepted mass of $3.6 \times 10^5 M_\odot$ I found a reasonable agreement with zero spin, an outer truncation radius of $\sim 1700 R_g$ and the temperature-dependent disc colour correction factor recommended by [Done et al. \(2012\)](#). For a smaller mass of $4 \times 10^4 M_\odot$ there is passable agreement but only if a very high, probably unphysical, correction factor of 2.4, is applied over the whole disc. A truncation radius of $\sim 14000 R_g$ is then required. Disc self gravity is probably unimportant in this low mass AGN. The dust sublimation radius is similar to the truncation radius for the small mass, so a dusty wind might shield the outer disc and provide an edge. However for the large mass the sublimation radius is further out so a line-driven wind might be more applicable. These observations provide the first good evidence for a truncated accretion disc in an AGN and caution that disc truncation must be included in modelling of AGN reverberation lags.

6.1.2 X-ray to UV/Optical Lags of NGC 4395

In Chapter 3, I presented extended ($4 \times \sim 117$ ks) simultaneous observations of NGC 4395 using XMM-Newton's EPIC camera in the X-ray, its Optical Monitor for UV, and 6 ground-based observatories for SDSS g-band coverage. Using a combination of Javelin and FR/RSS Cross-Correlation methodology, I determined from Javelin X-ray to UVM2 time lags of $511.6^{+8.3}_{-9.7}$ s and $414.8^{+57.4}_{-81.7}$ s, supported by FR/RSS CCF lags in agreement. I also found an X-ray to U-band Javelin time lag which is somewhat inconsistent with both models or a BLR-generated U-band excess, which requires further observation to investigate further. I also found X-ray to g-band Javelin lags of $696.4^{+3.8}_{-54.3}$ s, $828.6^{+9.7}_{-18.5}$ s, and $747.4^{+84.3}_{-97.5}$ s from three observations, the fourth's disagreement likely results from a lack of g-band data. I then fit these lags to two models, the function $\tau \propto \lambda^{4/3}$ as required by the Shakura-Sunyaev α -disc model, and simulated lag spectra generated by KYNreverb, a General Relativistic ray-tracing reverberation code. For the α -disc model, in contrast to the observations performed by [McHardy et al. \(2016\)](#), I found an X-ray to UV lag excess that indicates an inflated inner disc, in the range of ~ 200 -400 seconds. However, the KYNreverb simulations suggest that this turn off from the α -disc model is caused by a truncated disc-edge bending the lag spectrum.

6.1.3 Multiwaveband Power Spectra of NGC 4395

In Chapter 4, I presented multiple optical light curves in addition to the X-ray light curves from Chapter 3 in order to investigate the power spectra of NGC 4395 across multiple wavebands. For light curves from previous work, I presented two light curves from [Cameron et al. \(2012\)](#) taken with *Swift* UVOT in the B and V filters as well

as a white light curve taken using TESS from [Burke et al. \(2020\)](#), and a very high cadence g' -band light curve from HiPERCAM ([Dhillon and Marsh](#)) taken from Gran Telescopio Canarias. To provide very long-timescale g -band observations I then used a combination of Liverpool Telescope's IO:O ([Steele et al., 2004](#)), LCOGT's Sinestro and Spectral cameras ([Brown et al., 2013](#)), ZTF's Palomar P48 ([Masci et al., 2018](#)), and data from the Zowada Observatory [Carr et al. \(2022\)](#). I then measured power spectra of all these light curves to create an overall optical power spectrum for NGC 4395. From this I fit several power-law models, a straight power-law and a power-law with a single gentle bend, to our data using PSRESP ([Uttley et al., 2002](#)). I found that a straight power-law gave a very poor model probability of 0.11, while the power-law with a single bend at $\sim 10^{-6}$ Hz gave a very good model probability of 0.66. Then to test how well the data can be described by X-ray reprocessing, I created a synthetic g' -band light curve using the convolution of longer simulated X-ray light curves created by the method from [Emmanoulopoulos et al. \(2013\)](#) and a response function generated from the KYNreverb simulation ([Kammoun et al., 2021a](#)). These synthetic optical power spectra were not compatible with the optical data, and instead detected a bend two decades above the bend frequency of the data. This suggested a combination of reprocessing and another long-term process, which would lead to two bends. To examine this I fit a double-bending power-law model to the data and found a probability of 0.67 for bends at 0.8×10^{-6} Hz and 1.7×10^{-4} Hz. As the data could be described by a double-bending power-law, I generated a light curve with long-term variations using the method of [Timmer and König \(1995\)](#) and a power spectrum with a bend at 1×10^{-6} Hz and added it to the synthetic g' -band light curve. The double-bending power-law model found a good fit probability of 0.68 with a mid frequency slope reproducing a similar slope to that of the data, though the high frequency slope was steeper. This could be dependent on the precise measurement of the Poisson noise level to find the high frequency slope of the data. These fittings indicate that the optical variability of NGC 4395 can be described by a combination of X-ray reprocessing and another source of long-term variability, the physical origin of which could be intrinsic disc variability but is so far inconclusive.

6.1.4 Timescale-dependent Lags of NGC 4593

In Chapter 5, I presented a 140ks observation of NGC 4593 with *XMM-Newton* providing simultaneous and continuous PN X-ray and OM UV (UVW1 2910Å) lightcurves which sample short-timescale variations better than previous observations. These observations were simultaneous with 22d of *Swift* X-ray and UV/optical monitoring, reported previously, and 4d of *AstroSat* X-ray (SXT), far (FUV 1541Å), and near (NUV 2632Å) UV allowing lag measurements between them and the highly-sampled XMM. From the XMM we find that UVW1 lags behind the X-rays by 29.5 ± 1.3 ks, \sim half the lag previously determined from the *Swift* monitoring.

Re-examination of the *Swift* data reveals a bimodal lag distribution, with evidence for both the long and short lags. However if I detrend the *Swift* lightcurves by LOWESS filtering with a 5d width, only the shorter lag (23.8 ± 21.2 ks) remains. The NUV observations, compared to PN and SXT, confirm the ~ 30 ks lag found by XMM and, after 4d filtering is applied to remove the long-timescale component, the FUV shows a lag of ~ 23 ks. The resultant new UVW1, FUV, and NUV lag spectrum extends to the X-ray band without requiring additional X-ray to UV lag offset, which if the UV arises from reprocessing of X-rays, implies direct illumination of the reprocessor. By referencing previous *Swift* and HST lag measurements, I obtain an X-ray to optical lag spectrum which agrees with a model using the KYNreverb disc-reprocessing code, assuming the accepted mass of $7.63 \times 10^6 M_{\odot}$ and a spin approaching maximum. Previously noted lag contribution from the BLR in the Balmer and Paschen continua are still prominent.

