From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

Doctoral Dissertations, 2011-2026 Astrophysical Sciences

2024

Reverberation Mapping Active Galactic Nuclei Accretion Disks Via Observational and Theoretical Routes

Secunda, Amy

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp019019s586q

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Reverberation Mapping Active Galactic Nuclei Accretion Disks Via Observational and Theoretical Routes

Amy Secunda

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Astrophysical Sciences
Advisers: Professor Jenny Greene and Dr. Yan-Fei Jiang

September 2024

© Copyright by Amy Secunda, 2024.

All rights reserved.

Abstract

Due to their great distance it is generally not possible to spatially resolve an active galactic nuclei (AGN) accretion disk, making it difficult to model their structure. However, modeling AGN disks is crucial to our understanding of accretion physics, and the internal physics of these disks are important because they fuel powerful radiation and outflows that can have an enormous impact on galaxy evolution. Thankfully, then, the UV and optical light curves emitted by AGN disks vary over a range of timescales, and as a result, temporal resolution can often be substituted for spatial resolution. One of the most prominent techniques that takes advantage of this temporal resolution, and the focus of my dissertation, is reverberation mapping. Traditional reverberation mapping uses temporal lags between the variability in light curves observed in different wavebands to measure the radial temperature profile and radial extent of AGN disks. This method has been used for decades, however, the response function of this reprocessing and how it can be affected by local variability in the disk are not well understood. Multi-frequency radiation magnetohydrodynamic codes make it possible to directly simulate the reprocessing of light by an AGN disk. In this dissertation I present the results from new radiation Athena++ simulations and discuss how they can be directly compared to observations to better understand the source of variability in AGN light curves and what they can tell us about the internal physics of the AGN disk. I will also discuss a second type of lag, the "long negative lag," which unlike the lag measured in traditional continuum reverberation mapping, depends on the inflow rate. Because long negative lags depend on the inflow rate, we can use these lags to model the vertical structure of AGN disks, a fundamental test of the commonly used standard thin disk model.

Acknowledgements

This journey began nearly ten years ago when I decided to upend my life and go from being an actor to an astrophysicist. Here is the place where I get to thank everyone who has made this journey possible.

First, I would like to thank my advisers Dr. Jenny Greene and Dr. Yan-Fei Jiang. Within ten minutes of sitting down in Jenny's office to discuss potentially working with her for my thesis, they had both laid out a basic outline for the project that turned into the dissertation you are about to read. We held that rapid fire pace for the rest of my thesis so that I am now able to defend a mere two years later. Their hard work and dedication to me pursuing my dream extended seven days a week and twenty-four hours a day, from chats before I got my morning coffee at Rojo's to live editing my paper drafts back and forth for hours on days that they had to be submitted. Jenny and Yan-Fei had no reason to take me in, but I knew from the minute we started that I was going to get my PhD and I am infinitely grateful to them.

I am very thankful for the amazing faculty in Peyton Hall. Dr. Jim Stone (who was here in Peyton Hall when I started) was the first professor to take me on as a semester project student and the first to teach me the ins and outs of Athena++, one of the primary tools of my thesis. Dr. Renyue Cen was an excellent semester project advisor and co-mentor for Alexander Menegas and Elizabeth Berzin, who were both terrific to work with. I also had the privilege of working on an experimental project with Dr. Hantao Ji, which was one of the highlights of my time at Princeton. Finally, I learned an immense amount from Dr. Jeremy Goodman, and I am proud of the two papers we wrote together. I appreciate the way Jeremy looked out for me and treated me with great kindness when things were not going our way in research. I am deeply grateful for the advice I received from all of my research mentors at Princeton on not just science but my career.

A big thank you to my thesis committee including my thesis readers, Dr. Jeremy Goodman and Dr. Jim Stone, as well as Dr. Michael Strauss, Dr. Jenny Greene, and Dr. Yan-Fei

Jiang. I appreciate your thoughtful feedback throughout my thesis. Thank you also to directors of graduate study, Dr. Gaspar Bakos and Dr. Josh Winn, for helpful conversations, good advice, and kind words.

Thank you to the many postdoctoral fellows I have gotten to meet here in Princeton. Thank you to the Monday running group and the afternoon #just-coffee group. A special thank you to Dr. Neven Caplar, Dr. Yusra AlSaayad, Dr. Patrick Mullen, and Dr. Charlotte Ward. Thank you also to the very kind and helpful Peyton staff, Mindy Lipman, Mami Akiyama, Susan Duncan, Leigh Koven, Stephanie Reif, Greg Motta, Steve Oakley, and Polly Strauss. Polly is such a light in the dark basement of Peyton. I am honored to have gotten to spend so much time with her.

To all of the graduated students I have shared the basement of Peyton Hall with, thank you. You have genuinely all touched my experience over the past six years. A special and most heartfelt thank you to my officemates and close friends Heather Prince, Brianna Lacy, Mackenzie Moody, Roohi Dalal, and Lachlan Lancaster, my pandemic officemate that I would discuss astrophysics with around our kitchen counter. Thank you to past and present office neighbors, Christopher Bambic, Rodrigo Córdova, Roberto Tejada Arevalo, and Nora Linzer. Thank you to my hosts when I first came to visit Princeton, Goni Halevi and Erin "I like living in Princeton" Flowers, who have remained close friends. Thank you also to the four boys I studied for General Exams with in Spring 2020 Miles Cranmer, Samuel Yee, Oliver Philcox, and Eric Moseley. Thank you to the lower year graduate students whom I enjoyed mentoring and getting to know these past few years, Robel Geda, Eden Girma, Camryn Phillips, Amir Siraj, and Philippe Yao, who helped me launch my thesis and drove all the way to White Plains with me to watch the Survivor 45 Finale. Finally, thank you to Erin Kado-Fong and Kris Pardo without whom I never would have knocked on Jenny's door two and a half years ago.

Thank you to everyone who has been a huge influence to me along my journey. My first research advisors Dr. Kathryn Johnston and Dr. Robyn Sanderson were wonderful

and supportive advisors who taught me the foundation of everything I know about astrophysics research and have been huge advocates for me. Thank you to the lunch crew Andy Tzanidakis, Chris Lam, and Gabrielle Suissa, and to Briley Lewis and Evan Morris, who were great friends to me while I was first tepidly dipping a toe into astrophysics. Thank you also to Dr. Mordecai-Mark Mac Low, Dr. Barry McKernan, Dr. Saavik Ford, Dr. Nathan Leigh, Dr. Jillian Bellovary, Dr. Wladimir Lyra, Harrison Cook, Katherine Gonglewski, Syeda Nasim, Betsy Hernandez, Jose Adorno, Gaia Fabj, and all the members of the Ti-DYNYC group, where I have always felt at home and where I grew so much as a scientist. I am very grateful to the many mentors I have had in this group who have put my name forward for talks many times, written me numerous reference letters, and always had their door open to me for advice.

I now turn to thanking my community outside of science. In particular I need to thank my friends whose unwavering belief in and support of me even in the darkest times is unparalleled. Thank you to Anna Callahan, for the many long catch up chats, and Amy Callahan, the best sister's best friend and best friend's sister in the world. Thank you to my biggest cheerleader, Antonio Cavazos, and my family whenever life takes me to LA, Alison Ho. Thank you to my favorite prank co-conspirator and getting old bemoaner, Shirley Li and my dear game for anything friend, Laura Rosenfeld. Thank you to Josh Phillips whose support and encouragement were crucial at the beginning of my journey. Thank you to Adam Weinrib and Zach and Kendall Cobbs-Collins for always helping me take the weekend off to eat good food or watch an insane amount of sporting events, and thank you to Allison Bowman and Amanda Smith for heartfelt conversations on long car rides. Thank you to Brian Crowley for so many kind words over so many years and to my travel and adventure buddy Jill Paganelli. Thank you to Joseph Reidel for always being my fiercest defender and Amelia Coccaro for believing in me for thirty years.

I would also like to thank my family. Walter, Judi, Laurie, Michael, Diane, David, Josh, Eric, Jenn, Alec, Owen, Chris, and my science cousins, Dr. Jake and Melissa. Thank

you also to the Maher family, Jim, Mary, and Sydney, who have been very welcoming and supportive of me over the past six years, and never give me a hard time when I have to work on Christmas Eve. A special thank you to my wonderful grandmother Sylvia and my grandfather Leo, who I know would be very excited for me to get a PhD in Astrophysics from Princeton in particular. A big thank you to my sister and best friend, Michelle, who I can call any time with a story only she would find funny or a complaint only she would empathize with.

Thank you to my parents. My father was the first person I told I was thinking of going back to school and his encouragement started me on this journey. Even at a young age he encouraged my love of math and has always made a concerted effort to stay up to date on my research and career. Most PhDs are not lucky enough to have a father who can understand the details of their research like me and I am truly grateful for the effort he puts in and his enthusiasm. My mother, while less STEM inclined, taught me that any goal is achievable with enough determination and stubbornness. She has also been there for me every day of my life with her love, support, encouragement, and advice.

Finally I have to thank my Princeton crew, Michael Smith, Molly Brennan, Anne Stickells, and Charles Emmett Maher. You all kept me sane through the GC, the pandemic, an advisor switch, and thesis writing. Obviously, I would be remiss without a special shout out to my partner Charles Emmett Maher who was there for me practically from day one. I am not going to write anything cringe but we both know this PhD would not have been possible without you.

For the people who know everything about me and love me anyway

Relation to Published Work

Chapters 2, 3, and 4 have all been published as Yao et al. (2023), Secunda et al. (2023), and Secunda et al. (2024), respectively. I am the second author on Yao et al. (2023), which lays the foundation for the dissertation that follows. Chapters 5 and 6 are works in progress for which I will seek publication at a later date.

This research has been presented at several conferences: LSST AGN SC Summer Meeting, Infinito Planetarium di Torino (July 2024); Black Holes on Broadway, Center for Computational Astrophysics, (Dec 2023); LSST AGN SC Summer Meeting, NRAO (July 2023); The Restless Nature of AGN: 10 Years Later, Centro Congressi Federico II (June 2023); AGN Santa Fe: Where are the Objects in AGN Disks, LANL (March 2023); and LSST Active Galactic Nuclei Science Collaboration Summer Meeting, Virtual (June 2022).

Contents

Abstract . . . iii

Acknowledgements . . . iv

Relation to Published Work . . . ix

List of Tables . . . xiv

List of Figures . . . xv

1 Introduction **1**

1.1 Traditional Reverberation Mapping and the Lamp-Post Model . . . 2

1.2 Challenges to the Standard Thin Disk Model . . . 3

1.3 Simulating AGN Disks . . . 4

1.4 The Long Negative Lag Probes Disk Structure Beyond the Standard Model . . . 6

1.5 Preparing for the Vera C. Rubin Observatory . . . 7

1.6 Dissertation Overview . . . 9

2 A Long Negative Lag from the Optical to the UV Continuum in Fairall 9 **10**

2.1 Introduction . . . 10

2.2 High Cadence Monitoring of AGN . . . 14

2.3 Finding Lags . . . 15

2.3.1 Cross-correlation lags . . . 16

2.3.2 Frequency-resolved lags . . . 17

2.4 Lags in Fairall 9 Time-series . . . 22

2.5 Wavelength dependence of the long lag . . . 26
2.6 Summary . . . 29
2.7 Additional Materials: Fairall 9 Lags . . . 32
2.8 Additional Materials: Simulated light curves . . . 34
2.8.1 Simulations . . . 34

3 Negative Lags on the Viscous Timescale in Quasar Photometry and Prospects for Detecting More with LSST 55

3.1 Introduction . . . 55
3.2 What is the Long Lag Timescale? . . . 59
3.3 Generating Mock Light Curves . . . 65
3.3.1 Driving Light Curves . . . 65
3.3.2 Reprocessed Light Curves . . . 69
3.3.3 LSST Cadence . . . 70
3.4 Lag Detection Methods . . . 72
3.4.1 Cross-Correlation Methods and the Von-Neumann Estimator . . . 72
3.4.2 Javelin . . . 76
3.4.3 Maximum-Likelihood Fourier Method . . . 78
3.5 Using LSST to Detect Long Lags . . . 83
3.5.1 Detecting Long Lags Alone . . . 83
3.5.2 Detecting Short and Long Lags Together . . . 89
3.5.3 Improving the Baseline Cadence . . . 93
3.5.4 Beyond DRW Light Curves . . . 99
3.5.5 Dimmer Magnitude Quasars . . . 103
3.5.6 Summary of Results . . . 106
3.6 Number of Detectable Long Lags . . . 108
3.7 Discussion and Summary . . . 110
3.7.1 Improved Simulated Light Curves . . . 110

3.7.2 Improved Lag Detection Methods . . . 112
3.7.3 Summary . . . 114
3.8 Additional Materials: Full Tests . . . 116
3.8.1 Long and Short Lags . . . 117
3.8.2 Improving the Baseline Cadence . . . 117
3.9 Additional Materials: Null Tests . . . 118
3.9.1 Javelin . . . 119
3.9.2 Fourier Method . . . 120

4 Simulating X-ray Reverberation in the UV-Emitting Regions of Active Galactic Nuclei Accretion Disks with 3D Multi-Frequency Radiation Magnetohydrodynamic Simulations 131
4.1 Introduction . . . 131
4.2 Methods . . . 134
4.3 Results . . . 136
4.3.1 Disk Structure . . . 136
4.3.2 Intrinsic Variability in UV Light Curves . . . 139
4.3.3 Hard X-rays Have Little Impact on UV Emission . . . 142
4.4 Summary and Future Work . . . 145

5 Simulating X-ray Reverberation in the UV-Emitting Regions of Active Galactic Nuclei Accretion Disks with 3D Multi-Frequency Radiation Magnetohydrodynamic Simulations II: Updated X-ray Opacity Models and Modeling the Response Function 149
5.1 Introduction . . . 150
5.2 Methods . . . 153
5.3 Results . . . 158
5.3.1 Disk Structure . . . 158

5.3.2 Light Curves . . . 161
5.3.3 Response Functions . . . 167
5.4 Discussion . . . 173
5.4.1 Next Steps . . . 176

6 Finding Long Negative Lags in Light Curves from the Zwicky Transient Facility **178**

6.1 Introduction . . . 178
6.2 Methods . . . 182
6.2.1 Lag Recovery Methods . . . 182
6.2.2 Light Curve Selection . . . 186
6.2.3 Mock Light Curve Generation . . . 187
6.3 Mock Light Curves . . . 188
6.4 Results . . . 195
6.4.1 g- to i-Band Lags . . . 196
6.4.2 g- to r-Band Lags . . . 197
6.5 Predictions from Previous Work . . . 200
6.6 Discussion . . . 203

Bibliography **207**

List of Tables

2.1 Summary of lags recovered for Fairall 9 22

2.2 Fairall 9 best-fit parameters 26

3.1 Properties of the three redshift ranges we use to bin SDSS DR7 quasars . . . 65

3.2 83

3.3 Summary of cross-correlation and Fourier method performance on even cadence light curves 83

3.4 84

3.5 Summary of JAVELIN and Fourier method performance on mock LSST light curves 84

4.1 Chapter 4 simulation parameters 137

5.1 Parameter values for the Chapter 5 Simulations 154

5.2 Summary of Four Simulation Runs 157

5.3 Pearson r-values between different simulated light curves 161

5.4 Response function model fits 167

6.1 Summary of long negative lags detected for our ZTF light curves 195

6.2 AGN properties and predicted long lags 200

List of Figures

1.1 Schematic of traditional reverberation mapping 2

1.2 Schematic of the long negative lag 6

2.1 Mock light curve and recovered lags 39

2.2 Recovered Fairall 9 lags 40

2.3 Fairall 9 power-law fits 41

2.4 Detrended Fairall 9 light curves 42

2.5 Long and short lags recovered by JAVELIN 43

2.6 Short lags recovered by JAVELIN for detrended light curves 44

2.7 Correlation coefficients from PyMCCF 45

2.8 Short lags recovered in frequency space 46

2.9 Long lags recovered in frequency space 47

2.10 Long lag recovered by JAVELIN between the reference $UVW2$ and the $M2$ bands 48

2.11 Comparison between the frequency-resolved and cross-correlation approaches as a function of the amplitude ratio between the two lags' response functions 49

2.12 Comparison between the Fourier and JAVELIN lag recoveries for a mock light curve 50

2.13 Results of 300 simulated light curves with input lags of -120 and +5 days analyzed via the Fourier method 51

2.14 Results of 300 simulated light curves with input lags of -120 and +5 analyzed via JAVELIN 52

2.15 Results of 100 simulated light curves with input lags of -40 and +5 analyzed via the Fourier method 53
2.16 Lags detected by JAVELIN for mock light curves with different amplitude ratios 54

3.1 Viscous timescales as a function of Eddington ratio and SMBH mass for various scale heights 59
3.2 Long lags predicted by our models for SDSS DR7 AGN 60
3.3 Example mock light curve and its recovered long lag 66
3.4 Cross-correlation coefficients of mock even cadence light curves with input long and short lags with different aspect ratios 73
3.5 The mean of the Von-Neumann value as a function of lag for three mock long season LSST light curves 74
3.6 An example JAVELIN distribution for a mock light curve with the baseline cadence 76
3.7 The time lags as a function of frequency between two even cadence light curves found using the Fourier method 78
3.8 Histogram of the frequency of LSST observations for the long season cadence 81
3.9 The distribution of the long lags detected by JAVELIN for 100 mock long season cadence LSST light curves 85
3.10 The distribution of long lags recovered by the maximum-likelihood method for 100 mock long season cadence LSST light curves 86
3.11 Comparing the distributions of the long lags detected by JAVELIN and the maximum-likelihood method for 100 mock long season cadence LSST light curves with only a long lag and with both a long and short lag 87
3.12 The short lags detected by JAVELIN and the maximum-likelihood method for 100 mock long season cadence LSST light curves with different amplitude ratios 88

3.13 The fractional error in long lags recovered by JAVELIN and the Fourier method for 100 mock long season cadence LSST light curves as a function of amplitude ratio . . . 92
3.14 Comparing long lags recovered by JAVELIN and the maximum-likelihood method for the long season and baseline cadences . . . 93
3.15 Comparing long lags recovered by JAVELIN and the maximum-likelihood method for the baseline cadence with and without LCO observations . . . 94
3.16 Comparing short lags recovered by JAVELIN and the maximum-likelihood method for the baseline cadence with and without LCO observations . . . 95
3.17 Comparing long lags recovered by the Von-Neumann method for the long season and baseline cadences . . . 95
3.18 Long lags recovered by JAVELIN for mock light curves made using MHD simulation light curves . . . 100
3.19 Long lags recovered by the maximum-likelihood method for mock light curves made using MHD simulation light curves . . . 101
3.20 Short lags recovered by JAVELIN for mock light curves made using MHD simulation light curves . . . 102
3.21 Long lags recovered by JAVELIN for higher SNR mock light curves . . . 104
3.22 Long lags recovered by the maximum-likelihood method for higher SNR mock light curves . . . 105
3.23 Long Lags detected by JAVELIN for different lag durations and amplitude ratios 122
3.24 Long Lags detected by the maximum-likelihood method for different lag durations and amplitude ratios . . . 123
3.25 Short Lags detected by JAVELIN for different lag durations and amplitude ratios 123
3.26 Short Lags detected by the maximum-likelihood method for different lag durations and amplitude ratios . . . 124
3.27 Long Lags detected by JAVELIN for different lag durations and cadences . . . 124

3.28 Long Lags detected by the maximum-likelihood method for different lag durations and cadences . . . 124
3.29 Long Lags detected by JAVELIN for different lag durations for the baseline cadence with and without LCO observations . . . 125
3.30 Long Lags detected by JAVELIN for different lag durations for the baseline cadence with and without LCO observations . . . 125
3.31 Short Lags detected by the maximum-likelihood method for different lag durations for the baseline cadence with and without LCO observations . . . 126
3.32 Long lags recovered by the Von-Neumann method for different lag durations and cadences . . . 126
3.33 Results of JAVELIN null tests for 100 mock light curves . . . 127
3.34 Eight individual JAVELIN null tests distributions . . . 128
3.35 Results of maximum-likelihood method null tests for 100 mock light curves . 129
3.36 Results of JAVELIN and maximum-likelihood method null tests for simulation light curves . . . 130

4.1 Schematic of our simulation and vertical profile and space-time diagram of the gas temperature . . . 137
4.2 Vertical profiles and space-time diagrams of density, turbulent velocity, and radiation energy . . . 138
4.3 Light curves from our two runs . . . 140
4.4 PSDs of light curves from our two runs . . . 141

5.1 Opacity as a function of frequency from TOPS . . . 155
5.2 A space-time diagram of the horizontally-averaged gas temperature . . . 158
5.3 Averaged Vertical Profiles of density, gas temperature, turbulent velocity, X-ray flux, and UV flux . . . 159
5.4 Averaged vertical profiles of UV and X-ray thermal timescale . . . 160

5.5 Chapter 5 Simulated Light Curves 162
5.6 Fiducial run PSDs 163
5.7 Low Flux run PSDs 164
5.8 Power-Law run PSDs 165
5.9 Shallow Power-Law run PSDs 166
5.10 Fiducial run response functions 168
5.11 Low Flux run response functions 169
5.12 Power-law run response functions 170
5.13 Shallow Power-Law run response functions 171

6.1 The masses and redshift of the ZTF AGN in our parent sample 185
6.2 Lags recovered by JAVELIN for mock ZTF ight curves 189
6.3 Lags recovered by the Fourier methods for mock ZTF light curves 192
6.4 Lags recovered for subset of ZTF light curves 193
6.5 Lags recovered for subset of ZTF light curves (cont). 194
6.6 Linear interpolation Fourier method lags recovered for subset of ZTF light curves 198
6.7 The long negative lag as a function of mass and Eddington ratio calculated with $\beta = 3.59$ 201
6.8 The long negative lag as a function of mass and Eddington ratio assuming a constant scale height 203

Chapter 1

Introduction

Active galactic nuclei (AGN) are among the most luminous objects in the Universe and are thought to have a large influence on everything from the circumgalactic medium, to galaxy evolution, to the clustering of dark matter. The gas accretion disks that surround and fuel these AGN also provide an important laboratory to study accretion physics. However, with the exception of the massive efforts of the Event Horizon Telescope (Event Horizon Telescope Collaboration, 2019) and the Gravity instrument (Gravity Collaboration, 2018) Collaborations, who were able to image for the first time two nearby supermassive black holes, it is in general not possible to resolve the microarcsecond scales needed to image the accretion disk, the broad-line region surrounding it, or the event horizon. Therefore, we often substitute temporal resolution for spatial resolution, taking advantage of the fact that AGN light curves vary on a range of timescales. Unfortunately, the source of variability in AGN light curves and the extent of what it can teach us about the structure, internal physics, and accretion rates of AGN disk remain unclear. This thesis seeks to better understand the source of variability in AGN light curves in order to work towards better understanding of AGN accretion disks.

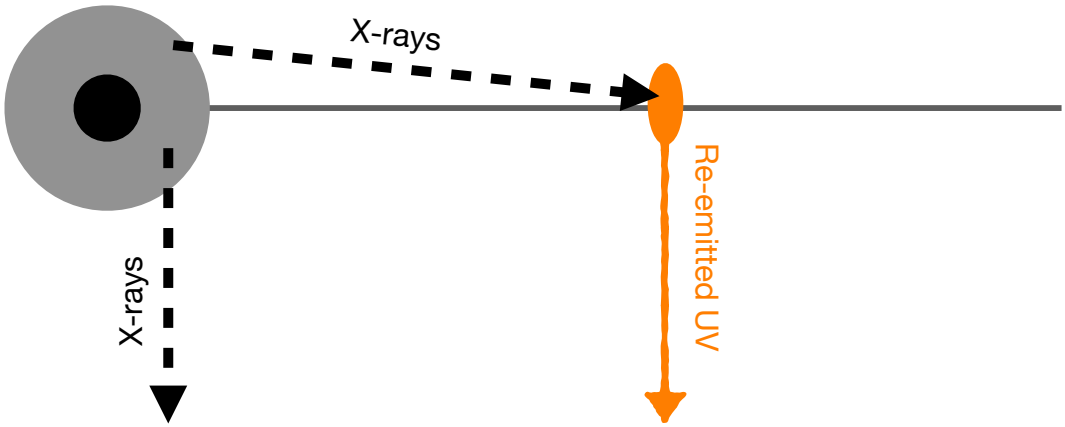


Figure 1.1: Schematic of traditional reverberation mapping, or the short lag. X-rays emitted near the black hole heat the disk further out, and the energy is re-emitted at UV wavelengths.

1.1 Traditional Reverberation Mapping and the Lamp-Post Model

One of the primary techniques that uses temporal resolution, or variability, to study AGN is reverberation mapping. Reverberation mapping was first proposed by Blandford & McKee (1982) as a way to measure distances to the broad-line region by finding a time lag between variability in continuum light curves and variability in broad emission lines (see also Kaspi et al., 2000; Peterson et al., 2004; Bentz & Katz, 2015; Grier et al., 2017). Disk continuum reverberation mapping, a schematic of which is shown in Figure 1.1, uses a similar technique to measure the radial extent and temperature profiles of AGN disks. Disk continuum reverberation mapping takes advantage of the fact that the accretion disks around AGN emit radiation over a wide range of wavelengths depending on the local temperature of the disk, which decreases with distance from the black hole. Traditionally, disk continuum reverberation mapping is based on the lamp-post model which says that as X-ray light emitted from the hottest inner regions of the AGN disk moves outwards along the disk it is absorbed and re-emitted by cooler and cooler regions of the disk at UV and then optical wavelengths (e.g. Sergeev et al., 2005; Cackett et al., 2007). Because this X-ray light varies on multiple

timescales, its variability is reverberated in these longer wavelengths. Traditional reverberation mapping measures the time lag, which is typically on the order of days, between these variability echoes, which corresponds to the light crossing timescale between the different temperature regions of the disk.

Over the past decade, significant investments in time domain surveys (e.g., Edelson et al., 2015; McHardy et al., 2018; Jiang et al., 2017; Mudd et al., 2018; Homayouni et al., 2022; Guo et al., 2022) and reverberation mapping campaigns with broad wavelength coverage (e.g., Edelson et al., 2019; Cackett et al., 2021) have improved the available cadence and duration of AGN light curves leading to disk continuum lag measurements, and therefore size measurements, for hundreds of AGN. In this thesis, I will often refer to these lags as "short" lags.

1.2 Challenges to the Standard Thin Disk Model

The empirical models derived from continuum reverberation mapping as well as simulations of AGN disks are typically compared to the standard Shakura & Sunyaev (1973) optically thick but geometrically thin disk model. However, a greater understanding of the underlying physics behind AGN light curves through simulations is urgently needed, because as simulations and observations of AGN disks improve several challenges to the standard disk model have arisen.

While the short lags measured from the high cadence observations of the last decade are generally consistent with the assumption from the standard thin disk model that $r \propto T^{4/3}$ (Shakura & Sunyaev, 1976), where r is the disk radius and T is the effective temperature, the observed timescales from these observations suggest that accretion disks are larger than predicted by the standard thin disk model (e.g., Jiang et al., 2017; Guo et al., 2022). Accretion disk sizes measured with micro-lensing also appear to be larger than predicted (e.g., Morgan et al., 2010). In addition, numerical simulations of AGN disks also suggest that

pressure support from magnetic fields and radiation can lead to puffed up, or slim not thin, disks (e.g., Jiang et al., 2016, 2019; Jiang & Blaes, 2020; Hopkins et al., 2023). This result is in direct contradiction of the standard thin disk model.

Finally, several recent reverberation mapping campaigns are finding that variability in UV-optical light curves and X-ray light curves from the same AGN are not well-correlated (e.g., Cackett et al., 2023; Kara et al., 2023) and that the variability in light curves emitted by the AGN disk cannot simply be modeled as reverberation of X-rays (e.g., Arévalo et al., 2009; Shappee et al., 2014; Hagen & Done, 2023). Both these results suggest that there is some intrinsic UV-optical variability originating from fluctuations in the UV-optical regions of AGN disks on a timescale shorter than predicted by the standard thin disk model (Lawrence, 2018). There is also evidence that these fluctuations are being propagated inwards on an inflow timescale that is much shorter than predicted by the standard disk model (e.g., Neustadt & Kochanek, 2022; Yao et al., 2023; Secunda et al., 2023).

1.3 Simulating AGN Disks

While we now have hundreds of measured continuum lags, numerically simulating the underlying physics behind reverberation mapping is a difficult problem, because it is crucial to self-consistently simulate both the magnetohydrodynamics (MHD) of the gas and radiation transport. Including the evolution of the magnetic fields is important in these simulations because the magneto-rotational instability (MRI) is now widely believed to be the physical mechanism responsible for the angular momentum transport in AGN disk (Balbus & Hawley, 1991), and also drives turbulent fluctuations that could impact the variability of AGN light curves (e.g., Lyubarskii, 1997). Calculating the radiation transport is necessary to determine the thermal properties of the disk and connect the accretion disk structures with the observed light curves directly. In addition, both radiation and magnetic pressure

play a role in supporting the AGN disk against gravity, with radiation pressure assumed to be the dominant force supporting the inner region of the AGN disk.

Early radiation MHD accretion disk simulations disagreed on whether AGN disks were subject to the well known thermal instability in the standard $\alpha-$disk model (e.g., Hirose et al., 2009; Jiang et al., 2013). Jiang et al. (2016) showed that including line opacities instead of only the electron scattering opacity when modeling the gas opacity in AGN disk simulations has a large impact on the stability and structure of the accretion disk. In particular, for the typical temperatures and densities of the UV-optical emission regions of AGN disks the bound-bound transition in iron can cause what is know as the iron opacity bump. This enhanced opacity can not only help stabilize the disk, it can also produce an additional source of variability by driving convection, as Jiang & Blaes (2020) showed using fully global 3D radiation MHD simulations of an AGN accretion disk. A recent improvement in radiation MHD simulations is the development of an implicit radiation method described in Jiang (2021), that does not require a closure scheme, such as the commonly used Eddington or variable Eddington tensors.

Most recently, the multi-frequency radiation MHD code described in (Jiang, 2022), provides for the first time an exciting opportunity to not just model the fluctuations in AGN disks, but to also directly simulate the reprocessing of high-frequency irradiation into lower frequency, UV or optical emission by the AGN disk from first principles using realistic opacity models for both the high and low frequency radiation. These simulations are the next step for understanding the drivers behind variability in AGN light curves and what these light curves can reveal about the structure and internal physics of AGN disks.

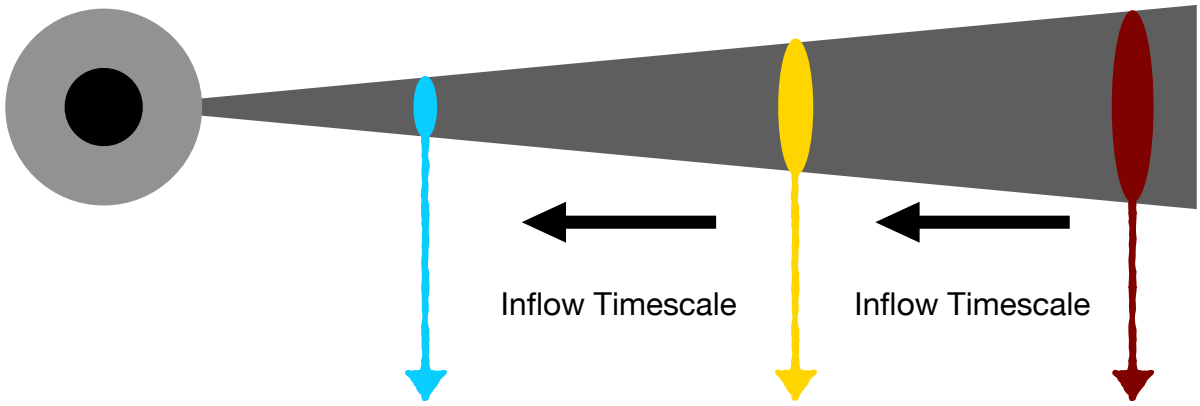


Figure 1.2: Schematic of the long negative lag. Fluctuations in the maroon oval lead to variability in light curves emitted at longer wavelengths from this cooler region. Over the inflow timescale this fluctuation is propagated inwards to the yellow then blue regions where it leads to variability in light curves emitted from gradually hotter regions at shorter wavelengths.

1.4 The Long Negative Lag Probes Disk Structure Beyond the Standard Model

One way to directly test the standard thin disk model is to study this inward propagation of fluctuations which lead to a new type of reverberation mapping signal that here we will call the "long negative lag." The long negative lag, illustrated in Figure 1.2, occurs when fluctuations in the outer UV-optical region produce variability in AGN light curves at longer wavelengths emitted from these cooler regions. This variability is then echoed at shorter wavelengths emitted by the hotter, inner regions that these fluctuations are propagated to. This echo, or lag, is in the opposite direction from traditional reverberation mapping (short wavelengths lag longer wavelengths) and on the much longer inflow timescale.

The standard thin disk model predicts that the inflow timescale is,

$$\tau_{\rm inflow} = 260 \text{ years } \frac{\rm M_{BH}}{10^8 \ \rm M_\odot} \frac{0.1}{\alpha} \left(\frac{\rm H/r}{0.01}\right)^{-2} \left(\frac{\rm r}{30 \ r_g}\right)^{3/2}, \tag{1.1}$$

where $M_{\rm BH}$ is the mass of the supermassive black hole, α is a dimensionless viscosity parameter, H is the scale height of the disk, r is the radial position in the disk, and $r_{\rm g}$ is gravitational radii. If the inflow timescale is on the order of hundreds of years it will not be possible to observe a lag on this timescale. However, as suggested by recent numerical simulations, if AGN disks are slim instead of thin, and their scale height is an order of magnitude larger, the inflow timescale is on the order of hundreds of days. A long negative lag with a timescale of around 100 days is possible to detect with long baseline light curves and can provide information on the vertical structure of AGN disks and intrinsic variability due to physical processes in AGN disks.

Note that in this dissertation we model the long negative lags we detect as due to this propagation of fluctuations, which is supported by evidence of the propagation of temperature fluctuations on monthly timescales observed by Neustadt & Kochanek (2022); Neustadt et al. (2023) and Stone & Shen (2023) in AGN light curves. However, there are alternative plausible explanations. Instead the long negative lags we detect could be the result of disk winds/columnated outflows, surface accretion that can proceed more rapidly in the optically thin regime, or the disk-instability model that is commonly used to explain dwarf novae (Hameury, 2020). Any of these alternative models would also challenge the standard thin disk model. As more long negative lags are detected it will be easier to differentiate between different models.

1.5 Preparing for the Vera C. Rubin Observatory

Now is the opportune moment to update how we model reverberation mapping because of the Vera C. Rubin Legacy Survey of Space and Time (LSST Ivezić et al., 2019). Beginning in 2025, LSST will survey half the sky for at least 10 years and will obtain light curves emitted by millions of AGN accretion disks in six different optical bands. The LSST observing strategy includes two main components, the Wide Fast Deep Survey and the Deep Drilling Fields.

The Wide Fast Deep Survey will be at least 18000 $\deg^2$ and each 9.6 $\deg^2$ field will have a median of 825 visits over ten years, or a visit every 3–5 nights. The Deep Drilling Fields are five relatively small (9.6 $\deg^2$) fields but will have deeper coverage and a higher cadence than the Wide Fast Deep Survey. The final cadence for LSST has not been selected yet, but for the Deep Drilling Fields the cadence could be as high as every 1 to 3 days. Kovacevic et al. (2022) showed that in the Deep Drilling Fields it will be possible to detect thousands of traditional disk continuum reverberation mapping lags with a confidence interval of at least 0.95. Chapter 3 outlines how it will be possible to detect up to hundreds of additional long negative lags in the LSST Deep Drilling Field.

Therefore, it is crucial to prepare for LSST with observations, forecasting, and simulations. Searching for long negative lags in already available UV and optical AGN light curves allows us to test and improve lag recovery methods. Detecting more long negative lags also enables us to learn about what timescales we expect these lags to be on, and improve models for AGN disk structure based on the lags we detect. Forecasting our ability to detect long negative lags in LSST light curves can also help us to tailor our lag recovery methods specifically toward detecting long negative lags in LSST light curves.

Finally, multi-frequency radiation MHD simulations of AGN disks will allow us to understand the source of variability in AGN light curves and understand how high-frequency irradiation is actually reprocessed by the disk by writing down a physically motivated response function that captures this reprocessing. This increased understanding will allow us to improve our traditional continuum reverberation mapping techniques and gain more information from AGN light curves about the internal physics of AGN disks than we are currently able to. These simulations are also a necessary first step to studying the propagation of fluctuations in these disks and simulating the long negative lag.

1.6 Dissertation Overview

The remainder of this dissertation is structured as follows. In Chapter 2, I introduce reverberation mapping in detail by describing the methods we used to detect both a short and one of the first ever long negative lags for the AGN Fairall 9. I use the best-fit models for the long negative lag as a function of wavelength from Chapter 2 to derive a model for the dependence of the vertical structure of the AGN disk on the long negative lag, and then use this model to forecast our ability to recover long negative lags in LSST light curves with different lag detection methods in Chapter 3. To take a closer look at the physics behind traditional continuum reverberation mapping, in Chapter 4, I perform radiation MHD simulations using ATHENA++ and the implicit multi-frequency radiation transfer scheme in Jiang (2022) to simulate the reprocessing of X-ray irradiation into UV emission by an AGN disk. I follow these simulations up in Chapter 5 with brand new simulations that use a more sophisticated model for the X-ray opacity. This new opacity model leads to more strongly-correlated X-ray and UV light curves that we can use to test the empirical response functions observers use to describe reprocessing. Finally, in Chapter 6 I describe my ongoing search for more long negative lags in light curves from the Zwicky Transient Facility, with a few promising long lag candidates and a discussion of the difficulties and pitfalls of detecting the long negative lag.

Chapter 2

A Long Negative Lag from the Optical to the UV Continuum in Fairall 9

2.1 Introduction

Supermassive black holes (SMBHs) are surrounded by an accretion disk that powers one of the most luminous sources in our universe. These central engines, termed active galactic nuclei (AGN), harbor precious information on accretion physics and galaxy evolution (Fabian, 2012; Kormendy & Ho, 2013). However, direct imaging of most accretion disks and surrounding structures is in general observationally inaccessible because of their relatively small sizes or large distance. Alternatively, temporal variability as a function of frequency can be used instead to map the structure of the accretion disks.

Reverberation mapping is a technique that traces delays between light curves at different wavelengths. By mapping the wavelengths to different temperature regions in the disk, the lags between bands can be used to probe the accretion disk structure (Blandford & McKee, 1982; Peterson et al., 2004). Classically, this technique looked at the lags between AGN continua and broad emission lines to measure the sizes of the Broad Line Regions (BLR), but the same principle applies to looking for continuum lags as a function of frequency. Since

the path lengths to observers on Earth are different from different regions of the source, we can use the light travel time to deduce the physical sizes and temperature profiles of AGNs. For example, radiation from the X-ray corona gets reprocessed and re-emitted earlier in the inner disk than the outer parts. Since the former possesses much higher temperature than the latter, the reprocessed light varies in wavelength with the ultraviolet light curve leading the optical. We will refer to this lag as the "short lag" (< 10 days), occurring in a light-crossing time, and this distinct signal has been observed in many active galaxies such as the Seyfert 1 galaxy Fairall 9 (See Cackett et al., 2021, and references there in). However, signals beyond reprocessing can also be embedded in the light curves from AGN.

In addition to the short lag, and the BLR lag, we also expect to see correlations between different bands on longer timescales, even propagating inwards in the disk. AGN accretion disks are believed to be turbulent due to magneto-rotational instability (MRI) (Balbus & Hawley, 1991). Turbulent fluctuations in the disk will naturally lead to stochastic variability from the locally emitted photons, which provide a natural explanation for the broadband photometric variability observed in AGNs. When gas flows towards the black hole via the accretion disk, it is expected that turbulence-driven perturbations at larger radii will propagate inward with the accretion flow itself and modulate the local variability at small radii (Lyubarskii, 1997).

Several characteristic timescales are expected from the accretion process (e.g. Stern et al., 2018). Turbulence will cause light curve variability on the local thermal time scale for optically thick disks (Jiang & Blaes, 2020; Burke et al., 2021). When perturbations propagate inward, fluctuations at different radii can correlate on the inflow time scale. Particularly, in AGN accretion disks, the UV/Optical emission region is expected to be strongly perturbed by convection due to enhanced opacity (Jiang & Blaes, 2020) in addition to the normal MRI turbulence, which can enhance the fluctuation propagation signals. Since the effective temperature at large radii should be lower than the effective temperature of the inner region of the disk, this kind of fluctuation propagation can cause long wavelength variability to lead

the variability at shorter wavelengths, which will be in the opposite direction as traditional reverberation signals. The two kinds of correlation will happen on quite different timescales but both of them (and potentially the BLR lag, which is not treated here) will exist in the light curve. Despite continuum lags, BLR lags also exist in AGN light curves, but are not the focus of this paper. Variability studies of X-ray binary light curves already find evidence that both lags can exist in the very inner region of the accretion disks (Uttley et al., 2014). Inward fluctuation propagation has also been proposed to explain the variability properties of many accretion systems (e.g. Arévalo & Uttley, 2006; Arévalo et al., 2008a,b; Hogg & Reynolds, 2016). Recent variability analysis of AGN has identified systems with two components with different timescales (e.g. Hernández Santisteban et al., 2020; Vincentelli et al., 2021), but the exact long-timescale lag remains uncertain. Inflow time scales in AGN accretion disks will depend on the detailed structure of the disk, and our a priori expectation is that these can be $\sim 100 - 1000$ days, but more empirical constraints are needed. Detecting such lags in AGN light curves will be very useful to probe internal structures of AGN accretion disks, and is a main goal of this work.

In this paper, we explore both cross-correlation (Gaskell & Peterson, 1987) and frequency resolved techniques to extract AGN time lags with specific interest in extracting a long lag in Fairall 9. In temporal space, we use the JAVELIN code (Zu et al. (2010), formerly known as SPEAR), which uses a damped random walk (DRW) model (Timmer & Koenig, 1995) to describe the AGN variability, and interpolates the light curves to increase sensitivity to the short-timescale variability. In frequency space, the Fourier method can provide excellent information on different timescales and periodicity. Fourier methods have seen many successes in X-ray astronomy, such as characterizing X-ray binaries. Beyond a single lag, frequency-resolved analysis yields lag as a function of frequency and can find the characteristic timescales of different processes. However, this technique can be challenging when irregular sampling is present. Both high and low frequency lags in X-ray reverberation of AGN have been detected through simple fast Fourier transform with continuously sampled

data (e.g. Papadakis et al., 2001; McHardy et al., 2007; Arévalo et al., 2008a; Fabian et al., 2009; Zoghbi et al., 2010; De Marco et al., 2013; Cackett et al., 2013; Uttley et al., 2014). Observations of AGN in the UV/optical require a maximum-likelihood approach in frequency space because of their non-uniform cadence. Hence, we use the method developed by Miller et al. (2010), which Zoghbi et al. (2013) and Cackett et al. (2022) later adapted to successfully recover lags from AGN light curves. We investigate the relative strengths of each method for extracting a range of lag times in a single spectrum. This analysis will be carried out both on simulated AGN light curves and on multi-wavelength observations of the local Seyfert galaxy Fairall 9, which hosts a 2.55×10^8 $M_\odot$ SMBH (Peterson et al., 2004).

Our discussion begins with presenting one of the best existing data sets for our studies from Hernández Santisteban et al. (2020) in Section 2.2. Then we move on to a introduction to both the cross-correlation and Fourier methods in Section 2.3. Next, in Section 2.4, we report our recoveries of both the long and short lags from the Fairall 9 data set, which informs how future long-term monitoring of AGN will improve our understanding of both signals. We then discuss how the wavelength dependence of the long lag informs the geometry of the inner accretion disk in Section 2.5. Finally, we provide a brief summary in Section 2.6. Details regarding the applicability of both methods in different scenarios are presented in Appendix 2.8, showing lag recovery from simulated light curves following a DRW model.

2.2 High Cadence Monitoring of AGN

With continuum reverberation mapping, significant progress has been made in understanding accretion physics, with pioneering observational projects such as the Space Telescope and Optical Reverberation Mapping Project (STORM) (e.g. De Rosa et al., 2015; Edelson et al., 2015; Fausnaugh et al., 2016; Starkey et al., 2017)), and other intensive disk-reverberation mapping campaigns (e.g. Edelson et al., 2017, 2019; Cackett et al., 2018; McHardy et al., 2018; Hernández Santisteban et al., 2020). Low-cadence sampling ($\sim$ 2-7 days) and short wavelength coverage are often the causes for failing to detect inter-band lags in analysis of AGN light curves (Hernández Santisteban et al., 2020). In a handful of AGN, high sampling, wide wavelength coverage, and long baseline are all available, allowing simultaneous studies of both short-timescale ($<$ 10 days) and long-timescale ($\sim$ 100 days) lags.

For our analysis, we focus on the intensive disk-reverberation mapping campaign of Fairall 9 carried out by Hernández Santisteban et al. (2020) using *Swift* and the Las Cumbres Observatory (LCO). Half of the data are already released to the public, hence we use light curves covering half of the 1.7-year duration between MJD 58250-58550. This timespan is still much longer than the timescale at which the short lag occurs. These data are also characterized by extended wavelength coverage from 1928 Å (*Swift* $UVW2$) to 8700 Å (LCO z_s). In the UV/optical regime, with $UVW2$, $UVM2$, $UVW1$, U, B, V broad-band filters from *Swift* and $uBVg'r'i'z_s$ filters from LCO, the wavelength dependence of both lags can be more clearly identified and associated with different regions of the accretion disk. Additionally, light curves from this campaign have a one-day cadence from *Swift* and a sub-day cadence from LCO, allowing more accurate derivation of the shorter lag.

In the intensive disk-reverberation mapping campaign of Fairall 9, Hernández Santisteban et al. (2020) were able to derive robust short lags using a cross-correlation method and provided a wavelength dependence consistent with $\tau \propto \lambda^{4/3}$ from classic predictions. For Fairall 9, with a Eddington ratio $L/L_{Edd} = 0.02$, the increase in lag from $UVW2$ to z_s band shows promising compatibility with accretion disk theory. Additionally, the authors

identified a low-frequency and slowly-varying component in Fairall 9 evidenced by a parabolic shape over $\sim$ 100 days. By measuring the shift in the minima of each band's parabolic fit, they derived a $\lesssim$ 10 day lag between the blue and red bands, which they conclude corresponds to the propagation time convolved with a weighted radial emission profile in different regions. Hernández Santisteban et al. (2020) associated this slower component with disk fluctuations, and left further analysis for future work when longer and higher-cadence monitoring of Fairall 9 becomes available. The reported lag ($\sim$ 10 days) for this slowly varying component is even smaller than the local thermal time scales in this region of the disk, which is one important motivation for us to reanalyze the light curves.

Frequency-resolved analysis could identify different processes taking place on different timescales (Uttley et al., 2014), and has not been carried out on this data set. With unevenly sampled light curves, Zoghbi et al. (2013) builds upon the Fourier method presented first in Miller et al. (2010), to successfully recover lags from NGC 4151 and MCG-5-23-16 that are consistent with standard cross-correlation methods. As a result, we are motivated to explore the timescale at which the slowly varying component seen in the intensive disk-reverberation mapping campaign of Fairall 9 occurs by assessing the ability of both the cross-correlation and Fourier methods in simultaneously constraining both varying components from the driving and the reprocessed light curves.

After demonstrating both methods in the following sections, we will report the first direct confirmation of the existence of a negative long-timescale lag (~ -70 days) in Fairall 9 in Section 2.4 obtained through cross correlating an interpolated light curve with JAVELIN. We have also confirmed this detection using the Fourier method.

2.3 Finding Lags

Continuum reverberation mapping aims at measuring the structure of AGN accretion disks, and hence determining properties of the black hole. Doing so relies on time-series in multiple

bands where time lags between light curves across a wavelength range can be measured. In recent years, both observational and theoretical work (e.g. Hernández Santisteban et al., 2020; Neustadt & Kochanek, 2022) has revealed the potential existence of the long lag inherent to accretion fluctuations in AGN UV/Optical light curves, especially in Fairall 9. As a result, in this section we build upon the methods of cross-correlation and Fourier to simultaneously constrain the short and long lags.

2.3.1 Cross-correlation lags

With an ideal uniformly-sampled time-series, cross-correlation analysis can be easily performed by measuring the cross-correlation coefficient at each delay-time τ following:

$$\mathrm{CC_{coeff}}(\tau) = \frac{\sum \mathrm{d(t)r(t+\tau)}}{\sqrt{\sum \mathrm{d(t)^2 r(t+\tau)^2}}}, \tag{2.1}$$

where d(t) and r(t) represent light curves from different bands used for cross-correlation; in this equation, r(t) is assumed to be delayed by time τ compared to d(t). The highest coefficient indicates the most probable time lag. Both positive and negative time lags can be measured in this manner by selecting a range of τ with both signs. The accuracy of this method is largely dependent on the discrete sample cadence. However, without the premise of uniform cadence, simple cross-correlation can no longer easily recover accurate lags, and interpolating to a uniform light curve becomes the natural solution.

JAVELIN [1], previously known as SPEAR, is an algorithm used to look for time lags in spectroscopic (Zu et al., 2011) or photometric (Zu et al., 2016) light curves. To perform the latter analysis, JAVELIN models quasar light curves DRW stochastic process. This has been shown to be an effective representation of quasar optical variability on timescales from days to years (e.g. Kelly et al., 2009; Kozłowski et al., 2010; Zu et al., 2013). By using Markov Chain Monte Carlo (MCMC) to estimate model parameters, JAVELIN averages over all

[1] As presented in Zu et al. (2010), code available at: https://github.com/nye17/javelin

possible continuous light curves consistent with the input data and the DRW model while simultaneously fitting a transfer function to maximize the likelihood function. To calculate the most probable time lags, JAVELIN takes into account multiple sources of error including those associated with interpolation. When applied to our simulated light curves and real data, JAVELIN takes the epoch, flux, and uncertainty as inputs, and returns the parameters for the interpolated DRW light curve, as well as the probability distribution for time lags. By selecting the range of lags we are looking for, both positive and negative, JAVELIN can constrain both lags simultaneously.

2.3.2 Frequency-resolved lags

Fourier transforms naturally separate variability at different frequencies. In the regime of X-ray binaries, frequency-resolved analysis distinguishes different processes taking place on different timescales (Uttley et al., 2014). Due to the short time scale of physical processes in this context, very uniform cadence can be achieved at X-ray wavelengths where the disk is most luminous. However, when applied to UV/optical light curves from AGNs, the relevant physical timescales are much longer, and the disk emission is spread out from UV through optical. As a result, uniform cadence is never fully realized, and more delicate treatment is necessary for an accurate lag determination.

2.3.2.1 Uniform cadence time-series

Reprocessing of the driving light curve $d(t)$, as shown in the top panel of Figure 2.1, can be represented as a response function $\psi(t)$ where the reprocessed light curve in a linearized reveberation mapping model is

$$\mathrm{r(t)} = \int_{-\infty}^{\infty} \psi(\tau)\mathrm{d(t-\tau)d\tau}. \tag{2.2}$$

In the Fourier domain, this integral becomes a multiplication in the form of

$$\mathrm{R}(\nu) = \Psi(\nu)\mathrm{D}(\nu), \tag{2.3}$$

where upper-case variables represent the Fourier transform of their lower-case counterparts at each frequency ν. The power-spectral density (PSD) of the driving light curve is simply equal to $|\mathrm{D}(\nu)|^2 = \mathrm{D}^*(\nu)\mathrm{D}(\nu)$, and the PSD for the reprocessed light curve follows in the same manner via $|\mathrm{R}(\nu)|^2 = |\Psi(\nu)|^2|\mathrm{D}(\nu)|^2$. Then, the cross-spectrum of the two light curves can be measured as the product of the complex conjugate of one Fourier-transformed light curve and the Fourier transform of the other:

$$\mathrm{C}(\nu) = \mathrm{D}^*(\nu)\mathrm{R}(\nu) = |\mathrm{D}(\nu)|^2\Psi(\nu). \tag{2.4}$$

Following Uttley et al. (2014) further, one can calculate the phase lag ϕ as the phase of the complex cross spectrum, and derive the time lag τ through

$$\tau = \phi/2\pi\nu. \tag{2.5}$$

As seen in Figure 2.1, the resulting distribution of τ in frequency space is largely dependent on the shape of the response function. Unlike the case for the single short lag, even when the light curve is reprocessed just with a single long lag, it can be difficult to discern by eye. But with the Fourier method, we can recover the lags in both cases indicated by the constant lag in the flat part of the lag distribution at lower frequencies. At higher frequencies, we see expected phase-wrapping as a form of aliasing. Phase wrapping depends on the centroid of the response function (τ_{cent}) at the frequency $\nu = 1/2\tau_{\mathrm{cent}}$. Because the waveform could be shifted by half a wave, the lag distribution in frequency space wraps around from π to $-\pi$, from positive to negative (Uttley et al., 2014).

2.3.2.2 Non-uniform cadence time-series

With UV/optical AGN light curves, the uniform sampling available from X-ray satellites is no longer feasible, at least at present. As a result, lag determination in frequency space cannot rely on simple fast Fourier transforms as presented above. Thus, a Fourier method with a maximum-likelihood approach, first presented in Miller et al. (2010), is more suitable. In this case, we model the power spectra and response functions as piecewise functions of Fourier frequency to construct a likelihood function, which is then used to estimate model parameters when maximized. For all of our analysis of simulated and real data, we follow the detailed outline of the method presented by Zoghbi et al. (2013).

We begin by creating N bins in frequency space with ranges and centroids depending on the simulated light curve length. The bins range from $f_{min} = 1/T$ to $f_{max} = 1/(2dt)$, where T is the longest separation between points in the light curves and dt is the shortest. Two additional bins are added at the start and end of the other bins to reduce bias in the first and last bins. We then create a model for the PSD and use a quadratic approximation to the likelihood to maximize the likelihood function and find best fit values of the PSD. Our uncertainties come from the Hessian matrix of the likelihood function, which places a lower limit. Future work can estimate the full uncertainty using the emcee package if needed. This step is repeated for each light curve and both PSDs are then used as inputs for modeling the cross spectrum to calculate the phase delay.

Here, the autocorrelation function can be calculated via:

$$\mathcal{A}(\tau) = \int_{-\infty}^{\infty} |D(\nu)|^2 \cos(2\pi\nu\tau) d\nu, \tag{2.6}$$

where $|D(\nu)|^2 = \sum_i D_i$ is the PSD as a piecewise sum over a range of frequency bins. Then, the time lag can be estimated with the cross-covariance, whose relationship with the cross

spectrum and phase lag are embedded in the cross-correlation function:

$$\begin{aligned}\chi(\tau) &= \int_{-\infty}^{\infty} \mathrm{C}(\nu)\cos(2\pi\nu\tau - \phi(\nu))d\nu \\ &= \int_{-\infty}^{\infty} |\mathrm{D}(\nu)|^2|\Psi(\nu)|\cos(2\pi\nu\tau - \phi(\nu))\mathrm{d}\nu,\end{aligned} \tag{2.7}$$

where $\Psi(\nu) = \sum_i \Psi_i$ and $\phi(\nu) = \sum_i \phi_i$ are the Fourier transform of the response function and its phase. Thus, the algorithm can construct and maximize a likelihood function to optimize for Ψ_i and ϕ_i at each frequency bin. The resultant ϕ_i distribution can be further divided by $2\pi\nu_i$ to calculate the time lag in frequency space, producing a lag distribution similar to the uniform cadence case with frequency-resolved analysis. The intervals of log(likelihood) informs the error: the 68% uncertainty associated with each parameter can be estimated as the value that changes $-2\log(\mathcal{L}/\mathcal{L}_{\max})$ by 1 while allowing other parameters to vary (Miller et al., 2010). In our case, with a small set of parameters to optimize, this is computationally inexpensive. However, with $\gtrsim$ 20 parameters, one can use MCMC to map the whole probability space.

2.3.2.3 Simulated light curves

We test both the Fourier (frequency space) method and the cross-correlation method (temporal space) with simulated DRW light curves, using JAVELIN for the latter. Our mock light curves for testing are inspired by the Fairall 9 data set in signal-to-noise ratio and cadence, but then we explore the impact of changes to the length of the campaign, as well as recovery success as a function of the shape of the response function. More details related to our simulated light curves and performance of both methods on extracting time lags with both uniform and non-uniform time-series are presented in Appendix 2.8.

In summary, both methods show significant promise in constraining long- and short-timescale lags simultaneously in non-uniform cadence light curves. With a cadence as high as the Fairall 9 data set ($\sim$ day to subday), both methods can easily recover the short lag

with low uncertainties. For the long lag, the Fourier method requires the time baseline of observations to be $\gtrsim 6$ times the long lag (~ 2 years in the case of a -120 day lag) to probe into the relevant time scale in frequency space. JAVELIN can recover the long lag given the length of current Fairall 9 campaign (~ 300 days), but because JAVELIN does not provide any information on the frequency-dependence of this lag, it lacks information regarding the physical timescale on which it acts. Since the actual shape of the response function for AGN is generally unknown, our test cases adopt several amplitude ratios between the responses of the short and long lags. As expected, when the response function for the long lag is stronger, the lags are more easily recovered, and vice versa.

Both JAVELIN long lags and Fourier short lags show no evidence of contamination from the other lag. However, JAVELIN short lags and Fourier long lags can be biased by the other lag to different degrees depending on the relative amplitudes of the response functions. This could provide a remedy to a well-known puzzle with continuum reverberation mapping that accretion disk sizes derived from cross-correlation lag spectrum are often $\sim 2-3$ times bigger than expected by the standard accretion theory in Shakura & Sunyaev (1973) (e.g. Edelson et al., 2015, 2017, 2019; Fausnaugh et al., 2016; Cackett et al., 2018; Pozo Nuñez et al., 2019). One possible explanation is that the value of the short lag can be biased by other embedded signals in AGN light curves. Figure 2.14 in Appendix 2.8 shows that lags recovered with cross-correlation can be notably longer than the input lag when the signal of the long lag is more significant, potentially resulting in inaccurate measurements of disk sizes. In this case, better modelling is needed to accurately deduce accretion disk sizes, as shown in Kammoun et al. (2021a). On the other hand, Figure 2.13 shows that short lags deduced with the Fourier method are all consistent with the input lag at 5 days, regardless of the relative amplitude ratio of the response functions. This corroborates our previous conclusion that the Fourier method tends to recover more unbiased and precise measurements of the short lag, which provides a hopeful outlook in accurately measuring accretion disk sizes across a large AGN population.