6.2 Future Work

Every investigation in this thesis could be expanded upon by further studies. Almost all of the uncertainty in the conclusions drawn from these studies originates in data volume and quality, and future observations will hopefully be able to improve on these aspects in order further constrain the lags and power spectra of the objects observed here.

The HiPERCAM data retrieved for the optical interband lags was incredible for its time resolution thanks to the enormous collecting area, but the unfortunate tracking issues described in Chapter 2 limited the utility of a large portion of this data. Extended observations of NGC 4395 with HiPERCAM on GTC without these technical issues would be of the utmost value; with such superb time resolution some phenomenally accurate measurements can be made. If some extended multiwaveband HiPERCAM optical observations were to be made simultaneously with high cadence *XMM-Newton* X-ray and UV observations, the ultra-high frequency variability relationship between X-ray/Optical and UV/Optical emissions of NGC 4395 could be measured with enormous precision.

For the *XMM-Newton* campaign, one of the biggest limiting factors was the sparseness of the ground-based data compared with the space-based data. While an incredible achievement of logistics to arrange in the first place, future observations with more sampled g-band data would improve the X-ray to optical lag relationship significantly. Additionally more UV measurements would not be missed, specifically UVW1 which was unfortunately missed due to technical issues in this campaign and could connect this campaign's lags to previous ones that have used that filter. More U-band data would also be appreciated to test if the anomalously small lag measured in this study

can be reliably measured, and if so, determine the origin of this unexpectedly small lag. To improve reliability of the lag modelling, further studies should also attempt to find reliable estimates for the colour correction factor, as this is a very limiting parameter for reverberation modelling.

The power spectral study had many surprising features, one that most definitely warrants further investigation is the apparent presence of long-term variability in addition to X-ray reprocessing. More sophisticated modelling should be employed to attempt to find a physical origin of this variability. In addition, as mentioned above, extended ultra-high frequency HiPERCAM observations could help with the high-frequency bend detection while also allowing a better connection between it and the mid-frequency data. Continuing long-term monitoring of NGC 4395 to take the power spectra to even lower frequencies could also probe the low frequency regime more thoroughly, though admittedly the logarithmic nature of power spectra and the rather linear nature of human life makes extending this too much farther impractical.

NGC 4593 displayed surprising features in its X-ray to UV lags that seem to indicate timescale-dependent lags. If this feature is something physically real, as appears in the results study, then follow-up observations are a necessity to determine its physical origin and if similar dependencies can be detected in other AGN.

6.3 Final Remarks

Something apparent from the results presented here is that AGN variability is an extremely complex phenomenon. Though in theory very simple models can be used to describe the observed features of AGN, it is clear that these models do not capture all the nuances of the real objects, and there are many complex physical processes underlying these accreting black holes that require equally complex models to accurately describe. Methods of observation also present their own challenges, many AGN are extremely dim when viewed from Earth and so equipment and statistics introduce their own uncertainties and biases that cloud our view of these extreme objects even more. The more components of a unified AGN model we can assemble and accurately describe, the closer we can get to retrieving a full true view of the underlying processes that drive these systems and come ever closer to have a complete description of our Universe and all it contains.

Appendix A

Comparison of KYNxiltr lags with our in-house lags

In Fig. A.1 we present, for comparison, lags from the most recent version of our in-house code (Veresvarska). This version uses the [Novikov and Thorne \(1973\)](#) disc temperature prescription but does not include any General Relativistic ray tracing to compute path lengths in the very inner disc. The parameters used are almost identical to that used in Figs. 2.22 and 2.23 except that here we apply no colour correction (although the code does allow colour correction) and, by eye, a slightly larger truncation radius of $2000 R_g$ rather than 1700 is closest to the data. An albedo of 0.2 is used here whereas for Fig. 2.22 the fraction of illuminating radiation absorbed is calculated using detailed X-ray reflection modelling within the KYNxiltr code. The two model lags are very similar.

We also present, in Fig. A.2, model lags from the code which used in previous publications (e.g. [McHardy et al., 2018](#)). This code uses the [Shakura and Sunyaev \(1973\)](#) disc temperature prescription. Here a slightly smaller truncation radius, $1400 R_g$ and maximum spin provides a closer ‘by eye’ fit. The lags, in this model, for zero spin are 20 percent larger, thus are not enormously different to those shown. Again the model lags are close to those of KYNxiltr.

The aim here is not to carry out detailed comparison between the KYNxiltr code and our original in-house code but simply to show that the main result presented here, i.e. that modelling the lags as resulting from disc reprocessing requires a truncated disc, does not depend significantly on any particular modelling code.




FIGURE A.1: As for Fig. 2.22 but the model lags are from our original in-house code with zero spin.




FIGURE A.2: As for Fig. 2.22 but the model lags are from our original in-house code with maximum spin.

Appendix B

Cross-Correlation and Javelin Methodology

B.1 Javelin

JAVELIN, or Just Another Vehicle for Estimating Lags In Nuclei (stylised in this report as 'Javelin'), is a Python 2 program based on the authors' older FORTRAN-based program SPEAR. The motivating idea behind both of these programs is that AGN variability on long time-scales relative to the time-scale of fluctuations in the corona/disc can be described well by a Damped Random Walk model. For most AGN this time-scale would be on the scale of days, but for one as low-mass as NGC 4395 this time-scale is more on the scale of hours, hence our ~ 117 ks observations are more than sufficient for this assumption to hold.

Using this method and looking at our theory of how X-ray and longer wavelength emissions should be dependent on one another, we can then model the X-ray with as a Damped Random Walk continuum and the longer wavelengths as smoothed, scaled, and lagged versions of this continuum. The DRW continuum is generated using three parameters: A mean flux, a time-scale of variation τ , and an amplitude of variation σ . The mean is mostly irrelevant by the end as the scaling takes that into account, and so the two values that are used to define the continuum model are τ and σ . To transform this continuum in order to fit it to the longer wavelength light curves, we need three other parameters. These are the width of the top hat function w used to smooth the light curve, the scaling factor s used to scale the mean of the light curve, and the time lag t . Therefore, by fitting these models to the light curves we can determine the time lag between them.