Filter	$\lambda_{\rm eff}$(Å)	Fourier short	JAVELIN short	Fourier long	JAVELIN long
Swift					
$M2$	2246	0.34 ± 1.33	0.23 ± 0.39	-2.78 ± 4.21	-29.6 ± 21.9
$W1$	2600	-1.20 ± 1.28	1.94 ± 1.16	-12.46 ± 1.79	-43.5 ± 15.9
U	3465	-0.18 ± 1.20	3.56 ± 1.45	-20.33 ± 6.92	-65.6 ± 3.4
B	4392	1.86 ± 0.69	6.02 ± 3.73	-10.02 ± 2.46	-58.5 ± 5.4
V	5468	0.39 ± 1.32	6.94 ± 2.25	-34.51 ± 19.77	-71.4 ± 8.5
LCO					
u'	3580	1.73 ± 0.40	4.41 ± 0.61	-25.97 ± 7.96	-58.6 ± 9.8
B	4392	1.26 ± 1.20	5.62 ± 0.63	-17.73 ± 2.21	-65.9 ± 4.6
g'	4770	1.04 ± 0.72	3.08 ± 0.96	-9.21 ± 1.31	-66.9 ± 7.4
V	5468	1.47 ± 4.71	4.62 ± 0.82	-6.26 ± 0.38	-69.1 ± 9.9
r'	6215	2.22 ± 0.71	4.44 ± 1.70	-12.72 ± 6.04	-59.2 ± 6.7
i'	7545	2.61 ± 0.66	8.77 ± 2.16	-29.57 ± 1.34	-77.3 ± 9.8
z_s	8700	3.30 ± 1.06	6.93 ± 1.53	-34.74 ± 8.95	-68.0 ± 6.2

Table 2.1: Table of lag values as plotted in Figure 2.2. Note that JAVELIN short lags come from detrended light curves. All lags are measured in units of days, relative to the reference $UVW2$ band (1928 Å).

2.4 Lags in Fairall 9 Time-series

Building on the work of Hernández Santisteban et al. (2020), we revisit the Fairall 9 UV/Optical light curves. Specifically, we hope to investigate whether the low-frequency component they uncovered may correspond with the long lag we expect. To make comparison easier, we show the lag calculations in each band relative to the $UVW2$ band (far-UV from *Switft*) as in Hernández Santisteban et al. (2020).

For analysis of short lags, detrending AGN light curves can remove variability on long timescales and hence produce more accurate measurements of the short timescale lags (McHardy et al., 2014, 2018; Pahari et al., 2020). Therefore, we perform quadratic detrending to analyzing the Fairall 9 data as in Hernández Santisteban et al. (2020) when measuring the short lag with JAVELIN. Our results also show short lags from undetrended light curves are about 2 times shorter than from detrended light curves. As seen in the top right panel of Figure 2.2, the short lag is increasing in magnitude with wavelength as predicted. In the bottom right panel, we also show the detected short lags from the Fourier

method. The values are an error weighted average of bins (orange points in Figure 2.8 in Appendix 2.7) in the relevant timescale of the short lag based on analysis of our simulated light curves around $10^{-2} - 10^{-1}\mathrm{day}^{-1}$. These lags come from undetrended light curves, since the Fourier approach should naturally separate variability on different timescales.

When compared to the cross-correlation results, we observe a recovery of the short-timescale lag that is better-fit with a power law and has a generally smaller error compared to the cross-correlation method. We note that the accuracy of the Fourier method here is sensitive to bin choices, the recovery of the short lag benefits from more bins in the higher frequency range. The success in this regime by the Fourier method can be attributed to the fact that the short lag occurs at a timescale significantly shorter than the full length of the available Fairall 9 data, and that interpolation is a noticeable source of error in JAVELIN. This leads to a conclusion that high-cadence (subday to $\sim$ 1-day) monitoring of AGN will significantly benefit the recovery of the short lag, also highlighting the strength of the Fourier method when the sampling and length of light curve is adequate.

For the long lag, when using JAVELIN on real data, we robustly detect long lags in almost all wavelengths except the $M2$ band immediately adjacent to the reference $UVW2$ band, as seen in the top panel of Figure 2.2. For the $M2$ band, we recover a long lag with relatively large uncertainties when limiting JAVELIN to search for only negative lags. Since the wavelength difference between $M2$ and $UVW2$ is small, the long lag is much shorter, and it is possible that the short lag overwhelms the long lag. While we nominally detect a signal, our simulations show that $\sim$ 10% of the time, JAVELIN can pick up an apparent long lag when there is only a positive short lag injected, while only focusing on the negative range of the cross-correlation function. Hence, we plot it as an open data point. Long lags in all other bands are observed with the full cross-correlation function and are easily distinguishable from the short lag.

When using the Fourier method, we can probe lower frequencies if we do not split the light curves into two segments.Therefore, we modify the Fourier approach by applying the

maximum-likelihood method to the full length of the light curve. With this modification, we probe into the timescale relevant to the long lag we recovered with JAVELIN and observe this lag in virtually all bands. We use the value of the frequency-space lag spectrum in the lowest frequency bin as the value of the long lag, which occurs at the longest timescale for fluctuation propagation in the accretion disk we can examine with our current Fairall 9 light curve.

The long lags recovered with the Fourier method are shorter than those recovered by JAVELIN. This difference is expected. In Figure 2.11 in Appendix 2.8, we show that as the amplitude of the long lag to the short lag decreases, so does the length of the long lag recovered with the Fourier method, even with uniform cadence light curves. For example, when the amplitude ratio is 0.05 the lag recovered is roughly a factor of 2-3 times smaller than the input lag. Hence, we scaled the JAVELIN fit by a factor of ~ 2.5 for comparison in Figure 2.2, which implies an amplitude ratio of $\sim$1:20 between the long and short lags. Detecting a long negative lag for Fairall 9 with the Fourier method, which adopts a completely independent approach from JAVELIN, is very encouraging and confirms that the JAVELIN long lags are not spurious. Both JAVELIN and Fourier lags are reported in Table 2.1. See cross-correlation functions and frequency-space lag distributions in Figure 2.5, 2.6, 2.8, 2.9, and 2.10 in Appendix 2.7.

As expected, the JAVELIN long lags are all negative, which indicates an opposite propagation direction to the short lag. The lag magnitudes range from ~ -30 to ~ -80 days, increasing with wavelength until 4392 Å and then gradually flattening. In Section 2.5, we will discuss the physical picture associated with this dramatic increase in the lag within regions closer to the black hole. This recovered long lag is a better match for theoretical predictions of the slowly varying component, which separates itself with both lag direction and magnitude from the short lag occurring at the light-crossing time. Figure 2.14 in Appendix 2.8 shows that regardless of the relative amplitudes of the long and short lag, the long lag can

be reliably recovered. Hence, here we do not expect the value of the long lag we recover to be significantly biased by the short lag with JAVELIN.

To confirm that this long lag is not spurious, we conduct the same analysis with JAVELIN using simulated inputs of a single short lag sampled at the same epochs as the real data. In this case, no long lag is detected in JAVELIN. Additionally, when the real data are shuffled, no lag is detected in the cross-correlation function. Further tests with DRW models without reprocessing are discussed in Appendix 2.8, along with an example of the original and detrended Fairall 9 light curves in Figure 2.4.

As is standard in the literature, we also carried out similar analysis with PyMCCF[2] code using its modified Interpolation Cross-Correlation Function method (ICCF) (Gaskell & Sparke, 1986; Gaskell & Peterson, 1987; Oknyanskii, 1993; Peterson et al., 1998; Sun et al., 2018; Oknyansky & Oknyansky, 2022). The original ICCF method relies on cross-correlating linear interpolation of unevenly sampled light curves, whereas PyMCCF optimizes for the number of interpolated points needed to reduce noise from interpolation errors. By passing Fairall 9 light curves through this algorithm, we identify an increase in the correlation coefficient near lags of ~ -50 days, enhanced when the wavelength differences get bigger (see Figure 2.7 in Appendix 2.7). As a result, the cross-correlation function as a whole is more symmetric when the band is close to the reference band, similar to an auto-correlation function being symmetric about $x = 0$ days. This matches our expectation from the long lags detected with both the Fouier method and JAVELIN, providing another evidence of the inward-propagating long-timescale variability using a simple method.

Similar to our test case when we create a simulated Fairall 9 data set that extends for only $\sim$ 300 days with a 120-day lag injected (see Appendix 2.8), our result is limited by the length of the Fairall 9 light curves and we expect improvements when the full 1.7-year Fairall 9 light curves become publicly available. This is further bolstered by the fact that our simulations in Appendix 2.8 show that light curves with $\gtrsim$ triple the time span of the

[2] https://bitbucket.org/oknyansky/python_mccf

$\sim$ 300-day light curves allow for more accurate long lags from the Fourier method. Hence, long-baseline observations are especially important for detecting long lags in Fourier space.

To compare, Hernández Santisteban et al. (2020) reported the long lags with a $\sim$ 100-day timescale of variability but with a negative time lag of only $\lesssim$ 10 days in magnitude between the $UVW2$ and z_s bands. This analysis is carried out by performing a quadratic fit to the light curves, and then using the time shift among minima of corresponding parabolas at each band as an indication of the time lag. The fitting is based on this single peak, which introduces significant error. For Fairall 9, our reported long lag values at ~ -70 days differ by more than an order of magnitude from the previous method. The lag value we find is also more consistent with the expected inflow timescale if this lag is produced by inward propagation of disk fluctuations. Further discussion regarding the power law fitting and how this result informs the characteristics of the accretion disk are included in Section 2.5.

2.5 Wavelength dependence of the long lag

Lag	Model Description	τ_0	γ	Reduced χ^2
Long	JAVELIN	-43.1 ± 4.9	0.39 ± 0.11	75
	JAVELIN (forced to zero, w/ M2)	66.8 ± 2.3	-4.90 ± 1.4	0.57
	JAVELIN (forced to zero, w/o M2)	66.3 ± 2.5	-5.18 ± 1.9	0.68
	Fourier	-20.0 ± 2.9	0.94 ± 0.43	88
Short	detrended JAVELIN	1.99 ± 0.59	0.94 ± 0.26	2.2
	Fourier	0.19 ± 0.14	1.90 ± 0.57	0.64
	Fourier (forced to zero)	0.90 ± 0.64	0.88 ± 0.47	0.80

Table 2.2: Best-fitting parameters of simple power law models for the long and short lags in Fairall 9.

At first glance, the steep wavelength dependence of the long lag seems suspicious. Here we will argue that in principle there is a physical explanation for it, namely radial variation in the disk thickness. To better understand the implications of our lag distributions as a function of wavelength, first we use the simple functional form of a power law to fit the lags

from both methods:

$$\tau = \tau_0 \left(\frac{\lambda}{\lambda_0} \right)^{\gamma}, \tag{2.8}$$

where τ is the recovered lag, τ_0 is the normalization constant, λ_0 is the reference wavelength band at 1928Å, and γ is the power law slope. For the detrended short lag, we obtain $\tau = (1.99 \pm 0.59) \times (\lambda/\lambda_0)^{0.94 \pm 0.26}$ from JAVELIN and $\tau = (0.19 \pm 0.14) \times (\lambda/\lambda_0)^{1.90 \pm 0.57}$ from the Fourier method, where wavelengths are measured in Angstrom. The former measurement is in good agreement with the lag measurement recovered by JAVELIN for the same detrended light curves in Hernández Santisteban et al. (2020). The latter has a significantly shorter normalization factor, a trend often found when comparing lags recovered for higher frequencies with the Fourier method to lags recovered with JAVELIN (e.g., Cackett et al., 2022)

Alternatively, we can use a fiducial wavelength $\lambda_c \equiv 5468$Å at the center of our wavelength range, and force the lag at reference wavelength ($1928\mathring{A}$) to be zero (e.g. Pal et al., 2017) with

$$\tau = \tau_0 \left[\left(\frac{\lambda}{\lambda_c} \right)^{\gamma} - \left(\frac{\lambda_0}{\lambda_c} \right)^{\gamma} \right]. \tag{2.9}$$

For the long lag from the Fourier method, we recover a relationship of $\tau = (-20.0 \pm 2.9) \times (\lambda/\lambda_0)^{0.94 \pm 0.43}$ with the simple power law fit with a reduced $\chi^2 = 87.9$. For the long lag from JAVELIN, we recover $\tau = (-43.1 \pm 4.9) \times (\lambda/\lambda_0)^{0.39 \pm 0.11}$ with the simple power law fit with a reduced $\chi^2 = 75.12$. Since the Fourier result is likely a scaled version of the true long lag, we will focus on the JAVELIN result for the remainder of this section. From the reduced χ^2 values, we conclude that this simple power law is not a good fit to the long lag. On the other hand, we can force the fit to go to zero at the reference wavelength with:

$$\tau = \tau_0 \left[\left(\frac{\lambda}{\lambda_0} \right)^{\gamma} - 1 \right], \tag{2.10}$$

with $\lambda_0 = 1928$Å. In this case, we recover a relationship of $\tau = (66.8 \pm 2.3) \times [(\lambda/\lambda_0)^{-4.90 \pm 1.41} - 1]$ with a reduced $\chi^2 = 0.57$ as illustrated in Figure 2.3. Note that, in Table 2.2, including the $M2$ data point here or not does not significantly affect the fitting results.

Given our knowledge of accretion disk theory, we can derive a more sophisticated functional form for the long lag to interpret the meaning of the power law slope. Long lags are thought to be the result of fluctuations originating at large radii and propagating inward with the accretion flow at the inflow timescale (e.g. Davis & Tchekhovskoy, 2020),

$$\tau_{\rm in} = \frac{1}{\alpha\Omega}\frac{r^2}{h^2}, \tag{2.11}$$

where α is the Shakura & Sunyaev (1973) stress parameter, $\Omega = (GM/r^3)^{1/2}$ is the orbital frequency, M is the mass of the SMBH, G is the gravitational constant, h is the scale height of the disk, and r is the radial position in the disk.

To derive the timescale of the long lag, we convert $\tau_{\rm in}$, to the inflow velocity, $v_{\rm in} = r/\tau_{\rm in} = \alpha\Omega h^2/r$, and integrate between the radius corresponding to band x, r_x, and the reference radius, r_0,

$$\begin{aligned} \tau_{\rm lag} &= \int_{r_0}^{r_x} \frac{dr}{v_{\rm in}} = \frac{r_0^{2\beta}}{h_0^2\alpha\sqrt{GM}} \int_{r_0}^{r_x} r^{5/2-2\beta} dr \\ &= \frac{r_0^{2\beta}}{(7/2-2\beta)h_0^2\alpha\sqrt{GM}} \left[r_x^{7/2-2\beta} - r_0^{7/2-2\beta} \right], \end{aligned} \tag{2.12}$$

where we have added an r dependence to h, by setting $h(r) = h_0\,(r/r_0)^{\beta}$. This choice is informed by the fact that, even though h is constant in the classic thin disk model, numerical simulations suggest that the scale height possesses a radial dependence (Jiang & Blaes, 2020). This function form is also more general as a constant scale height is just the special case with $\beta = 0$. Without assuming a radial temperature dependence, putting Equation 2.12 in terms of λ gives,

$$\tau \propto \left[\left(\frac{\lambda}{\lambda_0} \right)^{\alpha(7/2-2\beta)} - 1 \right], \tag{2.13}$$

where $r \propto \lambda^{\alpha}$.

With this relationship, we can use the fitting result $\gamma = -4.96 \pm 1.49$ from Equation 10 to derive the value of β. If we assume the radial temperature dependence of the Shakura & Sunyaev (1973) thin disk model, $T_{\rm eff}^4 \propto 1/r^3$, and use the characteristic emission wavelength $T \propto 1/\lambda$, we get $r \propto \lambda^{4/3}$. Following Equation 13, when keeping α fixed at 4/3, we obtain $\beta = 3.59 \pm 1.22$. Given the relationship between the scale height and radius, this slope implies a scaling relationship where the disk puffs up at larger radii. This interpretation of a puffed-up geometrically thick disk contradicts the traditional picture of a thin accretion disk, but is often invoked to explain the line driving broad absorption line quasars (e.g. Murray et al., 1995; Leighly, 2004). The best-fitting parameters for all models discussed are shown in Table 2.2.

Here, we have not explored the possibility of a long lag extending to the X-ray bands in this work. Given the accretion rate of Fairall 9, we expect that the X-ray emission emerges from only a few to ten gravitational radii and is likely optically thin (Jiang et al., 2019). Therefore, the fluctuation propagation that drives the long lag may not translate to the X-ray emission region. Furthermore, the relationship between the X-ray variability and the short lag remains complicated (Edelson et al., 2019; Mahmoud & Done, 2020; Hagen & Done, 2023), causing us to defer more detailed analysis of the X-ray light curves. For future studies, since obtaining more wavelength coverage at higher frequencies is difficult, a more complete demographic analysis of the AGN population will help distinguish if this slope is ubiquitous in all systems. At the same time, numerical simulations of AGN disks can enhance our understanding of the possible distribution of observed lags.

2.6 Summary

Variability studies of AGN not only inform basic properties of their accretion disks, but also have the potential to illuminate the fundamentals of accretion physics. For example,

past work such as Cackett et al. (2022) show that multiple lags can be recovered from UV/optical continuum reverberation mapping due to distinct processes occurring in the AGN disk. For this purpose, continuous and high-cadence monitoring of AGN with multi-wavelength coverage, such as the intensive disk-reverberation campaign of galaxy Fairall 9 is required. We apply both the traditional cross-correlation method with JAVELIN and a Fourier method in frequency space to simulated and real data sets, in an attempt to recover both the well-known and rapidly varying short lag that propagates outward, and the more elusive and slowly varying long lag that propagates inwards, in AGN light curves.

In our simulated DRW light curves, with a length slightly longer than 2 years and ∼day cadence, JAVELIN produces a highly accurate recovery of the long lag while failing to accurately localize the short lag without detrending. The Fourier method can easily identify the timescale at which both processes are occurring, and can very accurately recover the shorter lag while constraining both lags simultaneously. However, JAVELIN short lags and Fourier long lags can be biased by the other lag, especially when the other lag has a strong response. On the other hand, long lags with JAVELIN and short lags with Fourier methods can almost always be accurately recovered. Note that all our tests are conducted with DRW light curves to simulate aperiodic AGN variability. However, discrepancies between this assumption and the real light curve behaviour have been observed (e.g. Mushotzky et al., 2011; Kasliwal et al., 2015; Smith et al., 2018).

When performing the same analysis on real Fairall 9 light curves, both methods successfully identify short and long lags that increase in length with respect to the wavelength of observation; the Fourier method measures the short lags with significantly higher precision. Detrended JAVELIN short lags are twice as long as they are from the full undetrended light curves, and better match the Hernández Santisteban et al. (2020) results. We also present the first recovery of a long lag in AGN measured through two independent methods, where JAVELIN constrains the long lag at ~ -70 days between bands with the largest wavelength difference, propagating inwards in the accretion disk. This is further supported by the re-

sults from the more straightforward PyMCCF method, where we see an enhanced correlation coefficient for a negative long lag at larger wavelength difference between cross-correlated bands.

The long lag detection by the Fourier method at ~ -30 days provides an independent confirmation that JAVELIN long lags are not spurious. In fact, based on our simulations in Figure 2.11 in Appendix 2.8, the Fourier method is expected to measure long lags significantly shorter than the real lag even in ideal situations. The factor by which the long lag measured by the Fourier method will be shorter than the real lag depends on the ratio between the amplitudes of the long and short lags. For example, in this case the JAVELIN and Fourier results differ by a factor of ~ 2.5, which if you go off of Figure 2.11 implies an amplitude ratio of $\sim 1:20$ between the long and short lags.

The Fourier method is sensitive to bin choices. Having more bins allows more precision in determining the relevant timescale of the short lag. However, for the long lag, given the current length of optical light curves, we can only probe into the relevant timescale when using the whole light curve as one segment and lower number of frequency bins. This setting will need to be tuned when much longer light curves become available. Future work is advised to use the Fourier method when the lengths of available high-cadence light curves are a few times longer than that of the longest lag. With a longer time span, the Fourier method performs significantly better, and can also reveal information that the cross-correlation method is incapable of, such as separating the timescales of distinct processes. Additionally, we suggest applying both methods to the same data set, so that the Fourier method can provide verification of recoveries and separate them from spurious signals in the cross-correlation method. A good opportunity for this test is with the upcoming release of the full 1.7-year intensive disk reverberation mapping of Fairall 9 with daily *Swift* and LCO monitoring. Having twice the length of the data we are currently using, this release could probe a larger frequency range with the Fourier method.

Overall, our methods show promising results that raise the exciting prospect of finding long lags in more AGN, and thus mapping accretion disk structure as a function of AGN parameters. Many other past and future observations can be potentially useful for the same purpose. Stone et al. (2022) compiled *gri* light curves of a sample of 190 quasars within SDSS Stripe 82 spanning $\sim 1998-2020$ from SDSS, PanSTARRS-1, the Dark Energy Survey, and follow-up monitoring with Blanco 4m/DECam. Other potentially useful data sets include the ASAS-SN survey in Yuk et al. (2022) and the upcoming decade-long Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST). Extracting the long lag from these data sets will require accurate analysis since the limited wavelength difference (e.g. from u to i in LSST) will lead to much shorter long lags. An added challenge here is, in some cases, the sparse sampling or large seasonal gaps. Because of LSST's uniquely high cadence and long time baseline, it will be suitable for searching for the long lag in a larger population of AGN. We will explore this possibility with a suite of simulated LSST light curves in a follow-up paper, with improved metrics to overcome inevitable seasonal gaps (Secunda et al., 2023). Future work can focus on distinguishing spurious signals from these sources in both frequency and temporal space, which may aid lag recovery. Additionally, testing our methods on a select sample of the these quasars, especially high-redshift galaxies where lags will become more prominent even with short wavelength coverage, will further verify the applicability of both the Fourier method and JAVELIN overall. The derived lag distribution from such a population of AGN will in turn deliver precious information regarding accretion physics from distant galactic central engines.

2.7 Additional Materials: Fairall 9 Lags

Here we show cross-correlation and Fourier lag distributions in temporal and frequency spaces, respectively. Figure 2.4 gives an example of detrending Fairall 9 light curves to remove long-timescale variability, making it more suitable and accurate for obtaining short

lags. Then we show original and detrended JAVELIN cross-correlation functions in Figure 2.5 and 2.6. The frequency-space lag distributions from the Fourier method are included in Figure 2.8 and 2.9. The former uses more bins and separates the light curve into two segments, focusing on extracting the short lag with an error-weighted mean of the orange points; the latter uses less bins and the whole light curve to focus on extracting the long lag, taking the value of the orange point (lowest frequency bin) as the value of the long lag. Finally, we measure the long lag in the $M2$ band (2246 Å, *Swift*) with JAVELIN separately with large uncertainty. This is done with JAVELIN only searching for negative lags to isolate the long lag from the short lag, with the cross-correlation function shown in Figure 2.10.

2.8 Additional Materials: Simulated light curves

In this section, we test both JAVELIN and the Fourier method with a variety of DRW light curves covering a range of parameters. Hence, we demonstrate here that the recovery of short and long lags in Section 2.4 is robust and reasonable. Given its sampling cadence and wavelength coverage, the intensive disk-reverberation mapping campaign of Fairall 9 is one of the most unique and extensive datasets currently in existence. This is also an improving set of light curves which is expected to double its current length in the foreseeable future. As a result, many of the test cases presented in this paper are built from the characteristics of the Fairall 9 data set, including the cadence, length, and error of our simulated light curves.

2.8.1 Simulations

To demonstrate the ability to constrain lags for each of our methods, we generate driving light curves following Timmer & Koenig (1995)'s DRW model assuming a PSD with a slope of $D(\nu) \propto \nu^{-2}$ and then transition to $D(\nu) \propto \nu^{-2}$ below some damping frequency. Then, we reprocess the light curve in frequency space with a response function following equation 2.3; when necessary, we apply an inverse-Fourier transform to $R(\nu)$ to obtain the reprocessed light curve in temporal space. Since we expect lags both with $\sim$ several-day and $\sim$ 100-day timescales in opposite directions, we construct our response function as a composite of two Gaussian functions each in the form of:

$$\psi(\mathrm{t}) = \frac{1}{\mathrm{St}\sqrt{2\pi}} \exp^{\frac{-(\mathrm{t}-\mathrm{M})^2}{2\mathrm{S}^2}}. \tag{2.14}$$

We then vary M and S to change the amplitude and width ratio between the two Gaussian functions. Unless otherwise specified, we use an amplitude ratio of 5:1 and a width ratio of 1:24 between the response functions of the short and long lags. We place our Gaussian centroids at +5 and -120 days respectively for test purposes such that they take place at different timescales and opposite directions, thus matching the physical motivation discussed

earlier. The resultant response function is presented in the first panel of figure 2.11 with a range of amplitude ratios from 1 to 0.01 between the -120-day and +5-day Gaussian functions. To create lags with opposite directions, the negative lag is a Gaussian function with negative time. Intuitively, larger S means a broader distribution, and M is the mean of the Gaussian.

For uniform-cadence light curves, given current limitations of ground-based observatories, we adopt a one day cadence and a 1 year length that matches the best currently available light curves. For non-uniform cadence light curves, we regenerate the DRW light curve and sample the data points according to the Modified Julian Dates (MJD) in the data from Hernández Santisteban et al. (2020). We use the signal-to-noise ratio of the same data sets to generate errors and perturb our simulated light curves.

2.8.1.1 Uniform cadence time-series

With uniformly-sampled light curves, we are able to use both the simple cross-correlation analysis presented in 2.3.1 and simple frequency-resolved analysis presented in 2.3.2. As seen in figure 2.11, the cross-correlation method identifies the long lag easily when its amplitude in the response function is comparable to its shorter lag counterpart, and the short lag is detectable at most ratios except when the amplitude of the long and short lags are comparable. On the other hand, frequency-resolved analysis can recover both lags at all frequencies, with outstanding accuracy and consistency for the short lag, but the accuracy for the long lag decays with the strength of its signal. We should also note that the cross-correlation method can fail occasionally when asked to constrain both the long and short lags simultaneously. By running 100 test runs of cross-correlation on random simulated light curves, the fraction that fails to identify both lags increases with decreasing signal strength from the longer lag. The cross-correlation example in figure 2.11 further elaborates on this by showing the inability for cross-correlation to identify the long lag when the amplitude ratio falls below 5%.

2.8.1.2 Non-uniform cadence time-series

Since UV/optical variations arise at a timescale much longer than a day, astrophysical observations naturally produce gaps in their time-series, hence obtaining uniformly sampled data becomes infeasible. JAVELIN and the Zoghbi et al. (2013) methods act as alternatives in this regime for lag determination without significantly compromising the results of the analysis.

The driving light curve in Figure 2.12 comes from the DRW model with a length of 829.86 days. This length is triple that of the $UVW2$ (far-UV from *Switft*) band light curve of Hernández Santisteban et al. (2020); the flux uncertainty comes from the S/N ratio in the Fairall 9 light curve as well. The Fourier method relies on parametrizing values of power and lags in a preset of frequency bins, but the choice of the number of bins can have a significant impact on the lags we can recover. The lowest frequency bin is set by the total length of the light curve. Having too many bins tends to distribute the power to separate bins while significantly increasing the uncertainty of the time lag with biases introduced. With an appropriate number of bins, the short lag is often visible with negligible errors independent of the light curve we use, usually confirmed by several of its adjacent bins in the frequency range between 10^{-2} and 10^{-1}. With input lags at $+5$ and -120 days, when adopting 13 bins, ~ 3 bins in this range measures an average short lag of 3.85 ± 2.73, but with 8 bins, only 2 of them show a short lag. In this case, 13 bins is more suitable because it also returns a more accurate long lag at -109 ± 15 days compared to having 7 bins which returns a long lag of -81 ± 46 days. In virtually all test cases, the long lag tends to lie within 1 sigma of the lag in the first or second frequency bin; with a sufficiently long light curve, we can see a peak that indicates the exact timescale of the long lag.

To compare the performance of this method in various astrophysical conditions with different physical properties, we show in figure 2.13 the results of 300 different simulated light curves with the same parameters as above but reprocessed with different response functions: the amplitude ratio between the long lag and the short lag is 20%, 10% and 5%

respectively. Because the length of the light curve is not significantly longer than the longest input lag, we take the longest lags in the lowest frequency bins as the long lag. We then take a weighted average for the short lag from bins within frequency ranges between 10^{-2} and 10^{-1}. Each lag is weighted by the inverse square of its error. These bin choices are informed by the fact that these two lags would almost always occur at their characteristic Fourier frequencies; additionally, these bins produce results that are most consistent with the input and cross-correlation lags in our test runs. The Fourier method continues its success in the frequency domain at recovering the short lag in all cases. The recovered long lag, on the other hand, becomes more accurate when its amplitude in the response function is high, and vice versa. However, we note that in about 70% of all test cases, the Fourier approach can confidently recover a long lag.

By interpolating the same light curve to uniform cadence and to 10 times more data points, JAVELIN is more consistent in recovering a long lag with even higher precision. In this case, the median value of the recovered long lag is -113.03 ± 2.68 days. However, interpolation at short timescales seem to introduce higher uncertainties to measuring the short lag: our input short lag of +5 days lies just outside the 1 sigma region of the recovered short lag at 20.94 ± 13.02 days. A lot more scatter is seen in recovering the short lag.

We then apply JAVELIN to the lag extraction of the 300 simulated light curves. Each long and short lag is taken from the median of the probabilistic distribution from each test run. We can observe from Figure 2.14 that the recovered long lag is very accurate and precise in all cases. On the other hand, by observing the scatter in the right panel, the recovered short lag becomes much more accurate when the signal from it is much stronger than that from the longer lag.

However, a 800-day high cadence light curve is not always available for targets of interest. Similar tests that we carried out with a $\sim$ 300-day light curve (3x shorter) shows that the Fourier method can still accurately and precisely recover the short lag at 5.6 ± 1.0 days, but the limited length shortens the frequency range and we cannot recover the full long lag

resulting in a measurement of -25.1 ± 26.1 days. On the other hand, JAVELIN still produces a scattered estimate of the short lag at 33.7 ± 9.2 days, but is consistent in recovering the long lag at -120.2 ± 1.9 days.

To confirm that our results do not come from artificial sources such as sample cadence and light curve length, we performed all previous tests on two randomly generated DRW light curves with no reprocessing nor any correlation between them. In these cases, neither method recovers any statistically significant lag.

Results from the two methods allow us to conclude that both methods have their strengths and weaknesses, especially among different physical regimes. For different light curve lengths and cadences, when recovering the long lag, the Fourier method requires the light curves to be a few times longer than the long lag while JAVELIN can provide accurate result with a much shorter one. When recovering the short lag, regardless of the length or cadence, the Fourier method performs much better even lacking an interpolation to uniform cadence. Both methods work best when the AGN is sampled with $\sim$day-cadence. One unique strength of the Fourier method JAVELIN lacks is that the former informs the precise timescales at which the physical processes take place. As a result, combining these two methods when analyzing light curves is the best way to constrain two lags with opposite directions simultaneously. To place the tightest constraint on both lags, future work is advised to use the results from the Fourier method as a prior for JAVELIN to carry out cross-correlation studies.

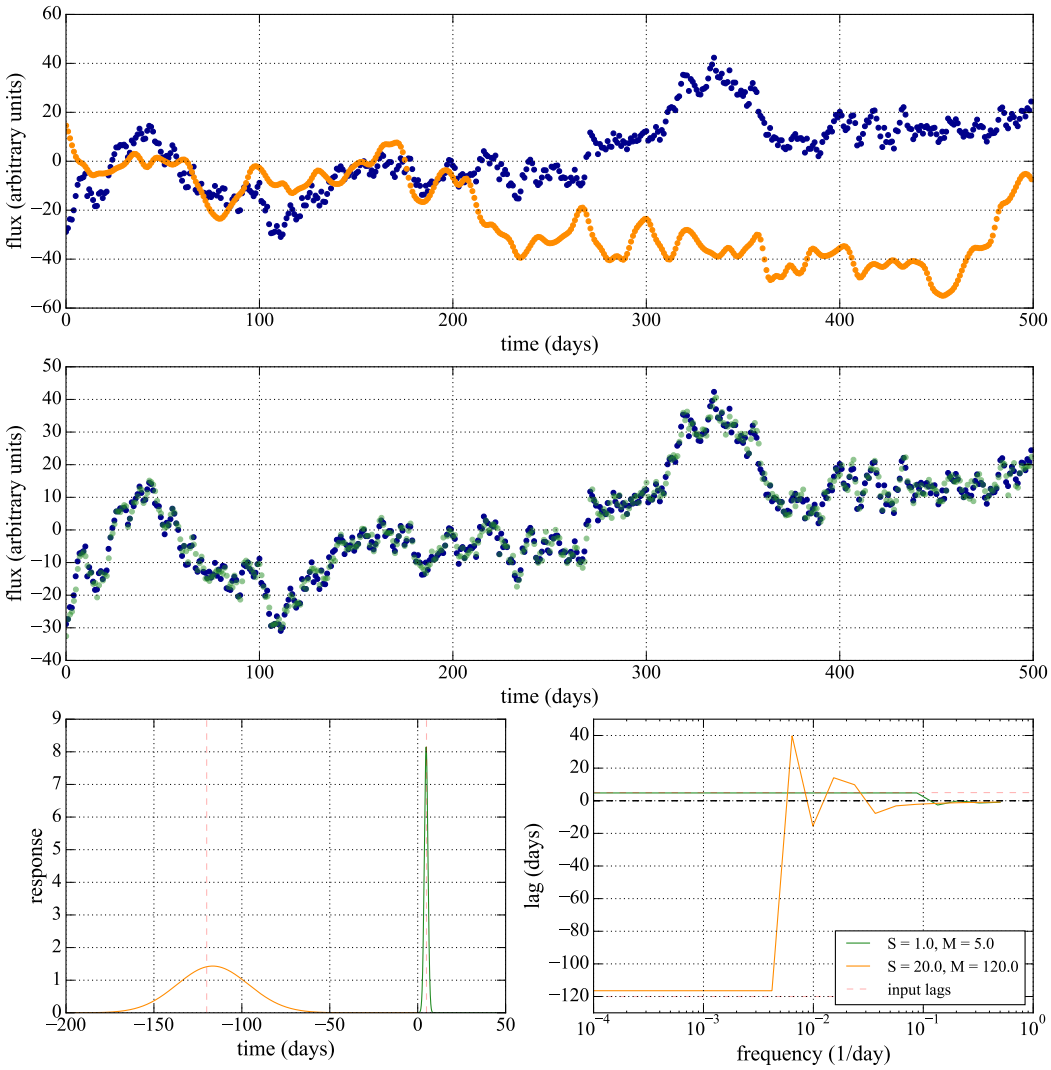


Figure 2.1: Top panel: simulated DRW light curve (blue) with uniform cadence, and reprocessed light curve with a long lag (orange). Middle panel: simulated DRW light curve (blue) with uniform cadence, and reprocessed light curve with a short lag (green). Each light curve has a length of 500 days and is sampled once per day. Lower left: response functions that contains either a long lag (orange) or short lag (green). The ratio of the amplitudes for each lag is set to be 1:5 for test purposes. Lower right: recovered long (orange) and short (green) lags with frequency-resolved lag analysis. In this ideal case, both lags can be easily recovered and identified by the flattening of lags at lower frequency.

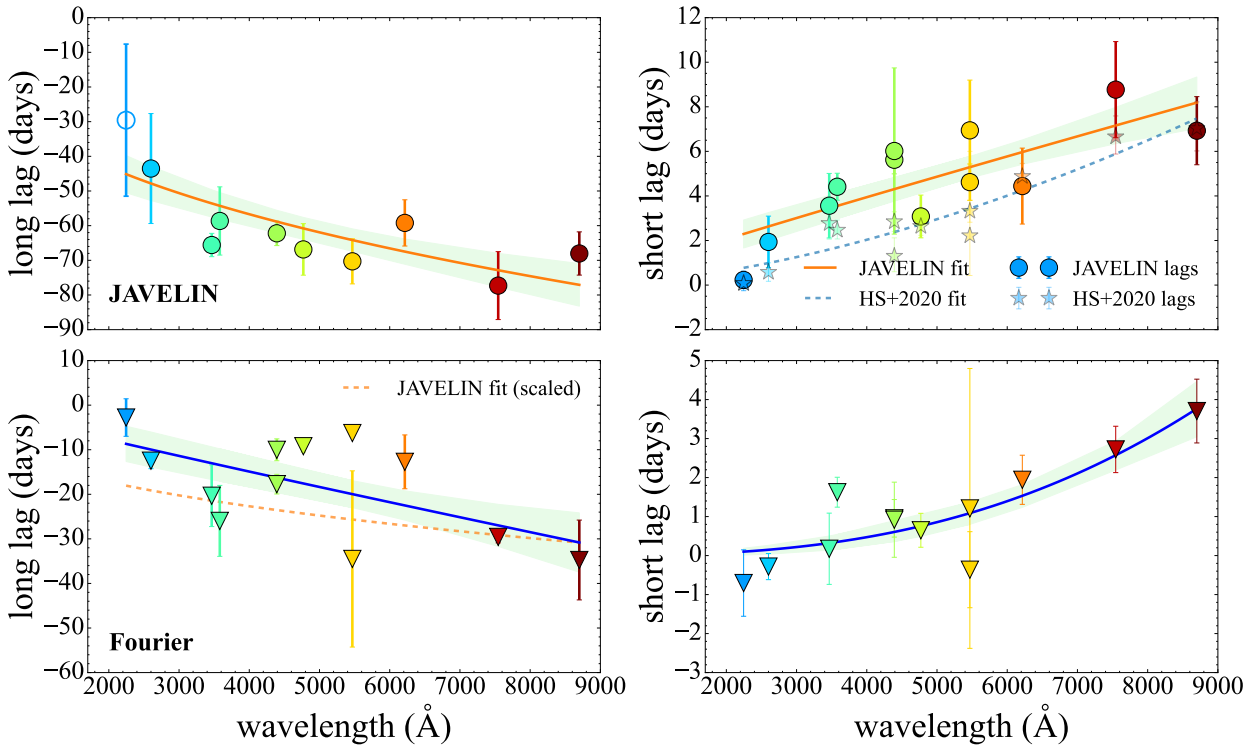


Figure 2.2: JAVELIN long lags (top left) and detrended short lags (top right), and Fourier method long (bottom left) and short lags (bottom right) recovered from the Fairall 9 data as a function of wavelength. Light curves are obtained from Hernández Santisteban et al. (2020). The $UVW2$ band is used as a reference band to perform cross-correlation against by JAVELIN. Long lags ~ -70 and ~ -30 days are detected with JAVELIN and the Fourier method, repsectively, in all bands with a sign opposite to the short lags. All lags are increasing in magnitude with wavelength. The simple power law fits with 1σ uncertainty are also shown as the blue solid line and the green shaded region. In the bottom left panel we show the JAVELIN fit scaled by a factor of 2.5 to compare with the Fourier results. Note that in JAVELIN the long lag in the $M2$ band (2246 Å, *Swift*) is measured separately with large uncertainty, so we plotted it as a open point. For easier comparison, cross-correlation short lags from Hernández Santisteban et al. (2020) (labeled as HS+2020) are overplotted with star-shaped points with the same power law fit in dashed lines, where circular points are JAVELIN short lags from detrended light curves. For the cross-correlation functions see Figure 2.5 & 2.6, and Figure 2.10 for the long lag in M2. For the frequency-space lag distribution see Figure 2.8 & 2.9.

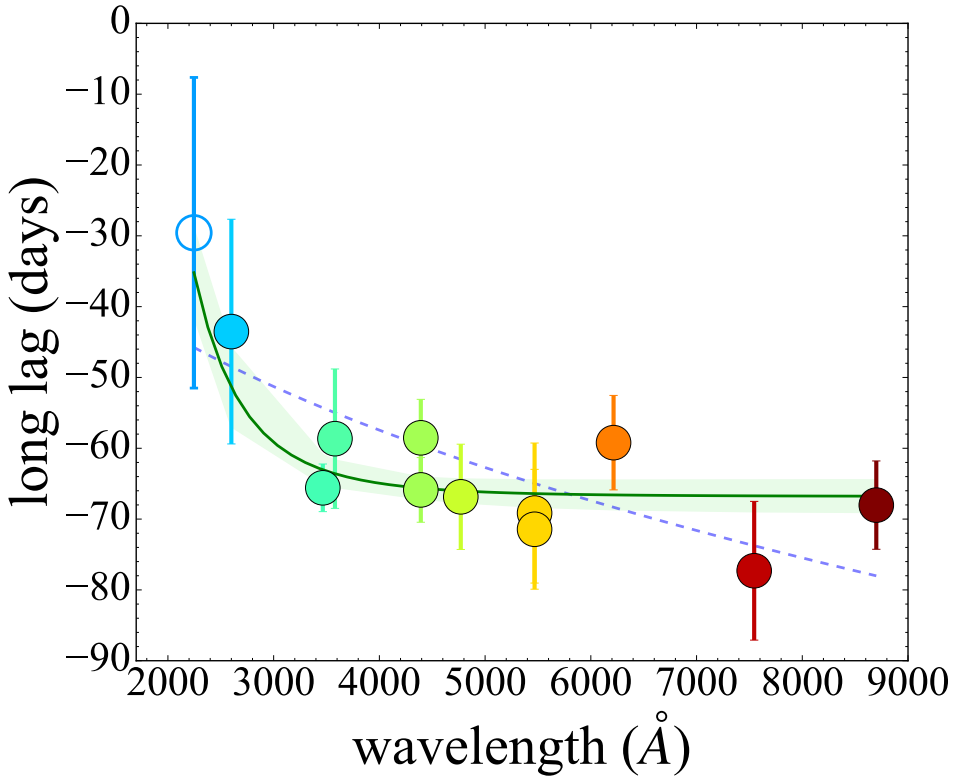


Figure 2.3: JAVELIN Long lags recovered from Fairall 9 as in Figure 2.2 but fitted by forcing the power law to go to zero at the reference wavelength. The green shaded region reflects the 1σ uncertainty of this fit. The blue dotted line shows the simple power law fit that resembles the fit in Figure 2.2 in the form of $\tau = \tau_0(\lambda/\lambda_0)^\gamma$. Note that in JAVELIN the long lag in the $M2$ band (2246 Å, *Swift*) is measured separately with large uncertainty, so we plotted it as a open point. Nevertheless, including this point or not does not significantly affect our fitting results presented in Table 2.2.

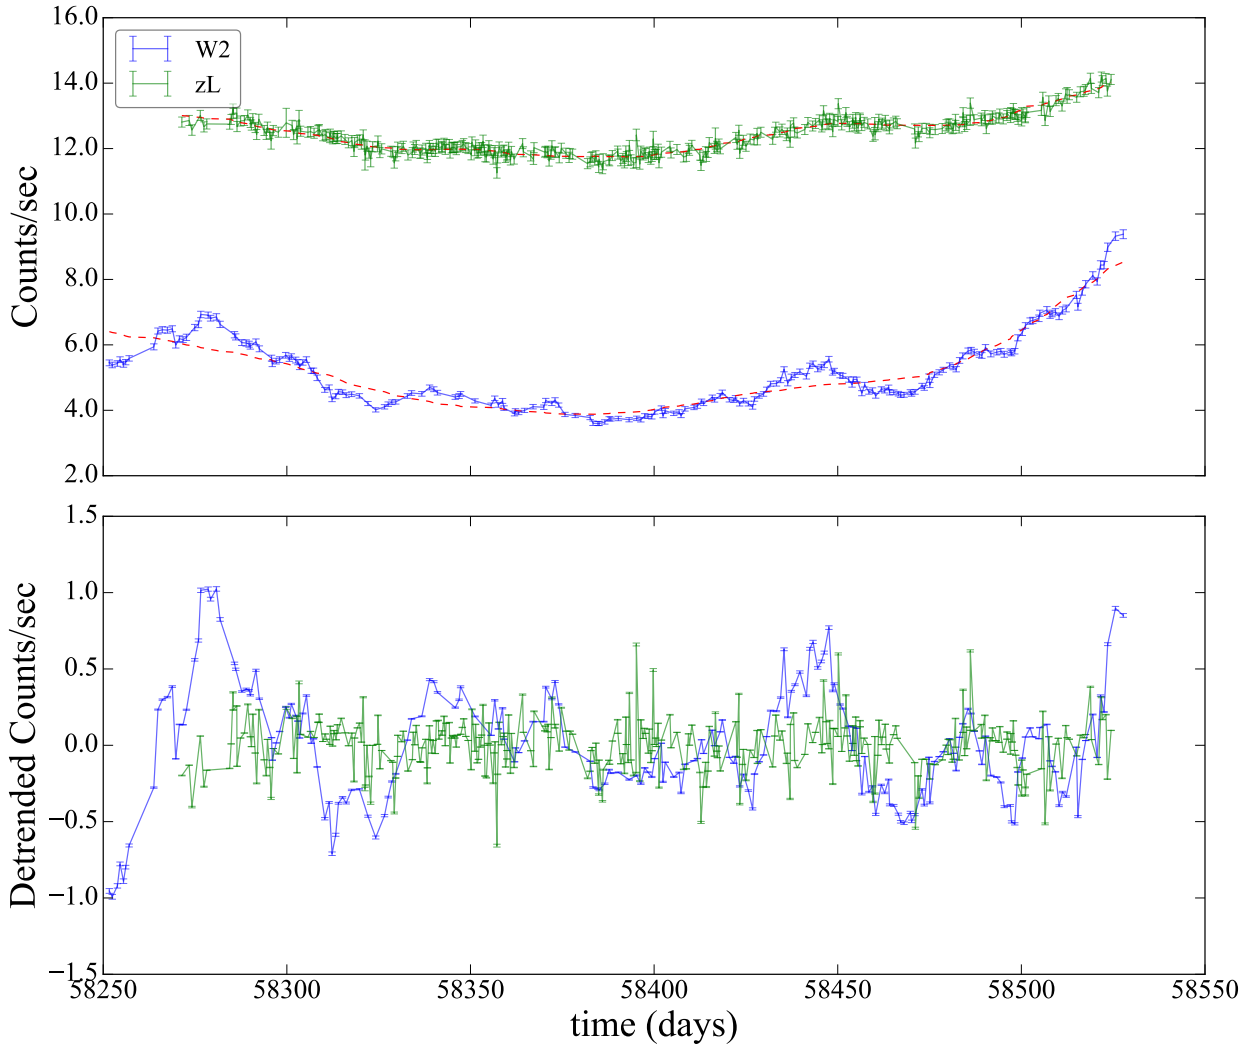


Figure 2.4: Top panel: example real light curves from Hernández Santisteban et al. (2020) for Fairall 9 in the reference $UVW2$ band (blue) and the longest wavelength band z_s (green). Parabolas for detrending are plotted in red. Bottom panel: light curves after detrending to obtain JAVELIN short lags.

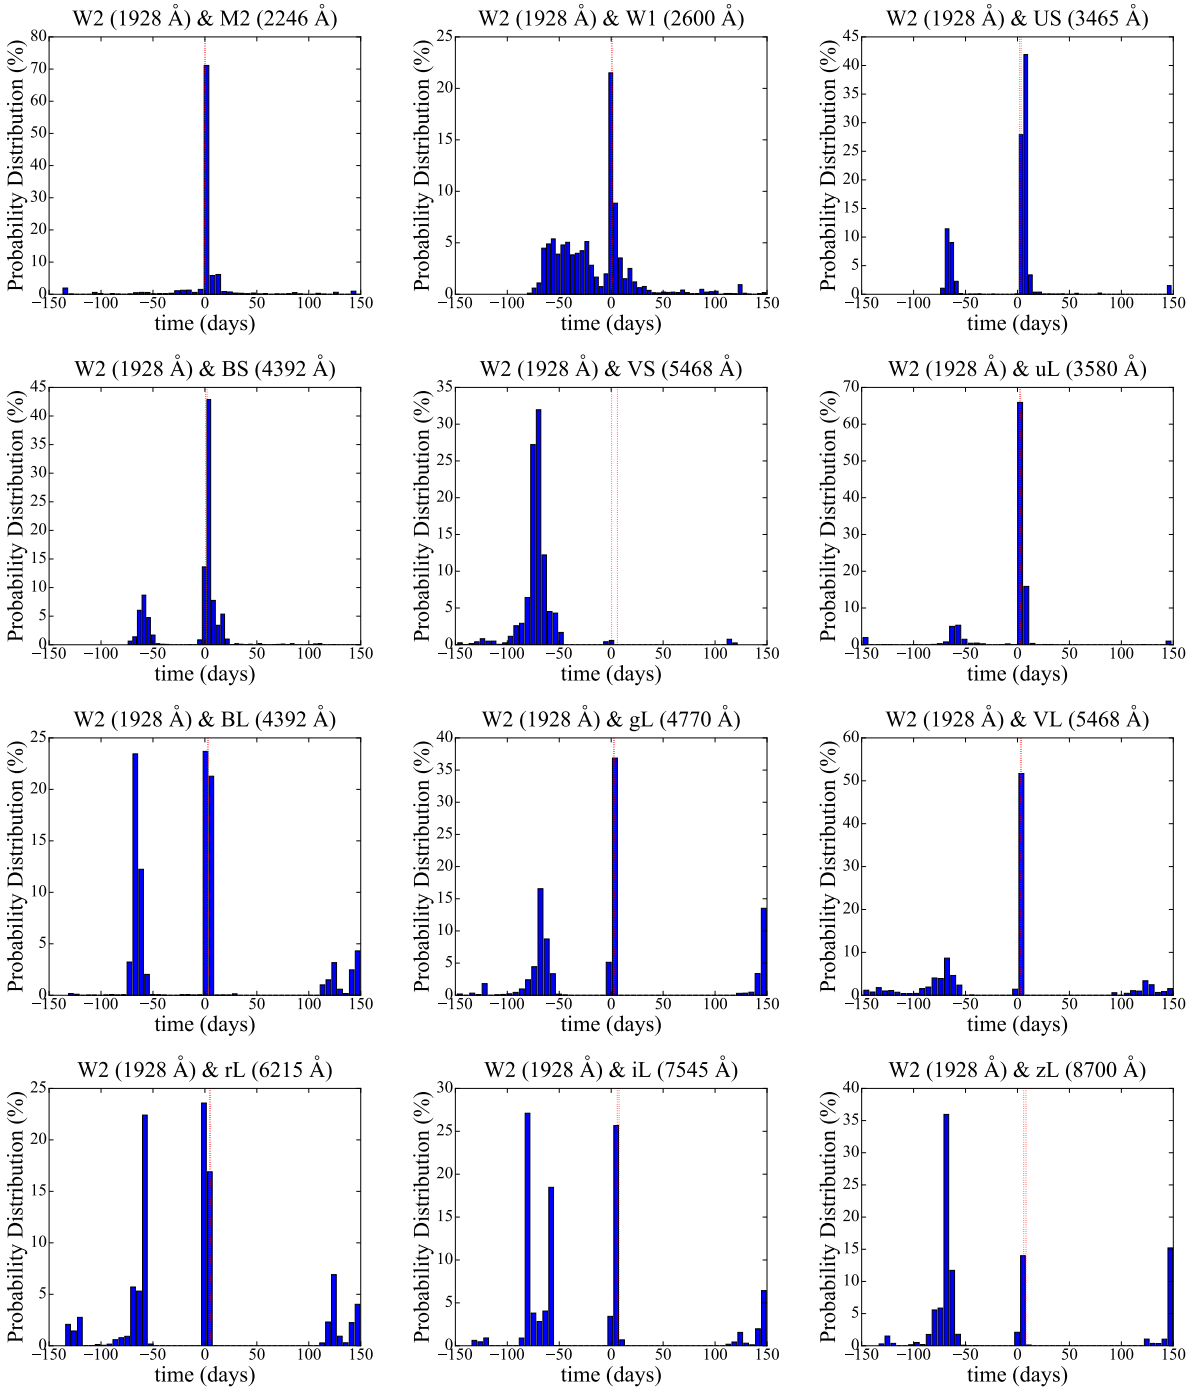


Figure 2.5: Cross-correlation functions from JAVELIN showing recovered long and short lags from Fairall 9 at each band relative to the $UVW2$ band at 1982 Å. The red dotted lines represents the 1σ range of the recovered short lags in Hernández Santisteban et al. (2020).

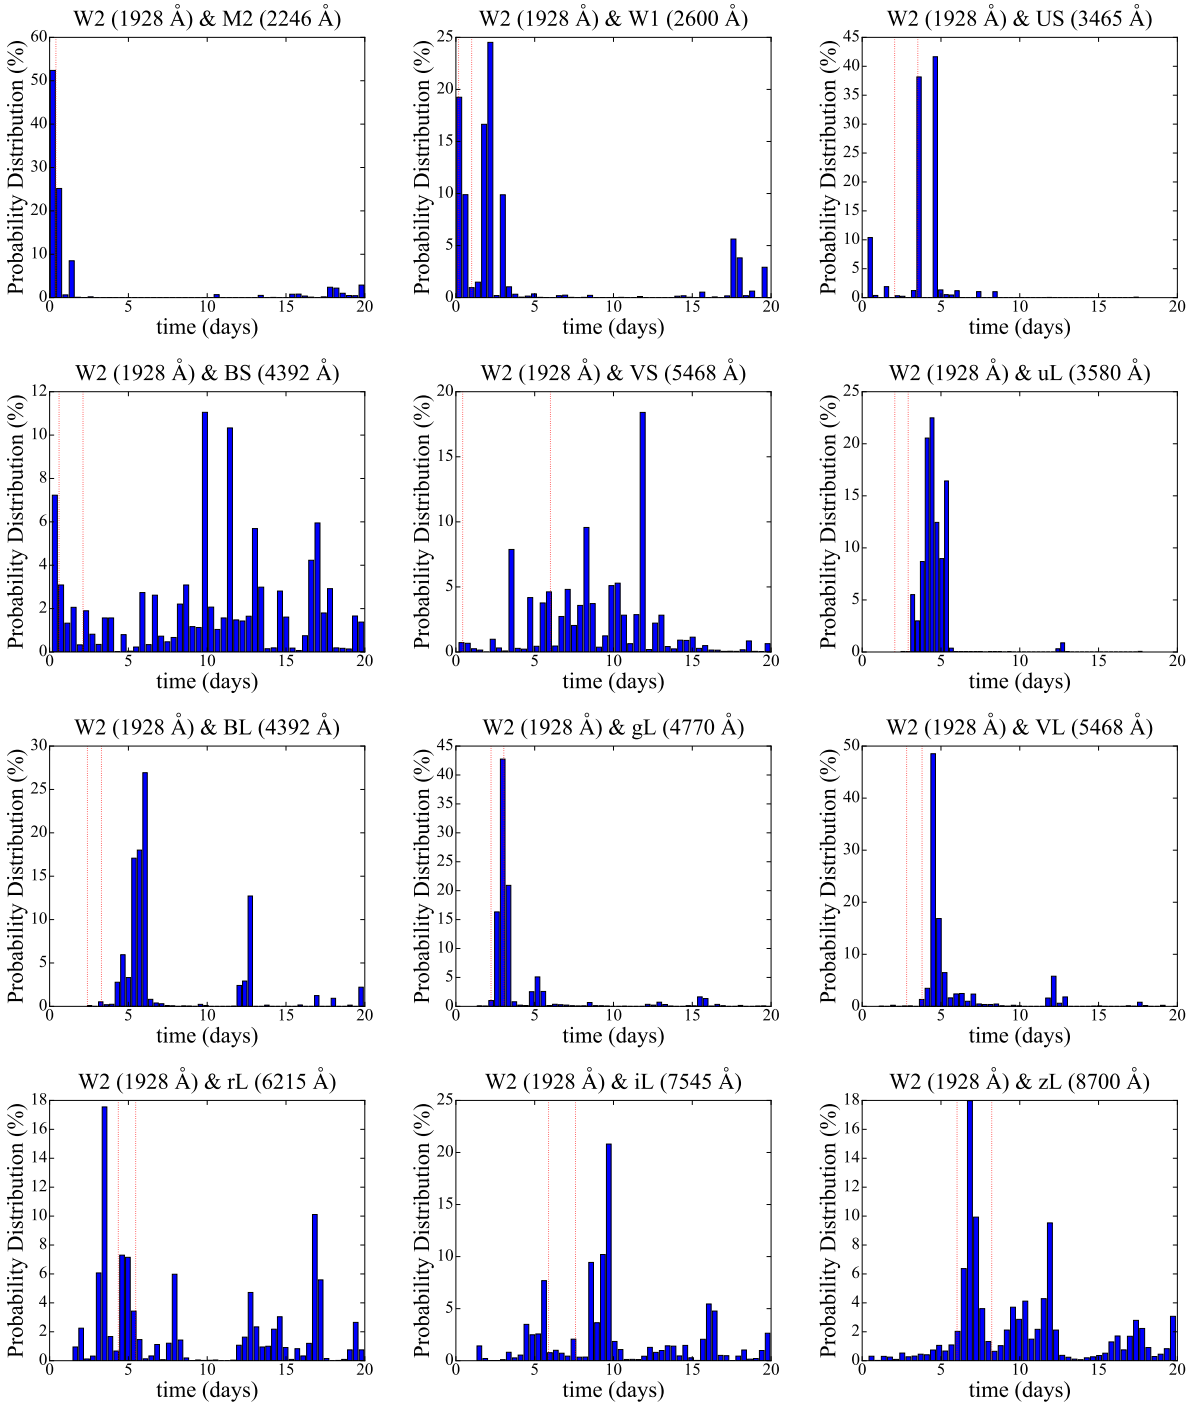


Figure 2.6: Cross-correlation functions from detrended JAVELIN light curves showing recovered short lags from Fairall 9 at each band relative to the $UVW2$ band at 1982 Å. The red dotted lines represents the 1σ range of the recovered short lags in Hernández Santisteban et al. (2020).