When actually running Javelin, one first generates the continuum model in order to create posterior distributions of τ and σ . These posteriors are then used in order to

interpolate the light curves across any gaps that may exist in the data. This now complete model then shifts, smooths and lags in order to fit to the wavebands using a Markov Chain Monte Carlo (mcmc) method. After doing this a large number of times during the mcmc run, Javelin determines updated posteriors for σ and τ as well as posteriors for s , w , and most useful for us, the time lag t .

One issue that can arise in this process is that it is possible there may be a feature on two light curves that gives a false lag value consistently which can generate a secondary peak. Luckily this can be mostly resolved by applying Javelin to the X-ray, UV, and Optical light curves all at once, as the presence of a third light curve tends to eliminate these discrepancies.

B.2 Cross-Correlation Function

Cross-Correlation Function methods of lag determination have been in use for a long time, and while generally they are now considered to be less precise than Javelin's method, it is still useful to be able to show our light curves can produce similar results from an independent method, even if the values it produces are less precise.

The particular method we're drawing from is from [Peterson et al. \(2005\)](#). This program utilises two methods in order to try to avoid some issues that can arise with a quick CCF. The first is Flux Randomisation which randomly varies each data point's flux value within its error bars, this ensures that a pattern made by chance in a sequence of large error bar data points will not throw off the lag determination. The second is Random Subset Selection, this process means each time the CCF is performed, it is not on the entirety of the light curves, but on randomly determined subsections of each. This ensures that the lags determined are a consistent feature across the light curve and are not created by a single by-chance-correlated piece of the light curve.

In addition, we generate Discrete Correlation Function Confidence Contours using our own code. This method uses the X-ray and UV data light curves as inputs. It then generates a Power Spectral Distribution (PSD) from the X-ray light curve, and attempts to fit a bending power law model to it. With this model PSD generated, the code can then create synthetic X-ray light curves from it, and then generate Discrete Correlation Functions from these synthetic light curves and the input UV light curves.

By doing this many times with many different synthetic light curves (we used 1000 for each observation period), the program can then determine the likelihood of a particular lag value appearing. What this essentially means is that the higher our peak is above the confidence contours, the more likely it is caused by specific features on the two data light curves (such as a lagged reverberation) instead of just generated by the background noise of the X-ray power spectrum.

Appendix C

NGC 4593 Residual Model Plot




FIGURE C.1: NGC 4593 Residual data lags from the Max Spin Model, and the BLR lag shape from Korista & Goad 2019.

Initially, one might think that the maximum lag model underestimates the lag in the data in the 6000-8000Å range as there is an apparent lag excess. This is in fact the

expected part of the Paschen component of the BLR lag excesses we observe in this lag spectrum, though it is less visually apparent than the Balmer excess. To demonstrate this we have plotted the residual data lags from the Maximum Spin KYNreverb model and plotted them alongside the expected shape of BLR lag excesses as seen in [Korista and Goad \(2019\)](#), as can be seen in Fig. C.1.

The lag values themselves are different as the [Korista and Goad \(2019\)](#) plot is not for NGC 4593, however it is, more importantly, demonstrative of the shape expected from the BLR contribution. While the Balmer excess at $\sim 4000\text{\AA}$ is much more visually obvious, the Paschen excess should also be present, and so a lower spin model that had the model spectrum intersecting the data points above 6000\AA would likely be fit incorrectly, as it would be removing the Paschen component of the BLR contribution.

References

- R. R. J. Antonucci and J. S. Miller. Spectropolarimetry and the nature of ngc 1068. *The Astrophysical Journal*, 297:621, 1985.
- P. Arévalo and P. Uttley. Investigating a fluctuating-accretion model for the spectral-timing properties of accreting black hole systems. *Monthly Notices of the Royal Astronomical Society*, 367:801–814, 04 2006.
- P. Arévalo, P. Uttley, S. Kaspi, E. Breedt, P. Lira, and I. M. McHardy. Correlated x-ray/optical variability in the quasar mr 2251-178. *Monthly Notices of the Royal Astronomical Society*, 389(3):1479–1488, 2008.
- P. Arévalo, P. Uttley, P. Lira, E. Breedt, I. M. McHardy, and E. Churazov. Correlation and time delays of the x-ray and optical emission of the seyfert galaxy ngc 3783. *Monthly Notices of the Royal Astronomical Society*, 397(4):2004–2014, 2009.
- D. I. Ashton and M. J. Middleton. Searching for energy-resolved quasi-periodic oscillations in AGN. *Monthly Notices of the Royal Astronomical Society*, 501(4):5478–5499, jan 2021.
- J. M. Bardeen. Kerr metric black holes. *Nature*, 226(5240):64–65, 1970.
- A. Baskin and A. Laor. Dust inflated accretion disc as the origin of the broad line region in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 474(2):1970–1994, 2017.
- M. W. J. Beard, I. M. McHardy, K. K., D. D. C., P. I., B. D., S. K. P., K. D., and P. M. Timescale-dependent x-ray to uv time lags of ngc 4593 using high-intensity XMM-Newton observations with *Swift* and *AstroSat*. *MNRAS*, 2022.
- M. C. Bentz, K. D. Denney, C. J. Grier, A. J. Barth, B. M. Peterson, M. Vestergaard, V. N. Bennert, G. Canalizo, G. De Rosa, and A. V. e. a. Filippenko. The low-luminosity end of the radius-luminosity relationship for active galactic nuclei. *The Astrophysical Journal*, 767(2):149, 2013.
- A. J. Berkley, D. Kazanas, and J. Ozik. Modeling the x-ray ultraviolet correlations in ngc 7469. *The Astrophysical Journal*, 535(2):712–720, 2000.

- R. D. Blandford and C. F. McKee. Reverberation mapping of the emission line regions of seyfert galaxies and quasars. *The Astrophysical Journal*, 255:419, 1982.
- H. Bondi. On spherically symmetrical accretion. *Monthly Notices of the Royal Astronomical Society*, 112(2):195–204, 1952.
- E. Breedt. PhD thesis, University of Southampton, 2010.
- E. Breedt, P. Arévalo, I. M. McHardy, P. Uttley, S. G. Sergeev, T. Minezaki, Y. Yoshii, C. M. Gaskell, E. M. Cackett, and K. e. a. Horne. Long-term optical and x-ray variability of the seyfert galaxy markarian 79. *Monthly Notices of the Royal Astronomical Society*, 394(1):427–437, 2009.
- E. Breedt, I. M. McHardy, P. Arévalo, P. Uttley, S. G. Sergeev, T. Minezaki, Y. Yoshii, Y. Sakata, P. Lira, and N. G. Chesnok. Twelve years of x-ray and optical variability in the seyfert galaxy ngc 4051. *Monthly Notices of the Royal Astronomical Society*, 403(2):605–619, 2010.
- M. d. Brok, A. C. Seth, A. J. Barth, D. J. Carson, N. Neumayer, M. Cappellari, V. P. Debattista, L. C. Ho, C. E. Hood, and R. M. McDermid. Measuring the mass of the central black hole in the bulgeless galaxy ngc 4395 from gas dynamical modeling. *The Astrophysical Journal*, 809(1):101, 2015.
- T. M. Brown, N. Baliber, F. B. Bianco, M. Bowman, B. Burleson, P. Conway, M. Crellin, E. Depagne, J. De Vera, and B. e. a. Dilday. Las cumbres observatory global telescope network. *Publications of the Astronomical Society of the Pacific*, 125(931):1031–1055, 2013.
- C. Brum, M. R. Diniz, R. A. Riffel, A. Rodríguez-Ardila, L. C. Ho, R. Riffel, R. Mason, L. Martins, A. Petric, and R. Sánchez-Janssen. A close look at the dwarf agn of ngc?4395: optical and near-ir integral field spectroscopy. *Monthly Notices of the Royal Astronomical Society*, 486(1):691–707, 2019.
- C. J. Burke, Y. Shen, Y.-C. Chen, S. Scaringi, C.-A. Faucher-Giguere, X. Liu, and Q. Yang. Optical variability of the dwarf agn ngc 4395 from the transiting exoplanet survey satellite. *The Astrophysical Journal*, 899(2):136, 2020.
- E. Cackett, C.-Y. Chiang, I. McHardy, R. Edelson, M. Goad, K. Horne, and K. Korista. Accretion disk reverberation with hubble space telescope observations of ngc 4593:evidence for diffuse continuum lags. *The Astrophysical Journal*, 857(1):53, 2018.
- E. M. Cackett, K. Horne, and H. Winkler. Testing thermal reprocessing in active galactic nuclei accretion discs. *Monthly Notices of the Royal Astronomical Society*, 380:669–682, 08 2007.
- E. M. Cackett, A. Zoghbi, C. Reynolds, A. C. Fabian, E. Kara, P. Uttley, and D. R. Wilkins. Modelling the broad fe ka reverberation in the agn ngc 4151. *Monthly Notices of the Royal Astronomical Society*, 438(4):2980–2994, 2014.

- E. M. Cackett, J. Gelbord, Y.-R. Li, K. Horne, J.-M. Wang, A. J. Barth, J.-M. Bai, W.-H. Bian, R. W. Carroll, and P. e. a. Du. Supermassive black holes with high accretion rates in active galactic nuclei. xi. accretion disk reverberation mapping of mrk 142. *The Astrophysical Journal*, 896(1):1, 2020.
- D. Cameron. PhD thesis, University of Southampton, 2014.
- D. T. Cameron, I. McHardy, T. Dwelly, E. Breedt, P. Uttley, P. Lira, and P. Arévalo. Correlated x-ray/ultraviolet/optical variability in the very low mass agn ngc 4395. *Monthly Notices of the Royal Astronomical Society*, 422:902–912, 03 2012.
- R. Carr, D. Cinabro, E. Cackett, D. Moutard, and R. Carroll. Wayne state university’s dan zowada memorial observatory: Characterization and pipeline of a 0.5 m robotic telescope. *Publications of the Astronomical Society of the Pacific*, 134(1034):045002, 2022.
- G. Chartas, C. S. Kochanek, X. Dai, D. Moore, A. M. Mosquera, and J. A. Blackburne. Revealing the structure of an accretion disk through energy-dependent x-ray microlensing. *The Astrophysical Journal*, 757(2):137, 2012.
- L. Coatman, P. C. Hewett, M. Banerji, and G. T. Richards. C^{iv} emission-line properties and systematic trends in quasar black hole mass estimates. *Monthly Notices of the Royal Astronomical Society*, 461(1):647–665, 2016.
- S. Collin-Souffrin and A. Dumont. Line emission from accretion disks in active galactic nuclei. ii. radial structure of the disc. *Astronomy & Astrophysics*, 229:292, 1990.
- S. Collins and J.-M. Huré. Influence of the metallicity and of the irradiation on the structure of accretion disks around massive black holes. *Astronomy & Astrophysics*, 341:385, 1999.
- S. Connolly. A python code for the emmanoulopoulos et al. [arxiv:1305.0304] light curve simulation algorithm. 2015. URL <https://arxiv.org/abs/1503.06676>.
- X. Dai, C. S. Kochanek, G. Chartas, S. Kozłowski, C. W. Morgan, G. Garmire, and E. Agol. The sizes of the x-ray and optical emission regions of rxj 1131-1231. *The Astrophysical Journal*, 709(1):278–285, 2009.
- S. W. Davis and S. El-Abd. Spectral hardening in black hole accretion: Giving spectral modelers an f. *The Astrophysical Journal*, 874(1):23, 2019.
- J. E. Deeter and P. E. Boynton. Techniques for the estimation of red power spectra. i - context and methodology. *The Astrophysical Journal*, 261:337, 1982.
- V. Dhillon and T. Marsh. Hipercam - resources. URL <http://www.vikdhillon.staff.shef.ac.uk/hipercam/resources.html>.
- T. di Matteo. Magnetic reconnection: flares and coronal heating in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 299(1):L15–L20, 1998.