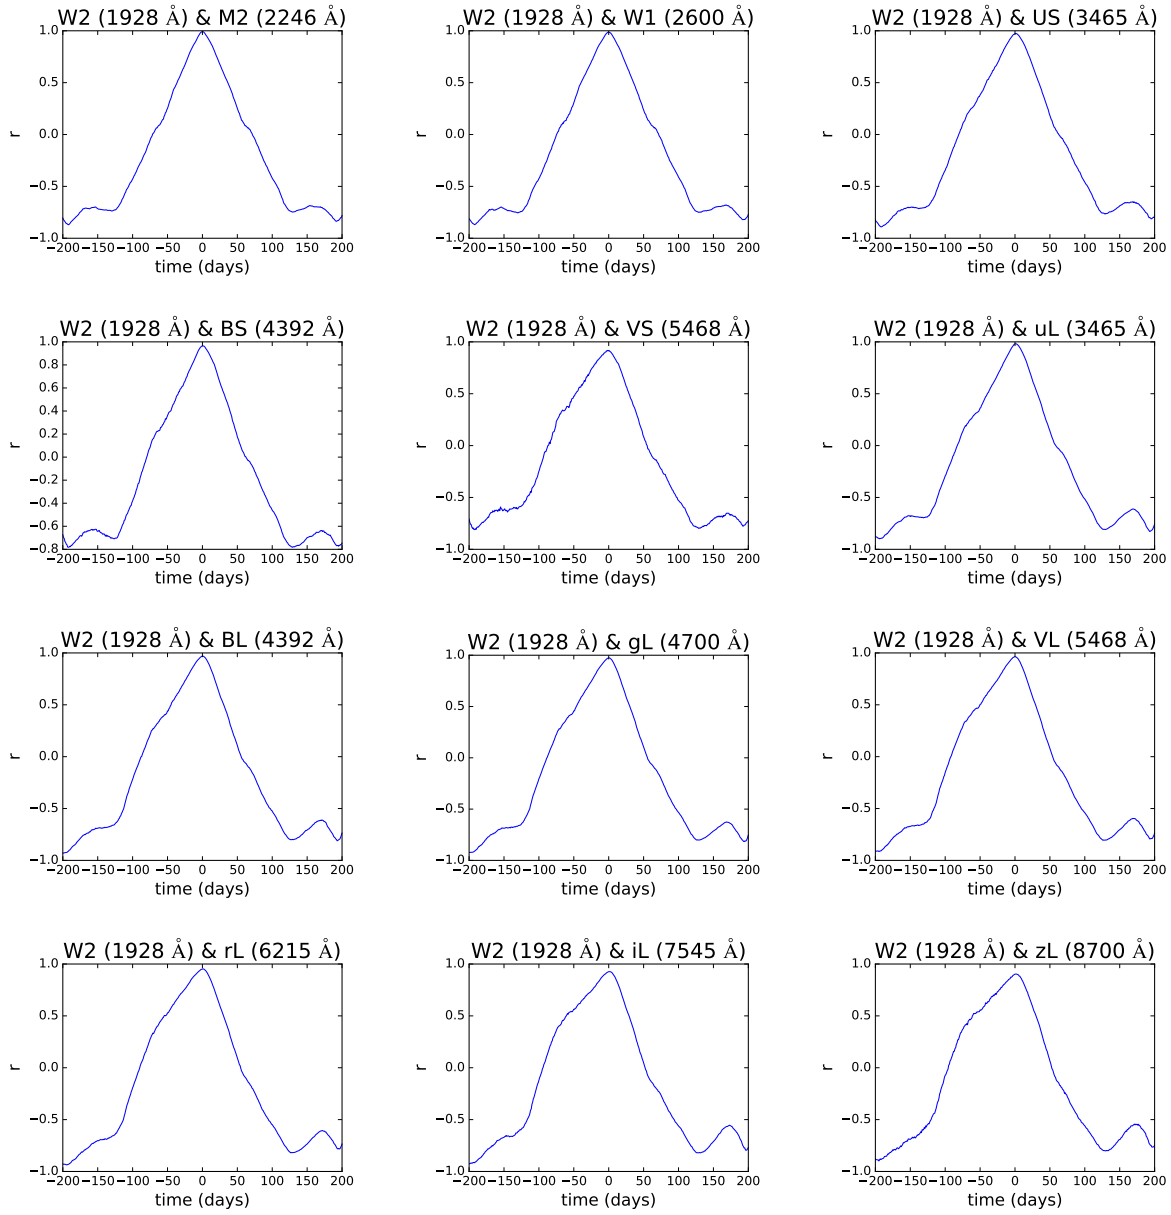


Figure 2.7: Cross-correlation functions using the ICCF method with PyMCCF showing the correlation coefficient for a range of lag times from Fairall 9 at each band relative to the $UVW2$ band at 1982 Å. The shape of the lag distribution skews towards negative lags as the wavelength difference increases, implying the existence of a long negative lag.

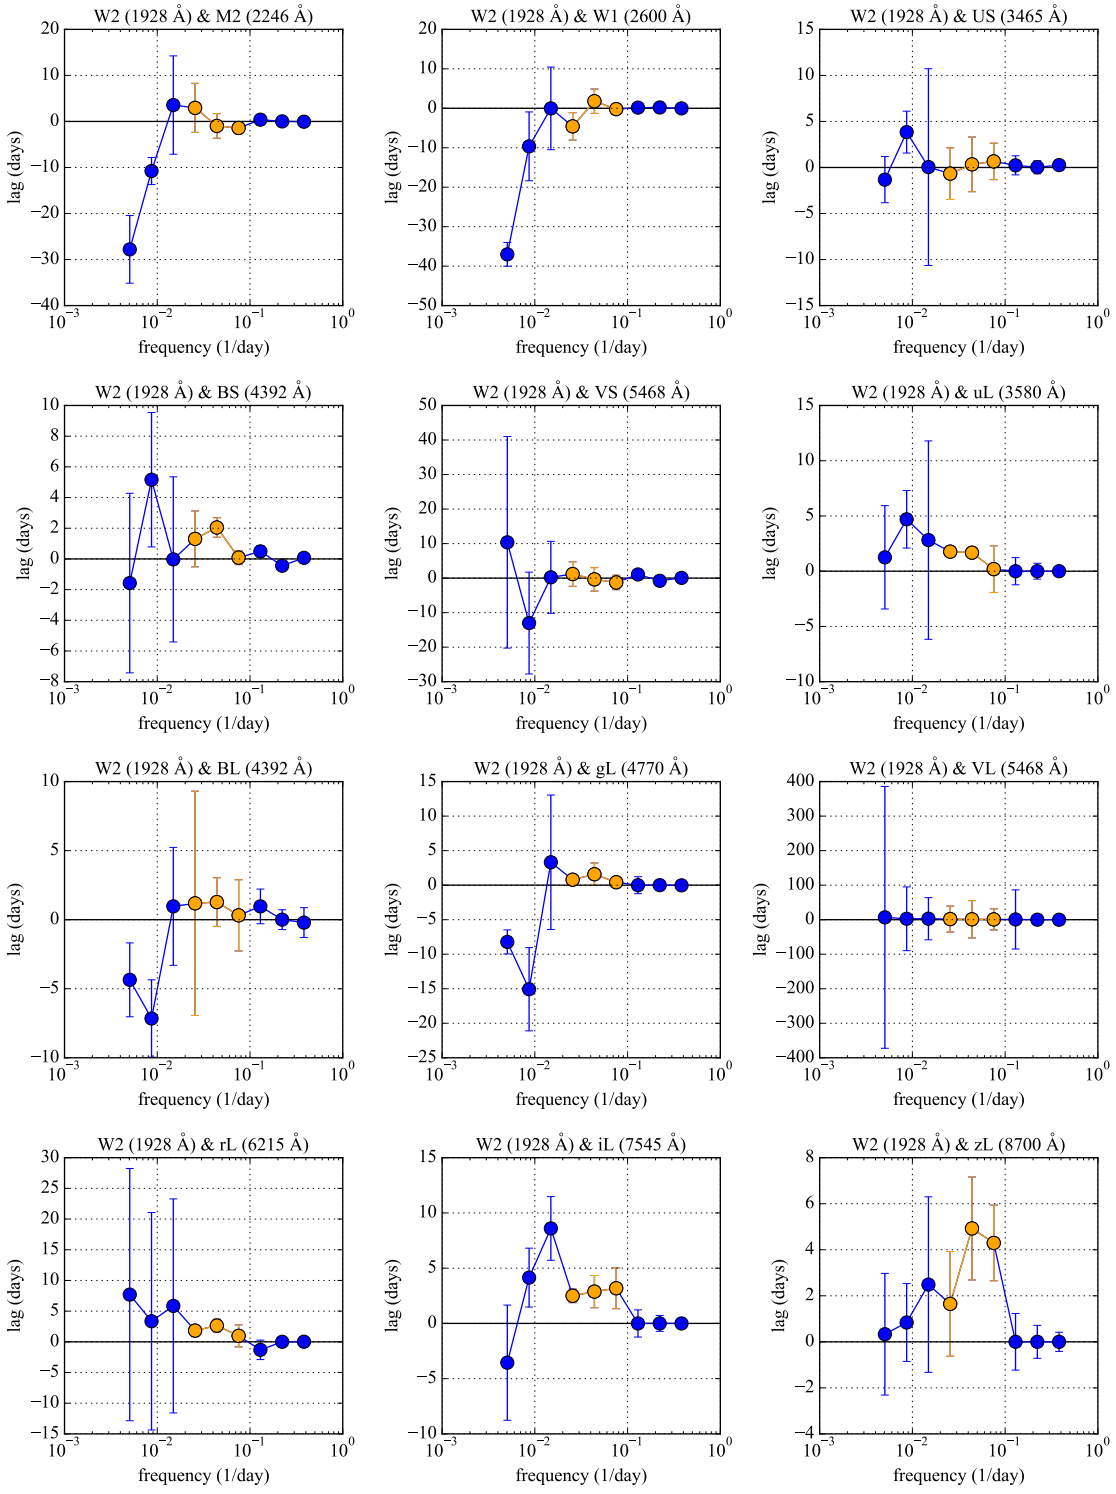


Figure 2.8: Lag distributions in frequency space with original light curve from Hernández Santisteban et al. (2020) separated into two segments and 9 bins, which is more efficient for detecting short lags. We take the the error weighted mean of the orange points as the value of the short lag.

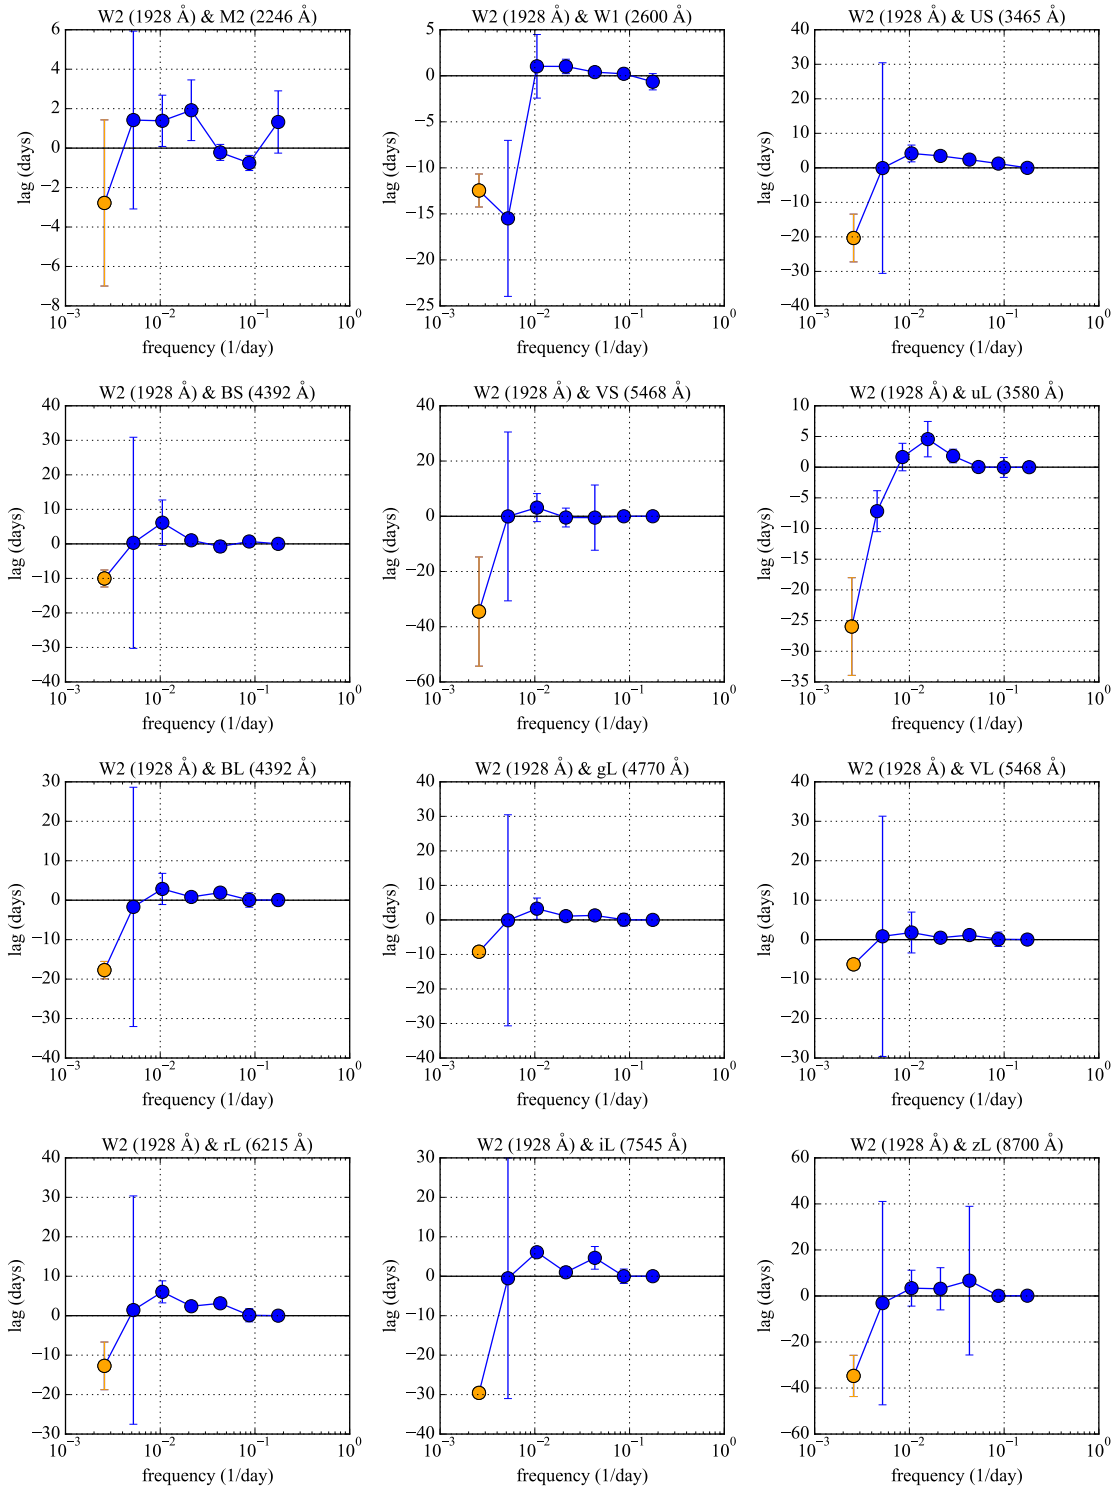


Figure 2.9: Lag distributions in frequency space with original light curve from Hernández Santisteban et al. (2020) with only one segment and 7 bins, which is more suitable for detecting long lags. We take the the first lag bin (orange) as the value of the long lag.

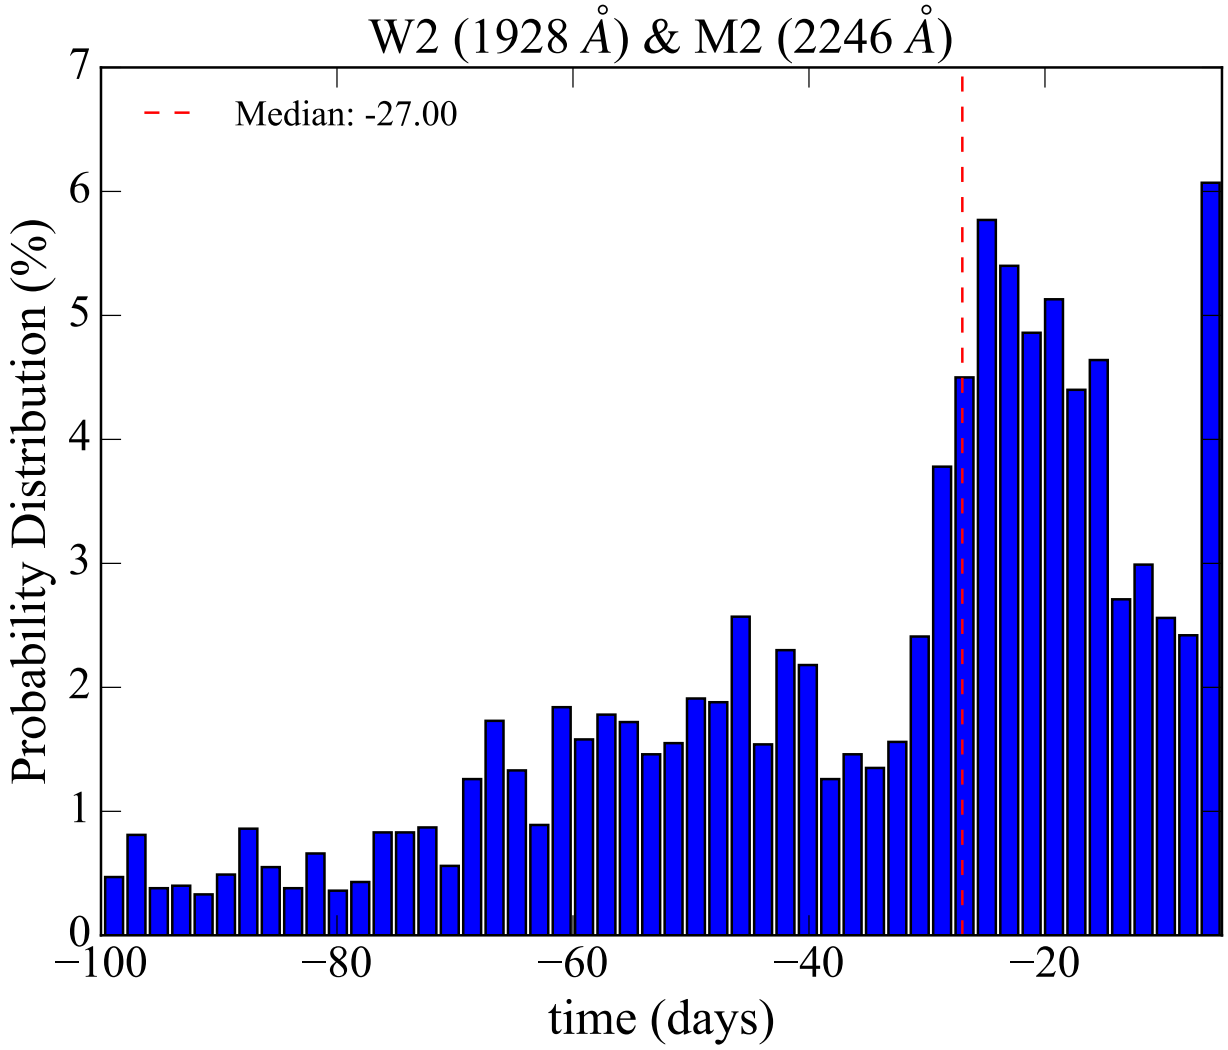


Figure 2.10: Cross-correlation functions from JAVELIN between the reference $UVW2$ band at 1982 Å and the the $M2$ band at 2246 Å. Here, we limit JAVELIN to only search for negative lags, because the power of the short lag significantly dominates over the long lag at this wavelength.

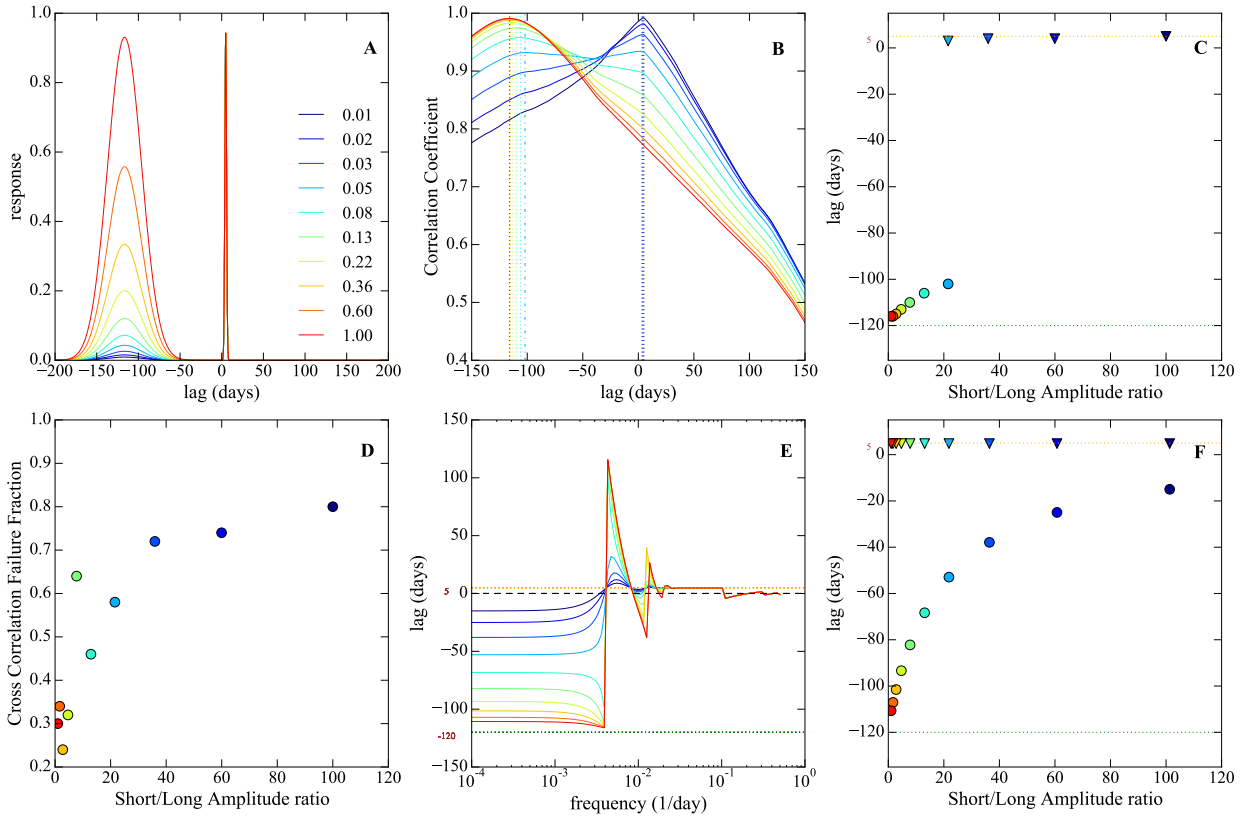


Figure 2.11: Comparison between the frequency-resolved and cross-correlation approaches as a function of the amplitude ratio between the two lags' response functions. Panel A show the concatenated response function with a varying long lag amplitude and constant short lag response. Panels B and C show the cross-correlation function and the recovered peaks in the functions respectively. Panel D demonstrates that cross-correlation can identify both lags easily when the amplitude of the response functions are comparable, with the failure rate increasing with diminishing long lag amplitude. Panels E and F show the recovered lag in frequency space and as a function of amplitude ratios respectively. We observe that as the long lag becomes less obvious, the recovered lag may not be an accurate determination of the input lag.

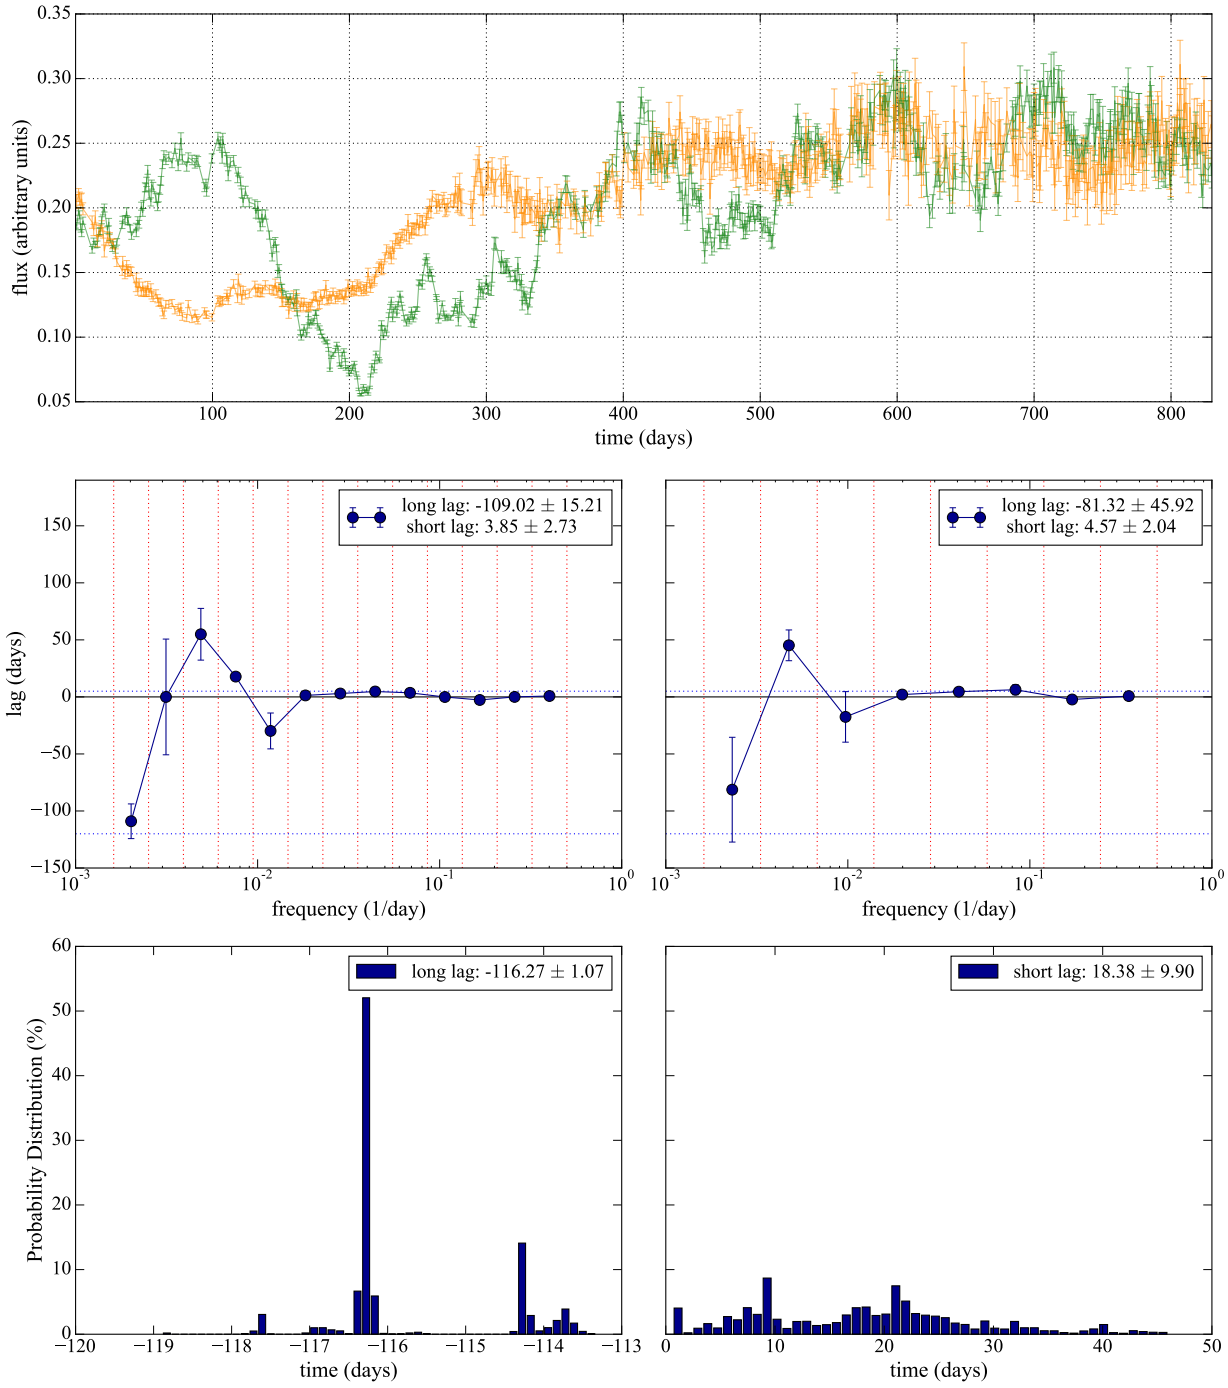


Figure 2.12: Comparison between the Fourier and JAVELIN lag recoveries. The top panel shows the simulated driving (green) and reprocessed (orange) light curves. The middle panels show the recovered lags by the Fourier method with two choices of binning (13 bins and 8 bins), which can have some impact on the accuracy of lag determination. The bottom panels show the recovered long (left) and short (right) lags by JAVELIN, where interpolation might have hindered a precise measurement of the short lag.

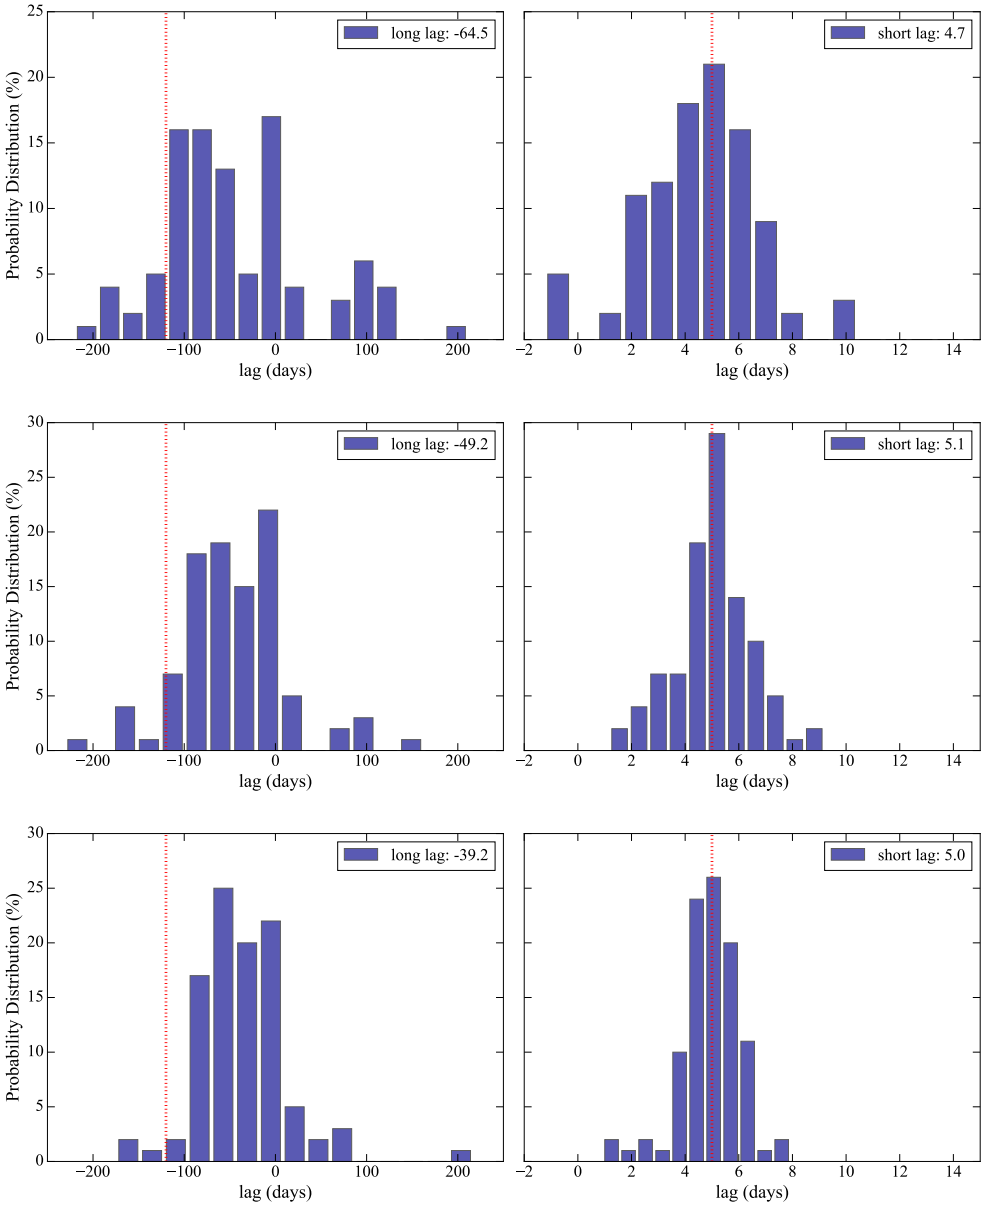


Figure 2.13: Results of 300 simulated light curves analyzed via the Fourier method. The top, middle, and bottom panels each represents 100 light curves reprocessed with response functions where the amplitude of the long lag is 20%, 10%, and 5% of that of the short lag respectively. The legend shows a median value of the recovered lags, and the red dotted lines are the input lag at -120 and +5 days. 78%, 85%, and 89% of the short lags recovered are significant ($> 3\sigma$ from 0). We see a clear trend where the median value approaches the input lag when the long lag signal is more prominent. In all three cases, the short lag can be easily recovered.

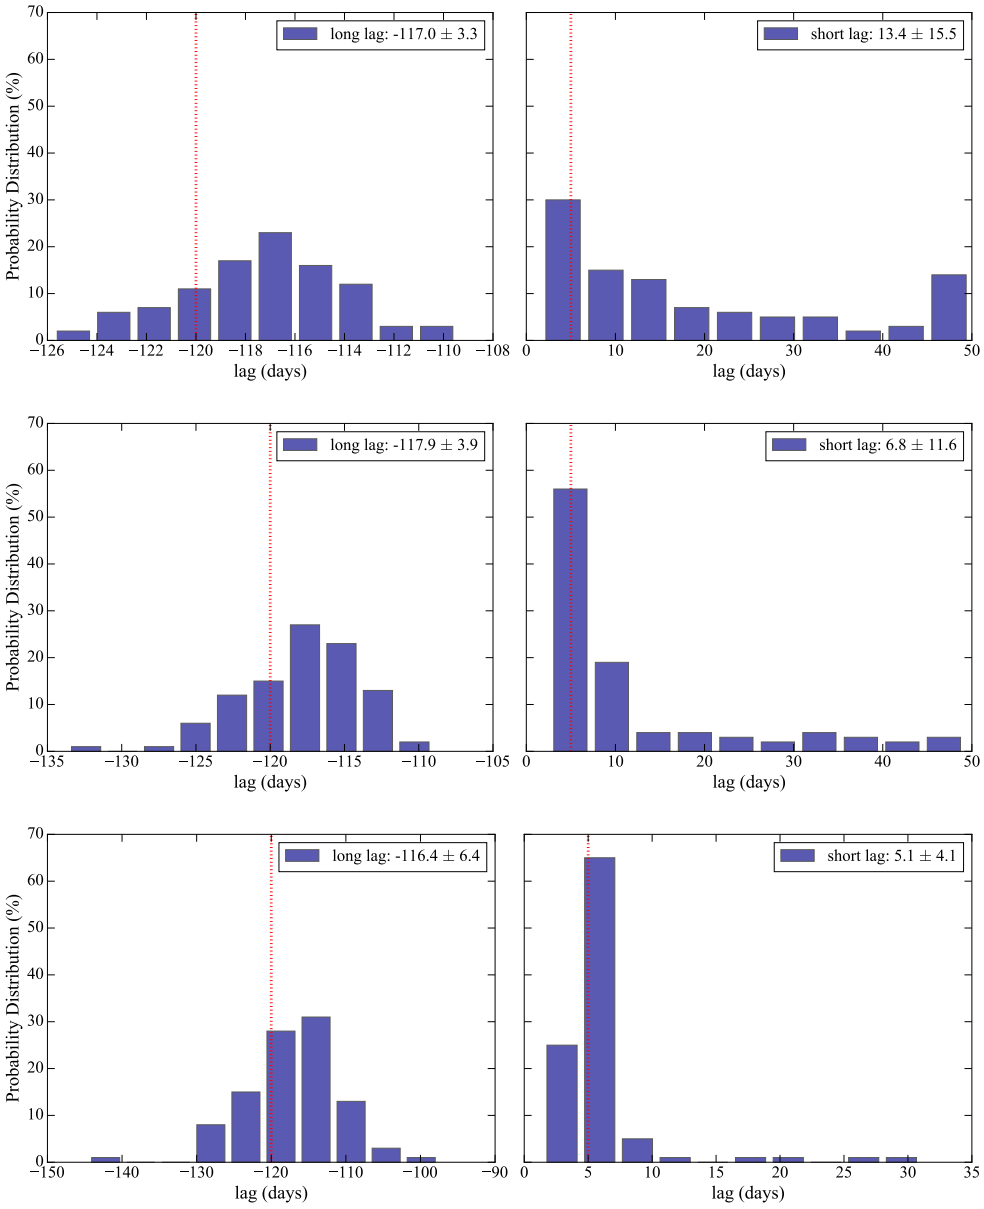


Figure 2.14: Results of 300 simulated light curves analyzed via JAVELIN. The top, middle, and bottom panels each represents 100 light curves reprocessed with response functions where the amplitude of the long lag is 20%, 10%, and 5% of that of the short lag respectively. The legend shows a median value of the recovered lags, and the red dotted lines are the input lag at -120 and +5 days.

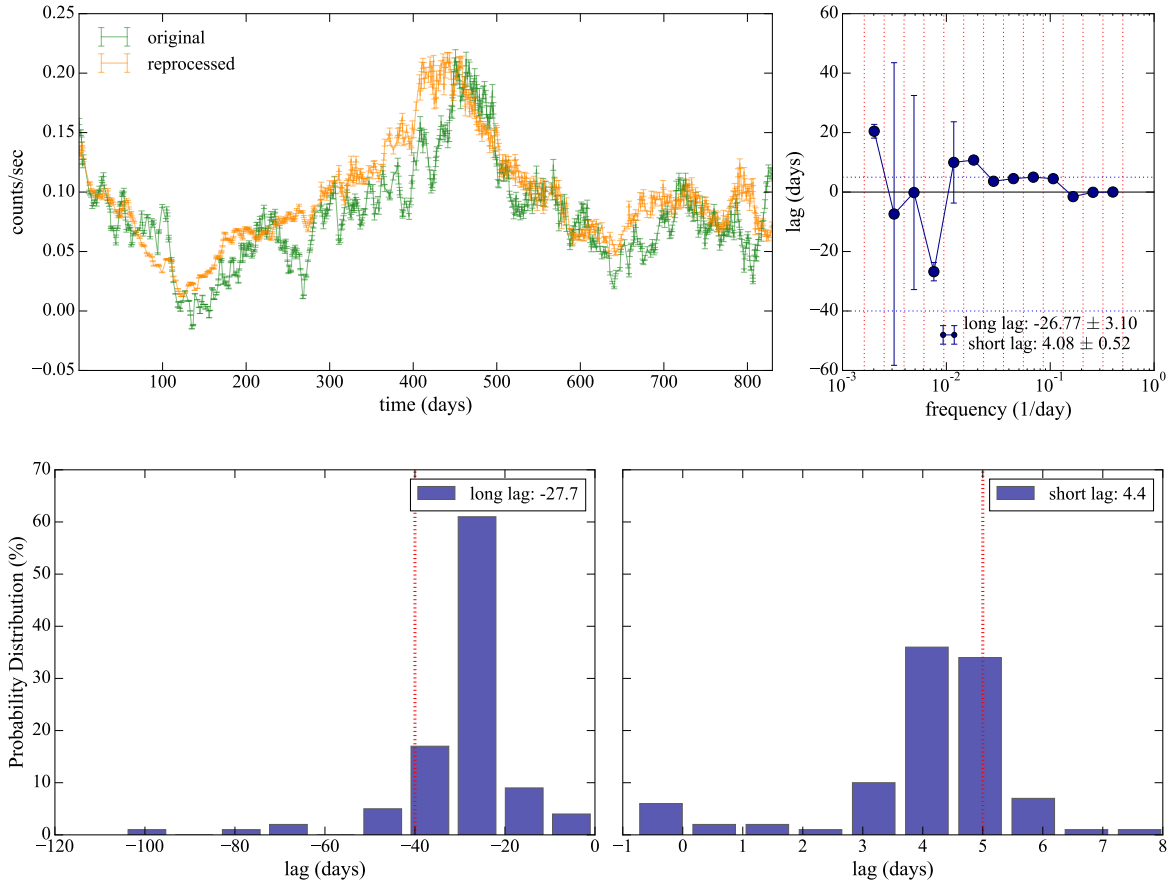


Figure 2.15: Results of 100 simulated light curves with input lags of -40 and +5 analyzed via the Fourier method. The top panel shows one of the simulations and its resultant lag distribution in frequency space. The input lags are -40 and +5 days with an amplitude ratio of 1:5. 80% of the short lags recovered are significant ($> 3\sigma$ from 0). This test demonstrates the ability of this method to clearly identify the long lag with a peak in frequency space when the length of the light curve is significantly longer than the longest lag.

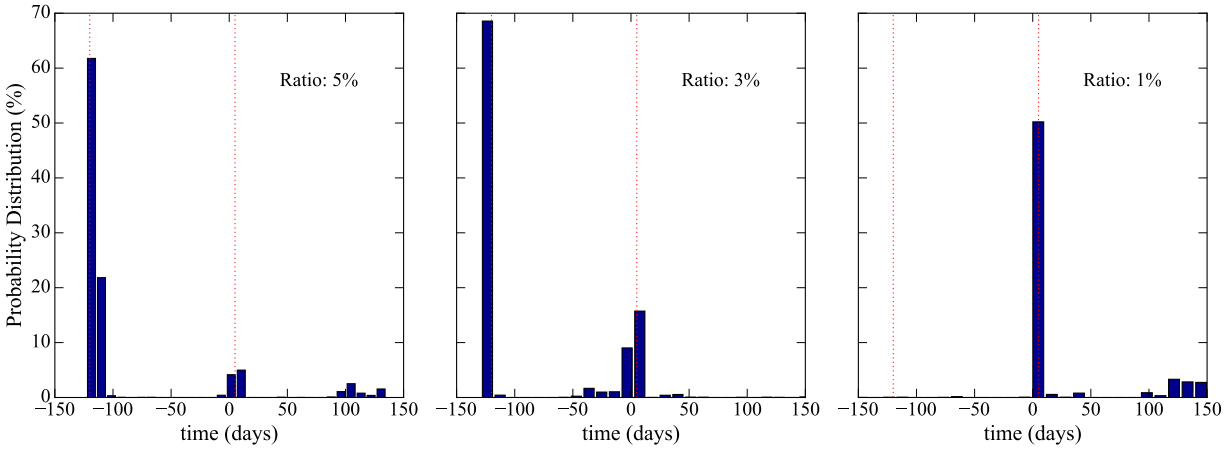


Figure 2.16: The cross-correlation function from JAVELIN indicating the most probable lags. The plots show results from different input response functions with amplitude ratios between the long lag and the short lag at 5%, 3% and 1% respectively. The red dotted lines show the input lags at -120 and 5 days respectively. The recovered short lag strengthens when the signal from the long lag weakens.

Chapter 3

Negative Lags on the Viscous Timescale in Quasar Photometry and Prospects for Detecting More with LSST

3.1 Introduction

Quasars are among the most luminous objects in the Universe. As a result, they are important to a variety of fields within astrophysics, from the study of accretion disks to galaxy evolution (Fabian, 2012; Kormendy & Ho, 2013). With the exception of the extraordinary efforts by the Event Horizon Telescope Collaboration (2019) and Gravity Collaboration (2018), it is generally not possible to resolve the microarcsecond scales of the accretion disk, broad line region (BLR), and event horizon. Since spatial resolution is so difficult to achieve and quasars vary over a range of frequencies, temporal resolution is instead often used to study the accretion disks of quasars.

In particular, Blandford & McKee (1982) proposed a technique known as reverberation mapping to measure distances to the BLR, by measuring time lags (of order 100 days) between the continuum emission and broad emission lines (see also Kaspi et al., 2000; Peterson et al., 2004; Bentz & Katz, 2015; Grier et al., 2017). A second type of reverberation mapping, first introduced by Collier et al. (1999) and known as disk continuum reverberation mapping, uses time lags between variability in different wavebands to measure the radial extent of the disk, using the lamppost model (e.g., Sergeev et al., 2005; Cackett et al., 2007, 2018; De Rosa et al., 2015; Edelson et al., 2015, 2017, 2019; Jiang et al., 2017; Fausnaugh et al., 2016; Starkey et al., 2017; Homayouni et al., 2022). The lamp-post model proposes that high-frequency (X-ray or hard-UV) emission from the corona will be reprocessed by the disk and re-emitted as UV/optical light. Because it takes light time to travel from the inner hotter regions of the disk, where it is re-emitted in shorter wavebands, to the outer cooler regions of the disk, where it is re-emitted in longer wavebands, there will be a time lag of order days between variability in short wavelength bands and long wavelength bands. Because this lag is on the order of days, we refer to it here as the "short" lag.

Using the Shakura & Sunyaev (1973) accretion disk model of an optically thick, geometrically thin disk to determine the temperature dependence of the disk, this time lag should scale as $\tau_c \propto \lambda^{4/3}$, and also depend on the mass of the supermassive black hole (SMBH) and accretion rate (see also, Lynden-Bell, 1969; Pringle & Rees, 1972; Novikov & Thorne, 1973; Friedjung, 1985). However, recent observations have suggested that the radial extent of the disk is larger than predicted by the standard thin disk model (Jiang et al., 2017; Cackett et al., 2018; Mudd et al., 2018; Guo et al., 2022), although this discrepancy may be due to nuclear extinction (Gaskell, 2017; Gaskell et al., 2023) or an underestimation of SMBH masses due to unknown BLR geometry (Pozo Nuñez et al., 2019). Other models suggest that AGN disks may appear more extended due to large azimuthal temperature fluctuations (Dexter & Quataert, 2012), disk winds (Laor & Davis, 2014), or magnetically elevated ac-

cretion (Dexter & Begelman, 2019), none of which are included in the standard thin disk model.

There is also growing evidence that the lamp-post model may be insufficient to fully explain the optical/UV variability in quasar light curves (e.g., Arévalo et al., 2009; Shappee et al., 2014; Edelson et al., 2019; Yu et al., 2022) Alternative sources of optical/UV light curve variability, which may occur on the thermal timescale, are magneto-rotational instability turbulence (Balbus & Hawley, 1991; Burke et al., 2021) and convection due to enhanced opacity in the optical/UV region of the disk (Jiang & Blaes, 2020). As these turbulent perturbations are accreted inwards on the viscous timescale, they can also modify the variability of light curves in inner regions of the disk (Lyubarskii, 1997). Because these perturbations are accreted inwards the lag will be "negative" compared to the short lag. In other words, light curve variability from disk turbulence in short wavelength bands will lag variability in long wavelength bands. The viscous timescale is (Davis & Tchekhovskoy, 2020),

$$\tau_{\rm visc} = \frac{1}{\alpha\Omega}\left(\frac{h}{r}\right)^{-2}, \tag{3.1}$$

where α is the Shakura & Sunyaev (1973) stress parameter, $\Omega = (GM/r^3)^{1/2}$ is the orbital frequency, M is the mass of the SMBH, G is the gravitational constant, h is the scale height of the disk, and r is the radial position in the disk.

In the standard thin disk model $\tau_{\rm visc}$ can be of order hundreds of years, because in the radiation-dominated region of the disk h is constant as a function of r, and $h/r \lesssim 0.01$. However, a recent simulation by Jiang & Blaes (2020) suggests that due to variations in the scale height of the disk, the viscous timescale could be on the order of 100 days. Furthermore, two recent independent works have suggested the detection of negative lags in the nearby quasar Fairall 9 (Hernández Santisteban et al., 2020; Yao et al., 2023). In particular, using the *SWIFT* and Las Cumbres Observatory (LCO) light curves with a baseline of ~ 270 days and sub-daily cadence from Hernández Santisteban et al. (2020), Yao et al. (2023) found a

long negative lag of ~ -70 days between the *W2*- and z-band. This sub-100 day negative lag suggests very significant disk puffing, even more so than in the simulation in Jiang & Blaes (2020), and that it may be possible to detect long lags in many more quasars.

In this paper, we derive a new empirical model for the "long" negative lag based on the model fit to the long lags detected for Fairall 9 by Yao et al. (2023). We then use this model to examine our ability to detect long lags in the future. Detecting more long lags will help to constrain our model and provide information on the vertical structure of disks. Additional lag measurements would also allow us to model how the vertical structure of disks change as a function of quasar properties such as SMBH mass and accretion rate.

Because long lags can last tens to hundreds of days, a long baseline of high cadence observations is necessary to detect them. The Vera Rubin Legacy Survey of Space and Time (LSST, Ivezić et al., 2019) will carry out a comprehensive survey of the southern sky and will contain over 100 million quasars. LSST will provide photometric light curves in *ugrizy* bands, allowing us to study the wavelength dependence of quasar lags. Kovacevic et al. (2022), Pozo Nuñez et al. (2023), and others found that LSST should be able to detect short lags in quasar light curves. Here, we generate mock LSST quasar light curves reprocessed with several different short and long lags. We then test our ability to detect long lags using several lag detection methods, primarily focusing on two: JAVELIN (Zu et al., 2011) and a maximum-likelihood Fourier method developed by Miller et al. (2010) and adapted for use on quasars by Zoghbi et al. (2013) and Cackett et al. (2022).

We derive our model in Section 3.2, describe in detail how we generate our mock LSST light curves in Section 3.3, and describe our lag detection methods in Section 3.4. In Section 3.5 we report the accuracy of our lag detection methods for light curves reprocessed with only long lags (Section 3.5.1), light curves reprocessed with long and short lags (Section 3.5.2), light curves observed with different proposed LSST cadences (Section 3.5.3), light curves from the simulation in Jiang & Blaes (2020) (Section 3.5.4), and light curves for dimmer magnitude quasars (Section 3.5.5) We provide a brief summary of the accuracy of our lag

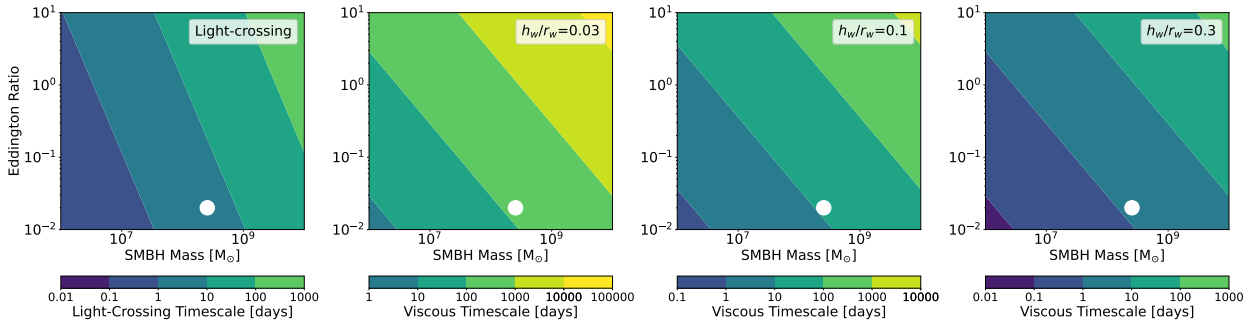


Figure 3.1: The leftmost panel shows the rest-frame short lag between the u- and y-band in days calculated using the standard thin disk model and the remaining panels show the long lag timescale between the u- and y-band in days for disks with aspect ratios from left to right of 0.03, 0.1, and 0.3 as a function of Eddington ratio and SMBH mass in $M_\odot$. The white point shows the approximate SMBH mass and Eddington ratio of Fairall 9 (Peterson et al., 2004; Hernández Santisteban et al., 2020).

detection methods for all light curves in Section 3.5.6. In Section 3.6 we make a prediction for the number of long lags detectable by LSST over its ten year baseline. We summarize our results in Section 3.7 and discuss the importance of improved light curve models (Section 3.7.1) and detection methods (Section 3.7.2).

3.2 What is the Long Lag Timescale?

The long lag timescale recovered by Yao et al. (2023) is much shorter (70 days versus 100 years) than the viscous timescale predicted by the standard Shakura & Sunyaev (1973) thin-disk model. However, this timescale is in good agreement with the recent simulation by Jiang & Blaes (2020) that shows $h/r > 0.01$ in the radiation pressure dominated region. Disks with larger aspect ratios are also often invoked to explain line driving broad absorption line quasars (e.g. Murray et al., 1995; Leighly, 2004). If this shorter timescale for the long lag is accurate, it should be possible to detect long lags with long baseline campaigns, like LSST, depending on the properties of the quasar and waveband coverage. We therefore use the results from Yao et al. (2023) to derive a new empirical model for the long lag timescale.

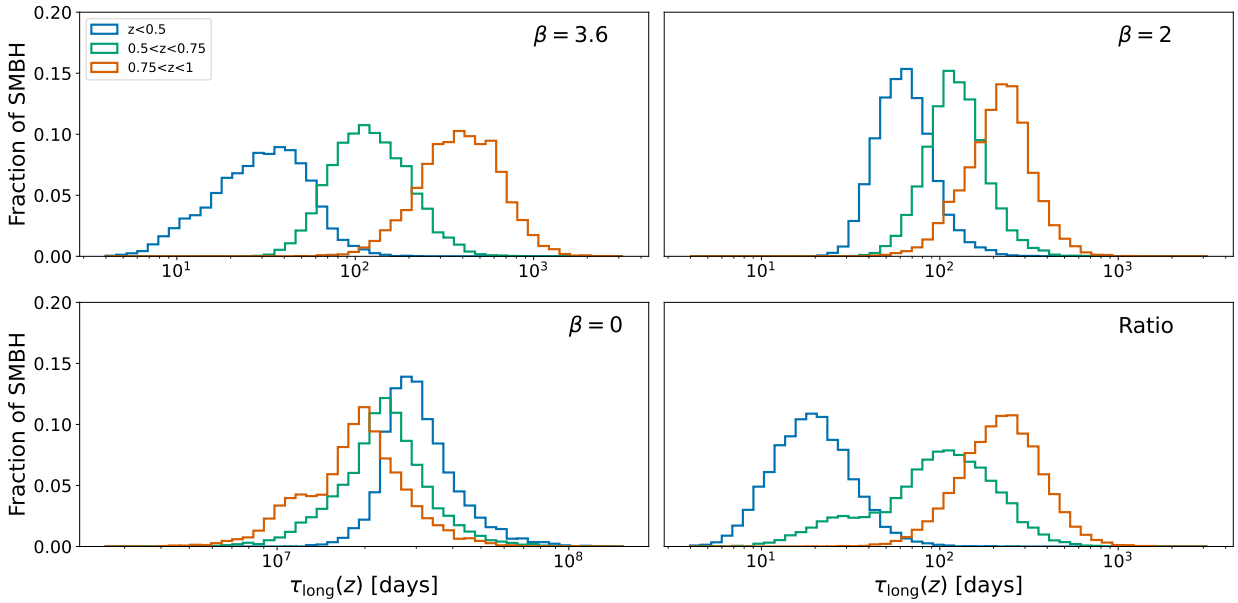


Figure 3.2: The normalized distribution of $\tau_{\rm long}(z)$ for quasars in SDSS DR7 (Shen et al., 2011) that fall within different redshift ranges. We assume $\alpha = 0.1$, and take the masses, Eddington ratios, and redshifts from Shen et al. (2011). In the top left panel we use $\beta = 3.6$ as in Yao et al. (2023); in the top right panel we use $\beta = 2$, which is a more conservative estimate that the aspect ratio increases linearly with radius; in the bottom left panel we use $\beta = 0$ and $h/r = 0.01$, which are typical values in the standard thin disk model and lead to much long lag timescales (see different x-axis range); and in the bottom right panel we use the ratios between the long and short lag for different wavebands in Yao et al. (2023)

.

As in Yao et al. (2023), we start by converting the viscous timescale to a velocity, $v_{\rm visc} = r/\tau_{\rm visc}$, and integrate from r_0, the radius corresponding to the reference band, to r_x, the radius corresponding to band x,

$$\begin{aligned}\tau_{\rm long} &= \int_{r_0}^{r_x} \frac{dr}{v_{\rm visc}} = \frac{r_0^{2\beta}}{h_0^2\alpha\sqrt{GM}} \int_{r_0}^{r_x} r^{5/2-2\beta} dr \\ &= \frac{r_0^{2\beta}}{(7/2-2\beta)h_0^2\alpha\sqrt{GM}} \left[r_x^{7/2-2\beta} - r_0^{7/2-2\beta}\right].\end{aligned} \tag{3.2}$$

Here we have written $h \equiv h_0(r/r_0)^\beta$ in order to capture the r-dependence of h. In the standard thin disk model, $\beta = 0$, as h is independent of radius. Note that at $\beta \to 7/4$, the β-dependent factors reduce to $r_0^{7/2}\ln(r_x/r_0)$, so there is no singularity in $\tau_{\rm long}$.