- C. Done, G. M. Madejski, R. F. Mushotzky, T. J. Turner, K. Koyama, and H. Kunieda. The x-ray variability of ngc 6814 - power spectrum. *The Astrophysical Journal*, 400: 138, 1992.
- C. Done, S. W. Davis, C. Jin, O. Blaes, and M. Ward. Intrinsic disc emission and the soft x-ray excess in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 420(3):1848–1860, 2012.
- M. Dovčiak. Stronggravity - kynreverb. URL <https://projects.asu.cas.cz/stronggravity/kynreverb>.
- R. Edelson and K. Nandra. A cutoff in the x-ray fluctuation power density spectrum of the seyfert 1 galaxy ngc 3516. *The Astrophysical Journal*, 514(2):682–690, 1999.
- R. Edelson, J. M. Gelbord, K. Horne, I. M. McHardy, B. M. Peterson, P. Arévalo, A. A. Breeveld, G. D. Rosa, P. A. Evans, M. R. Goad, G. A. Kriss, W. N. Brandt, N. Gehrels, D. Grupe, J. A. Kennea, C. S. Kochanek, J. A. Nousek, I. Papadakis, M. Siegel, D. Starkey, P. Uttley, S. Vaughan, S. Young, A. J. Barth, M. C. Bentz, B. J. Brewer, D. M. Crenshaw, E. Dalla Bontà, A. D. L. Cáceres, K. D. Denney, M. Dietrich, J. Ely, M. M. Fausnaugh, C. J. Grier, P. B. Hall, J. Kaastra, B. C. Kelly, K. T. Korista, P. Lira, S. Mathur, H. Netzer, A. Pancoast, L. Pei, R. W. Pogge, J. S. Schimoia, T. Treu, M. Vestergaard, C. Villforth, H. Yan, and Y. Zu. Space telescope and optical reverberation mapping project. ii.swiftandhstreverberation mapping of the accretion disk of ngc 5548. *The Astrophysical Journal*, 806:129, 06 2015.
- R. Edelson, J. Gelbord, E. Cackett, S. Connolly, C. Done, M. Fausnaugh, E. Gardner, N. Gehrels, M. Goad, K. Horne, I. McHardy, B. M. Peterson, S. Vaughan, M. Vestergaard, A. Breeveld, A. J. Barth, M. Bentz, M. Bottorff, W. N. Brandt, S. M. Crawford, E. D. Bontà, D. Emmanoulopoulos, P. Evans, R. F. Jaimes, A. V. Filippenko, G. Ferland, D. Grupe, M. Joner, J. Kennea, K. T. Korista, H. A. Krimm, G. Kriss, D. C. Leonard, S. Mathur, H. Netzer, J. Nousek, K. Page, E. Romero-Colmenero, M. Siegel, D. A. Starkey, T. Treu, H. A. Vogler, H. Winkler, and W. Zheng. Swiftmonitoring of ngc 4151: Evidence for a second x-ray /uv reprocessing. *The Astrophysical Journal*, 840:41, 05 2017.
- R. Edelson, J. Gelbord, E. Cackett, and B. M. Peterson. The first swift intensive agn accretion disk reverberation mapping survey. *ApJ*, 870(123), 2019.
- R. A. Edelson and J. H. Krolik. The discrete correlation function - a new method for analyzing unevenly sampled variability data. *The Astrophysical Journal*, 333:646, 1988.
- H. Edri, S. E. Rafter, D. Chelouche, S. Kaspi, and E. Behar. Broadband photometric reverberation mapping of ngc 4395. *The Astrophysical Journal*, 756(1):73, 2012.
- M. Elvis. A structure for quasars. *The Astrophysical Journal*, 545(1):63–76, 2000.

- D. Emmanoulopoulos, I. E. Papadakis, F. Nicastro, and I. M. McHardy. X-ray spectral analysis of the low-luminosity active galactic nucleus ngc 7213 using long xmm-newton observations? *Monthly Notices of the Royal Astronomical Society*, 429(4): 3439–3448, 2013.
- D. Emmanoulopoulos, I. E. Papadakis, M. Dovciak, and I. M. McHardy. General relativistic modelling of the negative reverberation x-ray time delays in agn? *Monthly Notices of the Royal Astronomical Society*, 439(4):3931–3950, 2014.
- B. L. Fanaroff and J. M. Riley. The morphology of extragalactic radio sources of high and low luminosity. *Monthly Notices of the Royal Astronomical Society*, 167(1):31P–36P, 1974.
- M. M. Fausnaugh, K. D. Denney, A. J. Barth, M. C. Bentz, M. C. Bottorff, M. T. Carini, K. V. Croxall, G. D. Rosa, M. R. Goad, and K. e. a. Horne. Space telescope and optical reverberation mapping project. iii. optical continuum emission and broadband time delays in ngc 5548. *The Astrophysical Journal*, 821(1):56, 2016.
- A. V. Filippenko, L. C. Ho, and W. L. W. Sargent. Hst observations of ngc 4395, the least luminous seyfert 1 nucleus - evidence against the starburst hypothesis for broad-lined active galactic nuclei. *The Astrophysical Journal*, 410:L75, 1993.
- C. F. Gammie. Nonlinear outcome of gravitational instability in cooling, gaseous disks. *The Astrophysical Journal*, 553(1):174–183, 2001.
- J. A. García, E. Kara, D. Walton, T. Beuchert, T. Dauser, E. Gatzert, M. Balokovic, J. F. Steiner, F. Tombesi, and R. M. T. e. a. Connors. Implications of the warm corona and relativistic reflection models for the soft excess in mrk 509. *The Astrophysical Journal*, 871(1):88, 2019.
- E. Gardner and C. Done. The origin of the uv/optical lags in ngc 5548. *Monthly Notices of the Royal Astronomical Society*, 470:3591–3605, 04 2017.
- A. R. Green, I. M. McHardy, and H. J. Lehto. On the nature of rapid x-ray variability in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 265(3): 664–680, 1993.
- J. V. Hernández Santisteban, R. Edelson, K. Horne, J. M. Gelbord, A. J. Barth, E. M. Cackett, M. R. Goad, H. Netzer, D. Starkey, and P. e. a. Uttley. Intensive disc-reverberation mapping of fairall 9: first year of swift and lco monitoring. *Monthly Notices of the Royal Astronomical Society*, 498(4):5399–5416, 2020.
- N. Higginbottom, D. Proga, C. Knigge, K. S. Long, J. H. Matthews, and S. A. Sim. Line-driven disk winds in active galactic nuclei: The critical importance of ionization and radiative transfer. *The Astrophysical Journal*, 789(1):19, 2014.