Following Fausnaugh et al. (2016) and Jiang et al. (2017) the radius at which the disk radiates at a wavelength λ is,

$$r(\lambda) = \left(\frac{X k_B \lambda}{\hat{h} c}\right)^{4/3} \left(\frac{f_i G M \dot{M}}{8\pi\sigma}\right)^{1/3}, \tag{3.3}$$

where k_B is Boltzmann's constant, $\hat{h}$ is Planck's constant, c is the speed of light, σ is the Stefan-Boltzmann constant, and $\dot{M}$ is the accretion rate of the SMBH. $X = 2.49$ is a factor to account for emission at this wavelength from a range of temperatures. f_i is a factor to account for the irradiation from an X-ray/UV source near the SMBH by changing the normalization of the temperature profile. We set $f_i = 1$, because three dimensional radiation magneto-hydrodynamic (MHD) simulations show that the UV/Optical region of the disk is very optically thick, making it difficult for X-ray photons to convert to UV photons, because the free-free absorption opacity for X-rays is very small and the Compton process is not important (Jiang et al., 2016, 2019; Jiang & Blaes, 2020, , although the bound-free opacity may be substantial for X-ray energies below 1 keV). As a result, the UV/Optical emission should be dominated by photons generated by the local disk ($f_i = 1$). $f_i = 1$ implies that X-ray light will not be directly reprocessed by absorption and re-emission in the disk. Instead in this model the X-ray short timescale variability is reprocessed by perturbing the local photosphere via scattering on the X-ray variability time scale.

While Equation (3.3) is derived for the standard thin disk model, $r \propto \lambda^{4/3}$ does not rely on the assumption that $h/r \ll 1$, only the assumption that the emission is local. Since we care most about the wavelength dependence of the long lag timescale we use Equation (3.3) here. However, the mass and accretion rate dependence of the long lag may need to be corrected when a more accurate model for the long lag can be derived either through simulations or additional observations. In particular, we note that h/r and thereby β may depend on the accretion rate and mass of the quasar. For simplicity, here we will define h/r and β independently of any quasar property.

Substituting Equation (3.3) into Equation (3.2) gives a viscous lag timescale of the form,

$$\tau_{\rm long} = \tau_0 \left[\left(\frac{\lambda}{\lambda_0} \right)^{4/3(7/2-2\beta)} - 1 \right], \tag{3.4}$$

where,

$$\tau_0 = 93 \text{ days} \left(\frac{-4}{7/2-2\beta} \right) \left(\frac{\lambda_0}{\lambda_{\rm W}} \right)^{4/3(7/2-2\beta)} \left(\frac{h_{\rm W}/r_{\rm W}}{0.2} \right)^{-2} \frac{0.1}{\alpha} \left(\frac{l/l_{\rm edd}}{0.1} \right)^{1/2} \left(\frac{M}{10^8 \text{ M}_\odot} \right)^{1/2}, \tag{3.5}$$

$\lambda_{\rm W} = 1928$ Å, and $h_{\rm W}/r_{\rm W}$ is the aspect ratio at the disk radius corresponding to the *W2*-band. The parameterizations and bands are based on the recent fitting that Yao et al. (2023) performed on the wavelength-dependent long lag they detected for Fairall 9 to this function for a reference band $\lambda_{\rm W}$. They found $\tau_0 = 66.8 \pm 2.3$ and $\beta = 3.59 \pm 1.22$. For Fairall 9, $M = 2.55 \pm 0.56 \times 10^8$ $\text{M}_\odot$ (Peterson et al., 2004) and the Eddington ratio $l/l_{\rm edd} = 0.02 \pm 0.004$ (Hernández Santisteban et al., 2020), which gives $\tau_0 = 66.8$ days for $\alpha = 0.1$ and $h_{\rm W}/r_{\rm W} = 0.2$.

The wavebands for LSST range from the u-band to the y-band. For $\lambda_{\rm u} = 3600$ Å we get,

$$\tau_0 = 3.6 \text{ days} \left(\frac{-4}{7/2-2\beta} \right) \left(\frac{h_{\rm W}/r_{\rm W}}{0.2} \right)^{-2} \frac{0.1}{\alpha} \left(\frac{l/l_{\rm edd}}{0.1} \right)^{1/2} \left(\frac{M}{10^8 \text{ M}_\odot} \right)^{1/2}, \tag{3.6}$$

for our fiducial values.

We show the lag between the u- and y-band at zero redshift for a range of quasar masses and Eddington ratios in Figure 3.1. For comparison we show the short lag timescale for the standard thin disk model in the left panel using Fausnaugh et al. (2016),

$$\tau_{\rm c} = \frac{r(\lambda_0)}{c} \left[\left(\frac{\lambda}{\lambda_0} \right)^{4/3} - 1 \right], \tag{3.7}$$

where $r(\lambda_0)$ is calculated from Equation (3.3). The remaining three panels are for different values of $h_{\rm W}/r_{\rm W}$: $h_{\rm W}/r_{\rm W} = 0.03$, which is similar to its proposed value in the standard thin disk model, and $h_{\rm W}/r_{\rm W} = 0.1, 0.3$, which are slightly smaller and larger, respectively, than the value found from fitting the lag found in Yao et al. (2023). Depending on $h_{\rm W}/r_{\rm W}$, at low redshift the long lag can range from a fraction of a day to dozens of years, for realistic quasar properties.

In addition to the mass and Eddington ratio of the quasar, the long lag also depends strongly on the redshift, z, of the quasar. First, time dilation due to increasing redshift should increase lag times by a factor of $(1+z)$. Second, the rest frame waveband of a quasar will change by a factor of $1/(1+z)$, i.e. λ_0 will decrease by a factor of $1/(1+z)$. Since τ_0 scales as $\lambda_0^{4/3(7/2-2\beta)}$ for $\beta \approx 3.6$ overall τ_0 scales as,

$$\tau_0 \sim (1+z)\left(\frac{1}{1+z}\right)^{4/3(7/2-2\beta)} \sim (1+z)^{5.9}. \tag{3.8}$$

The top left panel of Figure 3.2 shows the normalized distribution of $\tau_{\rm long}$ from the u- to y-band for quasars in the Sloan Digital Sky Survey Data Release 7 (SDSS DR7, Shen et al., 2011) for three different redshift ranges, using the $\beta \approx 3.6$ found in Yao et al. (2023). $z = 1$ is the highest redshift shown for two reasons. First, at $z > 1$ the observed-frame u-band corresponds to rest frame bands shorter than the shortest UV band in Yao et al. (2023) and we do not want to extrapolate our models beyond the data available. Second, for $z \gtrsim 1$ the long lag is over three years long for a majority of quasars if β is indeed as large as predicted by the model fit in Yao et al. (2023). Lags longer than three years would be difficult to detect even with the ten year baseline of LSST.

In the top right panel of Figure 3.2, instead of using the exact fit from Yao et al. (2023) we assume $\beta = 2$. For $\beta = 2$ the lag timescale would still flatten at longer wavelengths, as in Yao et al. (2023), but the aspect ratio would only increase linearly as a function of

radius. If the average value of β is closer to this value, lags in quasars out to $z \lesssim 2$ should be measurable.

We note that the dependence of the long lag on reference wavelength would be very different for $\beta < 7/4$. The term in the brackets in Equation (3.4) would become a positive factor and τ_0 would become negative. The increase in h/r becomes less steep than the other radial dependencies of the long lag. Therefore the long lag would increase with reference wavelength as is the case with short lags (see Equation (3.7)).

For example, for $\beta = 0$, as in the standard thin disk model, $\tau_{\rm long} \propto \lambda_0^{14/3}$, which is a stronger dependence on λ_0 than the $\tau_{\rm short} \propto \lambda_0^{4/3}$ for the short lag. As a result, the observed long lag from the u- to y-band would actually decrease as a function of redshift, because emission from the inner disk where the lag timescale is shorter gets redshifted into longer wavebands. The bottom left panel of Figure 3.2 shows the distribution of $\tau_{\rm long}$ for three different redshift ranges for $\beta = 0$ and $h/r = 0.01$, which are typical assumed values for the standard thin disk model. These distributions peak around $10^5 - 10^6$ years.

As an alternate way of scaling from Fairall 9 to other targets, we can assume that the ratio between short and long lags is preserved between Fairall 9 and other targets. We use this alternative way of scaling to examine how removing our various model assumptions changes the long lag distributions for our different redshift ranges. To calculate these lags we take the redshift of each quasar in SDSS DR7 with $z < 1$ and use that to determine the rest-frame wavebands that correspond to the observed frame u- and y-band for that quasar. We then take the ratio between the long and short lag detected in Yao et al. (2023) in the rest-frame and multiply it by the short lag (calculated using Equation (3.7)) for that quasar. In practice we use the error-weighted mean of this ratio in the rest-frame waveband corresponding to each quasar and both adjacent wavebands.

Calculated in this way the lags are shifted towards slightly shorter timescales, relative to our model with $\beta = 3.6$. This shift is especially the case for $z < 0.5$ quasars and $0.75 < z < 1$ quasars for which the distributions peak around 20 days instead of 40 days and 250 days

Redshift	$\tau_{\rm long}$	$\tau_{\rm c}$	i-mag	$\tau_{\rm damp}$
z <0.5	50 days	7 days	17.8 mag	145 days
$0.5 < z < 0.75$	130 days	9 days	18.6 mag	170 days
$0.75 < z < 1$	400 days	12 days	18.7 mag	204 days

Table 3.1: The three redshift ranges we use to bin SDSS DR7 quasars when generating our mock light curves and their corresponding median properties: long lag timescale ($\tau_{\rm long}$), short lag timescale ($\tau_{\rm c}$), i-band magnitude, and random walk damping time ($\tau_{\rm damp}$).

instead of 400 days, respectively. This shift towards shorter lags would enable the detection of long lags for light curves with shorter baselines than LSST and make it possible to detect long lags out to $z > 1$ with LSST. The main reason for the shift to shorter timescales is that this method removes the mass and luminosity dependence of the lag timescale, by scaling all lags to Fairall 9. In the bottom right panel of Figure 3.2 we show the long lag timescale calculated using this approach.

3.3 Generating Mock Light Curves

With a prescription for calculating lags in hand, we turn now to creating mock multi-band LSST-like light curves incorporating this negative lag. We outline our process for generating mock LSST light curves that have been reprocessed with lags. First, we either generate a driving light curve using a DRW model or use light curves from the radiation MHD simulation in Jiang & Blaes (2020). Then, we reprocess these light curves with a Gaussian response function to add lags. Finally, we sub-sample these light curves with proposed LSST cadences.

3.3.1 Driving Light Curves

3.3.1.1 Damped Random Walk Light Curves

For light curves analyzed in Sections 3.5.1, 3.5.2, and 3.5.3 we start by generating a highly sampled ($\Delta t = 0.001$ days) 2000 year-long light curve, using a damped random walk (DRW) model as in Kelly et al. (2009). A DRW is a stochastic process that goes from a ν^{-2}

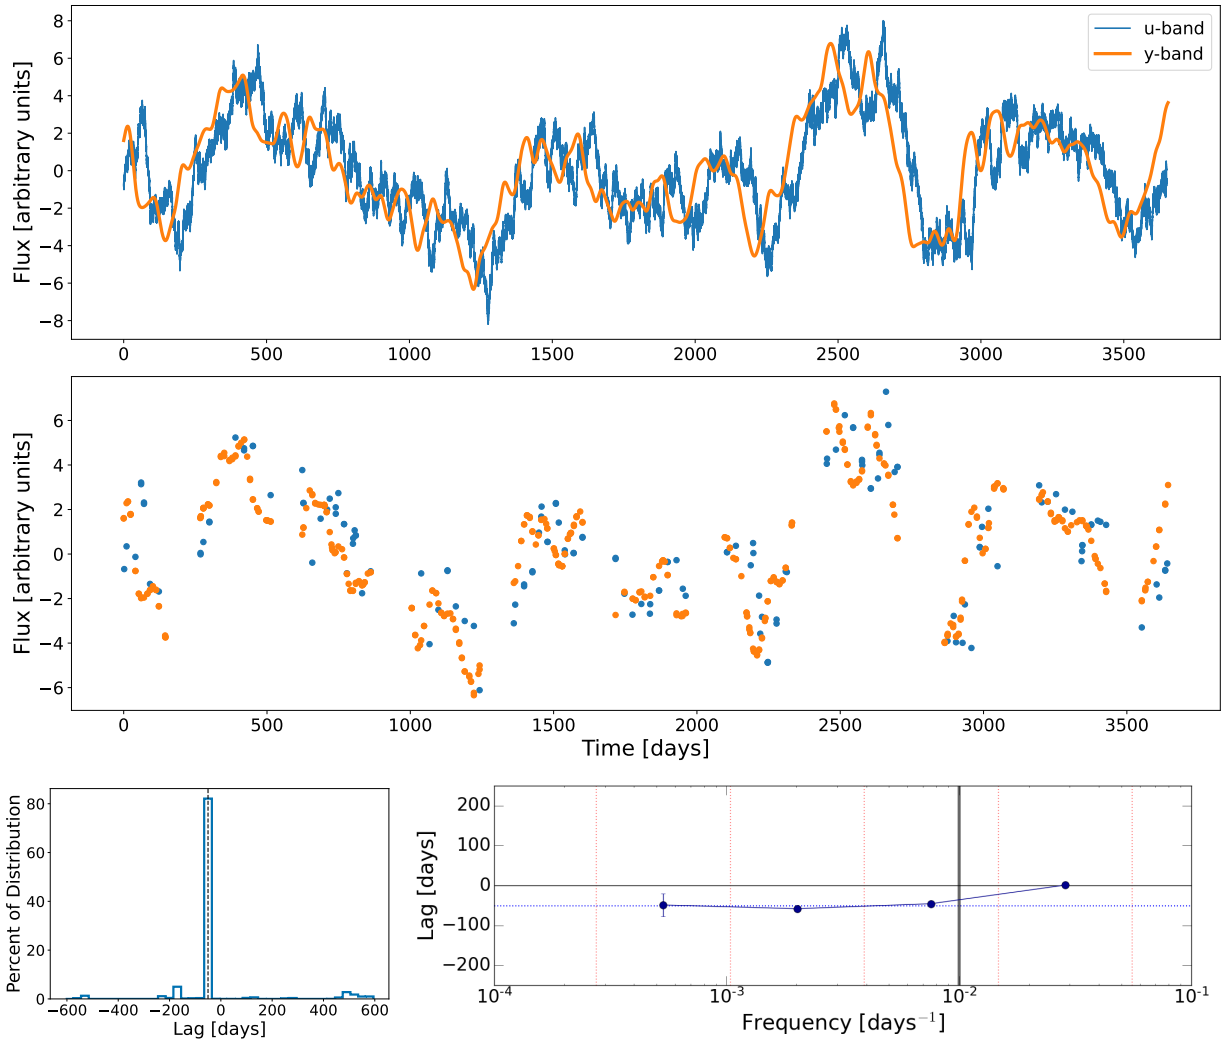


Figure 3.3: Top panel: Simulated driving u-band light curve, in blue, and y-band light curve, reprocessed with a long lag of -50 days, in orange. Middle panel: The same light curve as in the top panel sub-sampled using the long season LSST cadence. Bottom left: JAVELIN probability distribution for the long lag between the light curves above. The long lag is shown as a vertical dashed line. The median of the negative values of the distribution is -50 days. Bottom right: An example of the lags as a function of frequency recovered by the maximum-likelihood method for a mock light curve with an input long lag of -50 days (shown as the horizontal dashed blue line). The points are the lag detected in each bin (shown as the dashed vertical red lines). The vertical lines show the error bars. The lowest frequency bin has the largest error bar. The next two bins are roughly consistent with the input lag and have an error-weighted mean of -48 days. Because there is no short lag and it is at frequencies higher than the wrapping frequency (vertical grey line), the highest frequency bin is consistent with zero.

power-law, where ν is frequency, at high-frequency to white noise at low-frequency at some characteristic damping timescale $\tau_{\rm damp}$. The DRW model has been shown to be a decent representation of quasar light curves for timescales of days to years (e.g. Kelly et al., 2009; Kozłowski et al., 2010; MacLeod et al., 2012; Zu et al., 2013). The physical interpretation of this damping timescale is still debated, although Kelly et al. (2009) and Burke et al. (2021) found that it may be related to the thermal timescale of the disk. Burke et al. (2021) found that $\tau_{\rm damp}$ scales with the SMBH mass as,

$$\tau_{\rm damp} = 107 \left(\frac{M}{10^8 {\rm M}_\odot} \right)^{0.38}. \tag{3.9}$$

To determine the properties of our mock light curves we take quasars from SDSS DR7, and put them in three redshift bins: $z < 0.5$, $0.5 < z < 0.75$, and $0.75 < z < 1$ (see Table 3.1). In this paper we only investigate our ability to detect lags in quasars at $z < 1$, because, as can be seen in Figure 3.2, a majority of quasars at $z \approx 1$ have long lags over 500 days long, which would be difficult to detect with the LSST baseline of 10 years. We use the median quasar mass in each bin to determine $\tau_{\rm damp}$. We assume the structure function at infinity, ${\rm SF}_\infty = 2.5$ in flux units, which is a typical value for quasar light curves from the Dark Energy Survey (Yu et al., 2020) and describes the long term variability of the quasar.

We split our 2000 year light curve into 100 individual light curves which we use as underlying driving light curves, $d(t)$. We generate and subsample a 2000 year light curve instead of generating 100, 20 year-long driving light curves because it is important when generating a DRW light curve that the baseline is significantly longer than the damping timescale (Kozłowski, 2017; Stone et al., 2022). Done this way our baseline DRW light curve is over 3000 times longer than the largest damping timescale we use.

3.3.1.2 Simulated Light Curves

While the DRW model does a good job describing quasar variability over a range of timescales, it is not based on a physical model of the disk. The DRW model may also not model the PSD properly at timescales less than 1 month (Mushotzky et al., 2011; Caplar et al., 2017; Smith et al., 2018; Sánchez-Sáez et al., 2018), and be unreliable on the longest timescales no matter how long the modeled DRW is (Stone et al., 2022). Therefore, in Section 3.5.4 we test the robustness of our results by using light curves from a simulation performed by Jiang & Blaes (2020).

Jiang & Blaes (2020) used ATHENA ++ to perform a global radiation MHD simulation of a quasar disk between 30 $r_{\rm g}$ and 100 $r_{\rm g}$, where $r_{\rm g}$ is gravitational radii. This region of the disk, which roughly corresponds to the region that emits in the optical, is dominated by the Rosseland mean opacity, which increases around 1.8×10^5 K due to the iron opacity bump. Jiang & Blaes (2020) found that the iron opacity bump causes the disk to become convectively unstable leading to turbulence which puffs up and heats the disk. Eventually the puffing up of the disk decreases the opacity enough for the disk to become convectively stable and cool, until it becomes optically thick again. In Jiang & Blaes (2020) this cycle repeats itself every few years and leads to variability in the luminosity of the disk by factors of around 3-6.

Figure 14 in Jiang & Blaes (2020) shows the luminosity scaled by the Eddington luminosity emitted from $r \leq 60 r_{\rm g}$ and $r \leq 80 r_{\rm g}$ over roughly 57 years for their simulated 5×10^8 $M_\odot$ quasar. Over the 57 years the Eddington ratio can vary from $l/l_{\rm edd} \approx 0.01 - 1$. We split both of these light curves into eight chunks of 20 years each, and linearly interpolate them down to sub-daily cadence from an original sampling of roughly every three days. This interpolation should not have a large impact on our results because the LSST cadence we down-sample the light curve with has a cadence of roughly every 10–30 days (see Section 3.3.3).

3.3.2 Reprocessed Light Curves

Next we reprocess our driving light curves with a Gaussian response of the form,

$$\psi(\mathrm{t}) = \frac{1}{\mathrm{S}\sqrt{2\pi}} \exp^{\frac{-(\mathrm{t}-\mathrm{C})^2}{2\mathrm{S}^2}}, \tag{3.10}$$

where C is the center of the Gaussian (i.e. the input lag) and S is the width of the Gaussian, set to $S \equiv C/5$, for simplicity. For a negative lag ($C < 0$) we replace t with $-t$. We choose $S \equiv C/5$ such that there is some width to the response that scales with the lag duration. However, the exact shape of the response function is unknown, especially for the long lag.

In cases where we are simulating both a long and a short lag we add two response functions together. We scale the relative strength of the two response functions by fixing the amplitude of the long lag response function to some fraction of the amplitude of the short lag response function. As a result, the integrated Gaussian for the long lag will no longer sum to one, which means the amplitude of the reprocessed light curve will differ from the amplitude of the driving light curve.

Because we want to examine a broad range of long lag timescales, we use long lags of -50, -130, and -400 days. From Figure 3.2 we can see this range of timescales does a good job representing quasars from SDSS DR7 at $z \lesssim 1$. To match long lag timescales to other quasar light curve properties we estimate that each quasar bin $z < 0.5$, $0.5 < z < 0.75$, and $0.75 < z < 1$ corresponds to a long lag of -50, -130, and -400 days, respectively. We also calculate the short lag corresponding to each bin using Equation (3.7). The long and short lags for each redshift bin are in Table 3.1.

To get the reprocessed light curve we Fourier transform the driving light curve ($d(t) \rightarrow D(\nu)$) and response function ($\psi(t) \rightarrow \Psi(\nu)$). In a linear reprocessing model the reprocessed light curve is,

$$\mathrm{r(t)} = \int_{-\infty}^{\infty} \psi(\tau)\mathrm{d(t}-\tau)\mathrm{d}\tau. \tag{3.11}$$

which in Fourier space is equivalent to,

$$\mathrm{R}(\nu) = \Psi(\nu)\mathrm{D}(\nu). \tag{3.12}$$

We can then inverse Fourier transform $R(\nu) \rightarrow r(t)$ to acquire our reprocessed light curve.

3.3.3 LSST Cadence

Now that we have a driving and reprocessed light curve, we down-sample their $\Delta t = 0.001$ cadence to a variety of LSST cadence possibilities. The LSST cadence is not final. Over the past decade the LSST community has conducted a thorough investigation of prospective cadences (Connolly et al., 2014). Different proposed cadences are available through the LSST Operations Simulator (OpSim, Reuter et al., 2016)[1]. OpSim simulates the field selection and image acquisition for LSST observations over its proposed 10 year baseline. It takes into account sky brightness, bad weather predicted from localized weather models, seeing, telescope maintenance times, slew times, and filter changes.

The LSST observing strategy includes two main components, the Wide Fast Deep Survey and the Deep Drilling Fields (DDFs). The Wide Fast Deep Survey will be at least 18000 deg^2 where each 9.6 deg^2 field has a median of 825 visits summed over all six filters. The DDFs are five relatively small (9.6 deg^2) fields which will have deeper coverage and a higher cadence than the Wide Fast Deep Survey. In this paper, we focus on objects observed in the DDFs. Specifically, here we use the proposed cadences for the XMM-LSS DDF. Other DDFs have similar cadences.

In this work, we use the baseline v2.0 DDF cadence, and the DDF cadence with the longest available season length (ddf season length slf0.10 v2), which uses 80% of the available season, or a total of roughly 210 days. We primary use the u- and y-band cadences, to take advantage of the largest wavelength range available for LSST. Over ten years the baseline

[1]Code available at: https://github.com/lsst/sims_featureScheduler

v2.0 cadence has 1156 and 4515 samples for the u- and y-bands, respectively, while the long season cadence has 842 and 4434 samples for the u- and y-bands, respectively. We down-sample this number to eliminate visits that occur within 1/100 of a day of each other and end up with 186 (150) and 443 (443) samples for the baseline (long season) cadence, for the u- and y-bands, respectively. This down-sampling still leaves a large fraction (15–35% depending on the band and cadence) of visits that occur within roughly 15 minutes of each other. For the baseline cadence, most of the remaining samples are separated by roughly 1 to 3 days. For the long season cadence the separation is typically around 10 to 30 days, with more visit separations on the longer side for the u-band.

We find a longer season length, and as a result shorter season gaps, to be beneficial for detecting long lags in our mock light curves, even though the longer season has lower cadence during the season (see Figure 3.6 and discussion in Sections 3.4.2 and 3.5.3). Therefore, in Sections 3.5.1, 3.5.2, 3.5.4, and 3.5.5 we use the long season cadence. In Section 3.5.3, we discuss the relative performance of our lag detection methods on light curves with the baseline cadence and ways to improve our results if the baseline cadence is selected.

After down-sampling our driving light curve, we add band-dependent dust extinction and measurement error by calculating the signal to noise ratio (SNR) for each observation using the sky brightness and instrument noise from OpSim. In this calculation we use the average magnitudes in each SDSS quasar bin for quasars with $i < 20$ mag (see Table 3.1), to first test our ability to detect the long lag in the most luminous quasars. The median SNR for our mock quasars over all bins is ~ 150. We show that we should be able to detect long lags for fainter quasars as well in Section 3.5.5.

The top panel of Figure 3.3 shows an example DRW light curve. The driving light curve is shown in blue, and the same light curve reprocessed with a long lag of -50 days is shown in orange. The middle panel shows the same light curves sub-sampled with the long season LSST cadence for the u-band and y-band, for the driving and reprocessed light curves, respectively.

3.4 Lag Detection Methods

In this section, we outline several methods developed for detecting short lags and BLR lags that we will apply to the detection of long lags. First, we briefly discuss different cross-correlation methods (Gaskell & Peterson, 1987; White & Peterson, 1994; Peterson et al., 2004), and the Von-Neumann estimator presented by Chelouche et al. (2017). Next, we give a more thorough description of the two methods we explore in detail in this paper, JAVELIN (Zu et al., 2011) and a Fourier method (Zoghbi et al., 2013). We test a variety of methods on our mock light curves because each one has different strengths and drawbacks. These drawbacks mostly arise because these methods need to work for unevenly sampled UV/optical light curves, like those LSST will produce, which will not have an even cadence because of day-time gaps, bad weather, seasonal gaps, shared telescope time, etc.

3.4.1 Cross-Correlation Methods and the Von-Neumann Estimator

With uniformly-sampled light curves, time lags can be estimated by determining the delay time, τ, with the largest cross-correlation coefficient,

$$CC_{\rm coeff}(\tau) = \frac{\sum d(t) r(t+\tau)}{\sqrt{\sum d(t)^2 r(t+\tau)^2}}. \tag{3.13}$$

Multiple lags can in principle be detected by finding multiple peaks in the correlation coefficients.

We show the cross-correlation coefficients as a function of lag for mock DRW light curves with an evenly-sampled, daily cadence and ten-year baseline in Figure 3.4. The light curves have been reprocessed with long (short) lags of -50 (7), -130 (9), and -400 (12) days, in the top, center, and bottom panel, respectively. As the ratio between the amplitude of the long lag and the amplitude of the short lag decreases the peak around the long lag decreases.

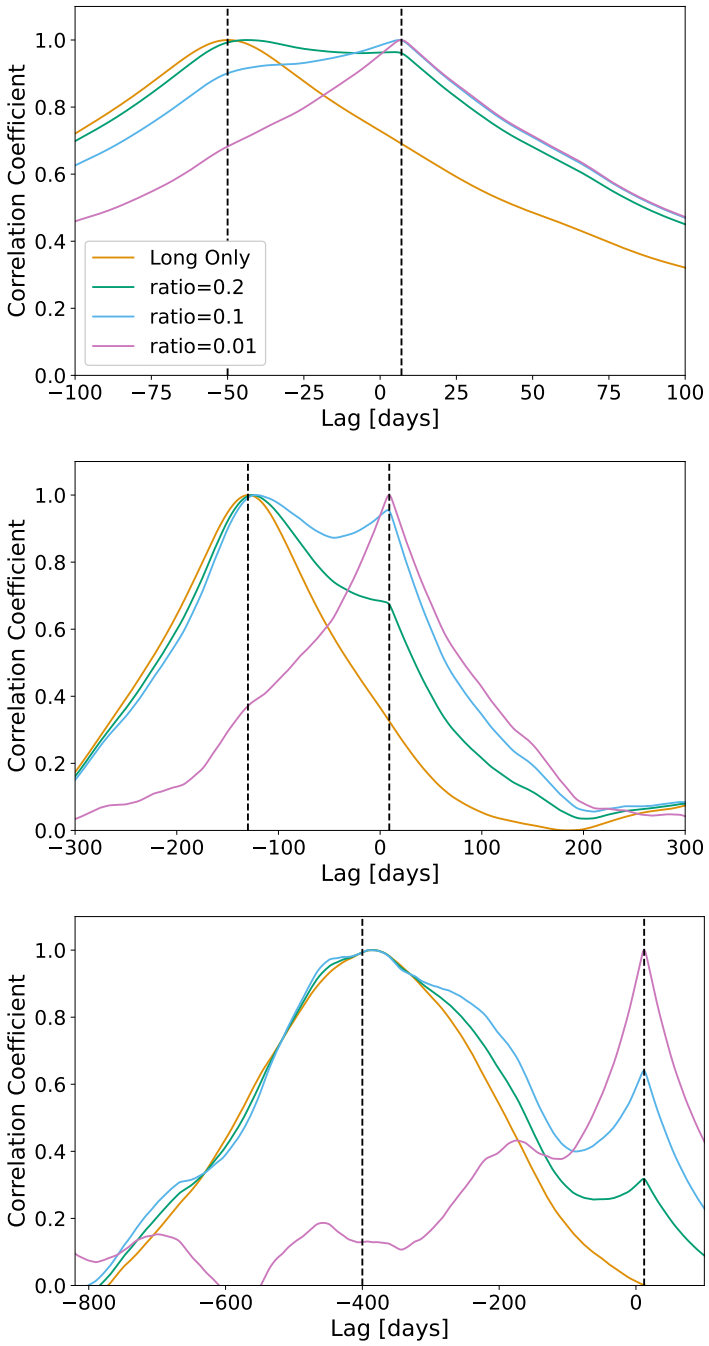


Figure 3.4: The cross-correlation coefficients of even daily cadence light curves that have been reprocessed with long (short) lags of −50 (7), −130 (9), and −400 (12) days, in the top, center, and bottom panel, respectively. Distributions shown in orange have been reprocessed with only a long lag. The ratios of the amplitude of the long to the amplitude of the short lag for the light curves with distributions shown in green, light blue, and pink, are 0.2, 0.1, and 0.01, respectively. The dashed vertical lines show the input lags.

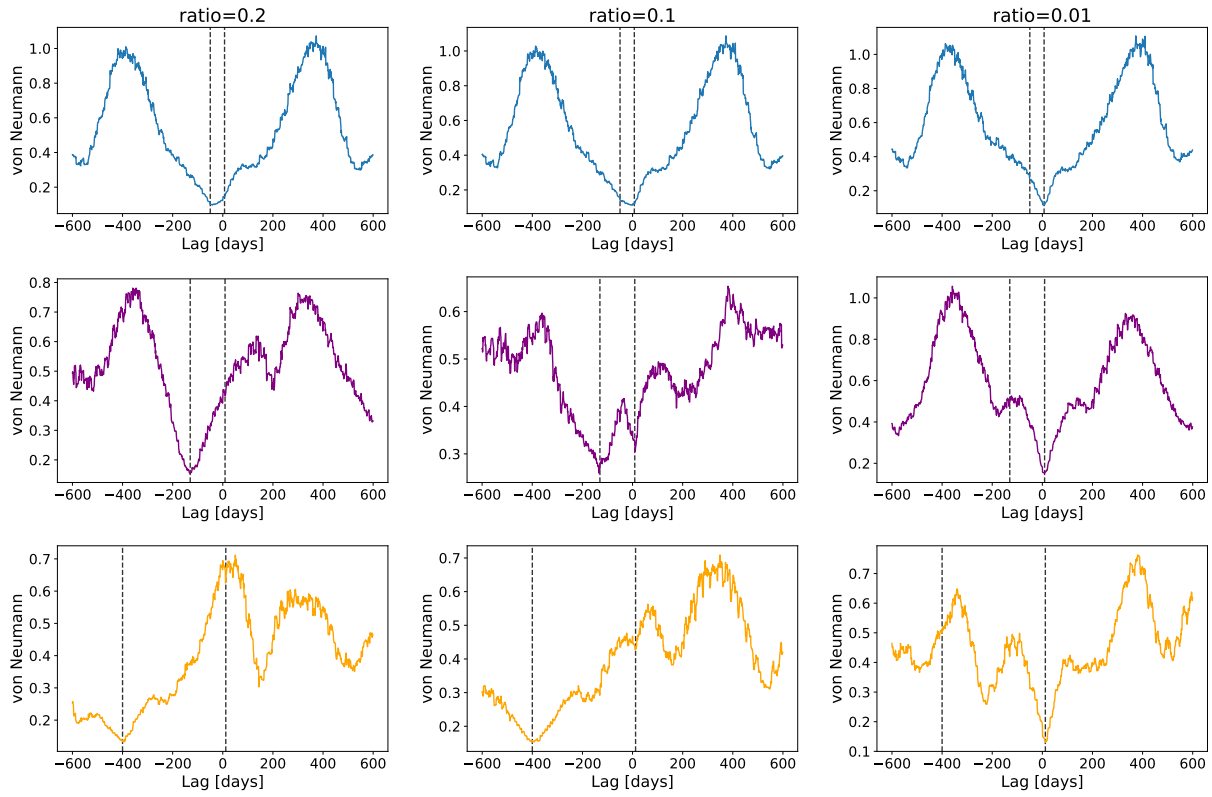


Figure 3.5: The mean of the Von-Neumann value or the mean-square successive difference for 1000 MCMCs as a function of lag for three mock long season LSST light curves with input long lags of −50, −130, and −400 days and short lags of 7, 9, and 12 days in the top, middle and bottom panel, respectively. The ratios of the amplitude of the long lag to the amplitude of the short lag are 0.2 (left panels), 0.1 (middle panels), and 0.01 (right panels). The dashed vertical lines show the input lags.

When the amplitude ratio is 0.01, the peak around the long lag nearly entirely disappears for input long lags of −50 and −400 days, although there remains a small peak around the long lag for the light curve with a −130 day input long lag.

On the other hand, as the amplitude ratio decreases, the peak around the short lag becomes more pronounced. For an input long lag of −50 days it becomes the largest peak when the amplitude ratio is 0.1. For the other two input long lags, the peak around the short lag becomes dominant when the amplitude ratio is 0.01. We show the positive and negative lag values with the highest correlation coefficients for these evenly-sampled light curves in Table 3.5.

To use this cross-correlation technique on unevenly sampled data, such as LSST light curves, we use the interpolated cross-correlation method (ICCF, Gaskell & Peterson, 1987), which performs a linear interpolation to add values to the gaps in the light curves to make them evenly sampled. We use the ICCF on the mock light curves described in Section 3.3 and find that the results are very consistent with what we find when we use cross-correlation on evenly sampled light curves. That is, the uneven cadence does not prevent us from being able to find an accurate lag, but as the amplitude ratio between the long and short lag decreases we are less (more) able to recover the input long (short) lag. We do not try other cross-correlation methods on our light curves, such as the z-transformed discrete correlation function (ZDCF, Alexander, 2013), but we expect these methods to perform similarly to the ICCF.

We do however try the Von-Neumann estimator presented by Chelouche et al. (2017). This Von-Neumann method shifts two light curves by some time τ and measures the mean-square successive-difference, which is a measure of the randomness, between the two light curves as a function of τ. The lag is therefore the τ with the smallest measure of randomness. This method is unique because it involves no model fitting, interpolations, or binning, and is therefore ideal for unevenly sampled light curves.

Figure 3.5 shows the mean of the mean-square successive differences for 1000 Monte Carlo Markov Chains (MCMCs) for an example mock light curve reprocessed with a long (short) lag of -50 (7) days in the top panel, -130 (9) days in the center panel, and -400 (12) days in the bottom panel. The ratio of the amplitude of the long lag to the amplitude of the short lag is 0.2, 0.1, and 0.01, in the left, center, and right panels, respectively. Because the Von-Neumann method assumes a lag that is constant with frequency it has trouble identifying two lags simultaneously. For the shortest duration, -50 day, lag this assumption means the Von-Neumann method combines the two lags into a single wider trough when the amplitude ratio is 0.2. That trough gets skewed towards the short lag when the amplitude ratio is 0.1. For the longer duration lags the assumption of only one lag results in the Von-Neumann

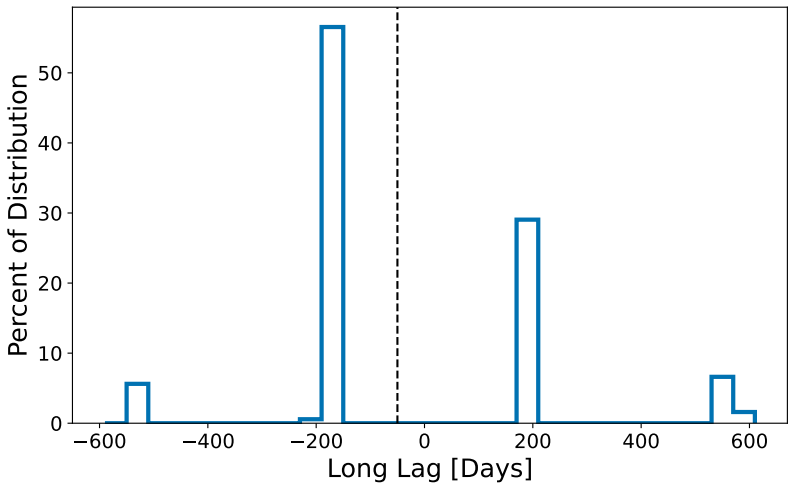


Figure 3.6: An example JAVELIN distribution for a mock light curve with the baseline cadence. While there is a very small peak around the input lag (shown as a dashed line) of −50 days, the peaks around ±180 days are larger.

method finding only one of either the long or the short lag for most amplitude ratios. Only for the light curve with an input long lag of −130 days and an amplitude ratio of 0.1 are both lags detected at all, with significant dips in the distribution near each lag.

As with the cross-correlation method, the single-lag assumption is a drawback of using the Von-Neumann method to detect a multi-lag signal, unless one signal can be filtered out in some way. However, for an amplitude ratio of 0.2, we find that the Von-Neumann method could be a helpful tool to detect long lags in LSST light curves with long seasonal gaps. We discuss this further in Section 3.5.3.

3.4.2 Javelin

Our first main lag detection method is JAVELIN [2], previously known as SPEAR, which is a stochastic process estimator used to look for time lags in spectroscopic (Zu et al., 2011) or photometric (Zu et al., 2016) light curves. Li et al. (2019) performed a comprehensive comparison of the performance of JAVELIN versus cross-correlation methods over a broad range of light-curve properties for mock SDSS light curves and found that JAVELIN performs

[2]As presented in Zu et al. (2010), code available at: `https://github.com/nye17/javelin`

best at accurately detecting BLR lags. As our lags are of a similar timescale we use JAVELIN as one of two main lag detection methods here.

JAVELIN works on unevenly sampled light curves by fitting a model to the available data and running an MCMC to find the posterior distributions for the DRW parameters ($\tau_{\rm damp}$ and SF_∞) of each light curve. JAVELIN then uses these posteriors to statistically interpolate each light curve as it shifts them temporally and uses an MCMC to simultaneously derive the posterior distributions for the DRW parameters and the response function between different light curves. For its default settings, which we use here, JAVELIN models the response function as a top hat with a finite width centered at time lag τ.

The DRW statistical interpolation works well for shorter gaps in unevenly sampled light curves but is unable to fit the large seasonal gaps in our data, which becomes a problem for the LSST baseline cadence. Figure 3.6 shows a JAVELIN distribution for an example mock baseline cadence light curve. The largest peaks are around $\pm$180 days, which is roughly the length of the gaps in between each observing season. We discuss solutions to this problem in case the baseline cadence is chosen in Section 3.5.3.

The bottom left panel of Figure 3.3 shows the JAVELIN distribution for the example light curves in the panels above, which have been sub-sampled with the long season cadence and reprocessed with an input long lag of -50 days. The probability distribution peaks very sharply around -50 days and the median of the negative values of the distribution is -50 days.

Because the long lag is negative and the short lag is positive, we take the long (short) lag as the median of the negative (positive) values of the JAVELIN distribution. This method could bias our values for the short lag towards slightly larger values, because the short lags are near zero and so the peak related to the short lag could extend below zero. However, we find that peaks in JAVELIN distributions around the input short lags rarely extend past zero.

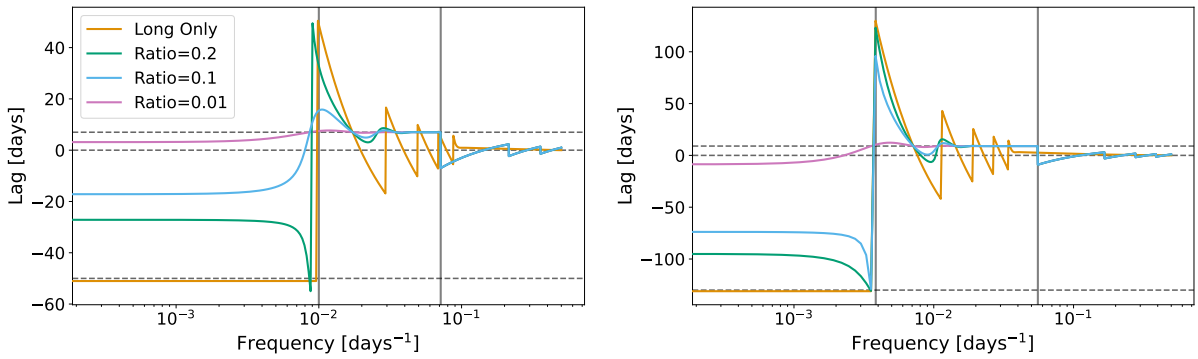


Figure 3.7: The time lags as a function of frequency between two evenly sampled, daily cadence light curves found using the phase of their complex cross-spectrum and Equation (3.15). The input long (short) lags are -50 (7) and -130 (9) days in the top and bottom panel, respectively. Distributions shown in orange have been reprocessed with only a long lag. The ratio between the amplitudes of the long and short lag for the light curves with distributions shown in green, light blue, and pink, are 0.2, 0.1, and 0.01, respectively. The dashed horizontal lines show the input lags and zero and the vertical grey lines show the wrapping frequencies.

While JAVELIN is a useful method for detecting lags in quasar light curves it does make assumptions about the light curve that are model dependent. Also, because the second light curve is modeled as a smoothed and shifted version of the first, the lag can depend on which light curve you model first. In addition, because JAVELIN uses a DRW model, which is what we use to generate our mock light curves, we could be biasing our comparison to favor JAVELIN. We address this potential bias in Section 3.5.4, where we test our methods on light curves from radiation MHD simulations instead of DRW mock light curves.

3.4.3 Maximum-Likelihood Fourier Method

Our second main lag detection method is a Fourier method. Fourier methods are useful because they allow us to separate lags as a function of frequency. They are straightforward as well because calculating the cross-spectrum is a simple multiplication. We can use Equation (3.12) and the power-spectral densities (PSD) of the reprocessed and driving light curve, $|\mathrm{D}(\nu)|^2 = \mathrm{D}^*(\nu)\mathrm{D}(\nu)$ and $|\mathrm{R}(\nu)|^2 = \mathrm{R}^*(\nu)\mathrm{R}(\nu)$, respectively, to derive the cross-spectrum of

the two light curves to be,

$$C(\nu) = D^*(\nu)R(\nu) = |D(\nu)|^2\Psi(\nu). \tag{3.14}$$

The phase lag, ϕ, is then just the phase of the complex cross-spectrum, which is related to the time lag as (Uttley et al., 2014),

$$\tau = \phi/2\pi\nu. \tag{3.15}$$

We show the lags as a function of frequency recovered by this Fourier method between two 10 year baseline, evenly-sampled, daily cadence light curves in Figure 3.7. The input long (short) lags are -50 (7) and -130 (9) days in the top and bottom panel, respectively. We do not use this method to detect long lags of -400 days because the LSST baseline is not long enough to recover enough data points at frequencies $< 10^{-3}$ days^{-1}. The vertical grey lines show the phase wrapping frequencies, which are equivalent to $\nu_w = 1/2C$ for a symmetric response function (Uttley et al., 2014). Phase wrapping, the sudden flip in the sign of the lag, occurs because phase lags are limited to between $\pm\pi$, so a positive or negative shift of half a wavelength cannot be distinguished.

For light curves with only an input long lag, at frequencies lower than the first phase wrapping the measured lag is -51 and -131 days for input long lags of -50 and -130 days, respectively. At frequencies above the long lag's phase wrapping, the lag values oscillate around and then approach zero. When a short lag is added and the ratio between the amplitude of the long lag and the amplitude of the short lag is 0.2, the lags measured at low frequencies become roughly -27 and -95 days for input long lags of -50 and -130 days, respectively. For an amplitude ratio of 0.1 the duration of these lags decrease again to -17 and -74 days for input long lags of -50 and -130 days, respectively. For an amplitude ratio of 0.01 the lag measured at the lowest frequencies becomes positive for an input long lag of -50 days and is only -8 days for an input long lag of -130 days.

The decrease in the duration of the recovered long lags as the amplitude ratio decreases is due to the growing influence of the short lag on lower frequencies. While the short lag has the strongest impact at frequencies around the inverse of the short lag timescale, it is also applied to all frequencies below that. Therefore it reduces the value of the long lag.

On the other hand, because the long lag only acts at lower frequencies, the short lag is easier to isolate in Figure 3.7. For light curves with two input lags, at frequencies higher than the long lag phase wrapping frequency the lag values oscillate before quickly converging on the short lag value for all amplitude ratios. The smaller the amplitude ratio the sooner they converge. The lag values stay at the input short lag, until the short lag wrapping frequency is reached, when the lag values again oscillate and then in this case approach zero. We summarize the long and short lag values in Figure 3.7 in Table 3.5.

This Fourier method is very useful for recovering the shape of a response function for evenly sampled X-ray data (e.g. Papadakis et al., 2001; McHardy et al., 2007; Arévalo et al., 2008a; Fabian et al., 2009; Zoghbi et al., 2010; De Marco et al., 2013; Cackett et al., 2013; Uttley et al., 2014), but requires a work-around for unevenly sampled data, which can not be fast Fourier transformed. Here we use a maximum-likelihood method described in Zoghbi et al. (2013) and first presented by Miller et al. (2010).

First, we create $N = 4$ bins in frequency space ranging from $f_{\rm min} = 1/T$, where $T =$ 10 years is the length of the light curve, to $f_{\rm max} = 0.055$ day^{-1}. As in Zoghbi et al. (2013) we also add a bin on either end to be discarded later to reduce bias in the minimum and maximum bins. $f_{\rm max}$ is chosen based on the maximum frequency for which there is significant power for an LSST cadence. Figure 3.8 shows the distribution of the inverse of the time between each observation for the LSST long season cadence. Because the length of the light curves is 10 years, frequencies less than 10^{-2} days^{-1} are common, but because the cadence is roughly every ten days, frequencies greater than a few times 10^{-2} days^{-1} are rarer.

It is easier to obtain accurate results with fewer bins. Previous testing shows that errors for each bin are smaller when the width of bins are at least a few times $1/T$. This requirement

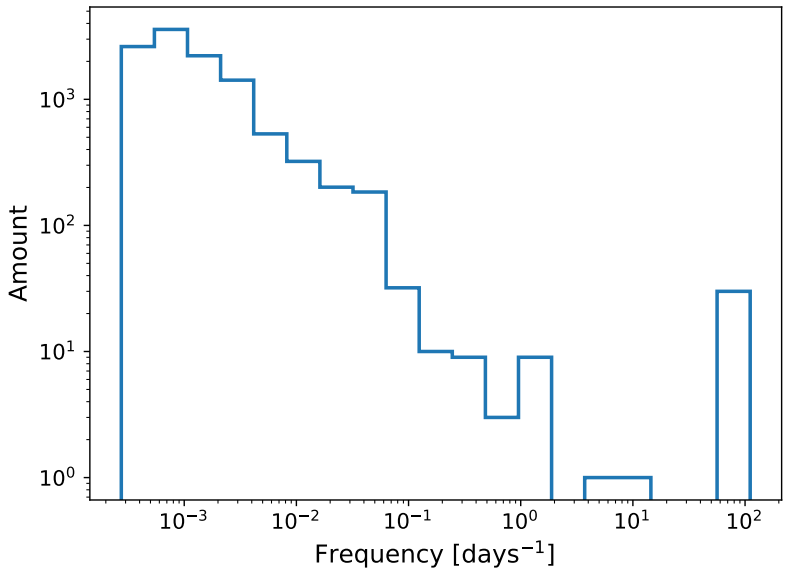


Figure 3.8: The distribution of the inverse of the time between each observation in the LSST long season cadence.

makes it difficult to detect -400 day lags, and so we do not attempt to detect that lag with this method. Otherwise, four bins are sufficient for our purpose here: detecting lags from two Gaussian response functions representing the long and short lag, which occur at order of magnitude different frequencies. However, additional bins would be useful for detecting the shape of more complicated response functions.

Using these bins we fit a piece-wise PSD to each light curve and calculate the likelihood using the autocovariance function,

$$\mathcal{A}(\tau) = \int_{-\infty}^{\infty} |\mathrm{D}(\nu)|^2 \cos(2\pi\nu\tau)\mathrm{d}\nu, \tag{3.16}$$

where $|D(\nu)|^2 = \sum_i D_i$ is the PSD as a piecewise sum over a range of frequency bins. We then maximize this likelihood to find the PSD for each light curve. We use these PSD as input parameters for an MCMC, which gives us a distribution for the power in each bin. We take the median of the distribution for each bin as the power for that bin and combine them to form a piecewise PSD over the full range of frequencies.

Next, we maximize the log-likelihood of the cross-covariance,

$$\begin{aligned}\chi(\tau) &= \int_{-\infty}^{\infty} C(\nu) cos(2\pi\nu\tau + \phi(\tau)) df \\ &= \int_{-\infty}^{\infty} |D(\nu)|^2 |\Psi(\nu)| cos(2\pi\nu\tau - \phi(\nu)) d\nu,\end{aligned} \tag{3.17}$$

to obtain the best fit values of $\Psi(\nu) = \sum_i \Psi_i$ and $\phi(\nu) = \sum_i \phi_i$ as a function of frequency using the PSD from the MCMC distributions. We again use the best fit parameters as input for an MCMC, which gives us the distribution of $\phi(\nu)$. For each bin we use the median of the distribution as that bin's $\phi(\nu)$ and the standard deviation as the error of $\phi(\nu)$. Finally, we can use $\phi(\nu)$ to determine the lag for each frequency bin using Equation (3.15).

We show an example of the lag as a function of frequency detected using this maximum-likelihood method in the bottom right panel of Figure 3.3, for the same light curve with an input long lag of -50 days shown above. The vertical grey line shows the phase wrapping frequency for -50 days. In Figure 3.3 the lowest frequency bin has the largest error-bar and is consistent with zero. In general for our mock light curves this bin tends to have the largest error and to be the least consistent with the input long lag. For these reasons we neglect the lag found in this bin, and then calculate the long lag as the error-weighted mean of all recovered lags < -10 days. We apply a cut-off of -10 days, because lags will approach zero at the wrapping frequency, and we only care about negative lag values when searching for the long lag. In Figure 3.3 the two bins below the wrapping frequency have an error-weighted mean value of -49 days, which is similar to the input lag. The bin above the wrapping frequency is consistent with zero for this light curve, because there is no input short lag. For light curves with input short lags we also use the maximum-likelihood method to detect the short lag. To avoid interference from a long lag signal, we follow Yao et al. (2023) and calculate a short lag by taking the lag value from the one bin centered between $0.01 < \nu < 0.1$.

Table 3.2.

	Cadence	Ratio	Input Lag		Cross-Correlation		Fourier	
			Long	Short	Long	Short	Long	Short
(1)	Even Daily	Long Only	−50	–	−50	1	−51	−1
			−130	–	−129	1	−131	3
			−400	–	−385	518	–	–
(2)	Even Daily	0.2	−50	7	−44	5	−27	7
			−130	9	−125	1	−96	9
			−400	12	−386	500	–	–
(3)	Even Daily	0.1	−50	7	−1	6	−17	7
			−130	9	−122	8	−74	9
			−400	12	−390	12	–	–
(4)	Even Daily	0.01	−50	7	−1	7	3	7
			−130	9	−1	9	−8	9
			−400	12	−1	12	–	–

Table 3.3. The peaks in the cross-correlation coefficient and the long and short lags detected by the Fourier method for the even-cadence light curves with different input lags and amplitude ratios (between the long and short lag) shown in Figures 3.4 and 3.7.

3.5 Using LSST to Detect Long Lags

3.5.1 Detecting Long Lags Alone

First we examine the performance of JAVELIN and the maximum-likelihood method on light curves with the long season LSST cadence (see Section 3.3.3) that have been reprocessed with a long lag only. Figure 3.9 shows the distributions of the medians of the negative values of the JAVELIN distributions for 100 mock light curves with input long lags of −50 days (top panel), −130 days (middle panel), and −400 days (bottom panel). The lag detected by JAVELIN agrees within 10% for all light curves with input lags of −50 days and −130 days and 88% of light curves with input lags of −400.

Table 3.4.

	Cadence	Ratio	Input Lag		JAVELIN		Fourier	
			Long	Short	Long	Short	Long	Short
(1)	Long Season	Long Only	−50	–	-50^{+1}_{-1}	460^{+60}_{-300}	-49^{+5}_{-4}	0^{+4}_{-4}
			−130	–	-130^{+0}_{-0}	490^{+50}_{-300}	-130^{+100}_{-30}	0^{+4}_{-4}
			−400	–	-410^{+10}_{-20}	270^{+260}_{-100}	–	–
(2)	Long Season	0.2	−50	7	-48^{+1}_{-4}	190^{+300}_{-90}	-25^{+7}_{-7}	6^{+2}_{-2}
			−130	9	-130^{+0}_{-0}	250^{+250}_{-120}	-98^{+43}_{-28}	8^{+2}_{-5}
			−400	12	-400^{+30}_{-40}	170^{+260}_{-60}	–	–
(3)	Long Season	0.1	−50	7	-48^{+3}_{-7}	14^{+120}_{-1}	-20^{+7}_{-4}	5^{+2}_{-2}
			−130	9	-130^{+0}_{-0}	240^{+250}_{-220}	-62^{+15}_{-27}	8^{+2}_{-4}
			−400	12	-410^{+40}_{-40}	160^{+310}_{-110}	–	–
(4)	Long Season	0.01	−50	7	-140^{+30}_{-170}	9^{+5}_{-1}	-26^{+3}_{-32}	5^{+2}_{-2}
			−130	9	-130^{+30}_{-80}	11^{+2}_{-2}	-53^{+38}_{-6}	8^{+2}_{-3}
			−400	12	-250^{+90}_{-230}	14^{+2}_{-1}	–	–
(5)	Baseline u-y	0.2	−50	7	-180^{+0}_{-10}	190^{+10}_{-10}	-31^{+12}_{-17}	9^{+2}_{-3}
			−130	9	-180^{+10}_{-0}	190^{+10}_{-10}	-64^{+21}_{-34}	9^{+2}_{-14}
			−400	12	-180^{+10}_{-350}	190^{+10}_{-0}	–	–
(6)	Baseline g-z	0.2	−50	7	-180^{+30}_{-10}	180^{+350}_{-10}	-27^{+8}_{-12}	7^{+2}_{-3}
			−130	9	-170^{+20}_{-10}	190^{+350}_{-10}	-98^{+50}_{-30}	9^{+2}_{-14}
			−400	12	-190^{+10}_{-340}	190^{+350}_{-10}	–	–
(7)	Baseline + LCO	0.2	−50	7	-51^{+2}_{-9}	450^{+100}_{-280}	-22^{+4}_{-5}	6^{+1}_{-1}
			−130	9	-130^{+0}_{-0}	370^{+170}_{-200}	-85^{+35}_{-28}	8^{+1}_{-1}
			−400	12	-390^{+3}_{-20}	170^{+280}_{-80}	–	–
(8)	Long Season (Non-DRW)	0.2	−50	7	-53^{+3}_{-1}	150^{+370}_{-50}	-21^{+1}_{-3}	6^{+0}_{-0}
			−130	9	-130^{+0}_{-10}	130^{+350}_{-20}	-95^{+6}_{-14}	9^{+1}_{-2}
			−400	12	-430^{+40}_{-20}	190^{+30}_{-20}	–	–
(9)	Long Season (i=21 mag)	0.2	−50	7	-50^{+2}_{-4}	480^{+40}_{-350}	-26^{+7}_{-7}	5^{+2}_{-2}
			−130	9	-130^{+10}_{-10}	480^{+40}_{-350}	-98^{+34}_{-30}	8^{+2}_{-8}
			−400	12	-410^{+30}_{-30}	170^{+320}_{-40}	–	–

Table 3.5. The median long and short lags detected by JAVELIN and the maximum-likelihood method for 100 mock light curves with each cadence, amplitude ratio, and input long and short lag used throughout the paper. The error ranges represent one standard deviation.

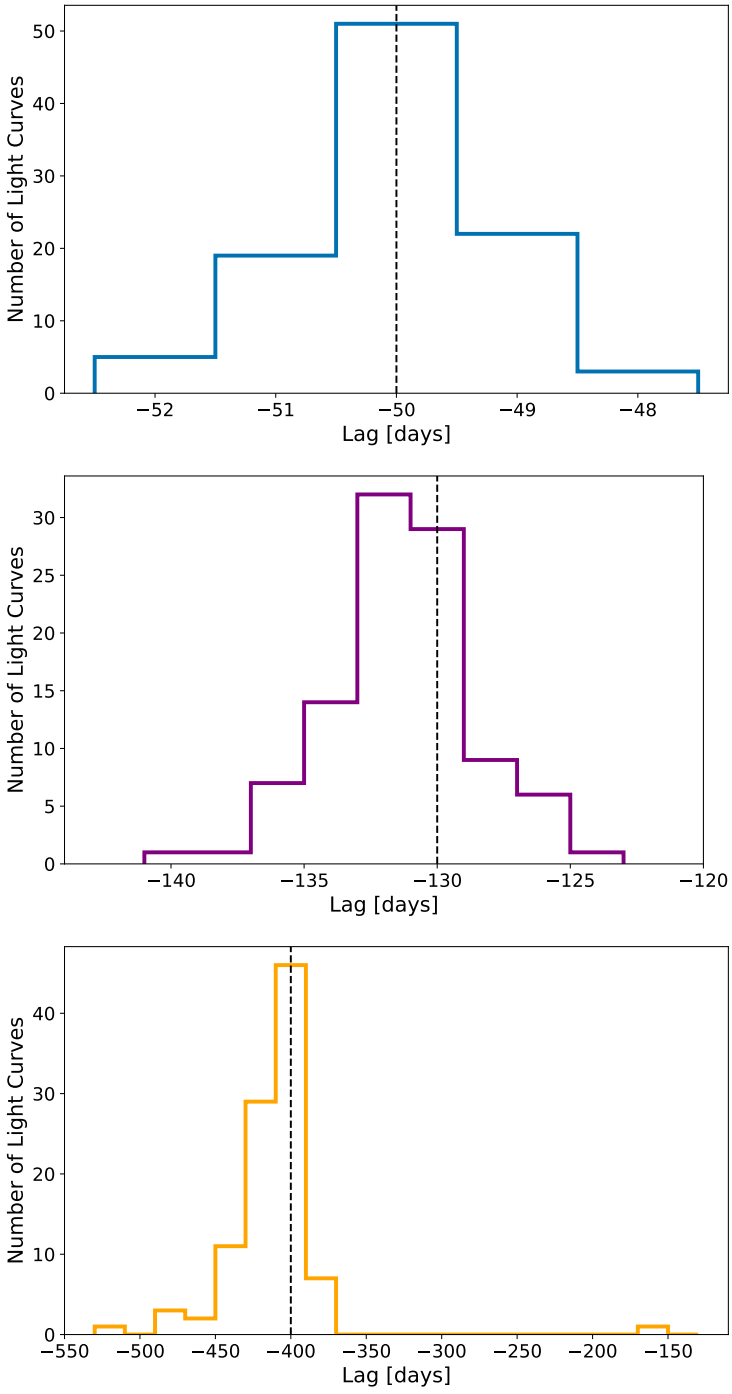


Figure 3.9: The distribution of the long lags detected by JAVELIN for 100 mock long season cadence LSST light curves with input long lags of −50, −130, and −400 days (shown as dashed black lines) in the top, center, and bottom panel, respectively.

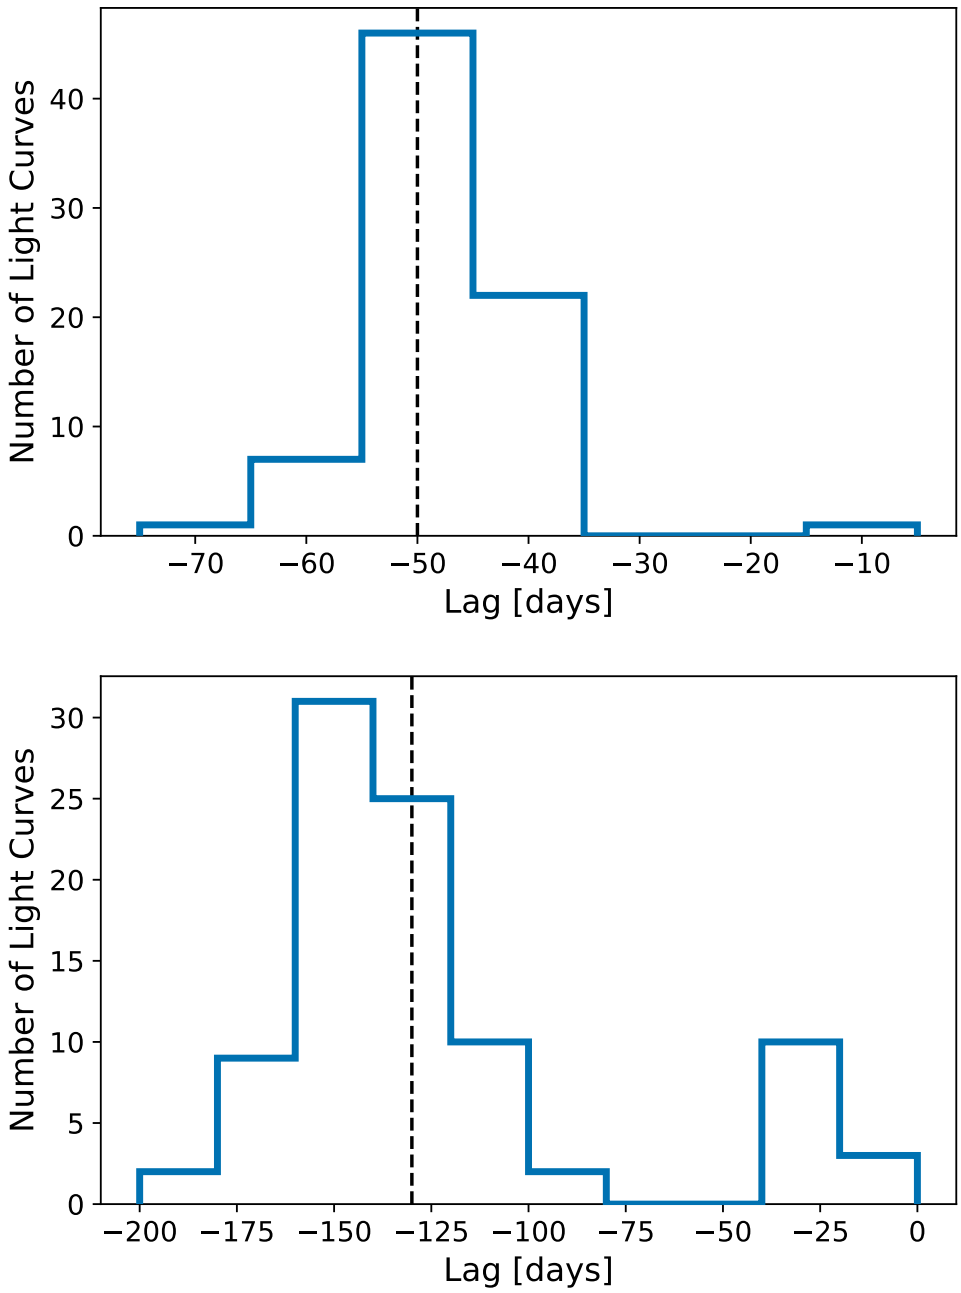


Figure 3.10: The distribution of long lags recovered by the maximum-likelihood method for 100 mock long season cadence LSST light curves with input long lags (shown as dashed black lines) of −50 days (top panel) and −130 days (bottom panel).

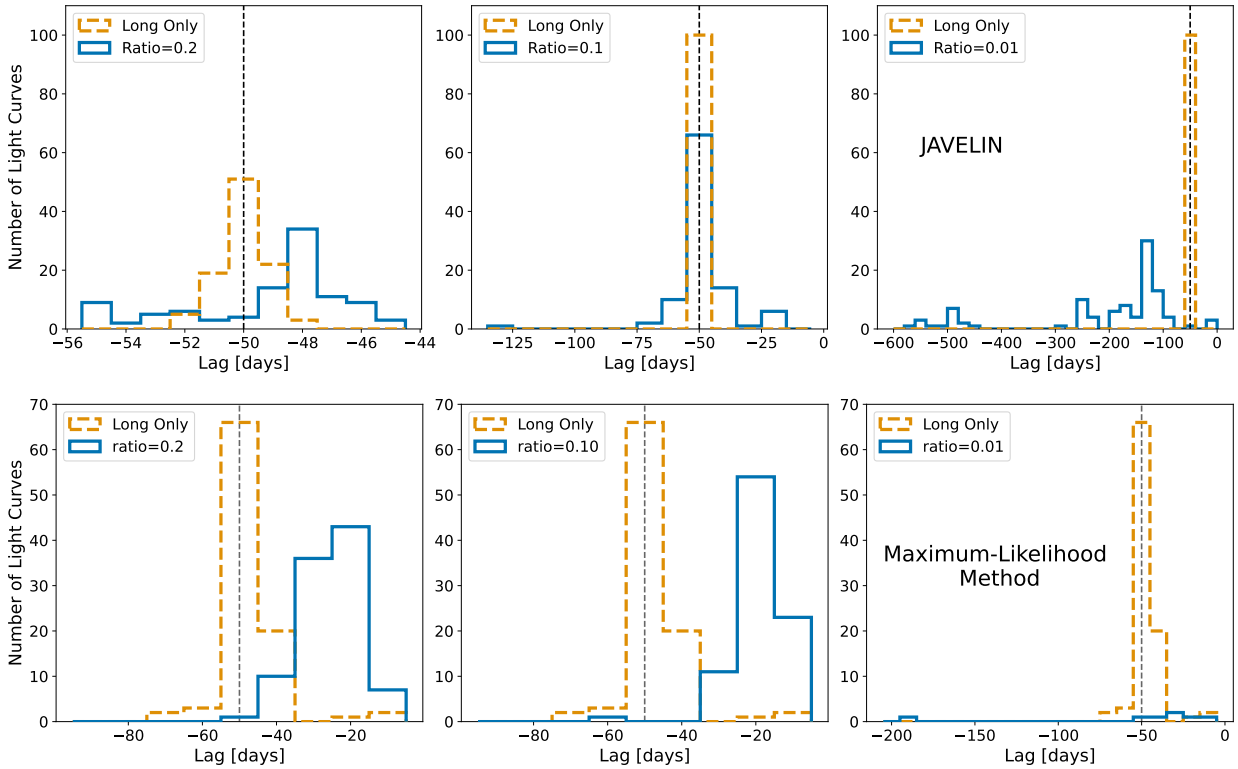


Figure 3.11: The distributions of the long lags detected by JAVELIN (top panel) and the maximum-likelihood method (bottom panel) for 100 mock long season cadence LSST light curves with an input long lag of −50 days (shown as a dashed black line). Distributions shown in blue are for light curves that also have a short lag of 7 days. The legends show the ratio of the amplitude of the long lag to the amplitude of the short lag for each blue distribution.

Figure 3.10 shows the distribution of lags recovered by the maximum-likelihood method for the same mock light curves excluding those with input long lags of −400 days. A long lag of −400 days is very difficult to detect with the maximum-likelihood method for light curves with the LSST baseline of ten years. For input lags of −50 and −130 days the maximum-likelihood method is less accurate than JAVELIN. However, the maximum-likelihood method is able to detect the long lag to within 30% of the input lag for around 70% of light curves. We show the medians and standard deviations of the JAVELIN and maximum-likelihood method distributions for these mock light curves in row (1) of Table 3.5.

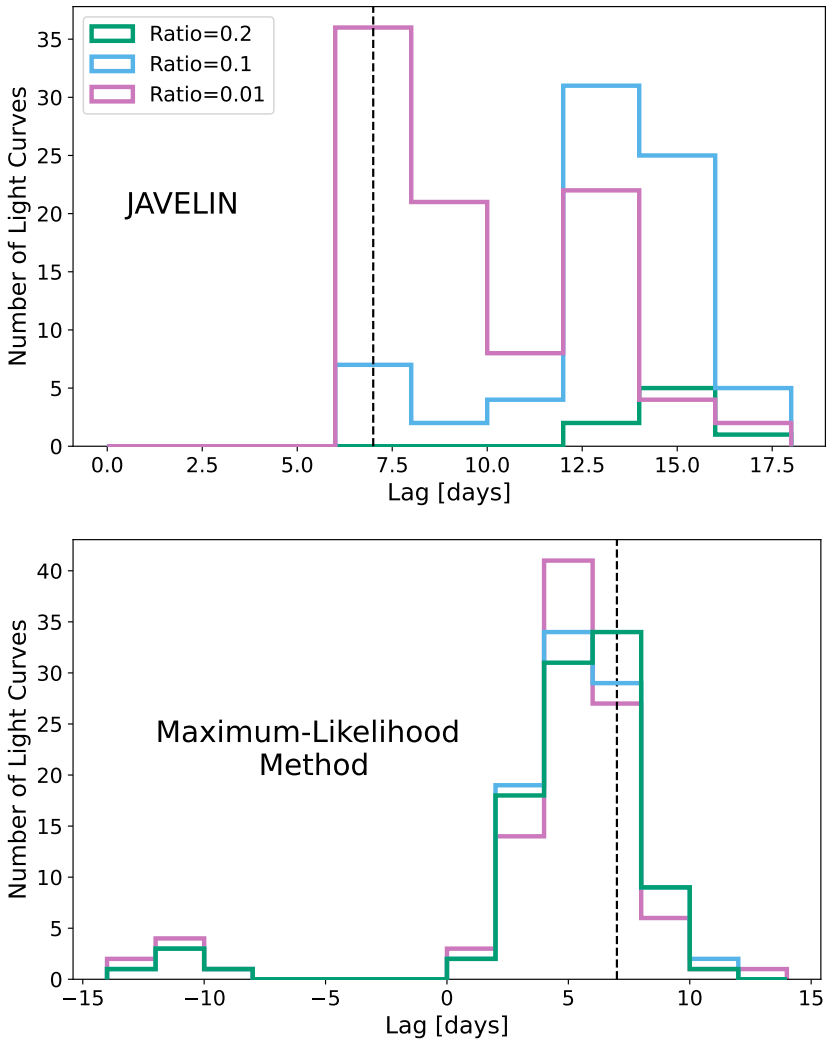


Figure 3.12: The distribution of the short lags detected by JAVELIN (top panel) and the maximum-likelihood method (bottom panel) for 100 mock long season cadence LSST light curves with different amplitude ratios (shown in the legend). In the top panel we zoom in on the portion of the distribution between 0 and 20 days to better show the accuracy of the short lag when detected. The input short lag, shown as the vertical dashed line, is 7 days.

3.5.2 Detecting Short and Long Lags Together

We now add a short lag to our mock light curves with the long season LSST cadence to examine how multiple lag signals can affect our ability to detect either lag. Our fiducial value for the ratio of the amplitude of the long lag to the amplitude of the short lag is 0.2, as we expect the short lag signal to be stronger. We also examine our ability to detect lags when this ratio is 0.1 and 0.01. For readability, in this section we only show figures for light curves with an input long lag of −50 days and an input short lag of 7 days. Unless otherwise noted, the input lag duration does not have an impact on the accuracy of our lag detection methods. We repeat the analysis in this section for light curves reprocessed with long (short) lags of −130 (9) and −400 (12) days in Appendix 3.8.1.

In the top panel of Figure 3.11 we compare the JAVELIN results for light curves reprocessed only with a long lag (in orange) to JAVELIN results for light curves reprocessed with a long and short lag (in blue). For light curves with both lags we take the median of the negative (positive) values in the JAVELIN distribution as the long (short) lag. When the amplitude ratio between the long and short lag is 0.2 or 0.1 the distributions of long lags detected by JAVELIN are consistent with the distribution of long lags detected for light curves without a short lag. Adding in a short lag only makes the distributions somewhat broader. Comparing these results to the correlation coefficients for evenly-sampled daily cadence light curves, shown in Figure 3.4, the long lag tends to be more prominent than we might expect, in particular for an amplitude ratio of 0.1. It may be that the roughly every 10–30 day cadence of the long season LSST cadence, helps us to detect the long lag, because it hides some of the signal from the short lag.

On the other hand, when the ratio is decreased to 0.01, we fail to detect the correct long lag for nearly all mock light curves with an input long lag of −50 days. In Figure 3.23 in Appendix 3.8.1, for an input long lag of −130 days there is still a peak in the distribution around −130 days. It is unclear whether this peak is tied to the cadence of the light curve or the injected long lag. For one, there is also a peak around −130 days in the lag distributions

for input long lags of −50 and −400 days when the amplitude ratio is 0.01. There is also some evidence for a peak around −130 days in our null tests using light curves that have not been reprocessed with a lag (see top left panel of Figure 3.33 in Appendix 3.9.1). JAVELIN detects a long lag of around −130 days for roughly 10% of these light curves with no input lag. However, in the right middle panel of Figure 3.23 JAVELIN accurately recovers a long lag of around −130 days for a significantly higher percentage ($\sim 50\%$) of light curves. Furthermore, in Figure 3.4 there is a slight peak in the cross-correlation coefficient for the evenly-sampled light curve with a −130 day input long lag and an amplitude ratio of 0.01, while there is no corresponding peak in the correlation coefficient for other input lags when the amplitude ratios are 0.01. Therefore it is unclear if the peak in the distribution in the right middle panel in Figure 3.23 is due to a recovered lag or not.

We show the medians of the positive values of the JAVELIN distribution for 100 mock light curves with input lags of −50 and 7 days in the top panel of Figure 3.12. We zoom in to the portion of the distribution falling between 0 and 20 days to better show the accuracy of the short lag when it is detected. When the ratio of the amplitude of the long lag to the amplitude of the short lag is 0.01 JAVELIN is able to detect the short lag, although the distribution is skewed towards larger values than the input short lag. In previous studies, JAVELIN has detected short lags that appear to be skewed towards longer frequencies due to diffuse BLR emission (e.g., Cackett et al., 2018). Something similar may be happening here due to the presence of both a long and short lag.

For an amplitude ratio of 0.1, JAVELIN detects a short lag of around 13–16 days for roughly 60% of light curves with input lags of −50 and 7 days. The null tests do not show a peak around 13 days for this cadence (see top right panel of Figure 3.33 in Appendix 3.9.1). In addition, for the even cadence light curve with an input long lag of −50 days shown in the top panel of Figure 3.4, the peak in the correlation coefficient around the short lag is larger than the peak around the long lag when the amplitude ratio is 0.1. The fact that this peak around the short lag is already larger at this amplitude ratio suggests that it might be

easier to detect the short lag with JAVELIN in this particular case. Therefore it is possible that the light blue peak in the top panel of Figure 3.12 is due to the input short lag, but heavily skewed towards longer lags because of the presence of the long lag. For other input long lags JAVELIN is unable to detect a short lag for most light curves unless the amplitude ratio is 0.01 (see Figure 3.25 in Appendix 3.8.1.)

The bottom panel of Figure 3.11 is the same as the top panel, but instead shows the long lag detected by the maximum-likelihood method. As with the even cadence light curves in Figure 3.7, when a short lag is added, the long lag is shifted towards shorter timescales. The median of the distribution for light curves with an input long lag of -50 days goes from -49 days for light curves with no short lag, to -25 days and -20 days for light curves with amplitude ratios of 0.2 and 0.1, respectively. When the amplitude ratio is 0.01, for light curves with input long lags of -50 days there are very few negative lags detected at all and for light curves with an input long lag of -130 days there are only about 20 negative lags detected, mostly between -10 and -65 days (see Figure 3.24 in Appendix 3.8). It is clear that like JAVELIN the maximum-likelihood method is unable to isolate the long lag when the amplitude ratio is 0.01.

The bottom panel of Figure 3.12 shows the distribution of short lags detected by the maximum-likelihood method. Unlike JAVELIN, the maximum-likelihood method is able to detect short lags that are accurate to within a few days of the input lag for over 60% of light curves with amplitude ratios of 0.2 and 0.1. Detecting short lags more accurately for a wider range of amplitude ratios is therefore one benefit to using the maximum-likelihood method in addition to JAVELIN.

The accuracy of the maximum-likelihood method on light curves where we include two input lags appears to be consistent with Figure 3.7, where we use the Fourier method on uniform 1-day cadence, 10 year-long, DRW light curves. It is promising that the maximum-likelihood method is able to perform nearly as well on light curves with the long season LSST cadence as the Fourier method performs on light curves with an even cadence. However, it

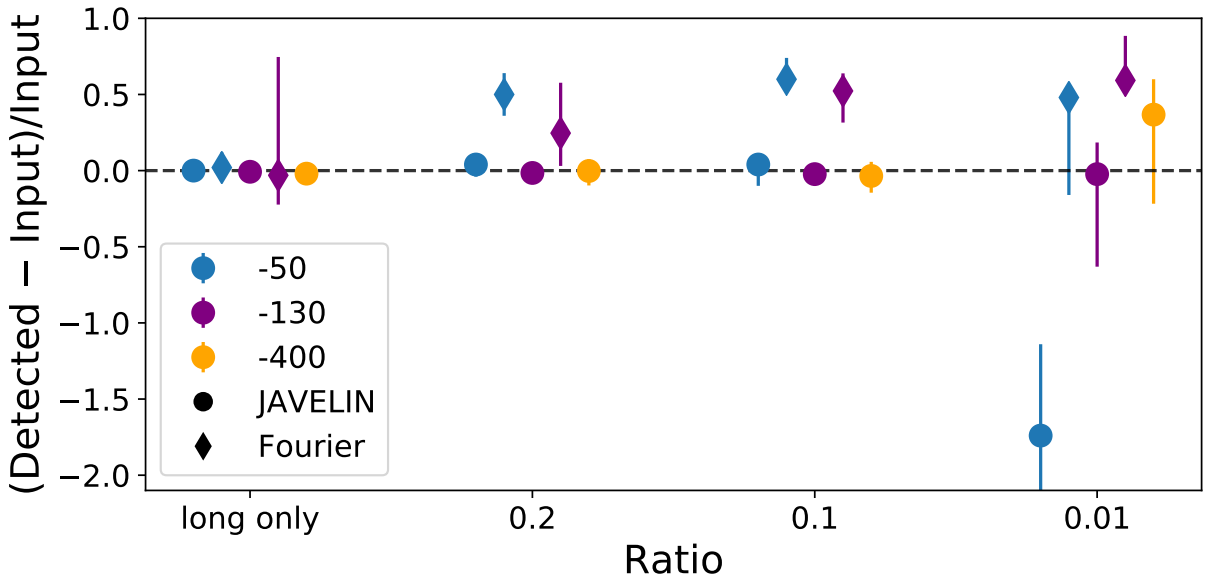


Figure 3.13: The fractional error of the median of the distribution of long lags recovered by JAVELIN (points) and the Fourier method (diamonds) for 100 mock long season cadence LSST light curves as a function of the ratio of the amplitude of the long lag to the amplitude of the short lag for input long lags of −50, −130, and −400 days in blue, purple, and orange, respectively. The error bars represent the fractional error of one standard deviation for each distribution and the horizontal dashed line shows zero. A slight scatter is added in the x-direction to make the individual points easier to view.

will be difficult with this method to disentangle the length of the long lag versus the amplitude ratio between the long and short lag.

We show the medians and standard deviations of the distributions of lags recovered by JAVELIN and the maximum-likelihood method for our 100 mock light curves in rows (2), (3), and (4) in Table 3.5 for light curves with amplitude ratios between their input long and short lags of 0.2, 0.1, and 0.01, respectively. In Figure 3.13 we show the fractional error of the medians of the distributions of lags recovered by JAVELIN and the maximum-likelihood method for the 100 mock light curves as a function of these amplitude ratios for all input long lags. The error-bars show the fractional error of one standard deviation for each distribution. The accuracy of the recovered lags is similar regardless of the durations of the input lags.

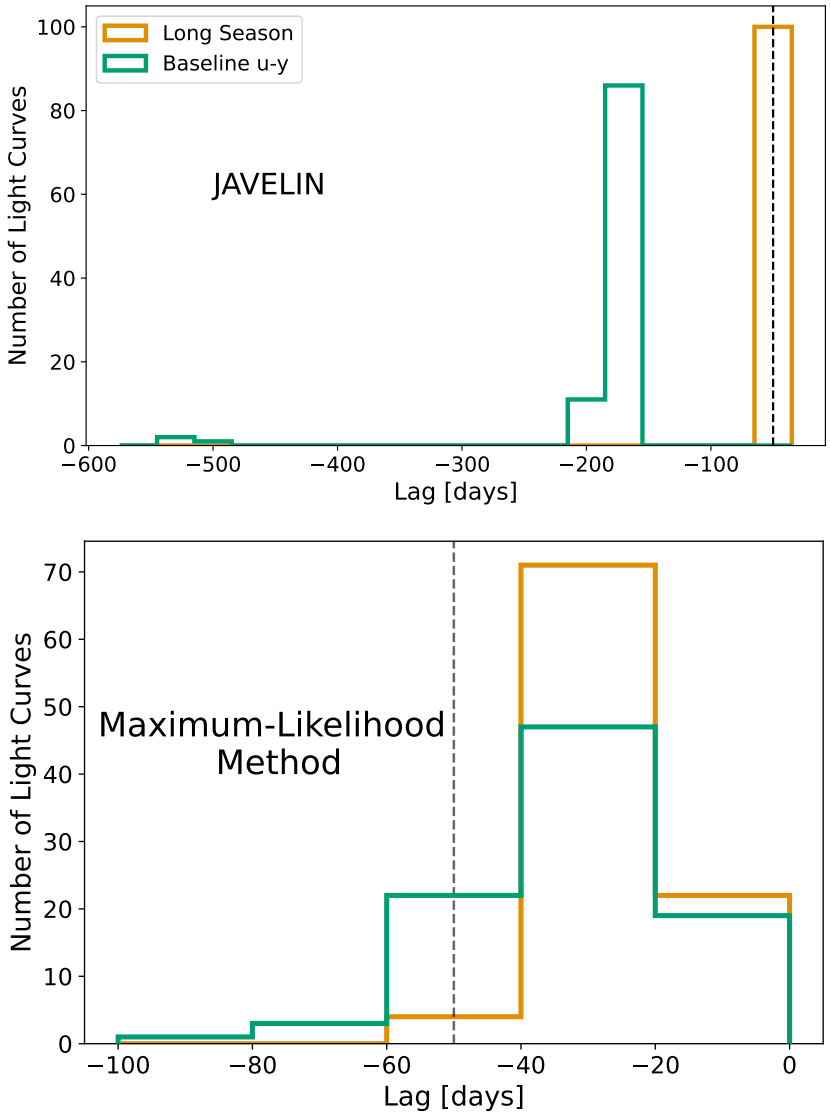


Figure 3.14: The distributions of the long lags recovered by JAVELIN (top panel) and the maximum-likelihood method (bottom panel) for 100 light curves with input lags of −50 and 7 days. The green and orange distributions are for mock light curves with the long season and baseline v2.0 LSST cadences, respectively, for the u- and y-band. The dashed vertical line shows the input long lag.

3.5.3 Improving the Baseline Cadence

The LSST collaboration is still evaluating different possible cadences. We have evaluated and compared the ability of our lag detection methods to accurately pick out long lags for several cadences. Above we asserted that the most significant factor for improving our ability to detect lags is increased season length. In the previous sections we use the long season LSST

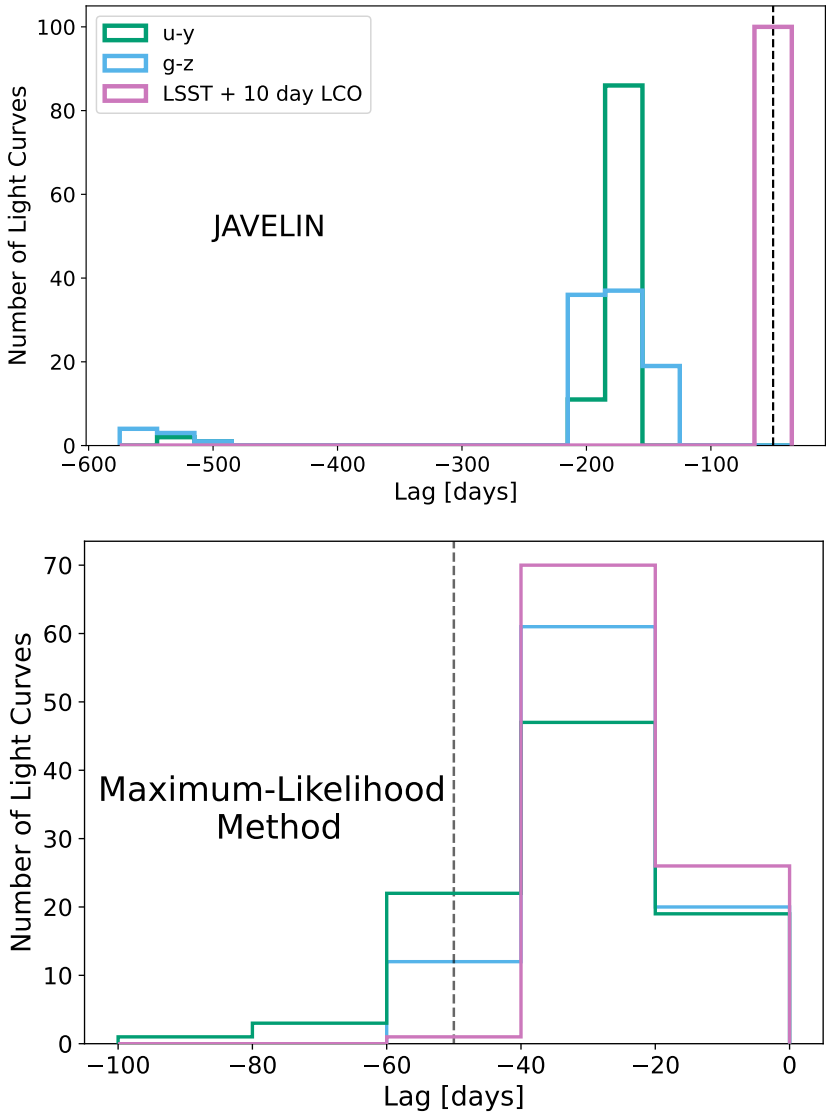


Figure 3.15: Distribution of the long lags recovered by JAVELIN (top panel) and the maximum-likelihood method (bottom panel) for 100 mock light curves with input lags of -50 and 7 days (top panel). The dashed vertical line shows the input long lag. The green and light blue distributions are for mock light curves with the baseline v2.0 LSST cadence for the u- and y-band and g- and z-band, respectively. The distributions shown in pink are for light curves with the baseline g- and z-band cadence plus an LCO observation taken roughly every ten days in each LSST observing gap. These LCO observations amount to around 18 observations in each gap, or per year, for a total of around 180 observations.

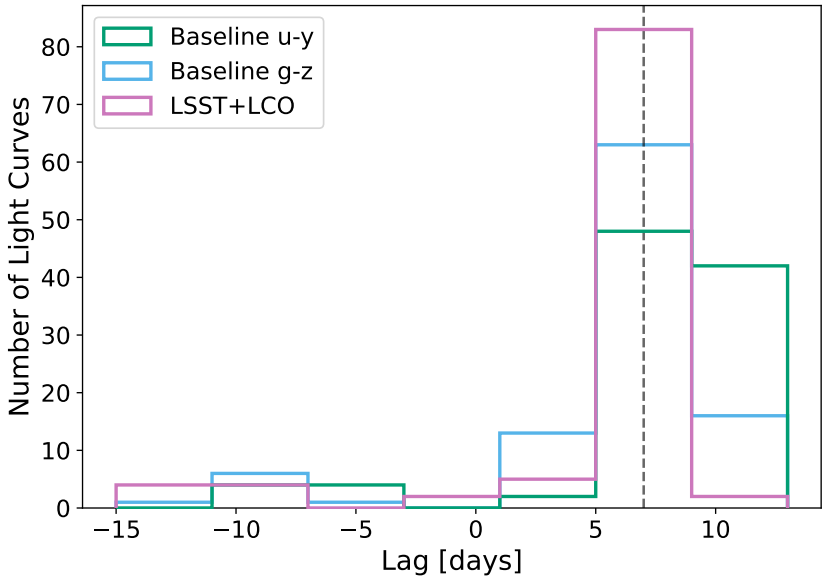


Figure 3.16: Distribution of the short lag recovered by the maximum-likelihood method for 100 mock light curves with input lags of −50 and 7 days. The dashed vertical line shows the input short lag. The green and light blue distributions are for mock light curves with the baseline v2.0 LSST cadence for the u- and y-band and g- and z-band, respectively. The distributions shown in pink are for light curves with the baseline g- and z-band cadence plus an LCO observation taken roughly every ten days in each LSST observing gap. These LCO observations amount to around 18 observations in each gap, or per year, for a total of around 180 observations.

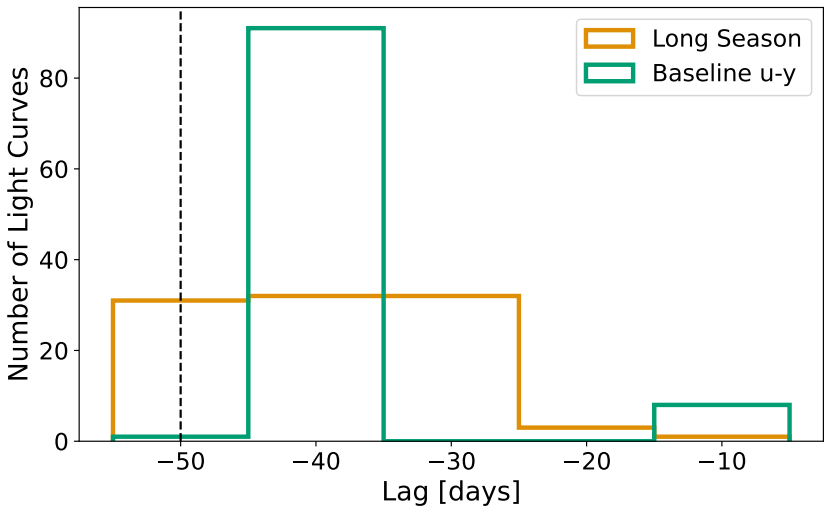


Figure 3.17: Distribution of the long lag recovered by the Von-Neumann method for 100 mock light curves with input lags of −50 and 7 days. The dashed vertical line shows the input long lag. The orange (green) distribution is for mock light curves with the long season (baseline u- and y-band) cadence.

cadence. In this section we compare the detected lags for light curves with this long season cadence to light curves with the baseline v2.0 cadence. We again only show results here for light curves with input lags of −50 and 7 days. We show results for all three lag bins in Table 3.1 in Appendix 3.8.2. Unless otherwise noted, our results for light curves reprocessed with a long lag of −50 days are consistent with results for light curves reprocessed with other input long lags.

The top panel of Figure 3.14 shows the long lags detected by JAVELIN for 100 mock light curves sub-sampled with the long season and baseline cadences, in orange and green, respectively. Unfortunately, when using the baseline cadence JAVELIN fails to accurately recover any long lags. Instead, JAVELIN primarily recovers long lags of around −180 days. The second panel from the top in Figure 3.33 in Appendix 3.9.1 shows the distributions of the medians of the negative and positive values of the JAVELIN distributions for 100 light curves with the baseline cadence but with no input lag, in the left and right panel, respectively. The positive and negative distributions peak strongly around ±180 days, like the distribution shown in green in top panel of Figure 3.14, suggesting that the peak around ±180 days is due to the cadence, not a lag.

The maximum-likelihood method, on the other hand, still performs reasonably well on the light curves sub-sampled with the baseline cadence. Because the maximum-likelihood method separates the signal out in frequency space it is easier to separate out the lag signal from the cadence signal. The bottom panel of Figure 3.14 shows the distribution of lags recovered by the maximum-likelihood method for light curves with the long season or baseline cadence in orange and green, respectively. The distribution of long lags recovered by the maximum-likelihood method is slightly less skewed towards shorter timescales, with a median of −31 days. In Figure 3.28 in Appendix 3.8 the distribution of long lags for light curves with an input long lag of −130 days is more skewed towards shorter timescales, with a median of −58 days. It could be that the presence of a gap at longer timescales actually helps to skew the long lag back towards longer timescales for light curves with input lags significantly

shorter than the observing gaps (~ -50 days), while longer lags (~ -130 days) closer to the gap length are affected differently.

In order to improve the performance of JAVELIN on light curves with large seasonal gaps, we can add observations from other telescopes to fill in the 180 day gap. Here, we explore the effectiveness of filling these gaps with observations from LCO. LCO has been used for numerous RM campaigns and has wavebands ranging from the u–y-band. Here we experiment with adding observations in the g- and z-bands to each seasonal gap on roughly a ten day cadence. We use the g- and z-bands instead of the u- and y-bands to increase the SNR. We determine the additional LCO cadence in the gaps using the LCO g- and z-band cadences from the Hernández Santisteban et al. (2020) Fairall 9 observing campaign, which are roughly daily. We randomize these Fairall 9 observing times over a uniform distribution with a range of half the distance to the following observation, and then down-sample to every tenth observation.

For each lag bin in Table 3.1, we make 100 mock light curves sub-sampled at times corresponding to the LSST baseline cadence for the g- and z-band and these additional LCO observations. The LCO observations add roughly 18 observations per year or 180 total. We add measurement errors corresponding to the SNR for LSST and LCO. The SNR for our mock LCO observations were calculated using the LCO Exposure Time Calculator[3] and range from 19 to 77, depending on the band and redshift bin.

We show the long lags recovered by JAVELIN for 100 of these mock light curves with input lags of -50 and 7 days in pink in the top panel of Figure 3.15. For comparison we show the lags recovered by JAVELIN for 100 mock light curves with the same input lags and the u- and y-band LSST baseline cadences (in green) and g- and z-band LSST baseline cadences (in light blue). Light curves with the g- and z-band LSST baseline cadences experience the same cadence-related issues as light curves with the u- and y-band LSST baseline cadences (see also Appendix 3.9.1). However, adding LCO observations eliminates most of the false

[3] https://exposure-time-calculator.lco.global/

lag detections from the large gaps in the baseline cadence. JAVELIN recovers a long lag within 10% of the input lag for 74% (100% are within 30%) of light curves with input long lags of −50 days. The decline in false lag detections with this LSST+LCO cadence can also be seen in our null tests in Appendix 3.9.1. In the bottom panel of Figure 3.33, the distributions of lags detected by JAVELIN for light curves with this LSST+LCO cadence and no input lag peak around zero, as expected.

The bottom panel of Figure 3.15 is the same as the top panel except for the maximum-likelihood method. The distribution of long lags detected for light curves with the LSST+LCO cadence is similar to the distribution for light curves with the long season cadence. The median long lag detected for light curves with the LSST+LCO cadence is −22 days for an input long lag of −50 days. This median is similar to, although skewed slightly shorter than, the median found for light curves with the long season cadence (−25 days).

We show the distributions of short lags detected by the maximum-likelihood method for the same light curves as above in Figure 3.16. For all cadences the distributions peak around the input short lag. However, for the baseline cadences only around 50% of short lags detected are within 2 days of the input short lag, whereas when the LCO observations are added around 80% of the short lags detected are within 2 days of the input short lag, which is slightly more accurate than with the long season cadence. For both JAVELIN and the maximum-likelihood method it is clear that adding just 18 LCO observations a year in the gaps of the baseline cadence should be sufficient to erase the gap issue and detect lags at rates consistent with those for the long season cadence. We show the medians and standard deviations of the distributions of lags recovered by JAVELIN and the maximum-likelihood method for light curves with the baseline u- and y-band cadences, baseline g- and z-band cadences, and baseline+LCO cadence in rows (5), (6), and (7) of Table 3.5, respectively.

Finally, we find that the Von-Neumann method (see Section 3.4.1) is more effective at detecting long lags for mock light curves with the baseline LSST cadence than JAVELIN or

the Fourier method. The Von-Neumann method is more accurate than these other methods in this case because it does not involve any interpolation or binning.

Figure 3.17 shows the distributions of long lags recovered by the Von-Neumann method for 100 mock light curves with input lags of -50 and 7 days and either the long season cadence (in orange) or the baseline v2.0 cadence for the u- and y-band (in green). Unlike with JAVELIN, the Von-Neumann method is still able to detect an accurate long lag for light curves with the baseline cadence. The median long lag for both distributions are -40^{+8}_{-7} days and -40^{+1}_{-2} days for an input long lag of -50 days for the long season and baseline cadences, respectively. Therefore, the Von-Neumann method appears to be a very useful tool for recovering the long lag if the baseline cadence is selected.

3.5.4 Beyond DRW Light Curves

Although the DRW model has been shown to be a good representation of quasar light curves (e.g., MacLeod et al., 2012), the DRW model is not physically motivated and may be unreliable at the smallest and largest scales. Furthermore, since JAVELIN uses a DRW model to interpolate unevenly-sampled light curves, it is prudent to investigate its performance on light curves that are not generated using a DRW. Therefore, in this Section we use the light curves from the radiation MHD simulation in Jiang & Blaes (2020) (see Section 3.3.1.2). These mock light curves have the long season LSST cadence and the fiducial amplitude ratio of 0.2. The simulation in Jiang & Blaes (2020) is for a 5×10^8 $M_\odot$ SMBH and the light curves vary in Eddington ratio from $l/l_{\rm edd} \approx 0.01 - 1$. A DRW can be fit to the light curves with $\tau_{\rm damp} = 200$, although this fit is not unique.

Figure 3.18 shows the distribution of the median of the negative values of the JAVELIN distribution for the 16 mock light curves made using the simulated light curves from Jiang & Blaes (2020). For an input long lag of -50 and -130 days all medians fall within $\sim 10\%$ of the input lag. 14 out of 16 lags detected for light curves with input long lags of -400 days are within 20% of the input long lag, with only two outliers. Both of these outliers are from

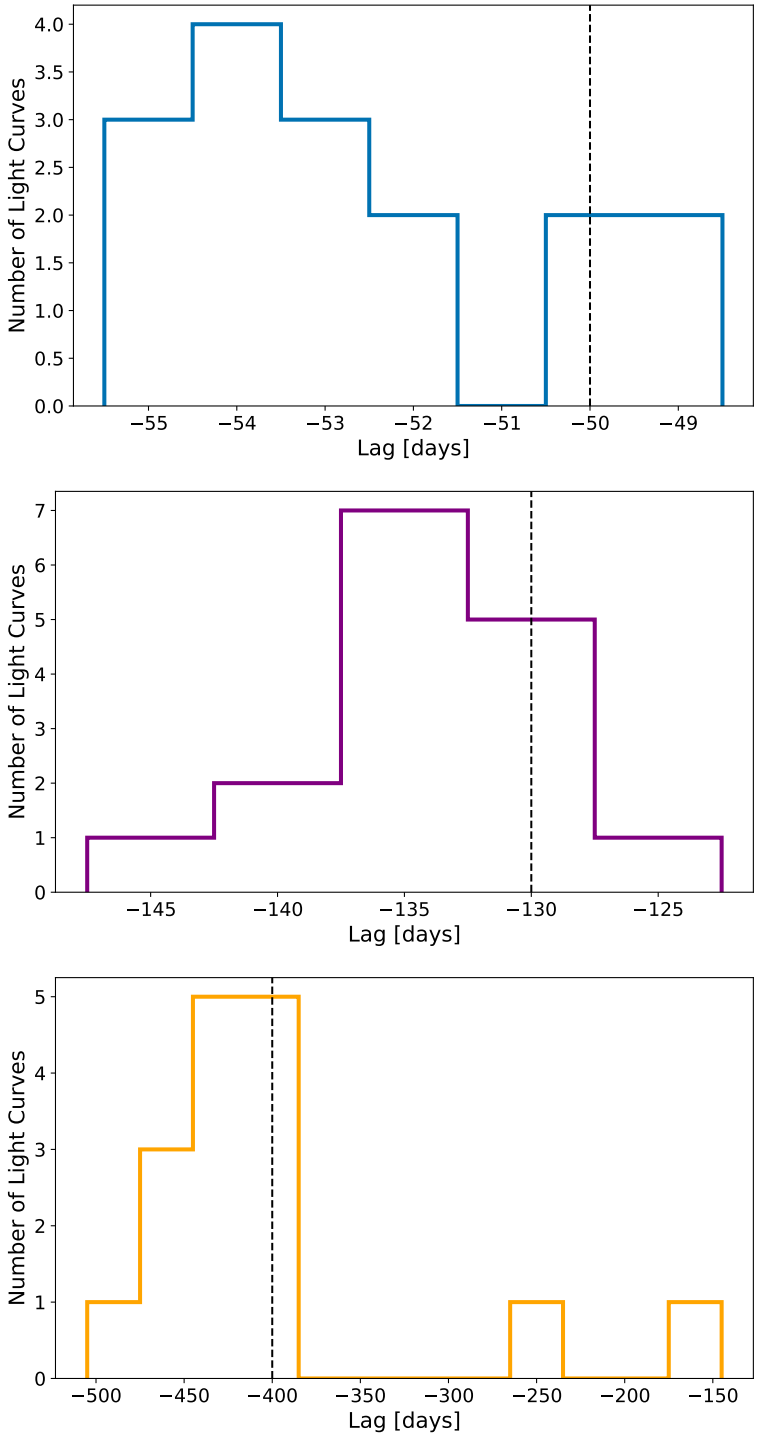


Figure 3.18: Distribution of the long lags recovered by JAVELIN for 16 mock long season cadence LSST light curves made using light curves from the simulation in Jiang & Blaes (2020) for input lags of −50 and 7 days (top panel), −130 and 9 days (middle panel), and −400 and 12 days (bottom panel). The long lags are shown as the dashed vertical line.

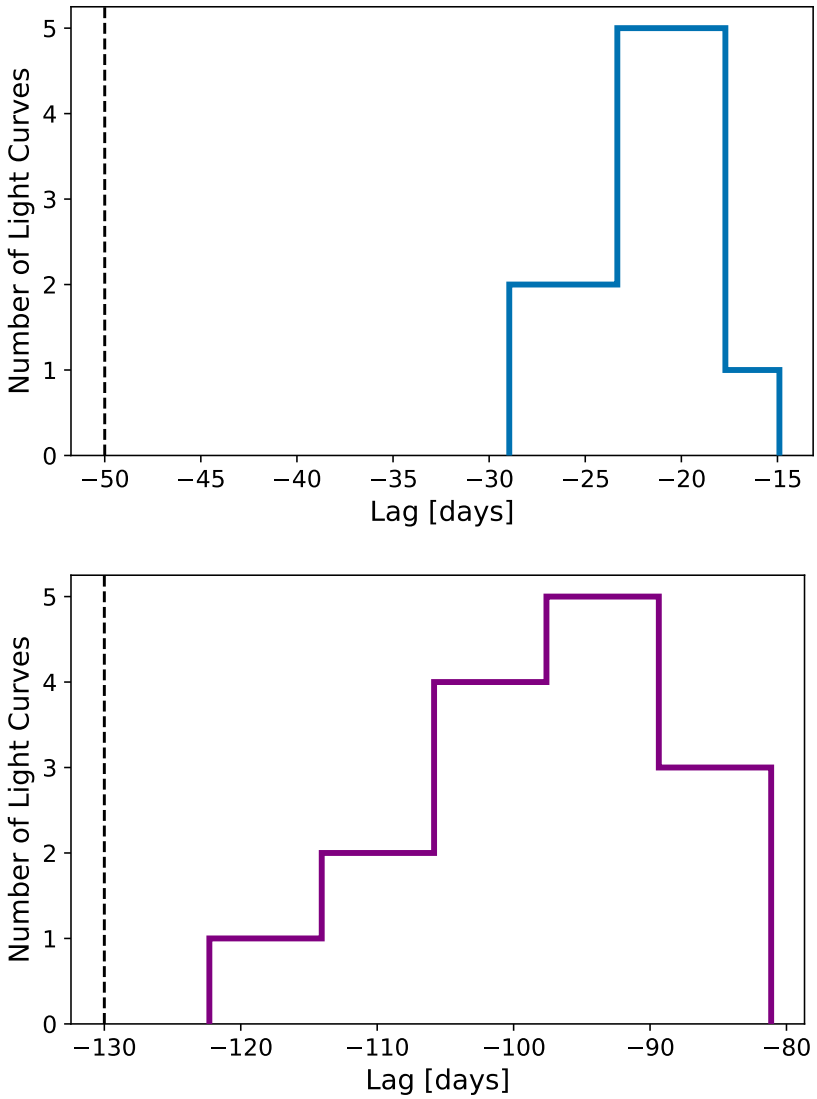


Figure 3.19: Distribution of the long lags detected by the maximum-likelihood method for 16 mock long season cadence LSST light curves made using light curves from the simulation in Jiang & Blaes (2020) for input lags of −50 and 7 days (top panel) and −130 and 9 days (bottom panel). The long lags are shown as the dashed vertical line.

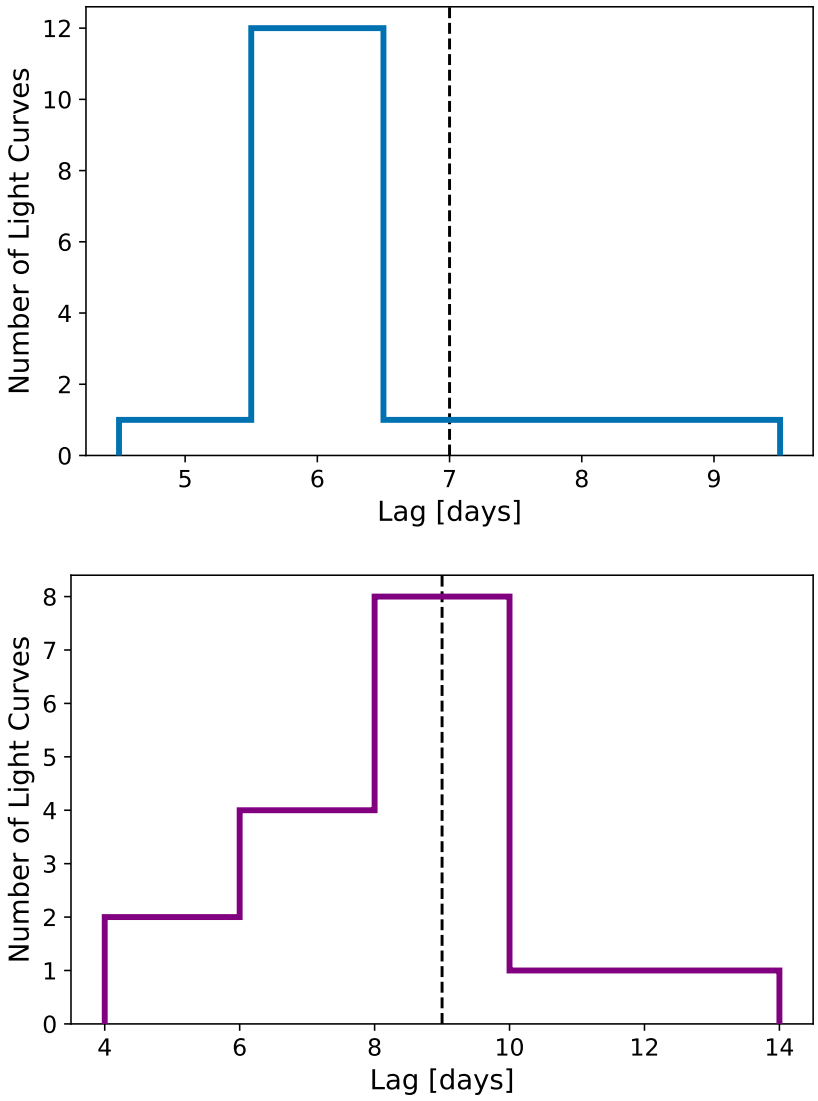


Figure 3.20: Distribution of the short lags detected by the maximum-likelihood method for 16 mock long season cadence LSST light curves made using light curves from the simulation in Jiang & Blaes (2020) for input lags of −50 and 7 days (top panel) and −130 and 9 days (bottom panel). The short lags are shown as the dashed vertical line.

light curves made using only the inner radius light curve (see Section 3.3.1.2), and may be from a portion of the light curve with lower variability, making it more difficult to detect the lag. Nonetheless, despite using a DRW model to interpolate light curves, JAVELIN still performs well on light curves that are not generated as a DRW.

We show the lags detected by the maximum-likelihood method for the 16 mock light curves with input long lags of -50 and -130 days in Figure 3.19. We again see that the distribution of long lags is skewed towards shorter timescales as was the case for light curves reprocessed with a long and short lag in Section 3.5.2. The shifts in these distributions are mildly larger. The maximum-likelihood method finds long lags of -20 days (40% of the input lag) versus -25 days (50% of the input lag) and -95 days (73% of the input lag) versus -98 days (75% of the input lag) for light curves with input long lags of -50 and -130, respectively.

As with DRW mock light curves where the ratio of the amplitude of the long lag to the amplitude of the short lag is 0.2, JAVELIN is unable to accurately detect the short lag. The maximum-likelihood method, however, is able to detect the short lag to within 50% accuracy for all light curves and to within a day for a majority of light curves.

We show the medians and standard deviations of the distributions of lags recovered by JAVELIN and the maximum-likelihood method for these non-DRW modeled light curves in row (8) of Table 3.5. Using light curves from the radiation MHD simulation in Jiang & Blaes (2020) instead of DRW light curves appears to have little affect on our ability to detect long and short lags. Because we derive similar success rates for more physically motivated light curves, we are hopeful that our results for light curves generated with a DRW model will hold for future observed LSST light curves.

3.5.5 Dimmer Magnitude Quasars

In the previous sections we test our lag recovery methods on light curves with $i < 19$ mag, which have SNR of around 150. To test if we can detect lags at lower SNR, we generate

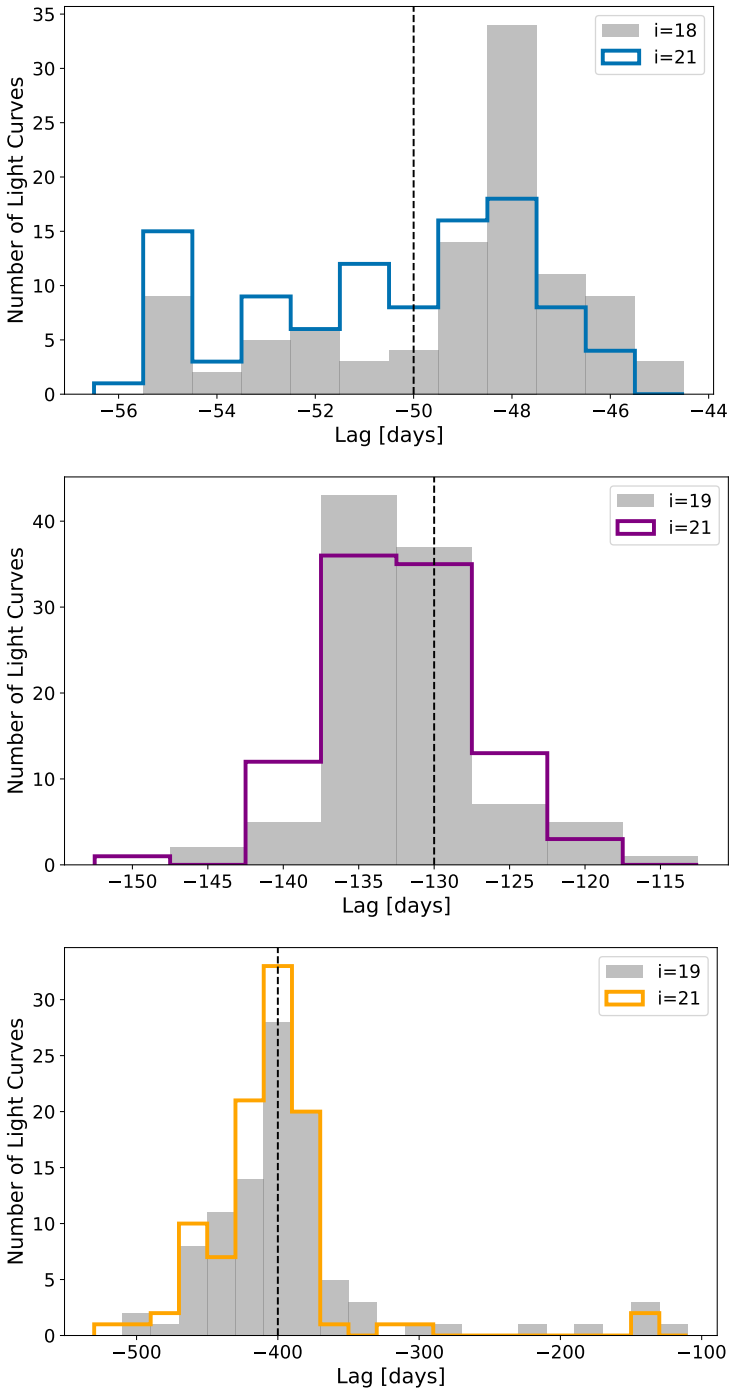


Figure 3.21: The distributions of the long lags recovered by JAVELIN for 100 mock light curves with the long season LSST cadence and input lags of -50 and 7 days (top panel), -130 and 9 days (middle panel), and -400 and 12 days (bottom panel). The vertical dashed lines show the input long lag. The colored lines are the distributions for light curves where we set $i = 21$ to explore how a lower SNR changes our ability to detect lags. The distributions of these lags are very similar to the distributions shown in gray which are for light curves with the i-band magnitudes in Table 3.1.

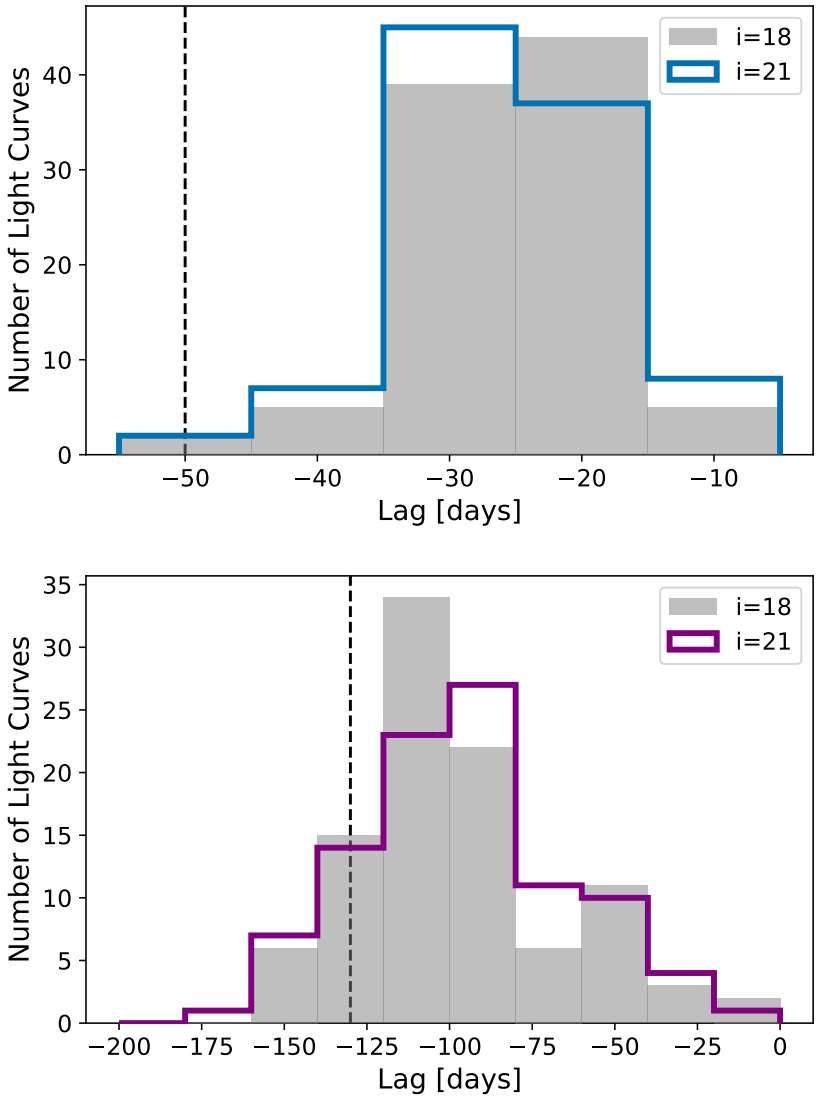


Figure 3.22: The distributions of the long lags recovered by the maximum-likelihood method for 100 mock light curves with the long season LSST cadence and input lags of −50 and 7 days (top panel) and −130 and 9 days (bottom panel). The vertical dashed lines show the input long lag. The colored lines are the distributions for light curves where we set $i = 21$ to explore how a lower SNR changes our ability to detect lags. The distributions of these lags are very similar to the distributions shown in gray which are for light curves with the i-band magnitudes in Table 3.1.

mock long season cadence LSST light curves with the same input long and short lags as in the previous sections and a ratio of 0.2 between the long and short lag amplitudes, but set the band magnitudes to $i = 21$ for all input lags. For LSST this band magnitude gives an SNR of around 30.