- L. C. Ho, A. V. Filippenko, and W. L. W. Sargent. A search for 'dwarf' seyfert nuclei. 2: an optical spectral atlas of the nuclei of nearby galaxies. *The Astrophysical Journal Supplement Series*, 98(2):555, 1995.
- L. C. Ho, A. V. Filippenko, and W. L. W. Sargent. A search for "dwarf" seyfert nuclei. iii. spectroscopic parameters and properties of the host galaxies. *The Astrophysical Journal Supplement Series*, 112(2):315–390, 1997.
- S. F. Hönig. Redefining the torus: A unifying view of agns in the infrared and submillimeter. *The Astrophysical Journal*, 884(2):171, 2019.
- K. Horne, B. M. Peterson, S. J. Collier, and H. Netzer. Observational requirements for high-fidelity reverberation mapping. *Publications of the Astronomical Society of the Pacific*, 116(819):465–476, 2004.
- A. R. Ingram and S. E. Motta. A review of quasi-periodic oscillations from black hole x-ray binaries: Observation and theory. *New Astronomy Reviews*, 85:101524, sep 2019.
- E. S. Kammoun, E. Nardini, A. Zoghbi, J. M. Miller, E. M. Cackett, E. Gallo, M. T. Reynolds, G. Risaliti, D. Barret, and W. N. e. a. Brandt. The nature of the broadband x-ray variability in the dwarf seyfert galaxy ngc 4395. *The Astrophysical Journal*, 886(2):145, 2019.
- E. S. Kammoun, M. Dovciak, I. E. Papadakis, M. D. Caballero-Garcia, and V. Karas. Uv/optical disk thermal reverberation in active galactic nuclei: An in-depth study with an analytic prescription for time-lag spectra. *The Astrophysical Journal*, 907(1):20, 2021a.
- E. S. Kammoun, I. E. Papadakis, and M. Dovciak. Modelling the uv/optical continuum time-lags in agn. *Monthly Notices of the Royal Astronomical Society*, 503(3):4163–4171, 2021b.
- E. Kara, M. Mehdipour, G. A. Kriss, E. M. Cackett, N. Arav, A. J. Barth, D. Byun, M. S. Brotherton, G. De Rosa, and J. e. a. Gelbord. Agn storm 2. i. first results: A change in the weather of mrk 817. *The Astrophysical Journal*, 922(2):151, 2021.
- A. L. King, J. M. Miller, M. T. Reynolds, K. Gültekin, E. Gallo, and D. Maitra. A distinctive disk-jet coupling in the lowest-mass seyfert, ngc 4395. *The Astrophysical Journal*, 774(2):L25, 2013.
- A. R. King, J. Frank, and D. J. Raine. *Accretion power in astrophysics*. Cambridge University Press, 2002.
- W. Kollatschny and K. J. Fricke. Influence of the metallicity and of the irradiation on the structure of accretion disks around massive black holes. *Astronomy & Astrophysics*, 143:396, 1985.

- K. T. Korista and M. R. Goad. The variable diffuse continuum emission of broad-line clouds. *The Astrophysical Journal*, 553(2):695–708, 2001.
- K. T. Korista and M. R. Goad. Quantifying the impact of variable blr diffuse continuum contributions on measured continuum interband delays. *Monthly Notices of the Royal Astronomical Society*, 489(4):5284–5300, 2019.
- A. Laor and H. Netzer. Massive thin accretion discs - i. calculated spectra. *Monthly Notices of the Royal Astronomical Society*, 238(3):897–916, 1989.
- A. Lawrence, M. G. Watson, K. A. Pounds, and M. Elvis. Low-frequency divergent x-ray variability in the seyfert galaxy ngc4051. *Nature*, 325(6106):694–696, 1987.
- D. Lawther, M. R. Goad, K. T. Korista, O. Ulrich, and M. Vestergaard. Quantifying the diffuse continuum contribution of blr clouds to agn continuum inter-band delays. *Monthly Notices of the Royal Astronomical Society*, 481(1):533–554, 2018.
- Y.-R. Li, J.-M. Wang, C. Hu, P. Du, and J.-M. Bai. A bayesian method for the intercalibration of spectra in reverberation mapping. *The Astrophysical Journal*, 786: L6, 04 2014.
- A. P. Lightman and A. A. Zdziarski. Pair production and compton scattering in compact sources and comparison to observations of active galactic nuclei. *The Astrophysical Journal*, 319:643, 1987.
- P. Lira, A. Lawrence, P. O’Brien, R. A. Johnson, R. Terlevich, and N. Bannister. Optical and x-ray variability in the least luminous active galactic nucleus, ngc 4395. *Monthly Notices of the Royal Astronomical Society*, 305(1):109–124, 1999.
- P. Lira, P. Arévalo, P. Uttley, I. McHardy, and E. Breedt. Optical and near-ir long-term monitoring of ngc 3783 and mr 2251-178: evidence for variable near-ir emission from thin accretion discs. *Monthly Notices of the Royal Astronomical Society*, 415(2): 1290–1303, 2011.
- A. Lobban and A. King. Agn light echoes and the accretion disc self-gravity limit. *Monthly Notices of the Royal Astronomical Society*, 511(2):1992–1998, 2022.
- Y. E. Lyubarskii. Flicker noise in accretion discs. *Monthly Notices of the Royal Astronomical Society*, 292(3):679–685, 1997.
- F. J. Masci, R. R. Laher, B. Rusholme, D. L. Shupe, S. Groom, J. Surace, E. Jackson, S. Monkevitz, R. Beck, and D. e. a. Flynn. The zwicky transient facility: Data processing, products, and archive. *Publications of the Astronomical Society of the Pacific*, 131(995):018003, 2018.
- J. McClintock and R. Remillard. Black hole binaries. *astro-ph/0306213*, 2003.

- I. McHardy and B. Czerny. Fractal x-ray time variability and spectral invariance of the seyfert galaxy ngc5506. *Nature*, 325(6106):696–698, 1987.
- I. M. McHardy. *On two topics in X-ray astronomy*. European Space Agency, 1989.
- I. M. McHardy, D. T. Cameron, T. Dwelly, S. Connolly, P. Lira, D. Emmanoulopoulos, J. Gelbord, E. Breedt, P. Ar’evalo, and P. Uttley. Swift monitoring of ngc 5548: X-ray reprocessing and short-term uv/optical variability. *Monthly Notices of the Royal Astronomical Society*, 444:1469–1474, 08 2014.
- I. M. McHardy, S. D. Connolly, B. M. Peterson, A. Bieryla, H. Chand, M. S. Elvis, D. Emmanoulopoulos, E. Falco, P. Gandhi, S. Kaspi, D. Latham, P. Lira, C. McCully, H. Netzer, and M. Uemura. The origin of uv-optical variability in agn and test of disc models: Xmm-newton and ground-based observations of ngc 4395. *Astronomische Nachrichten*, 337:500–506, 05 2016.
- I. M. McHardy, S. D. Connolly, K. Horne, E. M. Cackett, J. Gelbord, B. M. Peterson, M. Pahari, N. Gehrels, M. Goad, and P. e. a. Lira. X-ray/uv/optical variability of ngc 4593 with swift: reprocessing of x-rays by an extended reprocessor. *Monthly Notices of the Royal Astronomical Society*, 480(3):2881–2897, 2018.
- I. M. McHardy, M. W. J. Beard, E. Breedt, J. H. Knapen, and F. M. Vincentelli. First detection of the outer edge of an agn accretion disc: Very fast multiband optical variability of ngc 4395 with gtc/hipercam and lt/io:0. *MNRAS*, 2022.
- E. C. Moran, A. V. Filippenko, L. C. Ho, J. C. Shields, T. Belloni, A. Comastri, S. L. Snowden, and R. A. Sramek. The nuclear spectral energy distribution of ngc 4395, the least luminous type 1 seyfert galaxy. *Publications of the Astronomical Society of the Pacific*, 111(761):801–808, 1999.
- C. W. Morgan, C. S. Kochanek, N. D. Morgan, and E. E. Falco. The quasar accretion disk size-black hole mass relation. *The Astrophysical Journal*, 712:1129–1136, 03 2010.
- A. M. Mosquera, C. S. Kochanek, B. Chen, X. Dai, J. A. Blackburne, and G. Chartas. The structure of the x-ray and optical emitting regions of the lensed quasar q 2237+0305. *The Astrophysical Journal*, 769(1):53, 2013.
- K. Nandra and I. E. Papadakis. Temporal characteristics of the x-ray emission of ngc 7469. *The Astrophysical Journal*, 554(2):710–724, 2001.
- K. Nandra, J. Clavel, R. A. Edelson, I. M. George, M. A. Malkan, R. F. Mushotzky, B. M. Peterson, and T. J. Turner. New constraints on the continuum emission mechanism of active galactic nuclei: Intensive monitoring of ngc 7469 in the x-ray and ultraviolet. *The Astrophysical Journal*, 505(2):594–606, 1998.
- R. Narayan and I. Yi. Advection-dominated accretion: A self-similar solution. *The Astrophysical Journal*, 428:L13, 1994.

- H. Netzer. Continuum reverberation mapping and a new lag-luminosity relationship for agn. *Monthly Notices of the Royal Astronomical Society*, 2021.
- I. D. Novikov and K. S. Thorne. *Black Holes (Les Astres Occlus)*. 1973.
- C. A. Onken, S. Lai, C. Wolf, A. B. Lucy, W. J. Hon, P. Tisserand, J. L. Sokoloski, G. J. M. Luna, R. Manick, X. Fan, and F. Bian. Discovery of the most luminous quasar of the last 9 gyr, 2022.
- D. E. Osterbrock. Spectrophotometry of seyfert 1 galaxies. *The Astrophysical Journal*, 215:733, 1977.
- M. Pahari, I. M. McHardy, F. Vincentelli, E. Cackett, B. M. Peterson, M. Goad, K. Gültekin, and K. Horne. Evidence for variability time-scale-dependent uv/x-ray delay in seyfert 1 agn ngc 7469. *Monthly Notices of the Royal Astronomical Society*, 494(3):4057–4068, 2020.
- I. E. Papadakis and A. Lawrence. Improved methods for power spectrum modelling of red noise. *Monthly Notices of the Royal Astronomical Society*, 261(3):612–624, 1993.
- B. M. Peterson, I. Wanders, K. Horne, S. Collier, T. Alexander, S. Kaspi, and D. Maoz. On uncertainties in cross-correlation lags and the reality of wavelength-dependent continuum lags in active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 110(748):660–670, 1998.
- B. M. Peterson, M. C. Bentz, L. Desroches, A. V. Filippenko, L. C. Ho, S. Kaspi, A. Laor, D. Maoz, E. C. Moran, R. W. Pogge, and A. C. Quillen. Multiwavelength monitoring of the dwarf seyfert 1 galaxy ngc 4395. i. a reverberation-based measurement of the black hole mass. *The Astrophysical Journal*, 632:799–808, 10 2005.
- P.-O. Petrucci, F. Ursini, A. De Rosa, S. Bianchi, M. Cappi, G. Matt, M. Dadina, and J. Malzac. Testing warm comptonization models for the origin of the soft x-ray excess in agns. *Astronomy & Astrophysics*, 611:A59, 2018.
- D. Proga and T. R. Kallman. Dynamics of line-driven disk winds in active galactic nuclei. ii. effects of disk radiation. *The Astrophysical Journal*, 616(2):688–695, 2004.
- C. Ricci. *Active Galactic Nuclei at hard X-ray energies: Absorption, Reflection and the Unified Model*. PhD thesis, Université de Genève, 2011.
- R. R. Ross, A. C. Fabian, and S. Mineshige. The spectra of accretion discs in active galactic nuclei. *Monthly Notices of the Royal Astronomical Society*, 258(1):189–197, 1992.
- S. G. Sergeev, V. T. Doroshenko, Y. V. Golubinskiy, N. I. Merkulova, and E. A. Sergeeva. Lag-luminosity relationship for interband lags between variations in b,v,r, and i bands in active galactic nuclei. *The Astrophysical Journal*, 622(1):129–135, 2005.