The colored lines in Figures 3.21 and 3.22 show the distributions of lags recovered by JAVELIN and the maximum-likelihood method, respectively, for these mock light curves. These distributions are very similar to the distributions, shown in gray, of lags recovered by both of these methods for light curves with the band magnitudes in Table 3.1. The median long lags recovered by JAVELIN are -50, -130, and -410 days for input long lags of -50, -130, and -400 days, respectively.

The median long lags recovered by the maximum-likelihood method for these light curves, -26 and -98 days for input long lags of -50 and -130 days, respectively, are also very similar to the median long lags recovered for light curves with brighter magnitudes. We show these median values, their standard deviations, and these values for the short lags in the bottom row of Table 3.5. Because the results for these dimmer magnitude quasars are nearly identical to the results in Section 3.5.2, we can expect to be able to detect long lags for quasars brighter than $i = 21$ mag, and suggest it might be possible to look for lags for even fainter quasars.

3.5.6 Summary of Results

We now provide a summary of Section 3.5. First, in Section 3.5.1 we find that JAVELIN and the maximum-likelihood method both detect long lags with high accuracy in mock light curves with the LSST long season cadence that only have an input long lag. When we also reprocess these mock light curves with a short lag in Section 3.5.2, JAVELIN is still able to accurately recover the long lag as long as the ratio of the amplitude of the long lag to the amplitude of the short lag is 0.2 or 0.1. The long lags recovered by the maximum-likelihood method get skewed to shorter and shorter durations as this amplitude ratio decreases, a

trend we also see in our tests of the Fourier method on even cadence light curves (see Figure 3.7). This trend could potentially help us to determine the amplitude ratio between the long and short lag in cases where we can make a robust long lag prediction from an alternative method. Both detection methods fail to recover the long lag for most light curves when the amplitude ratio is 0.01, although JAVELIN may in some cases be able to detect a long lag for this ratio when the input long lag is -130 days.

JAVELIN is unable to detect a short lag for an amplitude ratio of 0.2 for most light curves. For an amplitude ratio of 0.1, JAVELIN is only able to detect the short lag for a significant amount of light curves for light curves with input lags of -50 and 7 days. When the amplitude ratio is 0.01 JAVELIN is able to detect a short lag for a majority of light curves, although the short lags detected are often skewed towards longer durations. The maximum-likelihood method, on the other hand, is always able to detect the short lag, which is a real strength of the method.

In Section 3.5.3 we find that JAVELIN performs poorly on light curves with the LSST baseline v2.0 cadence, identifying the season gaps instead of the long or short lag. The maximum-likelihood method and the Von-Neumann method perform better on light curves with the baseline cadence, seeming to be only minorly affected by the change in cadence. We show that adding just 18 observations a year from LCO in the gaps in the baseline cadence fixes the cadence-issues for JAVELIN, resulting in highly accurate lag detections. Adding these additional LCO observations also improves the accuracy of the short lags recovered by the maximum-likelihood method.

In Section 3.5.4 we find that both JAVELIN and the maximum-likelihood method detect lags with roughly the same level of accuracy for the light curves from the radiation MHD simulation in Jiang & Blaes (2020) as for our DRW light curves. That our methods perform this consistently on more physically motivated light curves suggests that the results in this section should be applicable to real LSST light curves in the future. Finally, in Section 3.5.5 we find that when we generate dimmer mock light curves with a magnitude of $i = 21$, which

lowers their SNR, the distributions of lags detected by JAVELIN and the maximum-likelihood method are nearly identical to the distributions for the brighter magnitude mock light curves ($i < 19$) in Section 3.5.2.

3.6 Number of Detectable Long Lags

We now provide an estimate for the number of long lags we may expect to measure with LSST. To do this estimate we will use quasar properties from SDSS and metrics developed for the robustness of BLR and short lag detections.

First, we select only quasars at $z < 1$, because for our models based on Fairall 9 Figure 3.2 suggests that we should be able to detect the long lags for a majority of SDSS quasars at $z < 1$. Note that if $\beta = 0$ and $h/r = 0.01$, as is often assumed for the standard thin disk model, the long lag will be at least on the order of 100 years for bright SDSS quasars, and therefore not detectable using LSST. Next, we select only quasars with an i-mag brighter than 21. In Section 3.5.5 we show that our lag detection methods are accurate down to this magnitude.

Whether a long lag is detectable or not also depends on the variability of the quasar. When looking for BLR lags in SDSS quasar light curves, Grier et al. (2017) defined a criterion SNR2 based on the continuum and line light curve rms variability. The rms variability is the intrinsic variability of the light curve about a fitted linear trend, divided by the uncertainty of the estimated intrinsic variability.

To detect C IV lags, Grier et al. (2019) require an SNR2 > 20. They found that 348 out of 492 ($\sim 71\%$) quasars in their sample from the SDSS reverberation mapping project that fall within $1.35 < z < 4.32$ met that criterion. Homayouni et al. (2020) look for Mg II lags for 453 quasar light curves with redshifts more similar to our target redshifts here, $0.35 < z < 1.7$, and find that 198, or $\sim 44\%$, have an SNR2 > 20. The number of quasars with SNR2 > 20 is fewer in this case because Mg II tends to be less variable.

A second criterion for determining lag significance for both BLR and short continuum lags is setting a minimum threshold for $r_{\rm max}$, the maximum ICCF correlation coefficient (see Section 3.4.1). Setting a minimum threshold for $r_{\rm max}$ ensures that the light curves are well correlated. Guo et al. (2022) require $r_{\rm max} > 0.8$ for their sample of $z < 0.8$ quasar light curves from the Zwicky Transient Facility. Of 4255 quasars, 559 (13%) met that criterion.

SDSS DR14 has $\sim 5 \times 10^5$ quasars over 9376 deg^2 (Pâris et al., 2018). Limiting ourselves to $z < 1$ and $i < 21$ mag quasars we are left with $\sim 9.5 \times 10^4$ quasars. Accounting for the five planned 9.6 deg^2 DDF, we can expect to detect around 500 of these quasars within the DDFs. If we take the more conservative fraction of quasars with SNR2 > 20 from Homayouni et al. (2022) we find that around 200 quasars should have detectable lags. If we assume only 13% of light curves will have $r_{\rm max} > 0.8$ as in Guo et al. (2022) we would be able to detect long lags for 65 quasars within the DDFs.

Here, we only look at $z < 1$ quasars because the long lag gets very long at $z > 1$ and LSST has a ten year baseline. However, it is possible LSST will operate for 20 years. In addition, at higher redshifts we could potentially use a narrower range of wavebands (for example, g-i) to detect lags with lengths similar to those tested here for $z < 1$. Finally, we could also expand on the 10 year LSST baseline by using lower cadence archival data. Future work should examine if such a cadence enables lag detection. If we use the same i-band magnitude cut-off for SDSS DR14 quasars out to $z < 2$ we would roughly double the number of detectable long lags.

The Wide-field Infrared Survey Explorer and eROSITA provide other catalogues of quasars not included in SDSS. In particular these catalogues could be useful for studying lower luminosity, and thus lower mass, or redder, more dust-obscured quasars. Since $\tau_{\rm long} \propto M^{1/2}$, the lag timescales for these quasars will be shorter. Shorter lag timescales could help us detect a higher yield of longs lags for lower mass quasars, although this exact yield is difficult to estimate.

Even with our conservative estimates, we should be able to detect long lags for dozens to hundreds of new quasars over the ten year lifetime of LSST. In addition, Yao et al. (2023) were able to potentially detect a lag of $\sim$ 70 days using light curves with a less than 1 year baseline, so we speculate that we should be able to detect long lags in some fraction of nearby quasars with durations $\lesssim$ 50 days within the first 1–3 years of LSST. Roughly a third of SDSS quasars with $i < 21$ and $z < 1$ are within $z < 0.5$, where we expect lags to have durations of $\lesssim$ 50 days. Therefore it could be possible to detect long lags for roughly 20–100 quasars within the first couple of years of LSST.

3.7 Discussion and Summary

In this Section we discuss caveats related to our simulated light curves (Section 3.7.1) and lag detection methods (Section 3.7.2) that we plan to address in future work. We then provide a summary of the paper in Section 3.7.3.

3.7.1 Improved Simulated Light Curves

In this paper we create mock observations of our driving and reprocessed light curves with both the proposed LSST baseline v2.0 and longest season cadence from OpSim. While OpSim models SNR and observing gaps by taking into account sky brightness, bad weather predicted from localized weather models, seeing, telescope maintenance times, slew times, and filter changes, the cadences and SNR measurements may still be optimistic. However, small changes in cadence should not have a large impact on our ability to detect lags on the order of months. In addition, our tests in Section 3.5.5 using SNR corresponding to $i = 21$ mag light curves show that even if the SNR corresponding to the magnitudes in Table 3.1 is too high by an order of magnitude, we should still be able to recover lags with similar accuracy.

We use two types of light curves, self-generated DRW light curves and light curves from the simulation in Jiang & Blaes (2020). While often used to model quasar light curves, the DRW models are not physically motivated and may fail to model the quasar PSD properly at short and long timescales. In particular, PSD from highly sampled light curves from the Kepler mission appear to have a steeper slope at short timescales (Mushotzky et al., 2011; Kasliwal et al., 2015; Smith et al., 2018) and the length of the light curve may have an effect on the PSD at long timescales (Stone et al., 2022). On the other hand, the simulated light curves from Jiang & Blaes (2020) are not intended to capture quasar variability for a wide range of quasar parameters. Therefore, future work to simulate light curves for a wider range of quasar parameters is needed to better understand our ability to detect long lags in their variability.

We reprocess our light curves with a Gaussian response function. However, the shape of the response function for a long lag is unknown. For this paper we set the width of the Gaussian to $S \equiv C/5$ (see Equation (3.10)), but preliminary tests suggest that JAVELIN has more difficulty accurately recovering lags for response functions with larger widths ($S \gtrsim C/2$). In addition, using other functional forms for the response function, such as the log normal distribution, may make it more difficult to detect a long lag using JAVELIN. Therefore the width and shape of the response function are important parameters. Here, we use the default JAVELIN response function model, a top hat with no priors on the width. It is possible that adjusting the model or priors for the width of the response function could improve the accuracy of JAVELIN. Still, JAVELIN may not be the best tool if detecting long lags relies on finding lags in variability that is reduced or smeared as it moves inwards. The maximum-likelihood method is better designed to capture the shape of the response function and should therefore be able to handle broader or more complex response functions.

Our response function is centered on frequencies corresponding to the timescale of the long lag. However, if the perturbations that are accreted inward leading to the long lag are on the thermal time scale it may be that, unlike the short lag timescale, the frequencies most

affected by the long lag do not correspond exactly to the inverse of the long lag timescale. The thermal timescale is similar to the viscous timescale we use to determine the long lag timescale, for the large aspect ratios in our model. However, a future improvement would be to experiment with the recoverability of long lags depending on the frequencies most affected by the long lag.

Another simplification is that we model the radial dependence and value of the aspect ratio as independent of SMBH mass or Eddington ratio, when that is very likely not the case. Most likely, as the Eddington ratio of a quasar increases, its aspect ratio will increase, which will shorten its long lag timescale.

Ultimately, simulations of the reprocessing of high frequency light in the UV/optical emitting regions of turbulent quasar disks for a wide range of quasar parameters are needed to better model a more accurate response function for the long lag. To perform these simulations it will be necessary to use a multi-band radiation MHD code, such as the code developed by Jiang (2022) for ATHENA ++. Radiation MHD is necessary to capture the Maxwell stress, magnetic pressure, radiative viscosity, and radiation pressure, all of which may play a role in locally generated turbulence that can be accreted inwards on the viscous timescale (Jiang & Blaes, 2020). The multi-band capabilities will be useful because it will allow us to inject high-frequency radiation and compare it to the radiation reprocessed and emitted by the local disk.

3.7.2 Improved Lag Detection Methods

Here we use two main lag detection methods. The first, JAVELIN, very accurately detects long lags when the ratio between the amplitude of the long lag and the amplitude of the short lag is $\gtrsim 0.1$. On the other hand, it performs very poorly when this amplitude ratio is 0.01. JAVELIN has been shown previously to be biased in the presence of multiple lag signals (e.g., Cackett et al., 2018).

We can potentially improve the performance of JAVELIN by simultaneously fitting several reprocessed light curves, instead of just fitting one driving and one reprocessed light curve. Fitting several reprocessed light curves will be possible for LSST, which will have six different wavebands. With multiple bands, JAVELIN will be less likely to misidentify seasonal gaps as lags when the lags of the different wavebands are significantly different from each other (Zu et al., 2016), which would be especially helpful if the baseline v2.0 cadence is selected. However, in our model for the long lag, the lag timescales between longer wavelength bands becomes very short. As a result, fitting multiple wavebands simultaneously may not be as helpful as it is when fitting BLR lags, which differ significantly from band to band. Another potential improvement would be to use a custom model for the response function in JAVELIN that includes a positive and negative lag, instead of using the default model which is just a single top hat function.

The maximum-likelihood method is less accurate overall than JAVELIN at detecting the long lag. The long lags also become increasingly skewed towards zero as the amplitude of the short lag becomes stronger relative to the long lag. Figure 3.7 shows that this skewing towards zero is due to the shape of the combined long and short lag response function. However, because of this bias, it will be difficult to disentangle the exact value of the long lag versus the ratio of the long lag amplitude to the short lag amplitude. Utilizing the full frequency dependent information instead of summarizing the long lag with a single average value may help disentangle the two and reveal more details about the shape of the response function. More details could also be obtained by increasing the number of bins. In this paper, we only use four frequency bins for the Fourier method, because of the relatively low sampling rate of LSST, and because it is easier to detect a single lag with fewer bins.

There are several ways that using all of the frequency dependent lag information can provide additional information on the shape of the response function. For one, regardless of the ratios between the long and short lags, the wrapping frequency seems to consistently occur very near $1/2\tau_{\rm long}$ in Figure 3.7. Therefore, we may be able to more effectively exploit

the location of the wrapping frequency to derive the lag. In addition, as can be seen in Figure 3.7, for larger amplitude ratios even if the value of the long lag at lower frequencies is skewed towards zero, right before the lag wraps it drops to roughly the value of the input long lag. If, on the other hand, the lags increase towards zero or positive values as the lag approaches the wrapping frequency, as is the case for smaller amplitude ratios, that would provide information on the possible presence of a long lag with a small amplitude relative to the short lag.

Overall, there is room for improvement in the lag recovery techniques used here, especially if we want to detect long lags with amplitudes less than 10% of the amplitude of the short lag. We have mentioned several ways our employment of JAVELIN and the maximum-likelihood method could be improved. It is possible that alternative methods used for detecting lags in quasar light curves may be more accurate. The Von-Neumann method, which we also test on our mock light curves, seems like a particularly promising alternative, especially if the baseline cadence is chosen for LSST (see Section 3.5.3). Ultimately, the greatest improvement may be to create a new method, directly tailored to detecting the long lag.

3.7.3 Summary

Inspired by the recent potential detection of a long lag in Fairall 9, we have laid out a model for the timescale of negative long lags, which occur due to turbulence in the outer disk and propagate inward on the viscous timescale. Because the viscous timescale depends on the aspect ratio of the disk, measuring long lags can provide information on the vertical structure of the disk. Based on the measured long lags in Yao et al. (2023) for Fairall 9, we find that the scale height depends on the radial disk position as $h \propto r^{3.59}$. However, additional long lags need to be detected to properly constrain our initial model for the radial dependence of the scale height. We therefore hope to use LSST to detect these additional long lags.

To study whether LSST will be able to detect long lags, we simulate mock LSST light curves using both DRW modeled light curves and light curves from the simulation in Jiang

& Blaes (2020), which we then mock observe with several LSST cadences from OpSim. We reprocess the lagging light curve with a Gaussian response function. We then test several lag detection methods, including the ICCF (Gaskell & Peterson, 1987; White & Peterson, 1994; Peterson et al., 2004), a Von-Neumann estimator presented by Chelouche et al. (2017), JAVELIN (Zu et al., 2011), and a maximum-likelihood Fourier method (Zoghbi et al., 2013), on our mock light curves. We primarily focus on two main lag detection methods. The first is the commonly used JAVELIN code. We find that when only the long lag is included (Section 3.5.1), JAVELIN is accurate to within 10% of the input lag for all light curves with input long lags of -50 and -130 days, and accurate to within 10% of the input lag for 79% of light curves with input lags of -400 days in our DRW modeled light curves.

When we include a short lag as well (Section 3.5.2) JAVELIN remains very accurate, unless the ratio of the amplitude of the long lag to the amplitude of the short lag is around 0.01, at which point the accuracy decreases dramatically. Conversely, the short lag is not detected by JAVELIN unless this ratio ≤ 0.1 (for an input long lag of -50 days) or ≤ 0.01 (for input long lags of -130 and -400 days). JAVELIN detects long lags with a similar level of accuracy even when using a DRW model to interpolate the non-DRW modeled light curves from the simulation in Jiang & Blaes (2020) in Section 3.5.4.

Because JAVELIN is not aimed at detecting two lags simultaneously we also use a maximum-likelihood Fourier method, which separates lags in frequency space. Separating lags in frequency space should make it easier to detect multiple lags. However, as the amplitude of the long lag decreases relative to the amplitude of the short lag, the influence of the short lag on lower frequencies increases, skewing the recovered long lag towards shorter durations. When the amplitude ratio between the long and short lags is 0.01 the maximum-likelihood method is unable to identify the long lag at all for most light curves, just like JAVELIN. On the other hand, the maximum-likelihood method provides a significant advantage when it comes to detecting the short lag, which it is able to accurately detect for

all amplitude ratios. Finally, the performance of this method on the light curves from Jiang & Blaes (2020) is comparable to its performance on DRW light curves.

Based on the results in Section 3.5.3 and other cadence tests we performed, we believe a long season cadence[4] is helpful for detecting long lags. A long season cadence is particularly beneficial when using JAVELIN to detect long lags, whereas the Von-Neumann method and the maximum-likelihood method for shorter duration long lags (~ -50 days) are able to detect long lags with similar accuracy regardless of the cadence. However, adding just 18 LCO observations annually to mock light curves sub-sampled with the baseline v2.0 cadence, leads to a comparable or higher lag detection rate than for the long season cadence, for both our main methods for all lag durations. Additional LCO observations or the Von-Neumann method may also be effective tools for finding long lags in the Wide Fast Deep Survey, which will cover a much larger area of the sky at lower cadence.

Even with our conservative estimates in Section 3.6, we should be able to detect long lags for dozens to hundreds of new quasars over ten years. These detections will significantly improve our ability to model the long lag as a function of wavelength, SMBH mass, and Eddington ratio. These models will help us constrain the vertical structure of quasar disks for a range of quasar parameters.

3.8 Additional Materials: Full Tests

In Sections 3.5.2 and 3.5.3 we show the distributions of lags detected by our various methods only for light curves with input lags of -50 and 7 days, for the sake of brevity. Here we present the same results, but for all lag bins in Table 3.1.

[4]including strong rolling or weak accordion cadences which both have long seasons

3.8.1 Long and Short Lags

We show the distributions of the long lags detected by JAVELIN for 100 mock long season cadence LSST light curves with input long lags of −50 days (top panel), −130 days (middle panel), and −400 days (bottom panel) in Figure 3.23. Distributions shown in orange have only long lags, while distributions shown in blue have long and short lags with varying ratios between the amplitude of the long and short lag. The accuracy of the distributions for different input long lags is mostly consistent, with the exception of the distribution for light curves with input long lags of −130 days and an amplitude ratio of 0.01, which appears to be more accurate than the other distributions (see Section 3.5.2 for more discussion).

Figure 3.24 is the same as Figure 3.23 but for the maximum-likelihood method. The results for light curves with an input long lag of −130 days are mostly consistent with the results for an input long lag of −50 days. The maximum-likelihood does do slightly better at detecting some negative lags for an amplitude ratio of 0.01 for −130 day lags. This improvement is expected based on our tests of the Fourier method with even-cadence light curves shown in Figure 3.7. It appears that because the long lag is greater for light curves with −130 day input lags, it is not skewed all the way positive by the short lag.

We show the short lags detected by JAVELIN and the maximum-likelihood method in Figures 3.25 and 3.26, respectively, for different input lags and amplitude ratios. As discussed in Section 3.5.2, JAVELIN does a better job detecting the short lag for light curves with an input long lag of −50 days and an amplitude ratio of 0.1 than for light curves with other input lags. Otherwise, the accuracy of these distributions for different input lags is similar.

3.8.2 Improving the Baseline Cadence

We compare the accuracy of JAVELIN and the maximum-likelihood method for the long season LSST cadence and baseline v2.0 cadence in Figures 3.27 and 3.28, respectively. JAVELIN identifies the seasonal gaps of −180 days for light curves with the baseline cadence for all input lags. As mentioned in Section 3.5.3, the maximum-likelihood method does better

at still detecting the input lags for light curves with the baseline cadence than JAVELIN, although it does better on light curves with −50 day lags than −130 day lags, perhaps because −130 days is closer to the length of the season gaps.

In Figures 3.29 and 3.30, we compare light curves that have both the LSST baseline cadence and a roughly ten day cadence of LCO observations in the LSST season gaps (see Section 3.5.3 for details), to light curves with the LSST baseline cadence for various wavebands. We find that adding LCO observations in the LSST seasonal gaps is very useful for improving the accuracy of JAVELIN for all input long lags. The distributions of lags detected by the maximum-likelihood method for light curves with different cadences are consistent with each other for all input long lags. However, Figure 3.31 shows that the maximum-likelihood method does a better job detecting short lags with the added LCO cadence for both input lags.

As mentioned in Section 3.5.3, the Von-Neumann method performs better on light curves with the baseline cadence than JAVELIN. We show the distribution of lags detected by the Von-Neumann method for light curves with the long season (in orange) and baseline (in green) cadences in Figure 3.32. The Von-Neumann method detects long lags within 20 days of the input lag for a majority of light curves, except for light curves with an input long lag of −400 days, where the distribution peaks around −550 days. The median long lags are -40^{+8}_{-7} days and -40^{+1}_{-2} days for an input long lag of −50 days, -130^{+10}_{-0} days and -140^{+10}_{-20} days for an input long lag of −130 days, and -400^{+10}_{-10} days and -510^{+80}_{-50} days for an input long lag of −400 days, for the long season and baseline cadences, respectively.

3.9 Additional Materials: Null Tests

To verify that our results are due to the presence of lags we run our two lag detection methods on mock light curves with no input lag for each of the different cadences presented in this paper. To generate these mocks we perform all of the steps outlined in Section 3.3, but skip

Section 3.3.2, i.e. we do not reprocess either light curve with a response function. As a result, we have two identical light curves that are then sub-sampled with the cadence for two different wave bands (u and y or g and z, depending on the cadence). Below we discuss the signals recovered by JAVELIN (Section 3.9.1) and the maximum-likelihood method (Section 3.9.2). Any signals recovered could be due to the cadence/gaps in the cadence, or timescales inherent to the light curves themselves, such as the damping timescale.

3.9.1 Javelin

Figure 3.33 shows the distributions of the medians of the negative (left panel) and positive (right panel) values of the JAVELIN distributions for 100 mock DRW light curves with no re-processing. Light curves in the top panel have been mock observed with the long season LSST u- and y-band cadence used for the light curves in Sections 3.5.1, 3.5.2, 3.5.4, and 3.5.5. In the remaining panels light curves have been mock observed with the baseline u- and y-band cadence, the baseline g- and z-band cadence, and the baseline g- and z-band cadence with the additional LCO cadence, which are all used to sub-sample the light curves in Section 3.5.3. The distribution of the negative medians for light curves with the long season cadence peaks at zero. The duration of the long lag detected by JAVELIN for 70% of light curves is < 15 days, suggesting that the lags recovered by JAVELIN in the previous sections are due to the input lag signal. On the other hand, over 50% of positive lags recovered by JAVELIN are over 5 days, suggesting that our short lag results could be contaminated.

The individual JAVELIN distributions for light curves with the long season cadence that have not been reprocessed with an input lag, several of which are shown in Figure 3.34, peak most prominently at zero, but have noise. Most of this noise is either around $\pm 150 - 200$ days, because that is the length of the season gap, or around the edges of the range over which we run JAVELIN (± 600 days). When we take the median of only the positive or negative parts of the distributions, this noise skews lag values to > 0 days.

In Figure 3.34 each panel is labeled with the median of the positive and negative values of the distribution, most of which are very near zero, due to the sharp peak in the distribution around zero. There are several exceptions where the positive or negative median is ±132 days, due to noise in the distribution. These large median values only amount to 10–15% of the overall distribution in the top panel of Figure 3.33. However, these values are similar to our input lag of −130 days and so could contribute somewhat to our ability to detect −130 day lags with JAVELIN.

The mock light curves made from the simulation in Jiang & Blaes (2020) presented in Section 3.5.4 also have the long season cadence. The top panel of Figure 3.36 shows that for mock light curves without an input lag, JAVELIN only detects a long lag longer than 10 days and a short lag longer than one day for two light curves each.

On the other hand, we show in Section 3.5.3 that the baseline cadence leads to contamination from the ±180 day seasonal gap. For light curves sub-sampled with the baseline cadence that have not been reprocessed, the medians of the positive and negative values of the JAVELIN distributions peak around +180 and −180 days, respectively (see middle panels of Figure 3.33). When we add a mock LCO cadence of roughly every ten days to the gaps, however, the results are significantly better (see bottom panel of Figure 3.33). 90% of negative lags detected have a duration of less than 15 days. The same is true for $\sim 60\%$ of the short lags detected. Unsurprisingly, it appears for this cadence the JAVELIN results are very robust.

3.9.2 Fourier Method

Figure 3.35 shows the distribution of the lags recovered by the maximum-likelihood method for light curves sub-sampled with the same cadences used throughout the paper, but with no input lag. Roughly 90% of long and short lags detected by the maximum-likelihood method are shorter than 10 days for all cadences. For the long season cadence (top panel)

the maximum-likelihood method detects only 3 non-zero long lags and only seven short lags longer than 1 day.

For light curves with no reprocessing and the LSST+LCO cadence (bottom panel of Figure 3.35), the maximum-likelihood method detects lags shorter than 1 day for roughly 90% of light curves. The maximum-likelihood method performs less well on the two baseline cadences (middle two panels of Figure 3.35) without the added LCO observations, although 80% of long (short) lags detected are shorter than 10 (3) days. The lags detected by the maximum-likelihood method for mock light curves made from the light curves in the simulation in Jiang & Blaes (2020) with no input lag are all less than 1 day (see bottom panel of Figure 3.36). Overall, the lags recovered by the maximum-likelihood method seem unlikely to be contaminated by any of the cadences used in this paper or the light curve properties.

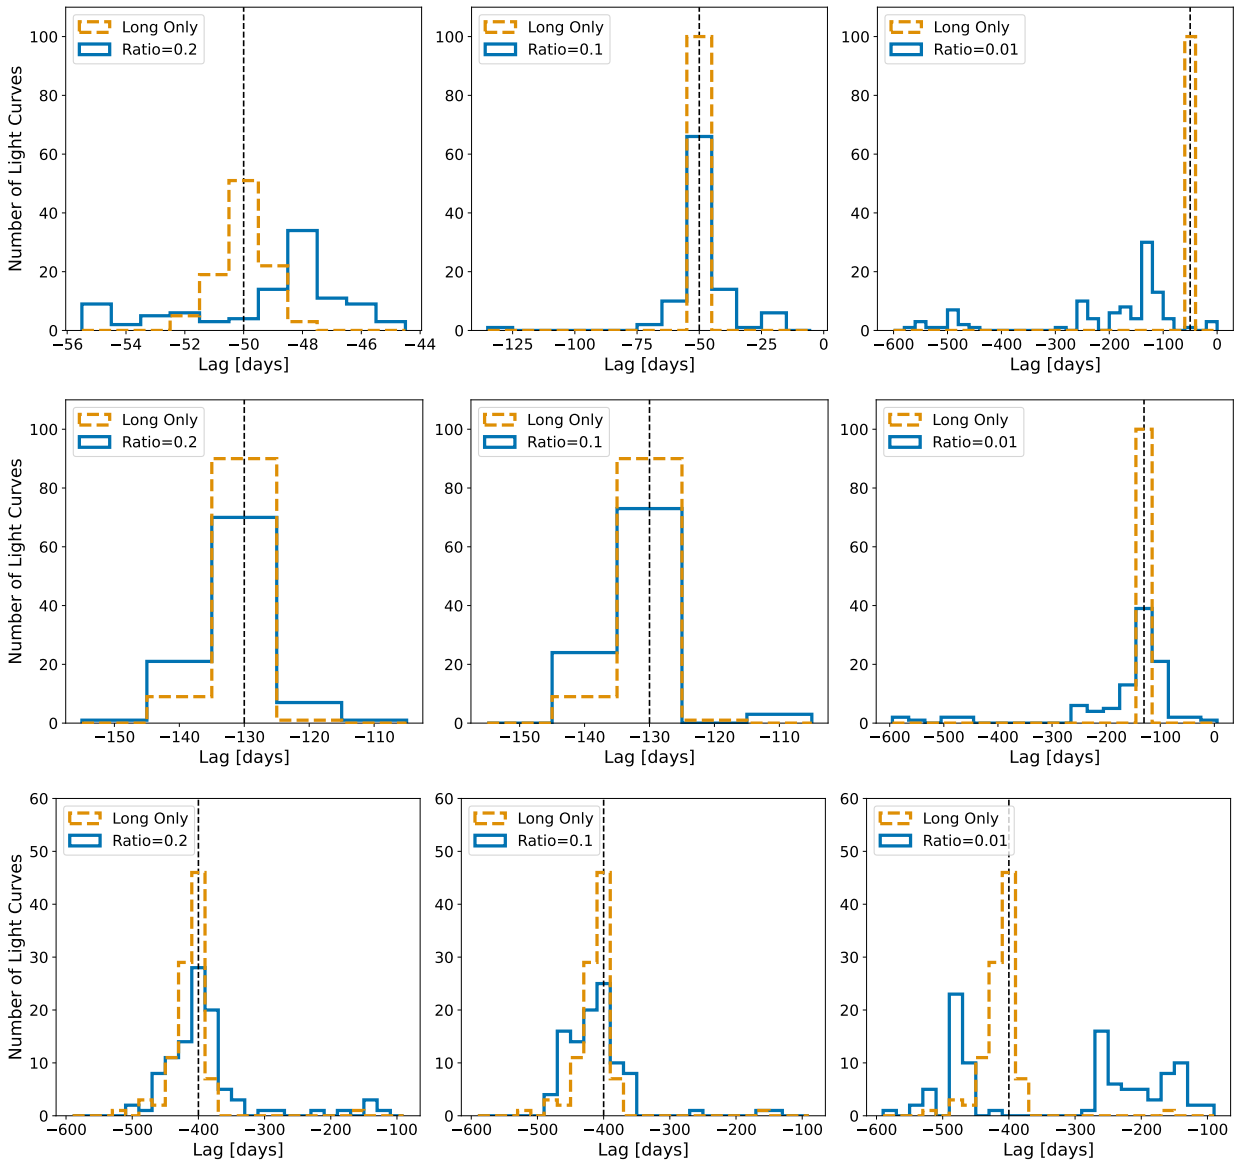


Figure 3.23: The distribution of the long lags detected by JAVELIN for 100 mock long season cadence LSST light curves with input long lags, shown by the dashed vertical lines, of −50 days (top panel), −130 days (middle panel), and −400 days (bottom panel). Distributions shown in blue also include a short lag of 7 days (top panel), 9 days (middle panel), and 14 days (bottom panel). The ratio of the amplitude of the long lag to the amplitude of the short lag is shown in the legends for each blue distribution.

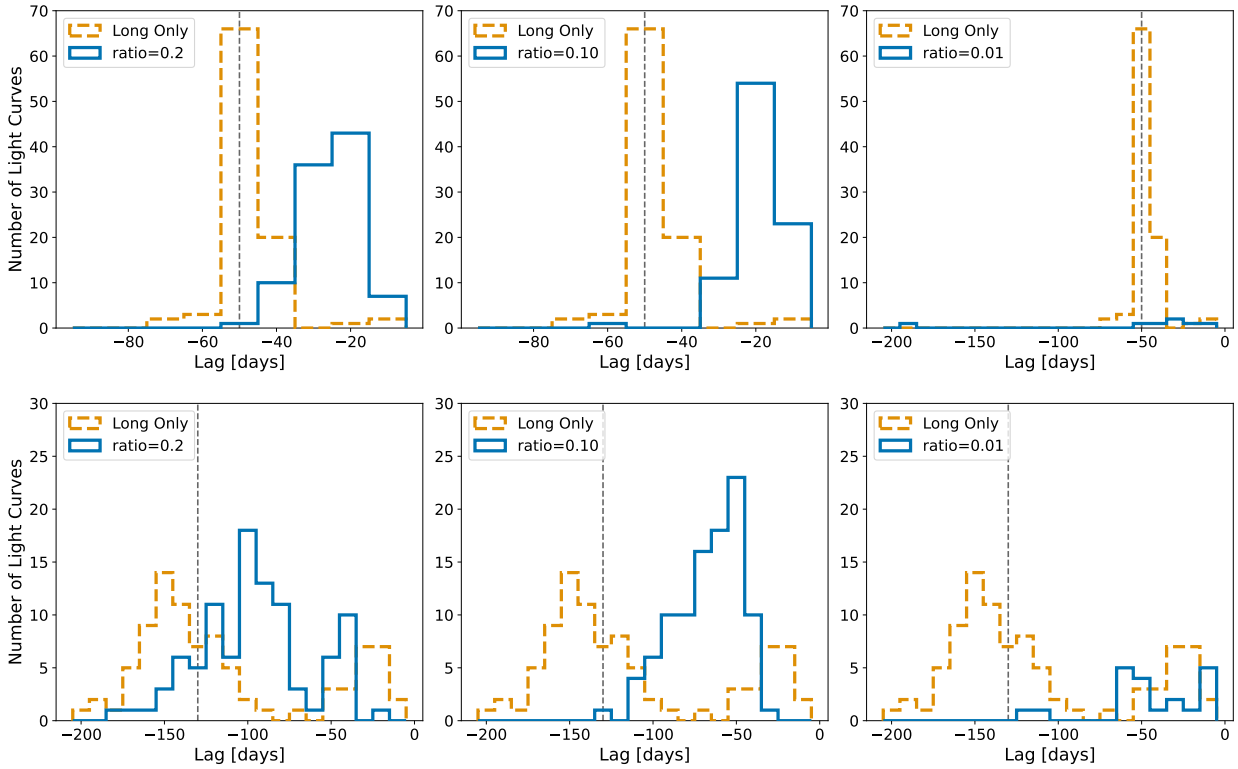


Figure 3.24: Distribution of the long lags recovered by the maximum-likelihood method for 100 mock long season cadence LSST light curves with input long lags, shown by the dashed vertical lines, of −50 days (top panel) and −130 days (bottom panel). Distributions shown in blue also include a short lag of 7 days (top panel) and 9 days (bottom panel). The ratio of the amplitude of the long lag to the amplitude of the short lag is shown in the legends for each blue distribution.

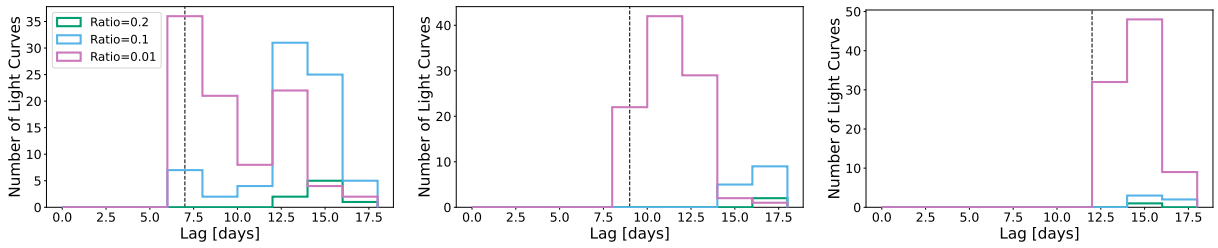


Figure 3.25: The distribution of the short lags detected by JAVELIN for 100 mock long season cadence LSST light curves with different amplitude ratios (shown in the legend). We zoom in on the portion of the distribution between 0 and 20 days to better show the accuracy of the short lag when detected. The input short lags, shown as the vertical dashed line, are 7, 9, and 12 days for the left, middle, and right panels, respectively.

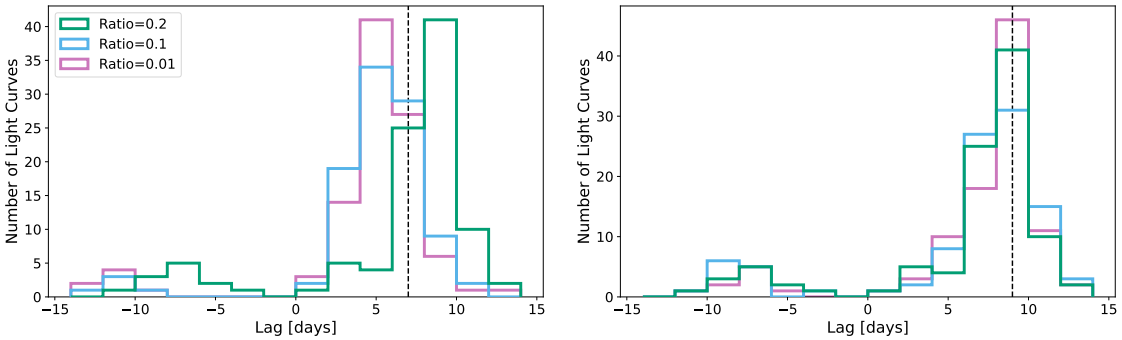


Figure 3.26: Distribution of the short lags recovered by the maximum-likelihood method for 100 mock long season cadence LSST light curves with different amplitude ratios (shown in the legend). The input short lags, shown as the vertical dashed line, are 7 days in the left panel and 9 days in the right panel.

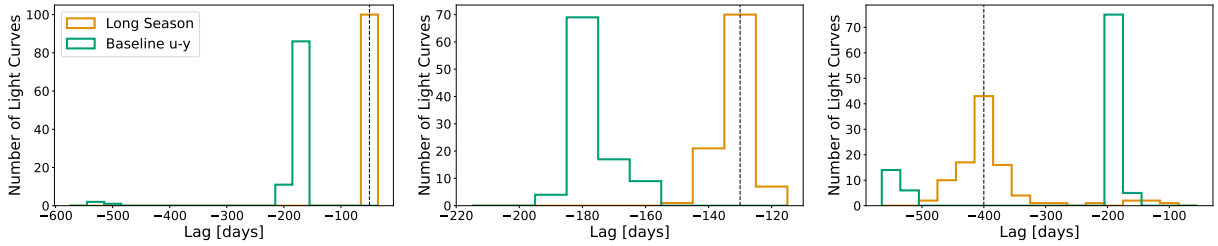


Figure 3.27: The distributions of the long lags recovered by JAVELIN for 100 light curves with input lags of -50 and 7 days (left panel), -130 and 9 days (middle panel), and -400 and 12 days (right panel). The orange and green distributions are for mock light curves with the long season and baseline v2.0 LSST cadences, respectively, for the u- and y-band. The dashed vertical line shows the input long lag.

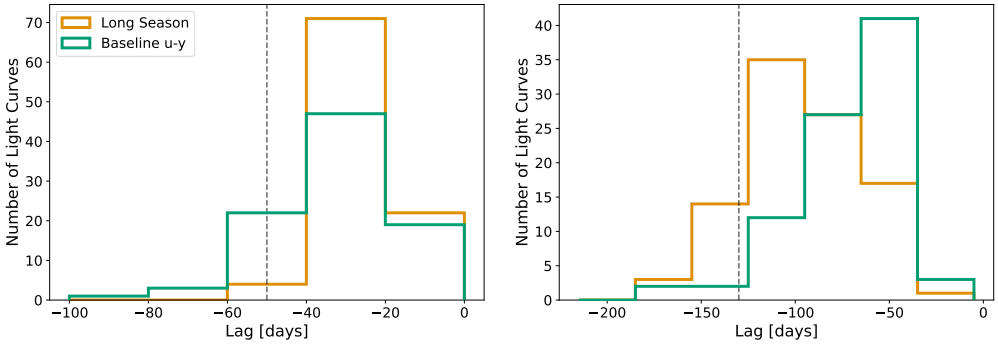


Figure 3.28: The distributions of the long lags recovered by the maximum-likelihood method for 100 light curves with input lags of -50 and 7 days (left panel) and -130 and 9 days (right panel). The orange and green distributions are for mock light curves with the long season and baseline v2.0 LSST cadences, respectively, for the u- and y-band. The dashed vertical line shows the input long lag.

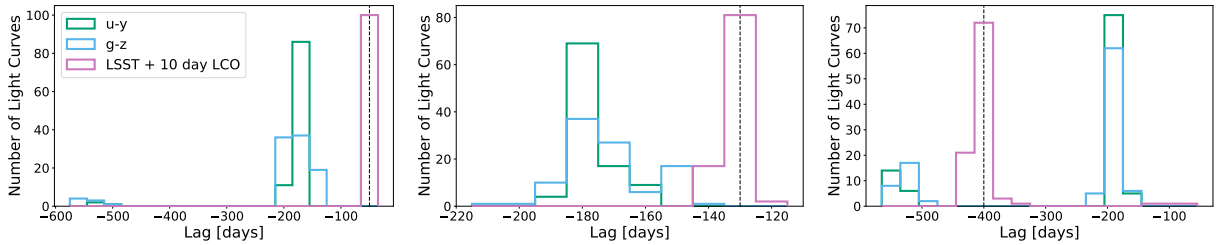


Figure 3.29: Distribution of the long lags recovered by JAVELIN for 100 mock light curves with input lags of −50 and 7 days (left panel), −130 and 9 days (middle panel), and −400 and 12 days (right panel). The dashed vertical line shows the input long lag. The green and light blue distributions are for mock light curves with the baseline v2.0 LSST cadence for the u- and y-band and g- and z-band, respectively. In addition to the baseline LSST cadence, the distributions shown in pink have an LCO observation taken roughly every ten days in each LSST observing gap. These LCO observations amount to around 18 observations in each gap, or per year, for a total of around 180 observations.

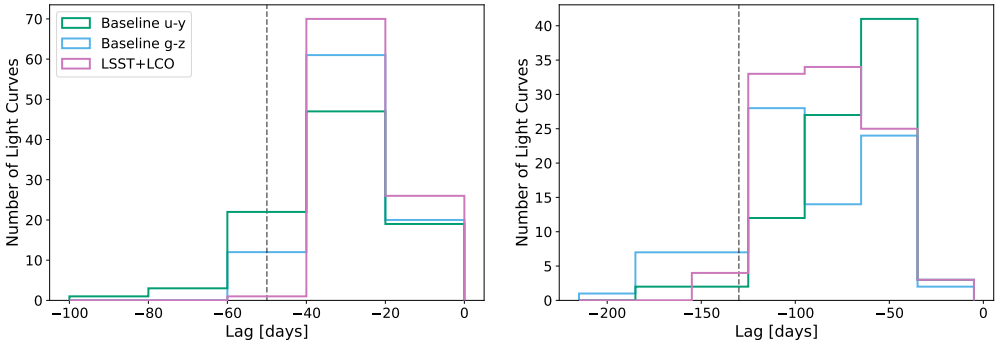


Figure 3.30: Distribution of the long lag recovered by the maximum-likelihood method for 100 mock light curves with input lags of −50 and 7 days (left panel) and −130 and 9 days (right panel). The dashed vertical line shows the input long lag. The green and light blue distributions are for mock light curves with the baseline v2.0 LSST cadence for the u- and y-band and g- and z-band, respectively. In addition to an LSST cadence, the distributions shown in pink have an additional LCO observation taken roughly every ten days in each LSST observing gap.

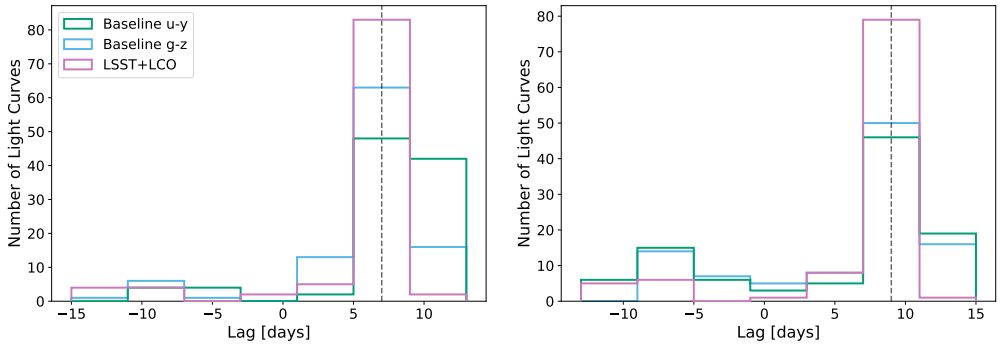


Figure 3.31: Distribution of the short lag recovered by the maximum-likelihood method for 100 mock light curves with input lags of −50 and 7 days (left panel) and −130 and 9 days (right panel). The dashed vertical line shows the input short lag. The green and light blue distributions are for mock light curves with the baseline v2.0 LSST cadence for the u- and y-band and g- and z-band, respectively. In addition to an LSST cadence, the distributions shown in pink have an additional LCO observation taken roughly every ten days in each LSST observing gap.

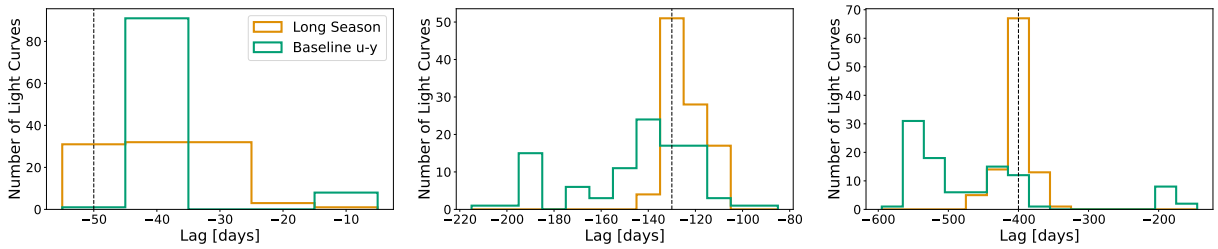


Figure 3.32: Distribution of the long lag recovered by the Von-Neumann method for 100 mock light curves with input lags of −50 and 7 days (left panel), −130 and 9 days (middle panel), and −400 and 12 days (right panel). The dashed vertical line shows the input long lag. The orange (green) distributions are for mock light curves with the long season (baseline) u- and y-band cadence.

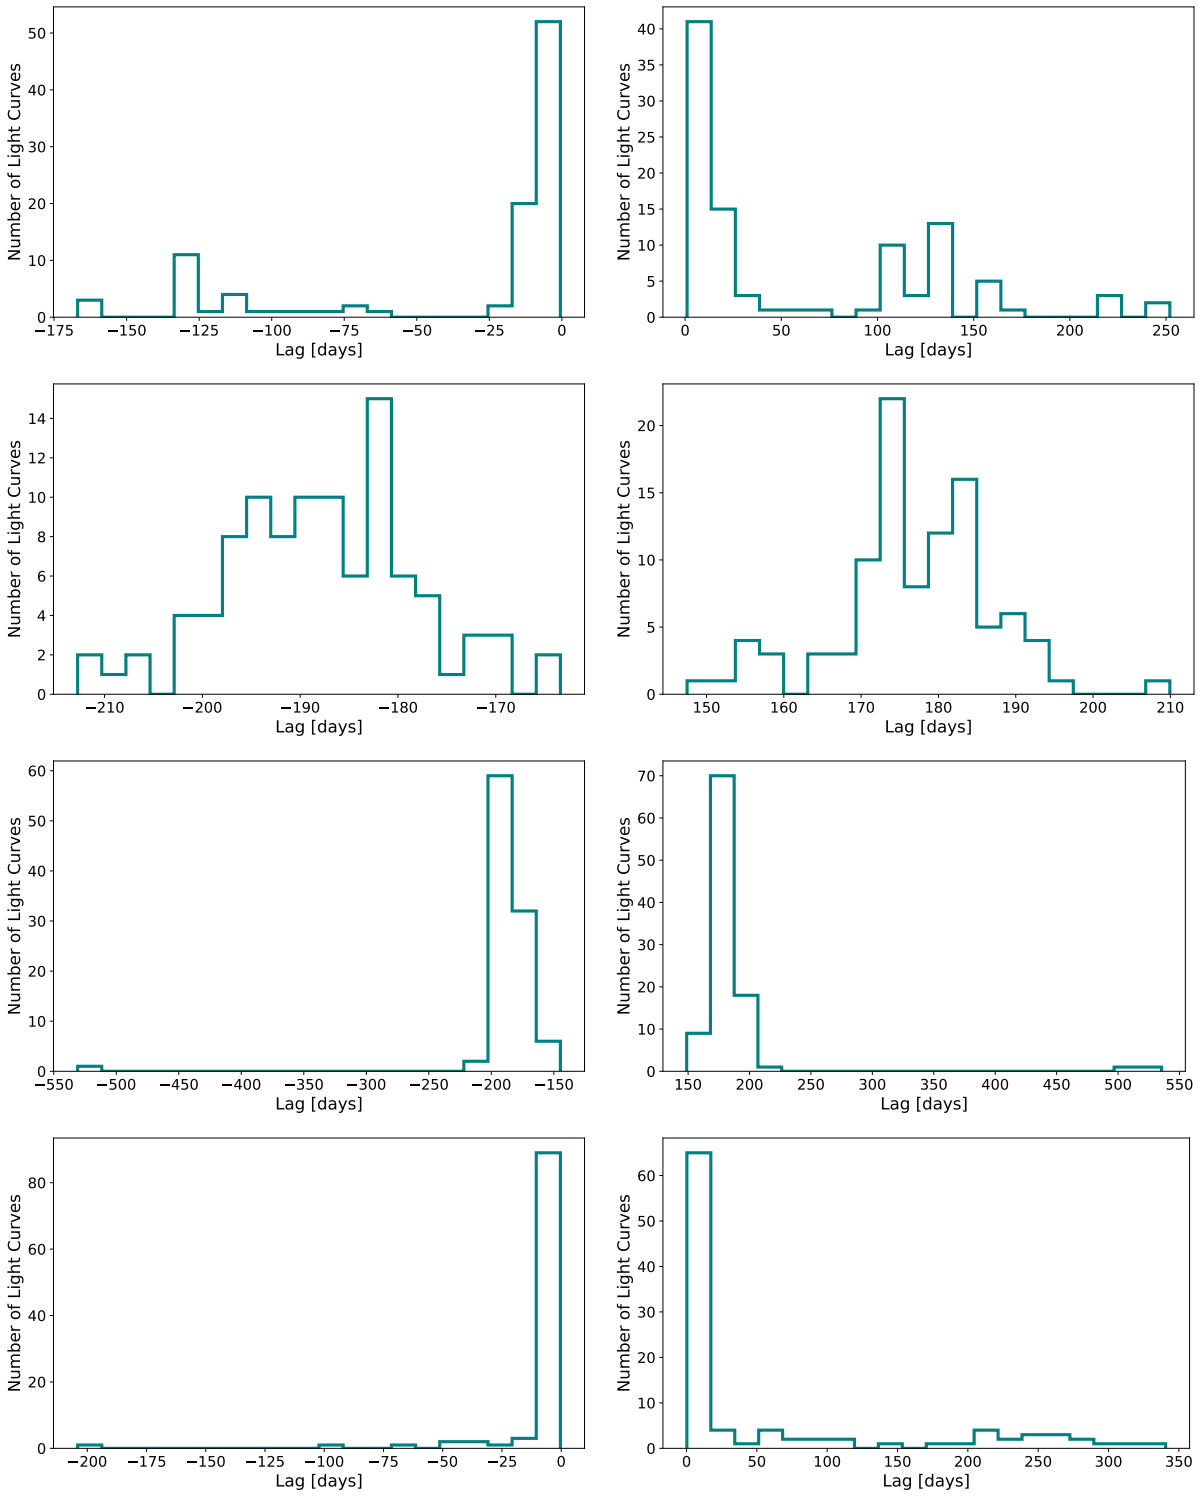


Figure 3.33: The distribution of the negative (left panels) and positive (right panels) medians of JAVELIN distributions for 100 mock light curves with no input lag. From top to bottom, the light curves have been sub-sampled with the long season LSST u- and y-band cadence, the baseline u- and y-band cadence, the baseline g- and z-band cadence, and the baseline g- and z-band cadence + LCO cadence.

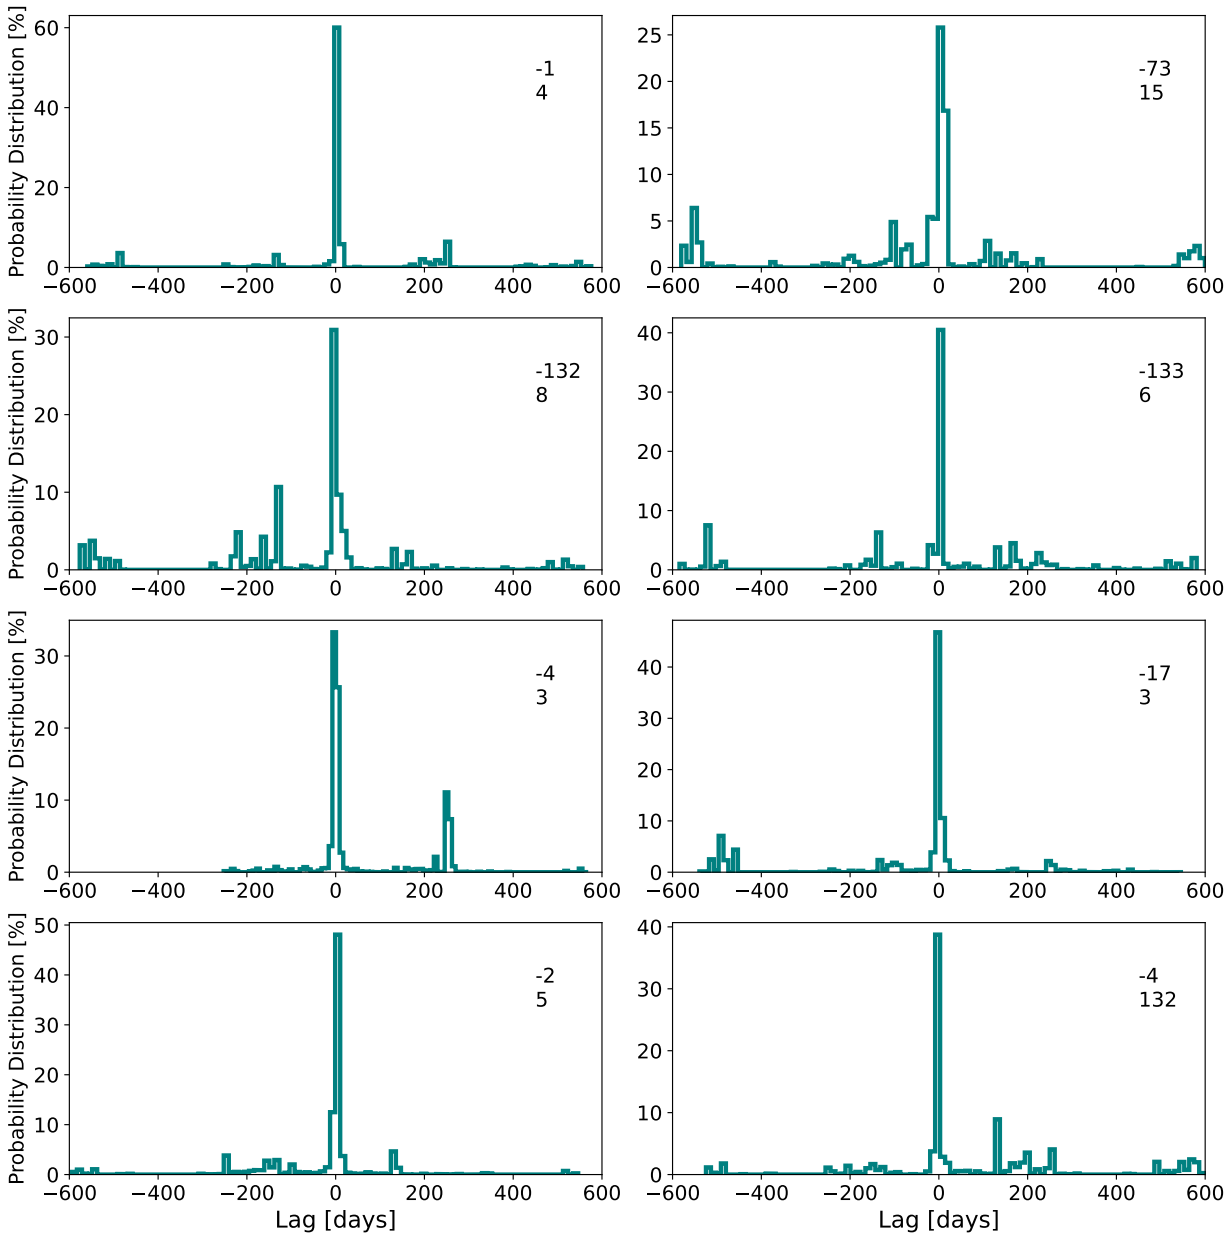


Figure 3.34: Eight individual JAVELIN distributions for light curves with no reprocessing and the long season LSST cadence. Each panel is labeled with the median of the positive and negative values of the distribution.

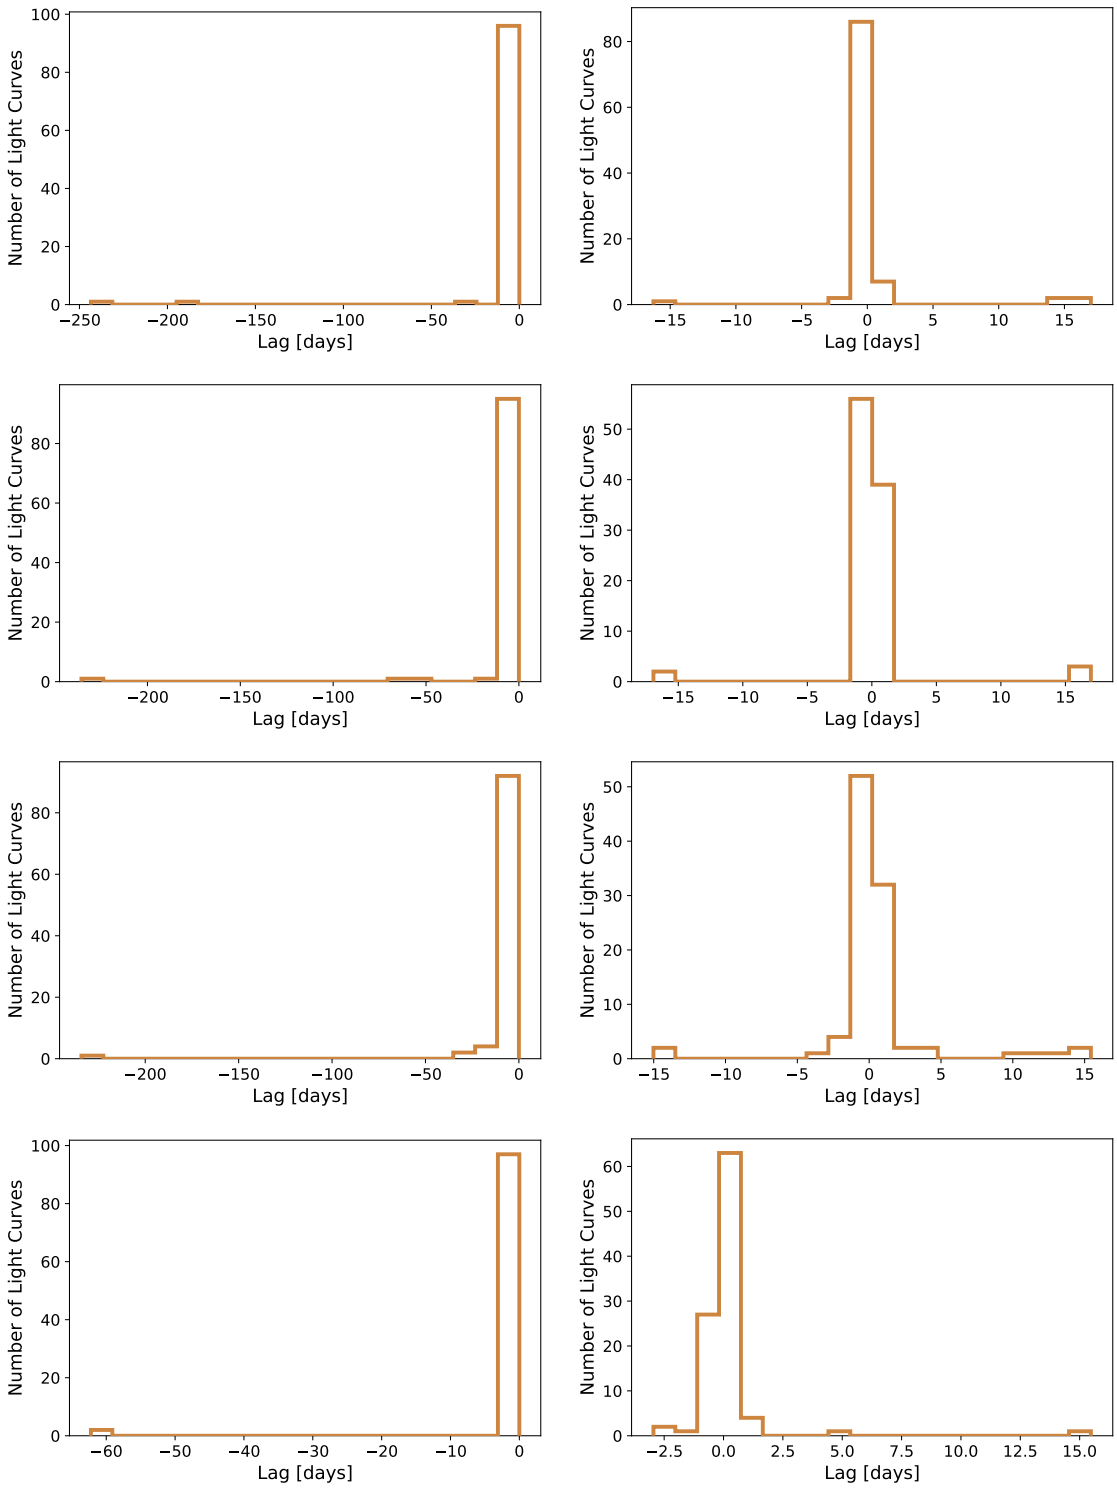


Figure 3.35: The distribution of long lags (left panel) and short lags (right panel) recovered by the maximum-likelihood method for light curves that have not been reprocessed with a lag for all cadences in the paper. From top to bottom, the light curves have been sub-sampled with the long season LSST u- and y-band cadence, the baseline u- and y-band cadence, the baseline g- and z-band cadence, and the baseline g- and z-band cadence + LCO cadence.

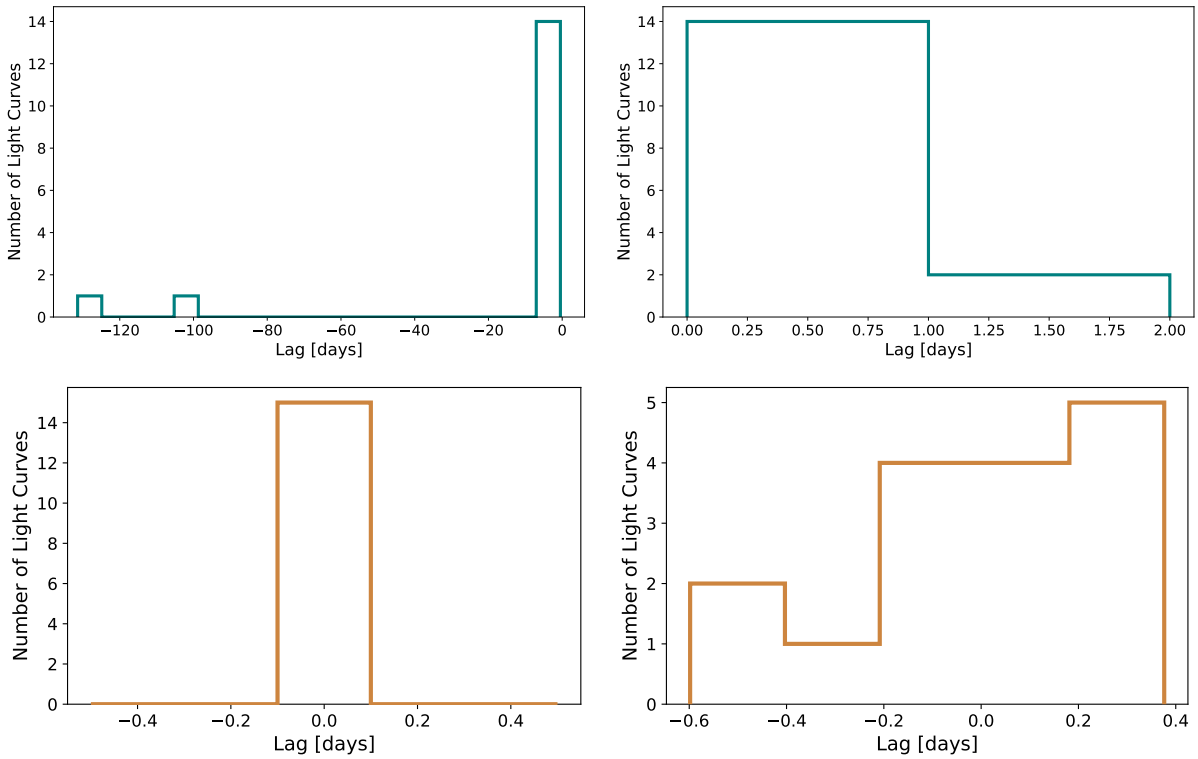


Figure 3.36: The top panel shows the lags recovered by JAVELIN and the bottom panel shows the lags recovered by the maximum-likelihood method for 16 mock light curves made from the simulation in Jiang & Blaes (2020) with no input lag and the long season LSST u- and y-band cadence.

Chapter 4

Simulating X-ray Reverberation in the UV-Emitting Regions of Active Galactic Nuclei Accretion Disks with 3D Multi-Frequency Radiation Magnetohydrodynamic Simulations

4.1 Introduction

Although active galactic nuclei (AGN) are among the most luminous objects in the Universe, due to their great distance it is generally not possible to resolve the microarcsecond scales of their accretion disks, broad line regions (BLRs), or event horizons. However, because AGN light curves vary over a wide range of frequencies, temporal resolution can be used in place of spatial resolution. A primary technique used to study AGN, called reverberation mapping, makes use of this temporal resolution (Peterson, 2014).

Reverberation mapping was first proposed by Blandford & McKee (1982) as a way to measure distances to the BLR by finding a time lag between variability in continuum light curves and variability in broad emission lines (see also Kaspi et al., 2000; Peterson et al., 2004; Bentz & Katz, 2015; Grier et al., 2017). Disk continuum reverberation mapping can similarly be used to measure the radial extent of the accretion disk itself, because X-rays emitted by the corona move outwards along the disk and are absorbed and re-emitted by different temperature regions of the disk at different wavelengths based on the local temperature (e.g., Sergeev et al., 2005; Cackett et al., 2007, 2018, 2021; De Rosa et al., 2015; Edelson et al., 2015, 2017, 2019; Jiang et al., 2017; Fausnaugh et al., 2016; Starkey et al., 2017; Homayouni et al., 2022). Therefore, there will be a lag on the light-crossing timescale between variability in light curves from short wavelength bands and light curves that have been reprocessed and re-emitted in longer wavelength bands. This lag is used to determine the distance between different temperature regions of the accretion disk.

Traditional models for reverberation mapping make the assumption that the variability in UV and optical AGN light curves is driven by the variability of the X-ray irradiation. However, there is significant evidence for additional variability in UV-optical light curves that does not appear to come from X-ray light curves (e.g., Uttley et al., 2003; Arévalo et al., 2009; Edelson et al., 2019; Yu et al., 2022; Hagen & Done, 2023; Cackett et al., 2023; Kara et al., 2023). The lack of correlation between X-ray and UV-optical light curves could be due to absorption by intervening material (e.g., Kara et al., 2021) or changes in the height and temperature of the corona on timescales shorter than the observing timescale (Panagiotou et al., 2022a). Another alternative is that this variability could come from fluctuations in the UV-optical regions of the accretion disk due to magnetorotational instability (MRI) driven turbulence (Balbus & Hawley, 1991) or convection resulting from the enhanced opacity in the UV-optical region of the disk (Jiang & Blaes, 2020).

While there has been considerable observational work done using reverberation mapping, theoretical work to model the underlying physics of the reprocessing of hard irradiation

into the softer radiation emitted by the disk has lagged behind. Most of the previous work modeling this reprocessing is (semi-)analytic and makes assumptions about the vertical profiles of various disk parameters, which often remain fixed over time or are not-calculated self-consistently to include the interaction of radiation and gas (e.g., Sun et al., 2020; Kammoun et al., 2021b; Salvesen, 2022; Panagiotou et al., 2022b). The recent development of accurate, high-speed, numerical methods to calculate radiative transfer coupled to magnetohydrodynamics (MHD) make it possible for the first time to examine how both radiation and magnetic fields have an impact on AGN disk structure and turbulence (Jiang et al., 2013; Jiang & Blaes, 2020; Jiang, 2021). Adding realistic Rosseland and Planck mean opacities that include line opacities instead of only electron scattering and free-free absorption also has an impact on driving convection in AGN disk models (Jiang & Blaes, 2020). However, all of these MHD simulations model radiation using a single frequency group, with the radiation integrated over all frequencies, and do not examine the impact of X-ray irradiation on the disk and the light curves emitted by the disk.

In this paper we provide an early application of a multi-frequency radiation MHD code, which instead of integrating radiation over all frequencies, integrates radiation over multiple frequency groups (Jiang, 2022). This multi-frequency capability allows us to simulate from first principles the reprocessing of high-frequency light incident on an accretion disk into light emitted in the UV based on the local disk temperature to better understand the drivers of variability in UV-optical AGN light curves. By first principles, we mean that the radiation energy and momentum are coupled to the gas evolution and that the temperature, opacity, and energy of the simulation depend self-consistently on the local gas density and pressure. In Section 4.2 we describe the methods we use for our simulations. In Section 4.3 we present, compare, and discuss the simulated disks and light curves from our simulations with and without X-ray irradiation. Finally, in Section 4.4 we summarize our results and present ideas for future work.

4.2 Methods

Using ATHENA++ we solve the ideal MHD equations coupled with time-dependent, implicit radiative transfer equations for intensities over discrete angles in the manner described in Jiang (2022), which builds on Jiang (2021) by allowing for multiple frequency groups. We use two frequency groups, one group covering the frequency range $[0, 1\mathrm{keV}]$ to model the radiation field emitted by the disk locally (at UV wavelengths) and the other group covering the frequency range $(1\mathrm{keV}, \infty)$ for the X-ray irradiation. Notice that we turn off the frequency shift due to Doppler effect as its effect is small and can cause numerical diffusion in the frequency space for the two groups we adopt.

We model a local patch of the accretion disk around a supermassive black hole with mass $M_{\rm BH} = 5 \times 10^7\ M_\odot$ under the local shearing box approximation with the box centered at $r_0 = 45 R_{\rm s}$, where $R_{\rm s}$ is the Schwarzschild radius. The corresponding Keplerian orbital frequency, $\Omega = 4.73 \times 10^{-6}\ \mathrm{s}^{-1}$ and shear parameter, $q = -d\ln\Omega/d\ln r = 3/2$. In the shearing box approximation the x- and y-coordinates correspond to the radial and azimuthal dimensions, while the z-coordinate corresponds to the height. We choose our simulation domain to be $(x, y, z) \in (-1.5H_g, 1.5H_g) \times (-6H_g, 6H_g) \times (-48H_g, 48H_g)$ with resolution $N_x = 48$, $N_y = 192$ and $N_z = 1536$. Here $H_{\rm g} = 1.11 \times 10^{12}$ cm is the gas pressure scale height of the disk, while the total pressure scale height (including radiation pressure) will be $H_{\rm r} = 4.06\ H_{\rm g} = 4.51 \times 10^{12}$ cm. We use the periodic and shear periodic ATHENA++ boundary conditions for the x- and y-boundaries, respectively. For the z-boundaries, we use outflow boundary conditions for the gas. Intensities for the low frequency group are allowed to leave the simulation domain but they are not allowed to enter. For the high frequency group, we set the incoming intensities based on a prescribed X-ray light curve. Outgoing intensities are copied from the last active zone to the ghost zone.

We first perform a one-frequency run to model only the locally emitted UV photons without X-ray irradiation. Then we restart the simulation with two frequency groups to model the locally generated UV photons and X-ray irradiation simultaneously. We use $N =$

48 discrete angles in each cell for each frequency group to resolve the angular distribution of the radiation field. Inside the simulation box we initialize the UV radiation isotropically based on the local disk temperature and the X-ray radiation isotropically based on a constant radiation energy, $E_{\rm r,1} = 10^{-4} E_{\rm r,0}^{\rm mid}$, where $E_{\rm r,0}^{\rm mid}$ is the midplane radiation energy for the UV frequency group.

We inject two different X-ray light curves for two otherwise initially identical two-frequency runs (Run A and Run B). The time-dependence of the intensity of this radiation is determined by a mock X-ray AGN light curve that we model as a damped random walk (DRW) as in Kelly et al. (2009), because the DRW model has been shown to be a decent representation of AGN light curves for timescales of days to years (e.g. Kelly et al., 2009; Kozłowski et al., 2010; MacLeod et al., 2012; Zu et al., 2013). For Run A, we use a damping timescale $\tau_{\rm damp} = 30$ days and a time-sampling of $dt = 1$ day, for Run B, we use $\tau_{\rm damp} = 0.1$ days and a time-sampling of $dt = 0.001$ days. For Run A, we choose a shorter damping timescale than the thermal timescale of our simulation box and a low time-sampling to avoid introducing very high frequency variability. For Run B, we choose a much shorter damping timescale that is closer to the dynamical timescale of the innermost disk where the X-rays originate, and a high time sampling rate to capture this short timescale variability. In summary, Run A will have more variability at low frequencies than Run B, while at high frequencies Run B will have more variability. The goal is to test and compare two very different injected light curves.

We found that if we normalized the intensity of the injected X-rays to the time-averaged total UV flux leaving the top and bottom of the simulation box in the one-frequency run, the X-rays do not have a large impact on our simulation, at least over a few thermal timescales. Since we are interested in X-ray reprocessing, we instead normalize the X-ray irradiation to eight times the average UV emission, which is the value we found sufficient to produce changes in the UV light curve. In the future, a parameter study is needed to determine how this normalization impacts the reprocessing. Finally, we average the injected flux over all

angles in our simulation and inject this averaged intensity along all angles pointing into the box. Note that because X-rays are injected directly into the simulation there will be no lag between the X-ray and UV light curves from the light travel time.

We determine the timestep-dependent opacity for the UV component in each cell with a bi-linear interpolation of the OPAL opacity tables from Iglesias & Rogers (1996) based on each cell's current gas temperature and density. For our X-ray component, we take the proper average of the free-free absorption opacity and electron scattering opacity, which ends up to be dominated by the scattering value.

We set the initial vertical profile of our simulations using the equations of radiation hydrostatic equilibrium, as in Jiang et al. (2016). Initial and normalization parameters are given in Table 4.1. For our initial conditions the ratio of the radiation to the gas pressure is $P_{\rm rat} = 4.4$. We initialize the magnetic field in our simulation as in Jiang et al. (2013) as two oppositely twisted flux tubes with zero net-flux and set the ratio of the gas pressure to the magnetic pressure to $\beta = 200$ in the midplane. We also add a purely vertical magnetic field with $\beta = 2 \times 10^4$ to increase the Maxwell stress in order to achieve a temporally averaged $\dot{M}/\dot{M}_{\rm edd} \approx 2 \times 10^{-2}$ once the simulation reaches steady state.

4.3 Results

4.3.1 Disk Structure

We show the horizontally-averaged vertical profiles as a function of time for the gas temperature, gas density, and radiation energy for the one-frequency run without X-ray irradiation (to the left of the grey line) and two-frequency Run A with X-ray irradiation (to the right of the grey line) in the right panels of Figures 4.1 and 4.2. In Figure 4.1 the red arrows represent the time varying X-ray irradiation we inject into the top and bottom of our simulation box in our two-frequency runs. The solid white lines show the location of the UV photosphere, which is very similar to the location of the X-ray photosphere, because both optical depths are

Parameter	Unit
$\boldsymbol{\Omega_0}$	4.73×10^{-6} s^{-1}
$\boldsymbol{\rho_0}$	10^{-7} g cm^{-3}
$\boldsymbol{T_0}$	2×10^5 K
$\boldsymbol{H_g}$	1.11×10^{12} cm
H_r	$4.06\ H_g = 4.51 \times 10^{12}$ cm
r_0	$r_0 = 45\ R_s = 6.68 \times 10^{14}$ cm
τ_0	6.26×10^4
Σ	12 g cm^{-2}
$P_{g,0}$	2.77×10^6 dyn cm^{-2}
$P_{rat} = P_{r,0}/P_{g,0}$	4.37
$c_{rat} = c/c_s$	5.69×10^3
τ_{therm}	80 days

Table 4.1: Ω_0, ρ_0, T_0, and H_g are fiducial units we use as normalization parameters in our simulation, for time, mass, temperature, and distance, respectively. H_r is the initial total (gas and radiation combined) pressure scale height. r_0 is the fixed location of the center of the shearing box and τ_0, $P_{g,0}$, and Σ are the initial midplane optical depth, initial midplane gas pressure, and constant surface density, respectively. P_{rat} and c_{rat} are the initial ratio of the radiation pressure to the gas pressure and the ratio of the speed of light to the sound speed, respectively, at the midplane. We also give the approximate thermal timescale (τ_{therm}) calculated from the average α-viscosity in our simulations once they reach steady-state.

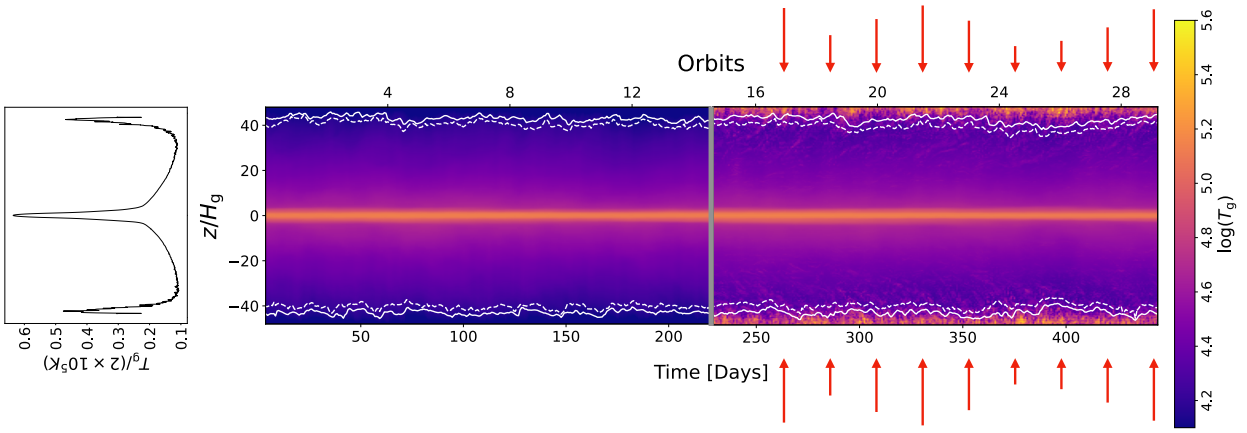


Figure 4.1: Right panel: Schematic of our simulation and space-time diagram of the gas temperature in the one- and two-frequency simulations in the region left and right of the grey line, respectively. The red arrows represent the time-varying X-ray radiation we inject in through the upper and lower boundaries in our two-frequency simulations. Left panel: Horizontally- and time-averaged vertical profiles of the gas temperature for our two-frequency simulation. All figures are for Run A.

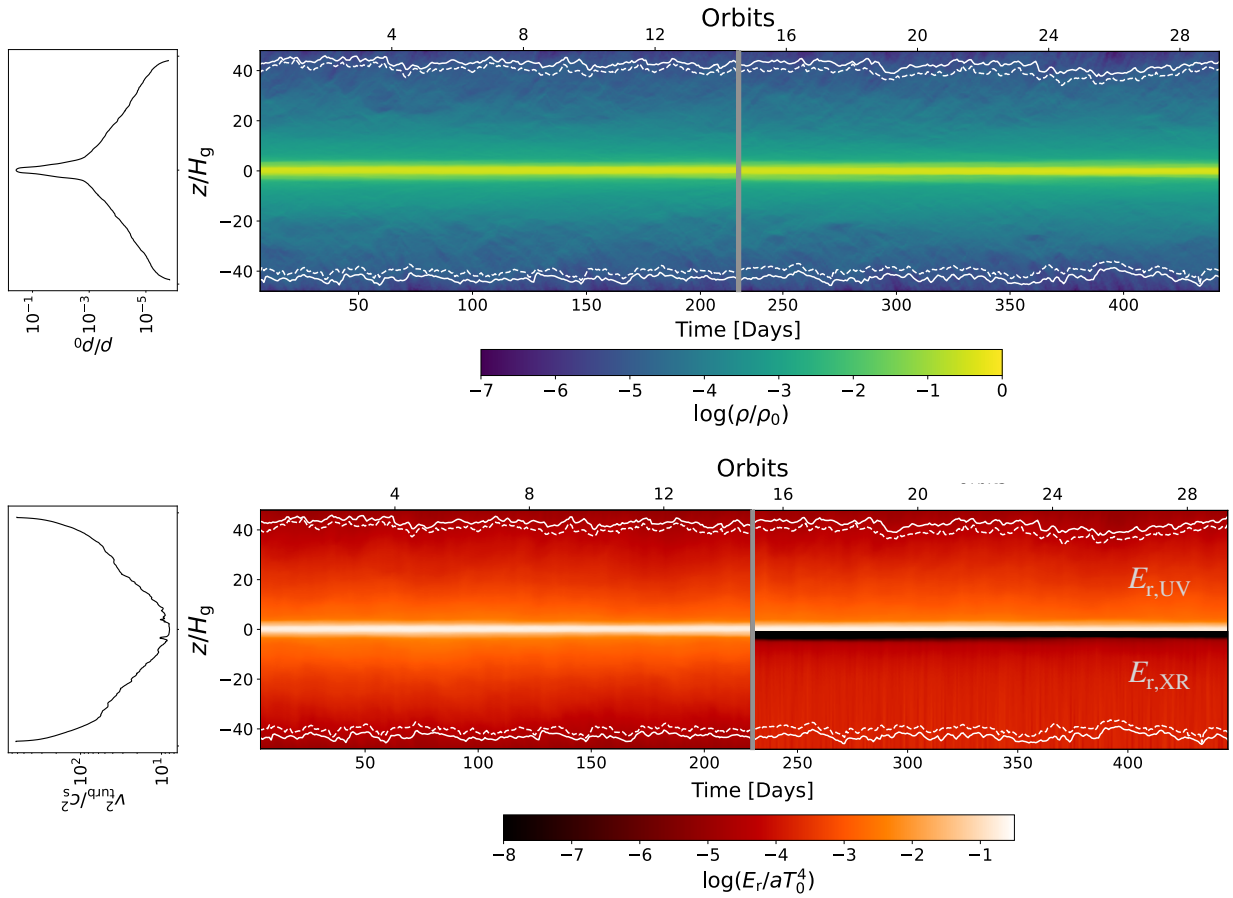


Figure 4.2: Top Right panel: Space-time diagram of the gas density in the one- and two-frequency simulations in the region left and right of the grey line, respectively. Bottom Right panel: Same as top panel, but for the radiation energy. We show the UV radiation energy above the midplane and the X-ray radiation energy below the midplane. In all panels the solid white line is $z(\tau_{\rm UV} = 1) \approx z(\tau_{\rm XR} = 1)$ and the dashed white line is $z(\tau_{a,\rm UV} = 1)$. $\tau_{a,\rm XR} < 1$ throughout our simulation. Left panels: Horizontally- and time-averaged vertical profiles of density, and turbulent velocity in the top and bottom panels, respectively, for our two-frequency simulation. All figures are for Run A.

dominated by the scattering opacity near the surface of the disk ($z(\tau_{\rm UV}=1) \approx z(\tau_{\rm XR}=1)$). The dashed white line shows the absorption photosphere for the UV component $z(\tau_{\rm a,UV}=1)$. $\tau_{\rm a,XR} < 1$ throughout the whole disk, but the effective absorption optical depth is $\tau_{\rm a,XR} \approx 1.3$. The left panels of Figures 4.1 and 4.2 show the horizontally- and time-averaged vertical profiles for the gas temperature, density, and turbulent velocity for the two-frequency Run A. The vertical profiles for two-frequency Run B, not shown here, are qualitatively the same as in Run A.

The vertical profiles are very similar between the one- and two-frequency runs, except for the profile for the gas temperature (top panel), which is higher throughout the simulation when we inject X-rays. Unsurprisingly, the temperature increases the most in the two-frequency run in the photosphere, where on short timescales the temperature can exceed the midplane temperature, and inwards an additional 10 $H_{\rm g}$, where the temperature can nearly reach $T_{\rm g} \sim 10^5$ K. In contrast to isothermal simulations, we do no see any sign of disk winds in our simulations (e.g., Bai & Stone, 2013).

In the bottom panel of Figure 4.2, we show the radiation energy for the low- and high-frequency group above and below the midplane, respectively, to the right of the grey line. The midplane radiation energy is dominated by the energy from the low-frequency group. However, at around ± 30 $H_{\rm g}$, the radiation energy from the high-frequency group becomes the dominant radiation energy. On average, the energy of the X-ray radiation manages to penetrate down to $\tau_{\rm XR} \approx 240$ or ± 6.5 $H_{\rm g}$ before it drops by an e-folding.

Overall, the α-viscosity parameter in our simulations ranges from around 0.025 to 0.04. In all runs, the Maxwell stress is about a factor of 6 larger than the Reynolds stress. The average thermal timescale is $\tau_{\rm therm} = 1/(\alpha\Omega) \approx 80$ days in all runs.

4.3.2 Intrinsic Variability in UV Light Curves

We show the intrinsic UV light curve from the one-frequency run without X-ray irradiation and the power spectral density (PSD) of this light curve in blue in the top panels of Figures

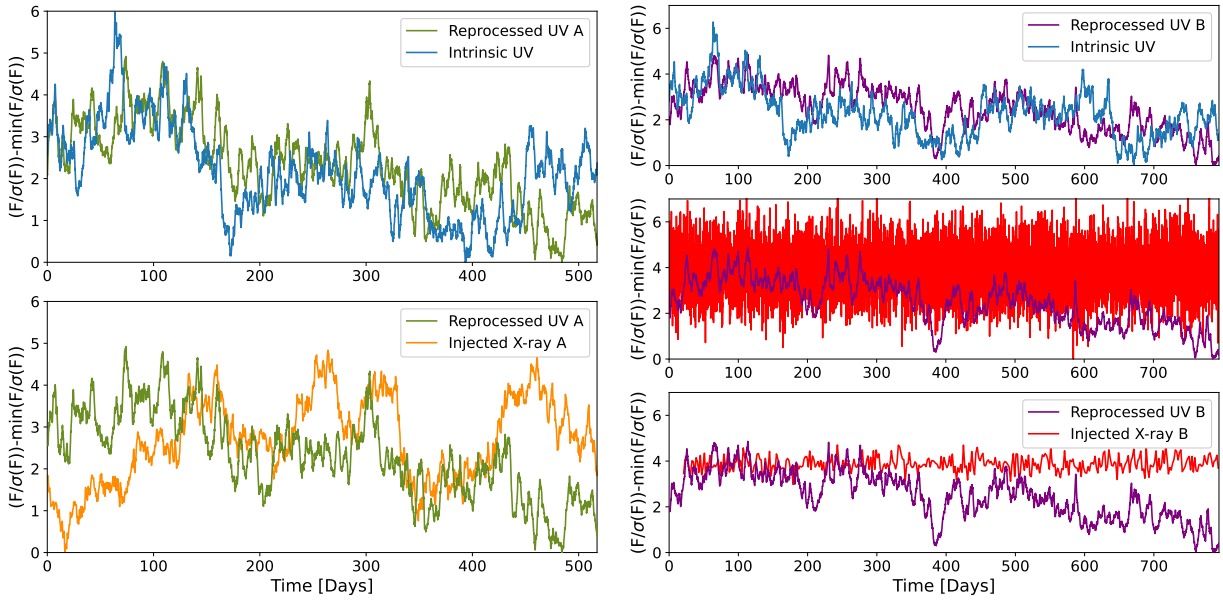


Figure 4.3: Top panels compare the intrinsic and reprocessed UV light curves which come from the one- and two-frequency runs, respectively. Middle and bottom panels compare the reprocessed UV light curves and injected X-rays from the two-frequency runs. $\tau_{\rm damp} = 30$ days for the injected X-ray light curve in the left panels (Run A) and $\tau_{\rm damp} = 0.1$ days for the injected X-ray light curve in the right panels (Run B). In the bottom right panel the X-ray light curve has been window-averaged over 2 days for easier comparison. To make each UV light curve, we sum the total outgoing flux from the low-frequency group over every cell at the top and bottom of the simulation box. For easier comparison, we scale the flux, F, by the standard deviation of the flux, $\sigma(F)$, and then subtract off the minimum flux for each light curve. Time zero corresponds to when we start injecting the X-rays. Note that because X-rays are injected directly into the simulation there will be no lag between the X-ray and UV light curves from the light travel time.

4.3 and 4.4, respectively. The intrinsic UV light curve has variability over timescales of monthly to sub-daily that is driven by MRI turbulence and convection. There is evidence that the UV PSD is damped at the lowest frequencies ($f \lesssim 1/(30\text{ days})$), and then goes as f^{-2} for $0.03 < f < 0.2$ days^{-1}, suggesting that a DRW model might be a good fit for the intrinsic UV light curve at these frequencies. Recent examinations of long baseline optical AGN light curves have suggested that the DRW damping timescale is related to the thermal timescale of the disk (Kelly et al., 2009; Burke et al., 2021; Stone et al., 2022). The damping timescale of around 30 days that we see evidence of in Figure 4.4, is very similar to the

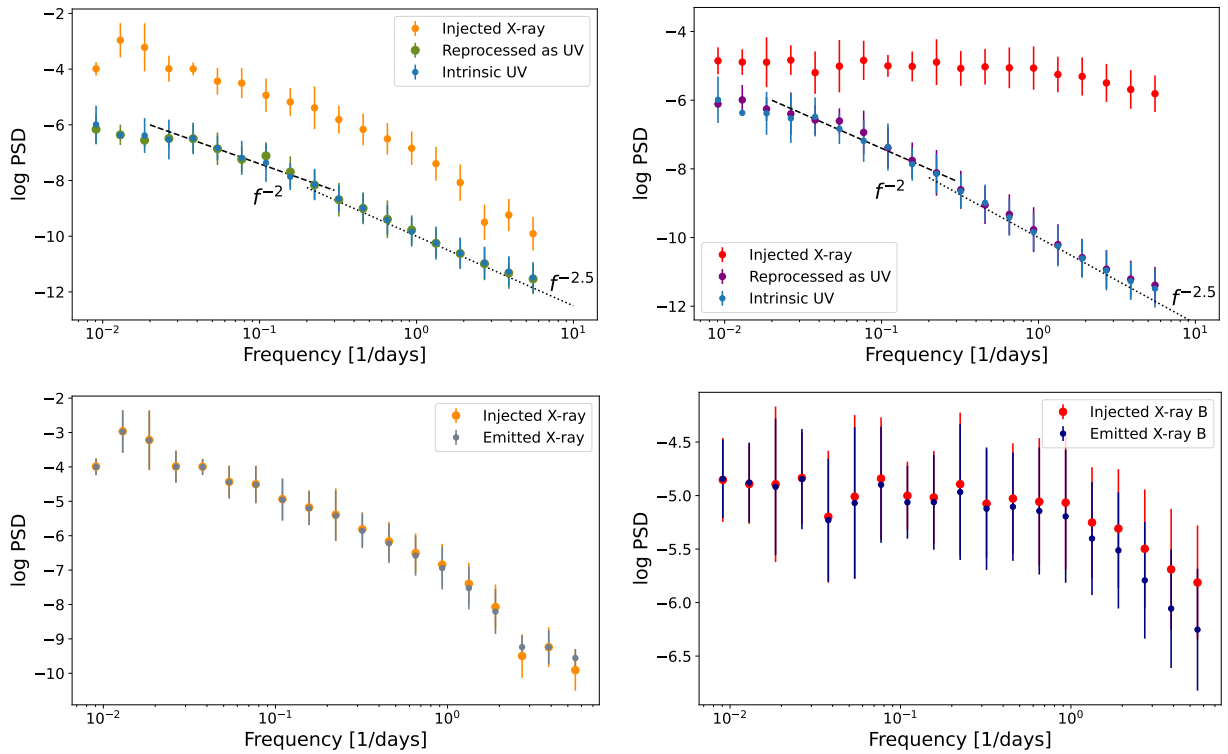


Figure 4.4: The top panels show the PSDs of the intrinsic UV, reprocessed UV, and injected X-ray light curves in Figure 4.3. The bottom panels compare the PSDs of the injected and emitted X-rays. The left (right) panels show the PSDs from Run A (Run B), where $\tau_{\rm damp} = 30$ days ($\tau_{\rm damp} = 0.1$ days) for the injected X-ray light curve. We see some evidence that the UV PSDs are damped at $f \lesssim 0.03$ days^{-1}. At $0.03 < f < 0.2$ days^{-1} the slope of the UV PSDs appear to be well fit by f^{-2}, shown as the dashed line, which is the DRW slope above the damping frequency. At $f > 0.2$ days^{-1} the UV PSDs are well fit by $f^{-2.5}$, shown as the dotted line.

thermal timescale near the surface of the disk in our simulations ($\sim$ 30 days), but longer light curves are needed to robustly model this damping timescale.

While the DRW model may be a good fit at lower frequencies, the slope of the UV light curve's PSD is steeper, $\sim f^{-2.5}$, than the DRW model at $f \gtrsim 1/(5$ days). In high cadence observations of AGN light curves from Kepler, on timescales around ten days there is evidence that the PSD slope is steeper than the f^{-2} power-law of the DRW (e.g., Smith et al., 2018). In our simulations, the break may be at a higher frequency because we are simulating the UV region as opposed to the optical region being probed by Kepler. Tachibana et al. (2020) proposed that this break comes from the fact that variability will be averaged over different

regions of the disk. Our simulations are highly localized which may be the reason why we do not see a break at frequencies as low as in Kepler light curves. On the other hand, naively, one might expect that if the turbulent heating occurs mainly at the midplane, then the spectrum would steepen at frequencies greater than $1/\tau_{\rm diff}$, where $1/\tau_{\rm diff} = \tau_0 H_r/c$ is the photon diffusion time from the midplane; about 100 days for our simulations. The shorter timescale for the drop in variability may be a clue that the turbulent dissipation providing the variability does not come from the midplane, but higher in the disk where the optical depth is lower.

Nonetheless, we find that our intrinsic UV light curve has variability on timescales of monthly to sub-daily, due to MRI- and convection-driven fluctuations. Previous global MHD simulations suggest that the fluctuations producing this variability would be accreted inwards (Hogg & Reynolds, 2016), leading to inward propagating lags (i.e. shorter wavelengths lagging longer wavelengths) on the inflow timescale, otherwise known as long negative lags. Evidence for these long negative lags have been observed for the AGN Fairall 9 (Hernández Santisteban et al., 2020; Yao et al., 2023; Secunda et al., 2023).

4.3.3 Hard X-rays Have Little Impact on UV Emission

We compare the injected X-ray and reprocessed UV light curves from the two-frequency runs which include X-ray irradiation in the bottom panels of Figure 4.3 and the PSDs of these light curves in the top panels of Figure 4.4. Note that because we are injecting X-rays directly into the simulation there will be no lag between the X-ray and UV light curves from the light travel time. For Run A, the Pearson r-coefficient between the injected X-ray light curve with damping timescale $\tau_{\rm damp} = 30$ days and the reprocessed UV light curve is $r = 0.34$, with a p-value, $p < 10^{-5}$. For Run B, where the damping timescale of the injected X-ray light curve is $\tau_{\rm damp} = 0.1$ days, we do not get a significant correlation between the X-ray and reprocessed UV light curves, $p = 0.64$. This lack of correlation is despite the fact that the injected X-ray flux is a factor of 8 larger than the emitted UV flux, which is larger

than typically observed. The lack of correlation between the X-ray and reprocessed UV light curves in Run B may be due to the short damping timescale of the X-ray light curve, which makes the X-ray light curve essentially white noise for all timescales of interest and therefore hard to cross-correlate.

For both runs, the injected X-ray light curves appear to have no impact on the PSDs of the reprocessed UV light curves, which are nearly identical to the intrinsic UV light curves from the one-frequency run without X-ray irradiation. This lack of impact is particularly surprising for Run B, where at high frequencies the X-ray irradiation has around 6 orders of magnitude more power than the intrinsic UV light curve. The lack of additional high frequency variability in the reprocessed UV light curve in Run B is apparent in the light curve itself, as well.

The reason the UV PSD from Run B is not enhanced at high frequencies may be related to the fact that the emitted X-ray light curve in Run B has less power at $f \gtrsim 1/(1\ \mathrm{day})$ than the injected X-ray light curve (see bottom right panel of Figure 4.4). The smoothing of high frequency variability in the injected X-ray light curve occurs because the X-ray irradiation gets scattered instead of absorbed by the disk. As a result, the injected and emitted light curves in Run B are less strongly correlated, $r = 0.89$, than in Run A, where $r = 0.998$ and there is less high frequency variability to be smoothed.

Due to the very low absorption opacity for the X-ray irradiation, there is very little absorption and direct heating, as is typically assumed for reverberation mapping. Instead, our results suggest that the scattering of the incident X-ray photons off of electrons leads to a momentum transfer that eventually will thermalize, producing the small amount of disk heating we see in our simulations (see Figure 4.1) and the correlation we see between the X-ray and UV light curves in Run A. This momentum transfer and disk heating does have some impact on the re-emitted UV light curve in Run B, as well, which in the top right panel of Figure 2 is clearly not identical to the intrinsic UV light curve. For both two-frequency runs, the reprocessed UV light curves are not strongly correlated with the intrinsic UV light

curve from the one-frequency run without X-ray irradiation, with $r = -0.28$ ($p < 10^{-5}$) and $r = 0.17$ ($p = 0.0032$), for Run A and Run B, respectively. The anti-correlation between the UV light curves in Run A results from modeling the injected X-ray light curve in this run as a DRW with a damping timescale on the order of the thermal timescale of the simulation, which introduces some incidental and non-physical correlation and anti-correlation between portions of the X-ray and intrinsic UV light curves. The overall correlation between the injected X-ray and intrinsic UV light curve in Run A is $r = -0.032$, $p = 0.0032$.

The scattering atmospheres in our simulations are similar to those modeled by Gardner & Done (2017), who found that UV emission from NGC 5548 was in good agreement with analytic models assuming a puffed up photosphere with low X-ray absorption opacities. In addition, AGN disks with scattering atmospheres may help account for the longer than anticipated lag timescales of recent AGN reverberation mapping campaigns (Hall et al., 2018). However, other models, such as those in Panagiotou et al. (2022a), suggest that the canonical reprocessing of X-rays through absorption can also produce longer lag timescales and poorly correlated X-ray and UV light curves, like those observed for NGC 5548. In addition, observations using X-ray spectroscopy suggest that at least the innermost regions of AGN disks may not be fully-ionized or scattering dominated (e.g., García et al., 2019).

The time-dependent opacities in our simulation are calculated self-consistently from the temperature and density determined by evolving the MHD equations using an initial vertical profile in radiation hydrostatic equilibrium. If our simulation box was located in a different temperature region of the disk or we used different AGN parameters, such as supermassive black hole (SMBH) mass or accretion rate, it is possible that we would have different initial conditions that would lead to an increased absorption or decreased scattering optical depth near the surface of the disk. We have also assumed that the opacity for the X-ray radiation is the proper average of the free-free absorption opacity and electron scattering opacity. This assumption is most accurate if the gas is fully ionized. The disk may not be fully ionized if, for example, the temperature near the disk surface drops below $T_{\rm g} \approx 10^4$ K, which it might

for different initial conditions. In general, a lower ionization fraction would decrease the scattering opacity and likely lead to increased absorption (García et al., 2013). Using a more sophisticated opacity model for high-energy photons could also lead to higher absorption opacities. Higher absorption opacities would result in a larger energy transfer from the X-ray irradiation to the gas and likely increase the correlation between the X-ray and UV light curves.

Salvesen (2022) found there was a non-zero thermalization time for the reprocessing of X-ray irradiation. We tested whether there is a lag between the X-ray and UV light curves and found that the strongest correlation between the X-ray and UV light curves in Run A is still at < 0.08 days. On the other hand, we found a very weak correlation for Run B, $r = 0.037$ and $p = 0.002$, occurs for a lag of 32 days, which is roughly the thermal timescale near the surface of the disk. However, a lag of around 30 days for reprocessing has not been observed in disk continuum reverberation mapping campaigns.

Overall, our simulations show that the low absorption opacity for the hard X-ray irradiation makes it unlikely that hard X-rays are the main driver of variability in emitted UV light curves. Several AGN reverberation mapping campaigns have found no correlation (Morales et al., 2019; Buisson et al., 2018, e.g.,) or only weak correlations (e.g., Schimoia et al., 2015; Edelson et al., 2019; Cackett et al., 2023) between X-ray and UV-optical light curves, which agrees with our results here. However, it is possible that this observed lack of correlation in these papers is simply due to absorption of intervening material (e.g., Kara et al., 2021).

4.4 Summary and Future Work

Using the multi-frequency radiation MHD code from Jiang (2022) allows us to simulate the reprocessing of X-ray irradiation into UV emission from first principles. This reprocessing is a key component of reverberation mapping, one of the primary tools used to study AGN disks. The traditional model for reverberation mapping assumes that X-rays are the main

driver of variability in UV and optical light curves emitted by the AGN disk. However, the results from our simulations do not agree with this assumption. Instead we find that there is significant intrinsic UV variability generated locally by the MRI and convection. This variability can be described by a DRW, with $\tau_{\rm damp} \approx 30$ days on timescales greater than five days, but shows less power than the DRW at higher frequency, in agreement with AGN light curves observed at high cadence by Kepler (Smith et al., 2018).

The X-rays we inject into our simulations have almost no impact on the PSDs of the reprocessed UV light curves in our two-frequency runs, because the short timescale variability of the X-ray irradiation gets smoothed when it is scattered in the outer regions of the disk. As a result, the X-ray and UV light curves from our two-frequency runs are not well-correlated if there is not sufficient X-ray variability on timescales greater than a day. Even with longer timescale variability, we only find a correlation of $r = 0.34$ between the X-ray and UV light curves, again because of the low absorption opacity for X-rays in our simulations. This correlation agrees with observational results that suggest that X-ray and UV light curves are not as well-correlated as we might expect if X-rays are the main driver of UV and optical variability as in traditional reprocessing models (e.g., Schimoia et al., 2015; Edelson et al., 2019; Cackett et al., 2023). However, canonical reprocessing models could produce lower X-ray to UV correlations if the X-ray light curve is non-stationary, for example due to changes in the height or temperature of the corona on timescales shorter than the observing campaign (Panagiotou et al., 2022a).

There are many radiation MHD simulations of AGN disks, but most are one-frequency simulations and none focus on the impact of X-ray irradiation on the UV-optical region of AGN disks and the light curves emitted by the disk. On the other hand, there are numerous sophisticated (semi-)analytical calculations for disk reprocessing of X-ray irradiation (e.g., Sun et al., 2020; Kammoun et al., 2021b; Salvesen, 2022; Panagiotou et al., 2022b). However, they all make assumptions about the vertical profiles of AGN disk parameters and do not evolve the gas turbulence in three dimensions nor the impact that radiation has on the gas

and vice versa as a function of time. On the other hand, the multi-frequency radiation MHD simulations presented here, provide detailed information on the nature of light curve variability produced by disk turbulence, and how X-ray light curves with different PSDs transfer energy and momentum to the AGN disk, affecting the variability in UV light curves emitted by the disk.

Increasing our understanding of the physical processes that influence reverberation mapping is crucial because there is growing evidence that the standard Shakura & Sunyaev (1973) thin disk model does not accurately model an AGN disk (e.g., Antonucci, 1988; Antonucci et al., 1989; Antonucci, 2023). For example, radiation MHD simulations show that pressure support from magnetic fields and radiation can lead to puffed up disks (e.g., Gaburov et al., 2012; Jiang et al., 2016, 2019; Jiang & Blaes, 2020; Hopkins et al., 2023), while recent observations provide additional evidence for thicker disks than predicted by the standard thin disk model (Yao et al., 2023; Secunda et al., 2023). In addition, both micro-lensing and reverberation mapping campaigns suggest that the radial extent of the disk is larger than predicted by the standard disk model (e.g., Morgan et al., 2010; Jiang et al., 2017; Guo et al., 2022).

Improved intuition about the physical processes involved in reverberation mapping is essential to improving our models for AGN disks. Future work should look at whether lower energy soft X-ray or far-UV irradiation, which will have larger absorption opacities compared to hard X-ray irradiation, is more likely to be absorbed and reprocessed by the disk producing stronger correlations between injected and reprocessed light curves. Future work should also examine a wider range of parameters for AGN disks in order to determine how the physics of reverberation mapping changes as a function of parameters such as SMBH mass and accretion rate. Finally, expanding these shearing box simulations to global simulations is necessary to understand how the fluctuations producing this intrinsic variability are propagated throughout the disk affecting the reverberation mapping signal. These simulations will allow us to make our reverberation mapping techniques more accurate and mine

more information out of AGN light curves about the structure and internal physics of AGN disks.

Chapter 5

Simulating X-ray Reverberation in the UV-Emitting Regions of Active Galactic Nuclei Accretion Disks with 3D Multi-Frequency Radiation Magnetohydrodynamic Simulations II: Updated X-ray Opacity Models and Modeling the Response Function

5.1 Introduction

Active galactic nuclei (AGN) accretion disks are important laboratories to study accretion physics and fuel the powerful radiation and outflows thought to have a large impact on galaxy formation (Fabian, 2012; Kormendy & Ho, 2013). Unfortunately, these AGN accretion disks are generally too distant to resolve spatially. However, because the light curves emitted by the AGN corona, accretion disk ,and broad line region vary with time, temporal resolution can be substituted for spatial resolution using a technique known as reverberation mapping.

Reverberation mapping was first proposed by Blandford & McKee (1982) as a way to measure distances to the broad line region by measuring the time lag between variability in continuum light curves and broad lines (see also, Kaspi et al., 2000; Peterson et al., 2004; Bentz & Katz, 2015; Grier et al., 2017). Collier et al. (1999) later introduced the idea of disk continuum reverberation mapping. Disk continuum reverberation mapping employs the lamp-post model, which proposes that time varying X-ray radiation emitted by the corona will move outwards to different temperature regions of the accretion disk where it will be absorbed, reprocessed, and re-emitted at different wavelengths based on the local temperature of the disk. Because the X-ray radiation takes a light-crossing timescale to travel from the inner hotter regions of the disk to the outer cooler regions, we can use the lag between variability in longer wavelength bands and variability in shorter wavelength bands to measure the radial temperature profile and radial extent of an AGN disk (e.g., Sergeev et al., 2005; Cackett et al., 2007, 2018; De Rosa et al., 2015; Edelson et al., 2015, 2017, 2019; Jiang et al., 2017; Fausnaugh et al., 2016; Starkey et al., 2017; Homayouni et al., 2022).

Traditionally, the lamp-post model assumes that the variability in UV-optical light curves emitted by the AGN disk is driven by the reprocessed variability from the absorbed X-ray irradiation. However, there are a growing number of examples of X-ray and UV-optical light curves that are only moderately (e.g., Edelson et al., 2019; Cackett et al., 2023; Kara et al., 2023) or weakly (Schimoia et al., 2015; Buisson et al., 2018) correlated. This lack or correlation could be due to absorption by intervening material (e.g., Kara et al., 2021) or changes in

the height and temperature of the corona on timescales shorter than the observing timescale (Panagiotou et al., 2022a). Alternatively, there is observational evidence for intrinsic variability in UV-optical light curves emitted by AGN disks (e.g., Arévalo et al., 2009; Neustadt & Kochanek, 2022; Hernández Santisteban et al., 2020). This variability could arise from fluctuations in the UV-optical regions of the accretion disk due to magnetorotational instability (MRI) driven turbulence (Balbus & Hawley, 1991) or convection resulting from the enhanced opacity in the UV-optical region of the disk (Jiang & Blaes, 2020).

The recent development of accurate, high-speed, numerical methods to calculate radiative transfer coupled to magnetohydrodynamics (MHD) makes it possible for the first time to examine how both radiation and magnetic fields have an impact on AGN disk structure and turbulence (Jiang et al., 2013; Jiang & Blaes, 2020; Jiang, 2021; Hopkins et al., 2023). However, these MHD simulations model radiation using a single frequency group, with the radiation integrated over all frequencies, and do not examine the impact of X-ray irradiation on the disk and the light curves emitted by the disk. Instead, reverberation mapping mostly employs empirical models to describe the reprocessing of X-ray irradiation into UV-optical emission.

The most common of these empirical models is a linear reprocessing model. In a linear reprocessing model the reprocessed light curve is,

$$\mathrm{r(t)} = \int_{-\infty}^{\infty} \psi(\tau_{\mathrm{L}})\mathrm{d(t} - \tau_{\mathrm{L}})\mathrm{d}\tau_{\mathrm{L}}, \tag{5.1}$$

where τ_{L} is the continuum lag on the light-crossing timescale, $d(t)$ is the driving light curve, and $\psi(t)$ is the response function, which is commonly modeled as a Gaussian or log-normal function. There have been efforts to develop more sophisticated theoretical models of reverberation mapping that can help to explain the discrepancies between the lamp-post model and observations. However, these models are primarily (semi-)analytic and make assumptions about the vertical profiles of various disk parameters, which often remain fixed over

time or are not-calculated self-consistently to include the interaction of radiation and gas (e.g., Sun et al., 2020; Kammoun et al., 2021b; Salvesen, 2022; Panagiotou et al., 2022b).

In Secunda et al. (2024), hereafter Paper I, we provided an early application of a multi-frequency radiation MHD code described in Jiang (2022), which instead of integrating over all frequencies of light, allows the user to integrate over different frequency bands. In Paper I we used two frequency groups to simulate the reprocessing of high-frequency or X-ray radiation into low-frequency or UV radiation by the AGN disk. Unlike in the previous semi-analytic models, in these simulations the radiation energy and momentum are coupled to the gas evolution and the temperature, opacity, and energy of the simulation depend self-consistently on the local gas density and pressure. We found that fluctuations in the disk due to the MRI and convection lead to intrinsic variability in the UV light curves emitted by the AGN disk, and the power spectral density (PSD) of this variability is similar to observed UV-optical AGN light curve PSDs. We also found that due to low absorption opacity for the X-ray radiation, the injected X-ray and reprocessed UV light curves from the simulations were only weakly correlated.

Paper I used tabulated Rosseland and Planck mean opacities that include line opacities instead of only electron scattering and free-free absorption for the low-frequency radiation group, because including line opacities for UV radiation has been shown to have a large impact on AGN disk models (Jiang & Blaes, 2020). However, the opacities for the high-frequency group were modeled simply using the proper average of the free-free absorption opacity and electron scattering opacity.

In this paper, we build on Paper I by using a more sophisticated opacity model for the X-ray component in order to see if with more realistic opacities the X-ray absorption will be enhanced, increasing the correlation between the injected X-ray and reprocessed UV light curves. In particular, a higher correlation between the injected X-ray and reprocessed UV light curves, would allow us to test the linear reprocessing model commonly used to describe reverberation mapping.

We summarize our simulation set-up, highlighting the differences between the simulations in Paper I and this paper in Section 5.2. We describe the disk structure of our simulations and present the light curves from our different runs in Sections 5.3.1 and 5.3.2, respectively. Because we find that our new opacity model for the X-ray radiation leads to strongly correlated X-ray and UV light curves in several of our runs, we are able to test different models of the response function for the reprocessing of X-ray irradiation into UV emission on our simulations for the first time in Section 5.3.3. Finally, we summarize and discuss these updated results in Section 5.4.

5.2 Methods

The initial set-up of our simulation is the same as in Paper I. In this section we provide a brief summary of the simulation set-up and refer the reader to Paper I for full details.

Using ATHENA++, we simulate a local patch of an AGN accretion disk initialized in radiation hydrostatic equilibrium using the local shearing box approximation. In the shearing box approximation the x- and y-coordinates correspond to the radial and azimuthal dimensions, while the z-coordinate corresponds to the height. We use the same parameters as in Paper I which are summarized in Table 5.1. We use periodic and shear periodic ATHENA++ boundary conditions for the x- and y-boundaries, respectively, and outflow boundary conditions for the gas at the z-boundary. We initialize the magnetic field in our simulation as in Paper I using two oppositely twisted flux tubes with zero net-flux and an additional purely vertical magnetic field to drive additional Maxwell stress.

In addition, we solve the implicit radiative transfer equations for intensities over discrete angles coupled to the MHD solver in ATHENA++ for two frequency groups in the manner described in Jiang (2021) and Jiang (2022). We use the same two frequency groups as in Paper I, one group covering the frequency range $[0, 1\text{keV}]$ to model the radiation field emitted by the disk locally (at UV wavelengths) and the other group covering the frequency range

Parameter	Unit
$\boldsymbol{\Omega_0}$	4.73×10^{-6} s^{-1}
$\boldsymbol{\rho_0}$	10^{-7} g cm^{-3}
$\boldsymbol{T_0}$	2×10^5 K
$\boldsymbol{H_g}$	1.11×10^{12} cm
H_r	$4.06\ H_g = 4.51 \times 10^{12}$ cm
r_0	$r_0 = 45\ R_s = 6.68 \times 10^{14}$ cm
τ_0	6.26×10^4
$P_{g,0}$	2.77×10^6 dyn cm^{-2}
$\beta_{x,y}$	200
β_z	2×10^4
$P_{rat} = P_{r,0}/P_{g,0}$	4.37
$c_{rat} = c/c_s$	5.69×10^3
N_x	48
N_y	192
N_z	1536
L_x	$3\ H_g$
L_y	$12\ H_g$
L_z	$96\ H_g$
N_θ	48
τ_{therm}	32 days

Table 5.1: Ω_0, ρ_0, T_0, and H_g are fiducial units we use as normalization parameters in our simulation, for time, mass, temperature, and distance, respectively. H_r is the initial total (gas and radiation combined) pressure scale height. r_0 is the fixed location of the center of the shearing box and τ_0 and $P_{g,0}$ are the initial midplane optical depth and gas pressure, respectively. $\beta_{x,y}$ and β_z are the ratio of the gas to magnetic pressure for the zero net-flux and vertical magnetic fields, respectively. P_{rat} and c_{rat} are the initial ratio of the radiation pressure to the gas pressure and the ratio of the speed of light to the sound speed, respectively, at the midplane. N_x, N_y, N_z are the number of cells in the x, y, and z direction, while L_x, L_y, L_z are the dimensions of the box. N_θ is the number of discrete angles in each cell for each frequency group. Finally, (τ_{therm}) is the approximate thermal timescale calculated from the average α-viscosity in our simulations once they reach steady-state.