- N. I. Shakura and R. A. Sunyaev. Black holes in binary systems: Observational appearances. *Symposium - International Astronomical Union*, 55:155–164, 1973.
- B. J. Shappee, J. L. Prieto, D. Grupe, C. S. Kochanek, K. Z. Stanek, G. De Rosa, S. Mathur, Y. Zu, B. M. Peterson, and R. W. e. a. Pogge. The man behind the curtain: X-rays drive the uv through nir variability in the 2013 active galactic nucleus outburst in ngc 2617. *The Astrophysical Journal*, 788(1):48, 2014.
- S. N. Shore and R. L. White. Self-gravitating accretion disk models for active galactic nuclei - self-consistent alpha-models for the broad emission-line region. *The Astrophysical Journal*, 256:390, 1982.
- K. P. Singh, G. C. Stewart, S. Chandra, K. Mukerjee, S. Kotak, A. P. Beardmore, V. Chitnis, G. C. Dewangan, S. Bhattacharyya, and I. e. a. Mirza. In-orbit performance of sxt aboard astrosat. *SPIE Proceedings*, 2016.
- K. P. Singh, G. C. Stewart, N. J. Westergaard, S. Bhattacharayya, S. Chandra, V. R. Chitnis, G. C. Dewangan, A. T. Kothare, I. M. Mirza, and K. e. a. Mukerjee. Soft x-ray focusing telescope aboard astrosat: Design, characteristics and performance. *Journal of Astrophysics and Astronomy*, 38(2), 2017.
- I. A. Steele, R. J. Smith, P. C. Rees, I. P. Baker, S. D. Bates, M. F. Bode, M. K. Bowman, D. Carter, J. Etherton, and M. J. e. a. Ford. The liverpool telescope: performance and first results. *SPIE Proceedings*, 2004.
- M. Suganuma, Y. Yoshii, Y. Kobayashi, T. Minezaki, K. Enya, H. Tomita, T. Aoki, S. Koshida, and B. A. Peterson. Reverberation measurements of the inner radius of the dust torus in nearby seyfert 1 galaxies. *The Astrophysical Journal*, 639(1):46–63, 2006.
- S. N. Tandon, A. Subramaniam, V. Girish, J. Postma, K. Sankarasubramanian, S. Sriram, C. S. Stalin, C. Mondal, S. Sahu, and P. e. a. Joseph. In-orbit calibrations of the ultraviolet imaging telescope. *The Astronomical Journal*, 154(3):128, 2017.
- S. N. Tandon, J. Postma, P. Joseph, A. Devaraj, A. Subramaniam, I. V. Barve, K. George, S. K. Ghosh, V. Girish, and J. B. e. a. Hutchings. Additional calibration of the ultraviolet imaging telescope on board astrosat. *The Astronomical Journal*, 159(4):158, 2020.
- J. Timmer and M. König. On generating power law noise. *Astronomy and Astrophysics*, 300:707, 1995.
- A. Treves, L. Maraschi, and M. Abramowicz. Basic elements of the theory of accretion. *Publications of the Astronomical Society of the Pacific*, 100:427, 1988.
- J. Troyer, D. Starkey, E. M. Cackett, M. C. Bentz, M. R. Goad, K. Horne, and J. E. Seals. Correlated x-ray/ultraviolet/optical variability in ngc 6814. *Monthly Notices of the Royal Astronomical Society*, 456(4):4040–4050, 2016.

- R. B. Tully, L. Rizzi, E. J. Shaya, H. M. Courtois, D. I. Makarov, and B. A. Jacobs. The extragalactic distance database. *The Astronomical Journal*, 138(2):323–331, 2009.
- F. Ursini, P.-O. Petrucci, S. Bianchi, G. Matt, R. Middei, G. Marcel, J. Ferreira, M. Cappi, B. De Marco, and A. e. a. De Rosa. Nustar/xmm-newton monitoring of the seyfert 1 galaxy he 1143-1810. *Astronomy & Astrophysics*, 634:A92, 2020.
- P. Uttley, I. M. McHardy, and I. E. Papadakis. Measuring the broad-band power spectra of active galactic nuclei with rxte. *Monthly Notices of the Royal Astronomical Society*, 332(1):231–250, 2002.
- P. Uttley, R. Edelson, I. M. McHardy, B. M. Peterson, and A. Markowitz. Correlated long-term optical and x-ray variations in ngc 5548. *The Astrophysical Journal*, 584(2):L53–L56, 2003.
- S. Vaughan, K. Iwasawa, A. C. Fabian, and K. Hayashida. The exceptional x-ray variability of the dwarf seyfert nucleus ngc 4395. *Monthly Notices of the Royal Astronomical Society*, 356(2):524–530, 2005.
- F. M. Vincentelli, I. McHardy, E. M. Cackett, A. J. Barth, K. Horne, M. Goad, K. Korista, J. Gelbord, W. Brandt, and R. e. a. Edelson. On the multiwavelength variability of mrk 110: two components acting at different time-scales. *Monthly Notices of the Royal Astronomical Society*, 504(3):4337–4353, 2021.
- F. M. Vincentelli, I. McHardy, J. V. Hernandez Santisteban, E. M. Cackett, J. Gelbord, K. Horne, J. A. Miller, and A. Lobban. The luminosity-dependent contribution from the broad-line region to the wavelength-dependent lags in mrk 110. *Monthly Notices of the Royal Astronomical Society: Letters*, 512(1):L33–L38, 2022.
- W. Welsh. On the reliability of cross-correlation function lag determinations in active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 111(765):1347–1366, 1999.
- J.-H. Woo, H. Cho, E. Gallo, E. Hodges-Kluck, H. A. N. Le, J. Shin, D. Son, and J. C. Horst. A 10,000-solar-mass black hole in the nucleus of a bulgeless dwarf galaxy. *Nature Astronomy*, 3(8):755–759, 2019.
- Y. Zu, C. S. Kochanek, and B. M. Peterson. An alternative approach to measuring reverberation lags in active galactic nuclei. *The Astrophysical Journal*, 735(2):80, 2011.
- Y. Zu, C. S. Kochanek, S. Kozłowski, and A. Udalski. Is quasar optical variability a damped random walk? *The Astrophysical Journal*, 765:106, 02 2013.