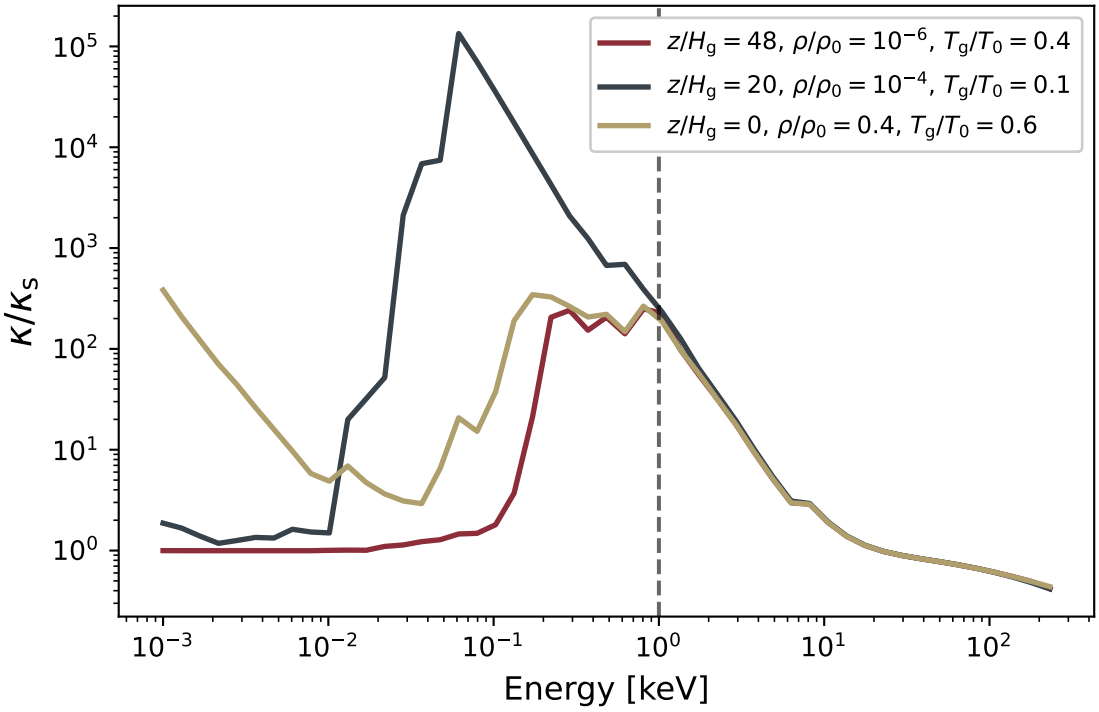


Figure 5.1: The opacity, κ, scaled by the scattering opacity, κ_s, as a function of frequency at densities and temperatures representative of three different heights in the simulation given in the legend. We show our cut off between our high- and low-frequency radiation groups at 1 keV. While the opacity is greater at lower frequencies, the mean opacities are weighted such that the opacity at the lowest energy in each group dominates. As a result, the mean opacity for the high-frequency group is larger throughout the disk.

$(1\text{keV}, \infty)$ for the X-ray irradiation. We initialize the UV radiation isotropically based on the local disk temperature. We initialize the X-ray radiation isotropically to a constant radiation density four orders of magnitude weaker than the midplane radiation energy for the UV frequency group, as in Paper I. At the z-boundaries, we allow intensities for both frequency groups to leave the simulation domain, with outgoing intensities copied from the last active zone to the ghost zone. For the low frequency group, intensities are not allowed to enter. For the high frequency group, we set the incoming intensities based on a prescribed X-ray light curve.

Our implicit radiative transfer scheme allows us to calculate the opacity for each frequency group based on the time-evolving gas density and temperature. As in Paper I, we determine

this opacity for the UV component in each cell with a bi-linear interpolation of the OPAL opacity tables from Iglesias & Rogers (1996) for each time-step based on each cell's current gas temperature and density. The largest change between our previous simulations in Paper I and this paper is the way we model the opacity for our X-ray component. In Paper I, we modeled the X-ray opacity as the proper average of the free-free absorption opacity and electron scattering opacity. We found that for the temperatures and densities of our simulation the proper average was dominated by the scattering opacity, which prevented the X-rays in our simulation from being absorbed and led to a low correlation between our injected X-ray and reprocessed UV light curves.

However, much like using tabulated OPAL opacities is crucial for modeling the structure of and fluctuations driving variability in AGN disks (e.g., Jiang et al., 2016), including a more sophisticated model for opacity for our higher frequency group could be the key to increasing X-ray absorption and the correlation between the driving X-ray and reprocessed UV light curves. We therefore use the online software TOPS (Abdallah & Clark, 1985) to calculate for solar composition the Rosseland and Planck mean opacities as a function of density and temperature for photon energies above 1 keV. We show the opacity as a function of frequency for densities and gas temperatures representative of three different heights in our simulation disk in Figure 5.1. We normalize the opacity by the scattering opacity $\kappa_{\rm s} = 0.3$ cm^2 g^{-1}, which was the dominant source of opacity in Paper I for the high-frequency component. It is clear from Figure 5.1 that even at photon energies greater than 1 keV, using mean opacities calculated from TOPS increases the X-ray absorption opacity significantly. In fact, because the Blackbody spectrum peaks at frequencies, $f < 1$ keV, for the temperature range in our simulation, calculated mean Rosseland and Planck opacities are weighted such that the opacity for the lowest photon energies dominate. As a result, the X-ray opacity is actually greater than the UV opacity throughout the simulated disk. We therefore expect this change in opacity model to have a large impact on our simulated light curves.

Run	X-ray Light Curve Parameters
Fiducial	DRW PSD with $\tau_{\rm damp} = 30$ days
Low Flux	DRW PSD with $\tau_{\rm damp} = 30$ days, 1/4 flux of Fiducial
Power-law	Power-law PSD with $P(f) \propto f^{-2}$
Shallow Power-law	Power-law PSD with $P(f) \propto f^{-1.5}$

Table 5.2: A summary of the different properties of the X-ray light curve injected into the simulation for the four different runs in this paper.

We perform four different runs which are summarized in Table 5.2. The only prescribed differences between these runs lies in the way we model the X-ray light curve we inject into the simulation. The first of these runs is our Fiducial run, which uses the same injected X-ray light curve as Run A in Paper I, a DRW light curve with $\tau_{\rm damp} = 30$ days. We do this in order to directly compare the impact of our new opacity model for the X-ray radiation. Next, we perform a run with the same injected X-ray light curve as the Fiducial run, but with 1/4 the flux, in order to study the impact of the amount of X-ray flux on reprocessed UV light curves. Our next two runs use a power-law instead of a DRW to model the PSD of the injected X-ray light curve. The first, our Power-Law run, has a power-law index of -2, which is the same as the power-law component of the DRW, while the second, our Shallow Power-Law run, has a slightly shallower power-law index of -1.5. These power-law indices were chosen based on empirical power-law fits to X-ray AGN light curves (e.g., Markowitz et al., 2003). All light curves are generated with a timestep of 10^{-3} days and a duration of 10^4 days. Finally, we use for comparison the results from our one-frequency run in Paper I, where we do not inject an X-ray light curve and integrate the radiation over all frequencies of light instead of using multiple frequencies groups.

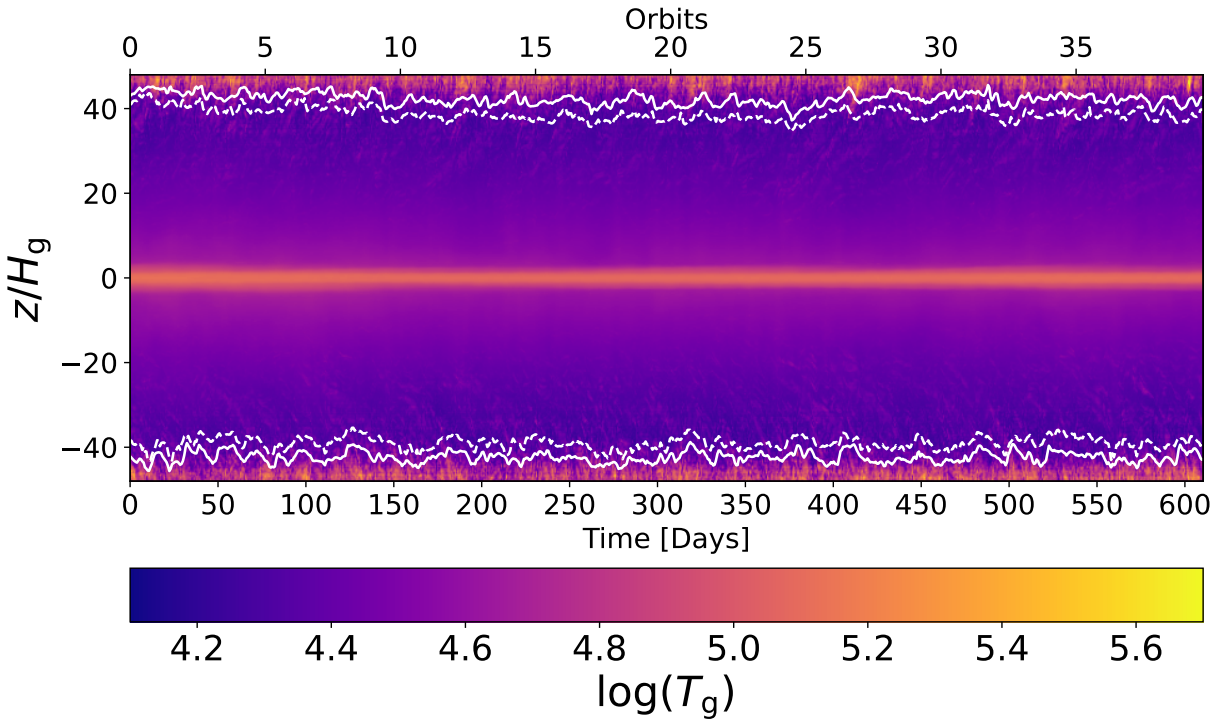


Figure 5.2: A space-time diagram of the horizontally-averaged gas temperature as a function of height and time. The solid white line shows the location of the UV photosphere (i.e $z(\tau_{\rm UV} = 1)$) and the dashed line shows $z(\tau_{\rm a,UV} = 1)$, where $\tau_{\rm a,UV}$ is the UV absorption optical depth.

5.3 Results

5.3.1 Disk Structure

Before presenting the light curves from our simulations, in this section we briefly describe the structure of the disk in our four runs. We show the space-time diagram for the gas temperature for the Fiducial run in Figure 5.2 and the time- and horizontally-averaged vertical profiles for the density, gas temperature, and turbulent velocity in the left panel of Figure 5.3. The properties of the disk are consistent between all four runs and with the disk structure found in Paper I for runs with injected X-rays.

In the right panel of Figure 5.3 we compare the X-ray and UV flux in the Fiducial run to the X-ray and UV flux in the other three runs. By design there is less X-ray flux in the Low Flux run. However, even though the amount of injected X-ray flux is equivalent in

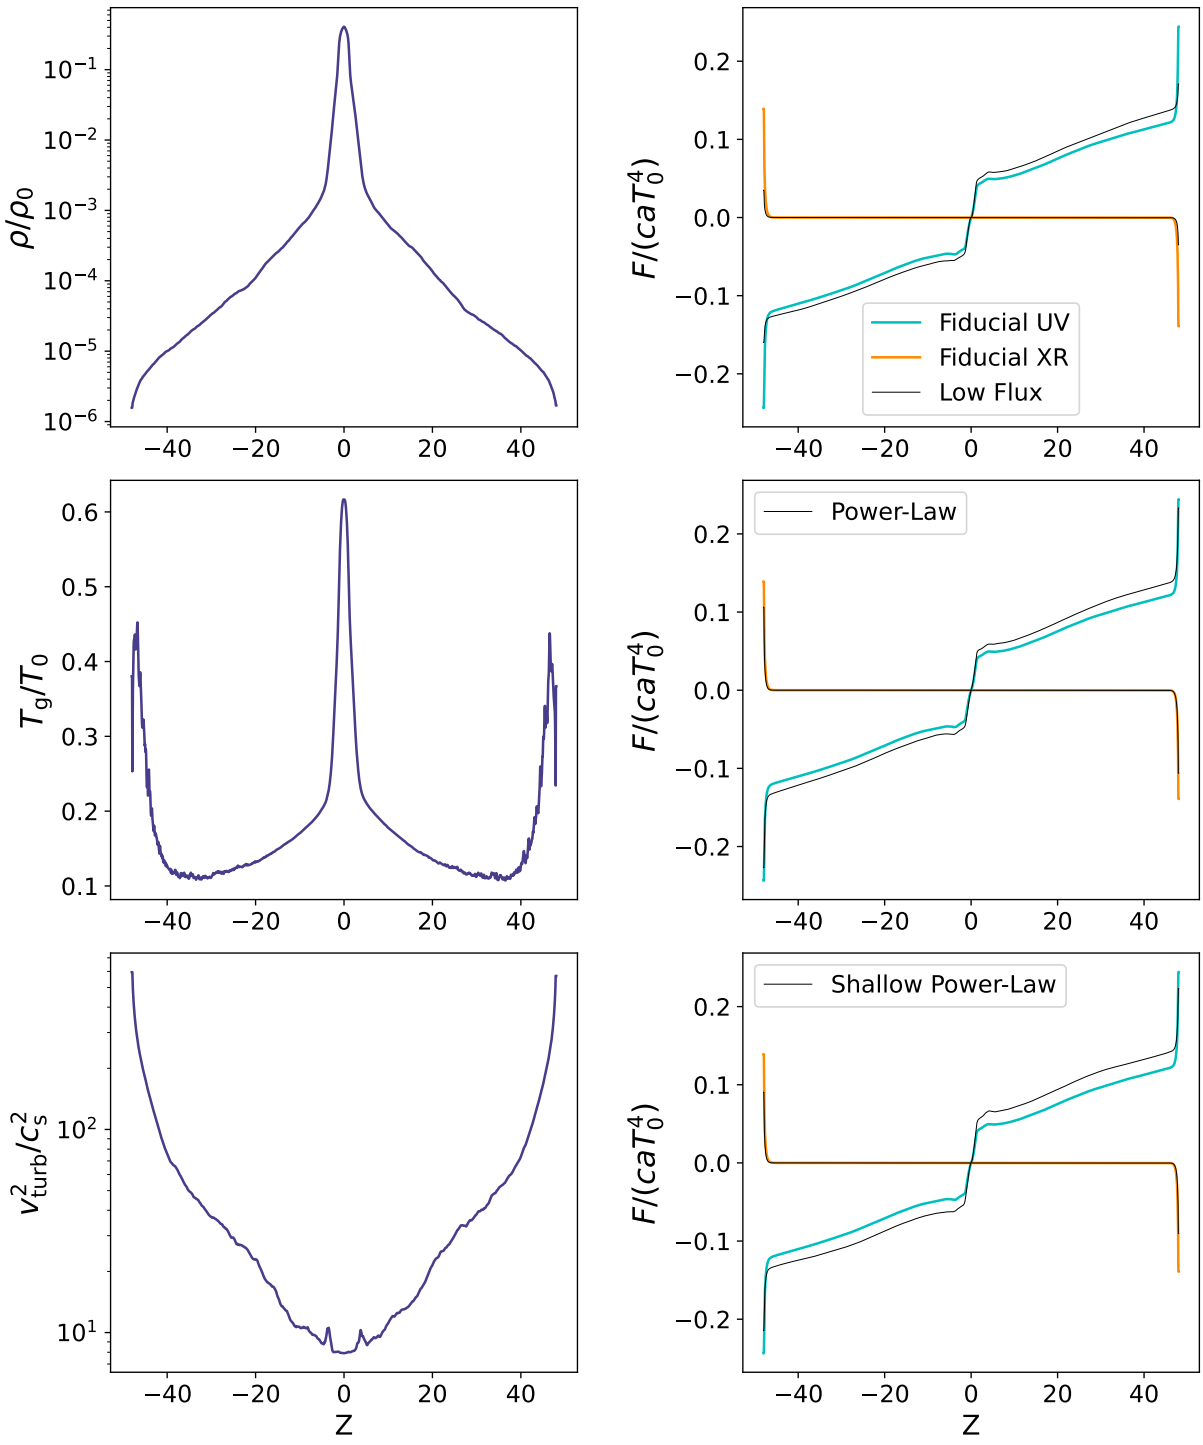


Figure 5.3: **Left panels:** The time- and horizontally-averaged vertical profiles for density (top panel), gas temperature (middle panel), and turbulent velocity (bottom panel) for the Fiducial run. The vertical profiles of these properties for the different runs are consistent with each other. **Right panels:** The UV and X-ray flux for the Fiducial runs are shown in blue and orange, respectively, while the fluxes for the three other runs are shown in black for comparison.

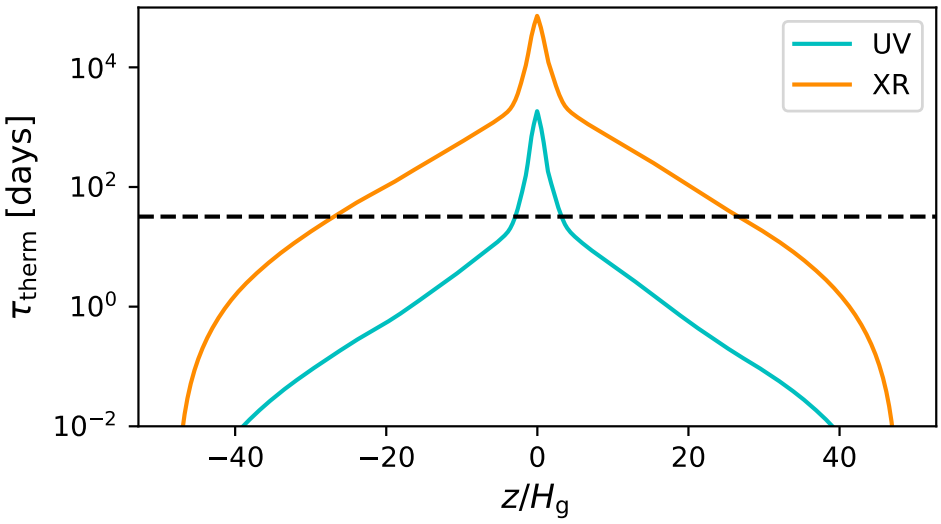


Figure 5.4: The horizontally- and time-averaged vertical profile for the UV and X-ray thermal timescales for the Fiducial run are shown in blue and orange, respectively. The dashed black line shows the average thermal timescale for the disk for the Fiducial model.

the Fiducial and the two Power-Law runs, the amount of flux that penetrates beyond the boundary cells and into the active cells is lower for the Power-Law Runs. Instead there is slightly more UV flux at regions below the surface of the disk, perhaps due to enhanced turbulent dissipation or convection in these runs.

We show the thermal timescale, $\tau_{\rm therm}(z) = \tau_{\rm UV,XR} L/c$, where L is the distance to the edge of the simulation box, $\tau_{\rm UV,XR}$ is the UV or X-ray optical depth, and c is the speed of light, as a function of height for the Fiducial run in Figure 5.4. The dashed black line shows the time-averaged thermal timescale $\langle\tau_{\rm therm}\rangle = 1/(\langle\alpha\rangle\Omega)$, where α is the Shakura & Sunyaev (1973) viscosity parameter and Ω is the orbital frequency, for the Fiducial run. Knowing the thermal timescale of our simulated disk is important because it sets the timescale of fluctuations in the disk. There is also observational evidence that the thermal timescale sets the damping timescale for DRW-modeled UV-optical light curves emitted by AGN disks (Burke et al., 2021).

Run	X-ray-reprocessed UV	inner UV-reprocessed UV	X-ray-inner UV
Fiducial	0.92	−0.03	−0.04
Low Flux	0.22	0.22	−0.20
Power-law	0.81	0.26	0.10
Shallow Power-law	0.41	0.0059	−0.31

Table 5.3: The Pearson r-values for the correlations between the X-ray and reprocessed UV light curves (second column), the inner UV and reprocessed UV light curves (third column), and the X-ray and inner UV light curves (fourth column), for the four runs shown in different rows. The inner UV is the outward flux at $z = \pm 30\ H_g$, below where the X-ray flux penetrates.

5.3.2 Light Curves

We present the simulated light curves from our four runs in Figure 5.5. For comparison, we show the outgoing UV flux at the upper and lower boundaries of the simulation box ($\pm 48 H_g$), which is the reprocessed UV light curve emitted from our simulation, in purple, and the UV flux at $\pm 30 H_g$ in blue, which is deeper than the X-ray flux directly penetrates (see Figure 5.3) and therefore is a decent representation of the intrinsic UV flux without X-ray driven variability. The fluxes have all been normalized for easier comparison. Figures 5.6 – 5.9 show the PSDs for the the injected X-ray light curves and the reprocessed UV light curves emitted from the upper and lower boundaries of our simulation box for each of our four runs. We also show the PSD of the UV light curve from the run in Paper I with no injected X-rays for comparison. Finally, the Pearson r-coefficients for the correlations between these light curves are given in Table 5.3.

In our Fiducial run the injected X-ray and the emitted UV light curves are highly correlated, which is clear both by eye in Figure 5.5 and from the Pearson r-coefficient between them of $r = 0.92$. In addition, the inner UV light curve is not well correlated with either the X-ray or emitted UV light curve, suggesting that the variability in the emitted UV light curve is nearly entirely from the reprocessing of the X-ray variability, despite the inner UV having its own intrinsic variability. This picture is supported by the PSDs of the different light curves in Figure 5.6, where the PSD of the reprocessed UV light curve is mostly con-

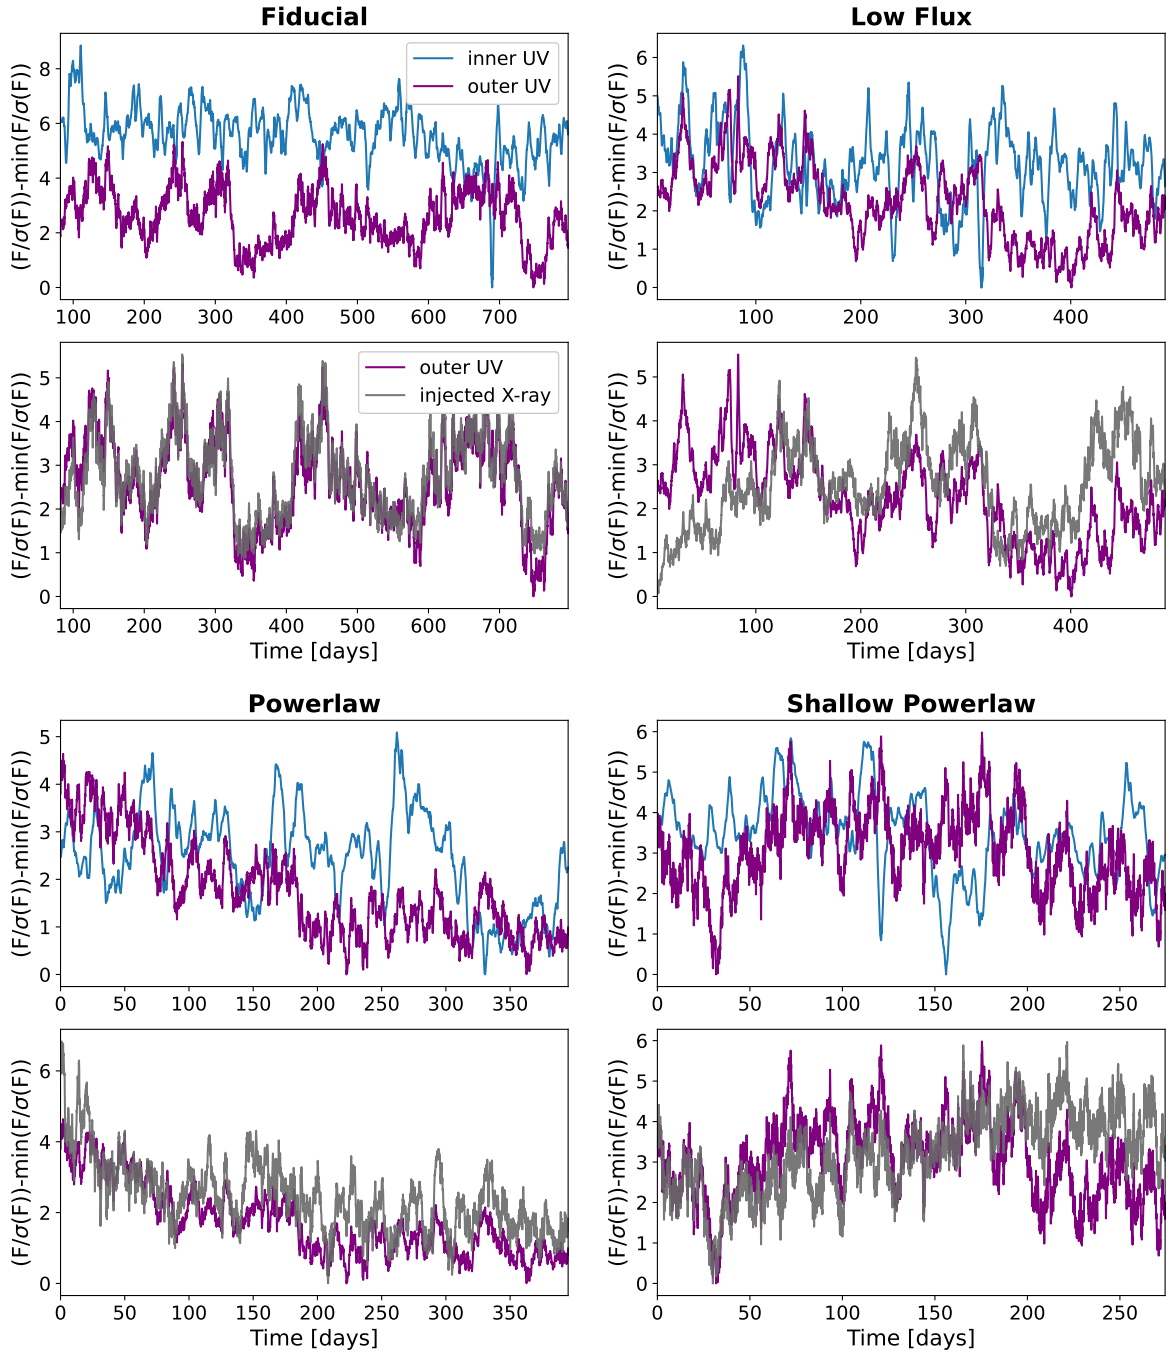


Figure 5.5: The light curves from our four runs. In purple we show the reprocessed UV flux emitted from our simulations, in blue we show the inner UV flux computed from the outward flux at ± 30 H_g below where the X-ray irradiation penetrates, and in grey we show the injected X-ray light curve. For easier comparison, we scale the flux, F, by the standard deviation of the flux, $\sigma(F)$, and then subtract off the minimum flux for each light curve. Time zero corresponds to when we start injecting the X-rays. Note that because X-rays are injected directly into the simulation there will be no lag between the X-ray and UV light curves from the light travel time.

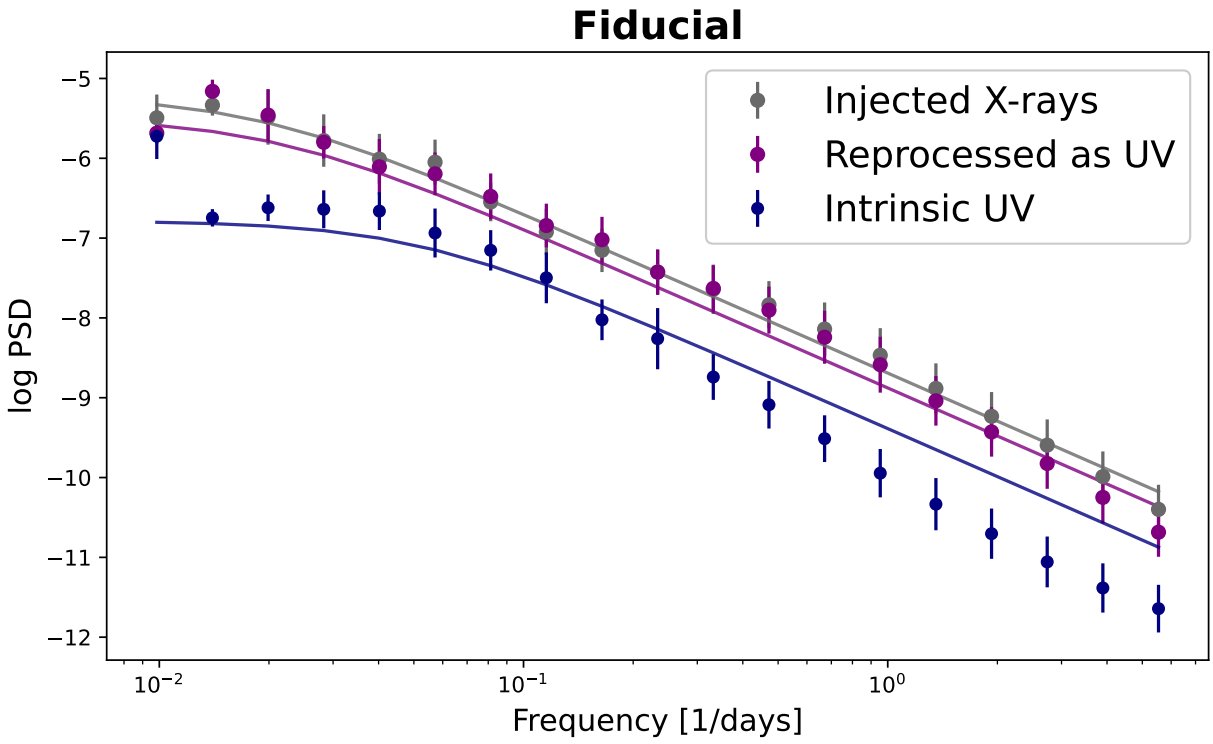


Figure 5.6: The PSDs for the injected X-rays in gray, reprocessed UV in purple, and intrinsic UV (from the one-frequency run with no injected X-rays) in navy for our Fiducial run. We also show here our best fit DRW-model for the intrinsic ($\tau_{\rm damp} = 3.18 \pm 1.5$ days) and reprocessed ($\tau_{\rm damp} = 7.81 \pm 1.4$ days) UV light curves and the X-ray light curve ($\tau_{\rm damp} = 8.67 \pm 2.9$ days).

sistent with the X-ray PSD, except at the highest frequency where the reprocessed UV has slightly less power.

In Figure 5.6 we also show our best fit DRW model for the reprocessed UV PSD, $\tau_{\rm damp} = 7.81 \pm 1.4$ days. Unsurprisingly, this $\tau_{\rm damp}$ is within one-σ of the best-fit damping timescale the PSD of the X-ray light curve that makes it through the boundaries to enter the simulation ($\tau_{\rm damp} = 8.67 \pm 2.9$ days). On the other hand, it is more than double the $\tau_{\rm damp} = 3.18 \pm 1.5$ days of our best-fit model to the intrinsic UV PSD at $f < 0.2$ days^{-1}. Interestingly, the damping timescale of the intrinsic UV light curve is an order of magnitude shorter than the average thermal timescale calculated in Section 5.3.1 for the Fiducial run. It is instead more similar to the thermal timescale at around ± 10 $H_{\rm g}$. The high correlation between the X-ray and reprocessed UV light curves may be because the injected X-ray light curve and

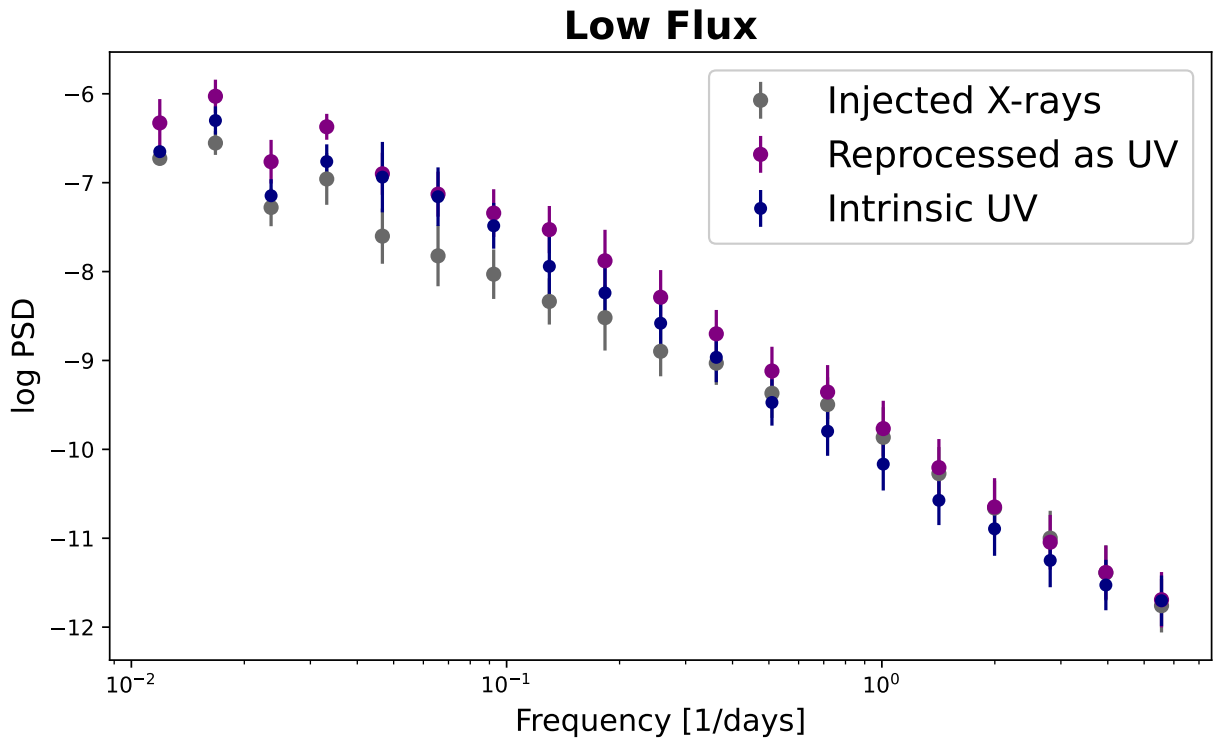


Figure 5.7: The PSDs for the injected X-rays in gray, reprocessed UV in purple, and intrinsic UV (from the one-frequency run with no injected X-rays) in navy for our run with 1/4 the injected X-ray flux as the Fiducial run. We are unable to fit a DRW or power-law to the reprocessed UV light curve PSD.

the inner UV light curve are both well-fit by a DRW with a damping timescale of around 5 days. However, the X-ray and reprocessed UV PSDs are significantly more similar than the inner UV PSD.

We fit our best-fit DRW model to the intrinsic UV PSD at $f < 0.2$ days^{-1}, because, as mentioned in Paper I, at higher frequencies the PSD of the intrinsic UV light curve steepens to $P(f) \propto f^{-2.5}$ from the DRW $P(f) \propto f^{-2}$. For the reprocessed UV PSD, the DRW model is a good fit up to the highest frequencies, because the reprocessed UV PSD never drops off significantly. A break to a steeper slope in the UV PSD has been observed in high cadence AGN light curves from Kepler (e.g., Smith et al., 2018). If as Tachibana et al. (2020) proposed this break comes from the averaging of variability over different regions of the disk, due to the fact that our simulation is highly local we would not expect to see this break.

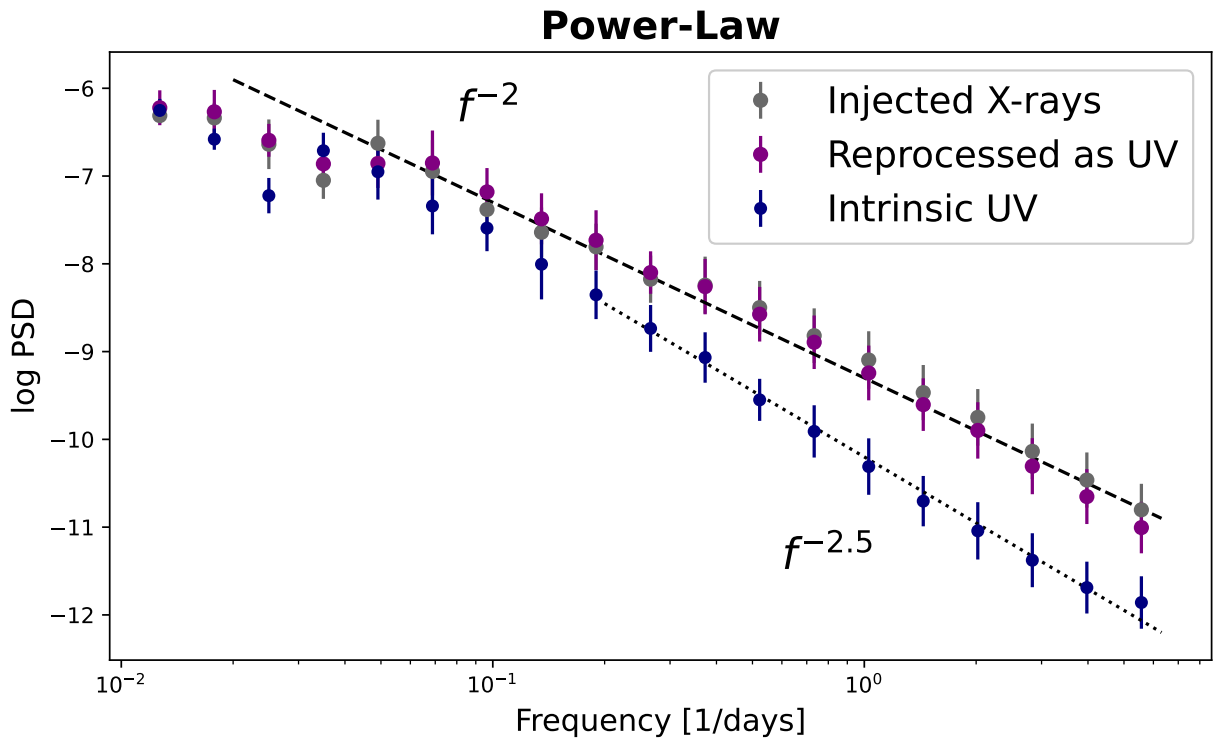


Figure 5.8: The PSDs for the injected X-rays in gray, reprocessed UV in purple, and intrinsic UV (from the one-frequency run with no injected X-rays) in navy for our Power-Law run. The dashed line shows $P(f) \propto -2$ and the dotted line shows $P(f) \propto -2.5$.

In our Low Flux run, we inject the same X-ray light curve as in our Fiducial run, but with 1/4 the X-ray flux. In this run the X-ray and reprocessed UV and the reprocessed and inner UV are both weakly correlated ($r = 0.22$). Some of the lack of correlation between the X-ray and the reprocessed UV light curve may be the result of a weak anti-correlation between the X-ray and inner UV flux, $r = -0.20$. In Figure 5.7, the PSDs of the reprocessed and intrinsic UV light curves from the runs with and without injected X-rays, respectively, are more similar than the injected X-ray and reprocessed UV PSDs. In particular both UV PSDs have more power at $f < 0.3$ days^{-1} than the injected X-ray PSD.

The correlation between the X-ray and reprocessed UV light curves for our Power-Law run, where we model the PSD of the injected X-ray light curve as a power-law that goes as f^{-2}, is nearly as high as in our Fiducial run, with $r = 0.81$. Interestingly, the inner and reprocessed UV are more correlated here than in any other run, $r = 0.26$, and more

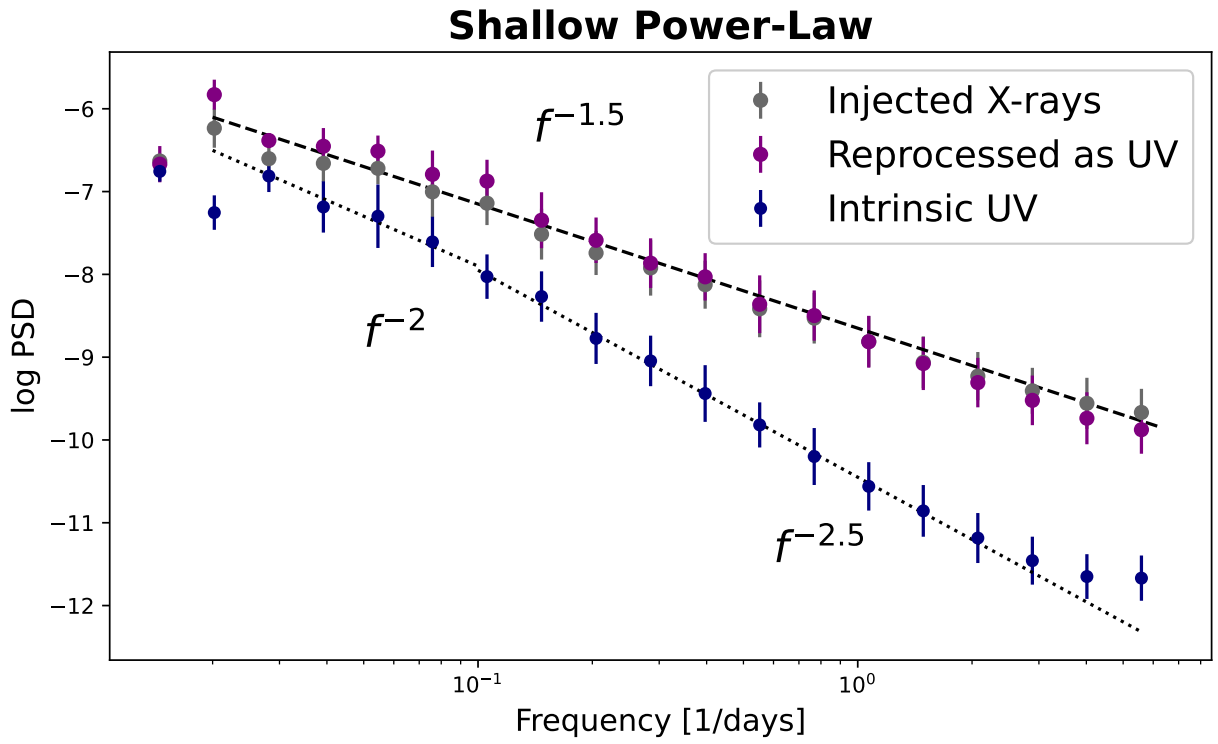


Figure 5.9: The PSDs for the injected X-rays in gray, reprocessed UV in purple, and intrinsic UV (from the one-frequency run with no injected X-rays) in navy for our Shallow Power-Law run. The dashed line shows $P(f) \propto -1.5$ and the dotted line shows $P(f) \propto -2$ at low frequencies and $P(f) \propto -2.5$ at high frequencies.

strongly correlated than the X-ray and inner UV light curves, $r = 0.10$. This UV correlation suggests that signals of the intrinsic UV variability from disk fluctuations may still exist in the emitted UV light curve, despite the X-ray and reprocessed UV light curve being strongly correlated. In Figure 5.8, the PSD of the reprocessed UV is similar to the PSD of the injected X-ray light curve, both appearing to approximately follow $P(f) \propto f^{-2}$ (dashed line) at all frequencies. On the other hand, as we found in Paper I, the intrinsic UV light curve steepens and is better fit by $P(f) \propto f^{-2}$ at $f > 0.2$ days$^{-}1$. As a result for the injected X-rays and reprocessed UV PSDs there is significantly more power at shorter timescales and similar power at longer timescales to the intrinsic UV PSD from the run without injected X-rays.

Finally, we also perform a run where we model the PSD of the injected X-ray light curve with a shallow power-law that goes as $f^{-1.5}$. The correlation between the X-rays and

Run	Fit	Parameters	Reduced-χ^2
Fiducial	Polynomial	$a = 480, b = -130, c = 9.8, \tau = 0.075$	0.11
	Gaussian	$b = 9.8, \tau = 0.033$	0.11
Low Flux	Polynomial	$a = 800, b = -180, c = 12, \tau = 0.070$	0.61
	Gaussian	$b = 12, \tau = 0.032$	0.62
Power-Law	Polynomial	$a = 540, b = -140, c = 12, \tau = 0.056$	0.17
	Gaussian	$b = 12, \tau = 0.034$	0.13
Shallow Power-Law	Polynomial	$a = 360, b = -130, c = 13, \tau = 0.10$	0.30
	Gaussian	$b = 13, \tau = 0.048$	0.22

Table 5.4: The best-fit parameters for polynomial and Gaussian modeled response functions fit to the reprocessed UV light curves from our four simulations and the reduced-χ^2 values for these models. The reduced-χ^2 values are very low because there is large variance in the reprocessed UV PSDs

reprocessed UV light curves in this run is half the correlation between the two in the Power-law run with f^{-2}. This weaker correlation may be due to the fact that the X-ray and inner UV are significantly anti-correlated with $r = -0.31$, especially because the PSDs shown in Figure 5.9 of the X-ray and reprocessed UV light curves are both better fit by a $P(f) \propto f^{-1.5}$ power-law (dashed line) than the broken power law we fit the intrinsic UV light curve with (dotted line). Overall, changing the PSD of the injected X-ray light curve has some effect on the correlations between the different light curves and the PSD of the reprocessed UV light curve, but lowering the amount of flux in the injected X-ray light curve seems to have the largest effect.

5.3.3 Response Functions

The reprocessing of X-ray emission from the corona by the AGN disk into UV-optical light is modeled using a response function, $\psi(t)$, which is typically assumed to be a Gaussian or log-normal function. Because our simulations produce evenly sampled driving X-ray and reprocessed UV light curves we can fast Fourier transform these light curves $r(t) \rightarrow R(f)$ and $d(t) \rightarrow D(f)$ and use the Fourier transform of Equation 5.1,

$$\mathrm{R(f)} = \Psi\mathrm{(f)D(f)}, \tag{5.2}$$

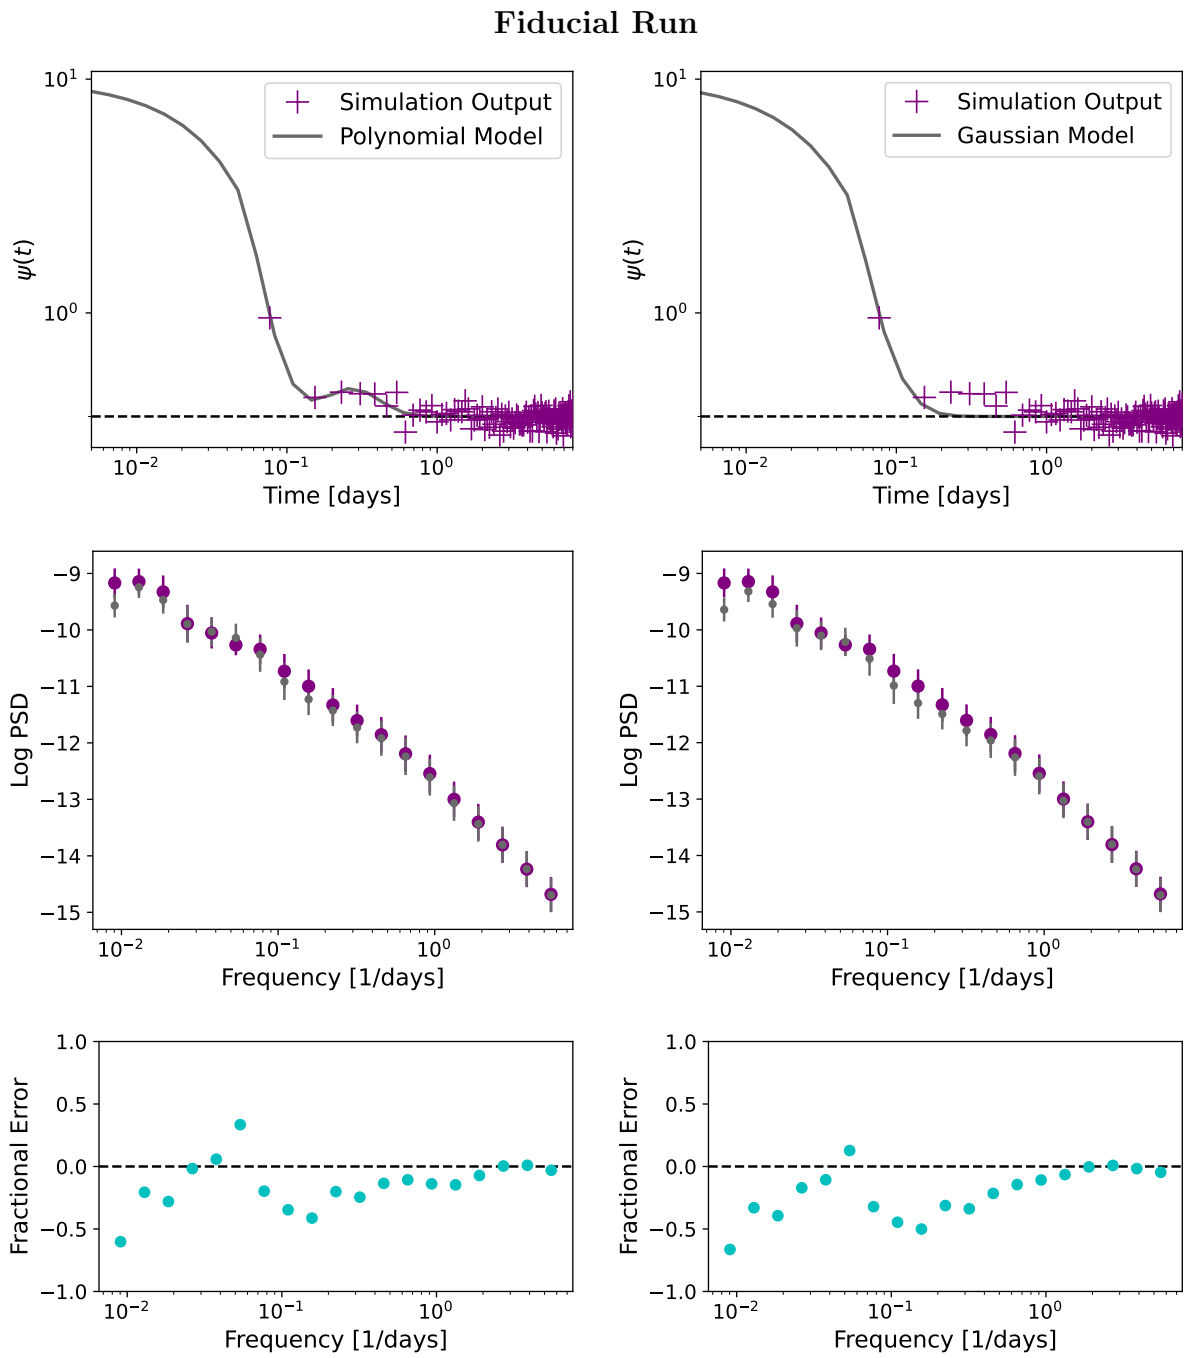


Figure 5.10: For the Fiducial run. **Top panels:** The response function in real space as a function of time, $\psi(t)$, calculated using our simulation data (in purple). The gray line shows our best-fit damped second-order polynomial (left panel) and Gaussian models (right panel) for the simulation output. **Middle panels:** The logarithm of the binned PSD for the simulation data is shown in purple, and the logarithm of the binned PSD of the X-ray light curve reprocessed with the response function shown in the upper panels in gray. **Bottom panels:** The fractional error, (modeled PSD - simulation output PSD)/simulation output PSD, as a function of frequency.

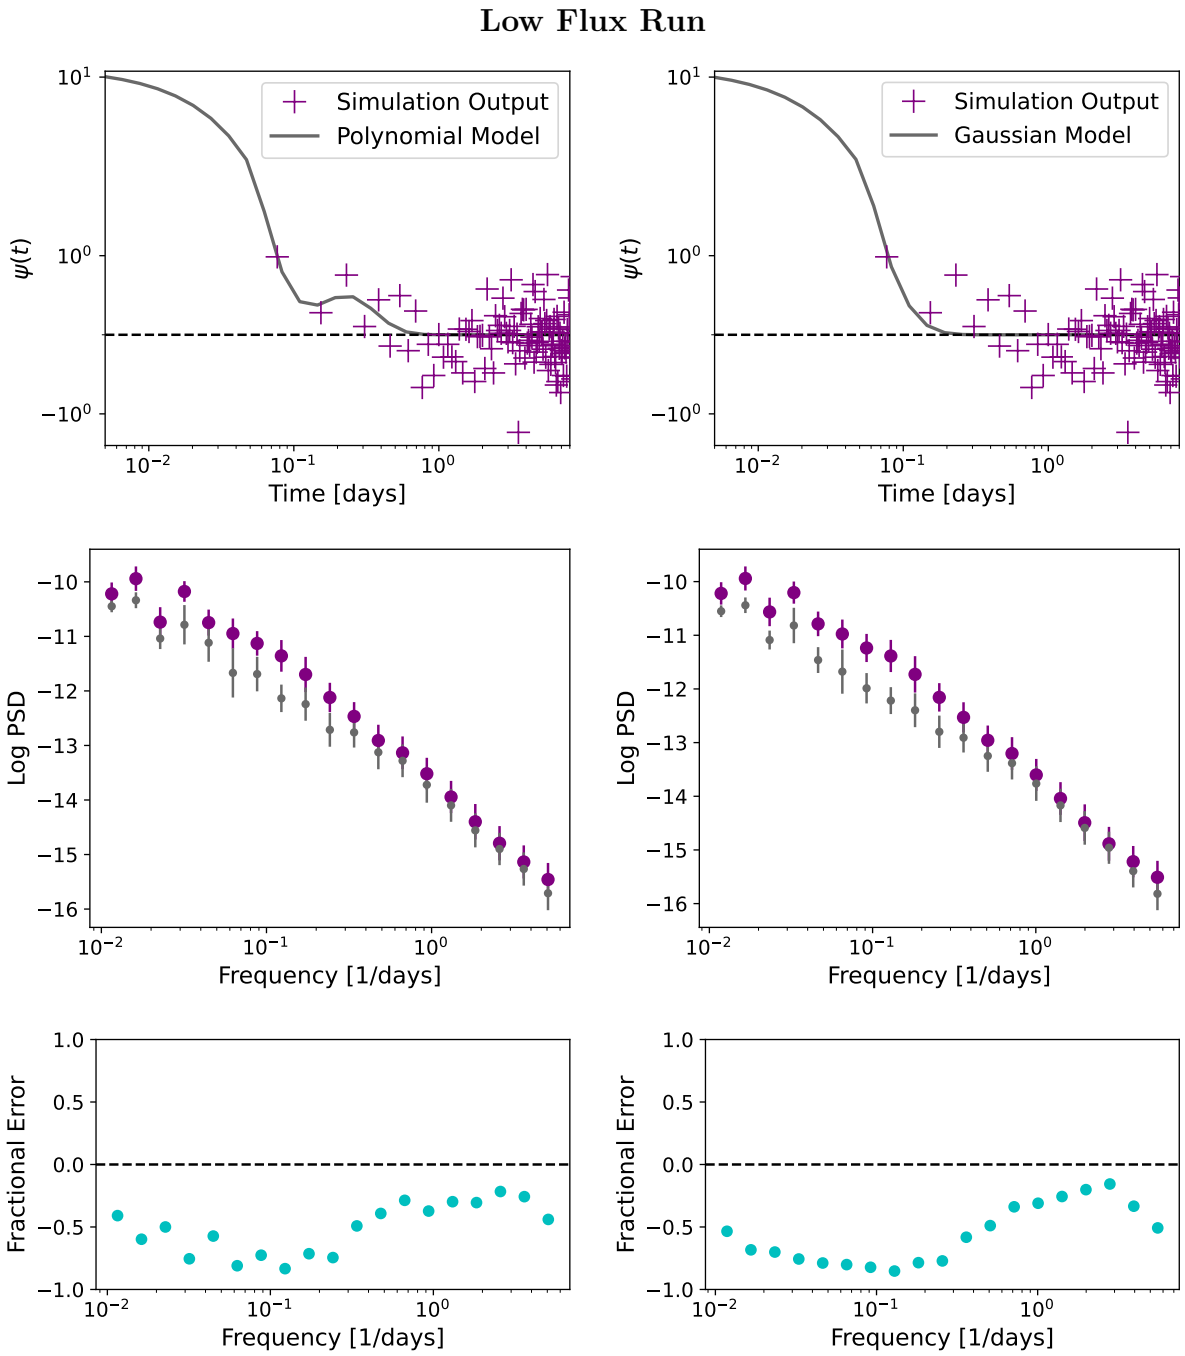


Figure 5.11: For the Low Flux run. **Top panels:** The response function in real space as a function of time, $\psi(t)$, calculated using our simulation data (in purple). The gray line shows our best-fit damped second-order polynomial (left panel) and Gaussian models (right panel) for the simulation output. **Middle panels:** The logarithm of the binned PSD for the simulation data is shown in purple, and the logarithm of the binned PSD of the X-ray light curve reprocessed with the response function shown in the upper panels in gray. **Bottom panels:** The fractional error, (modeled PSD - simulation output PSD)/simulation output PSD, as a function of frequency.

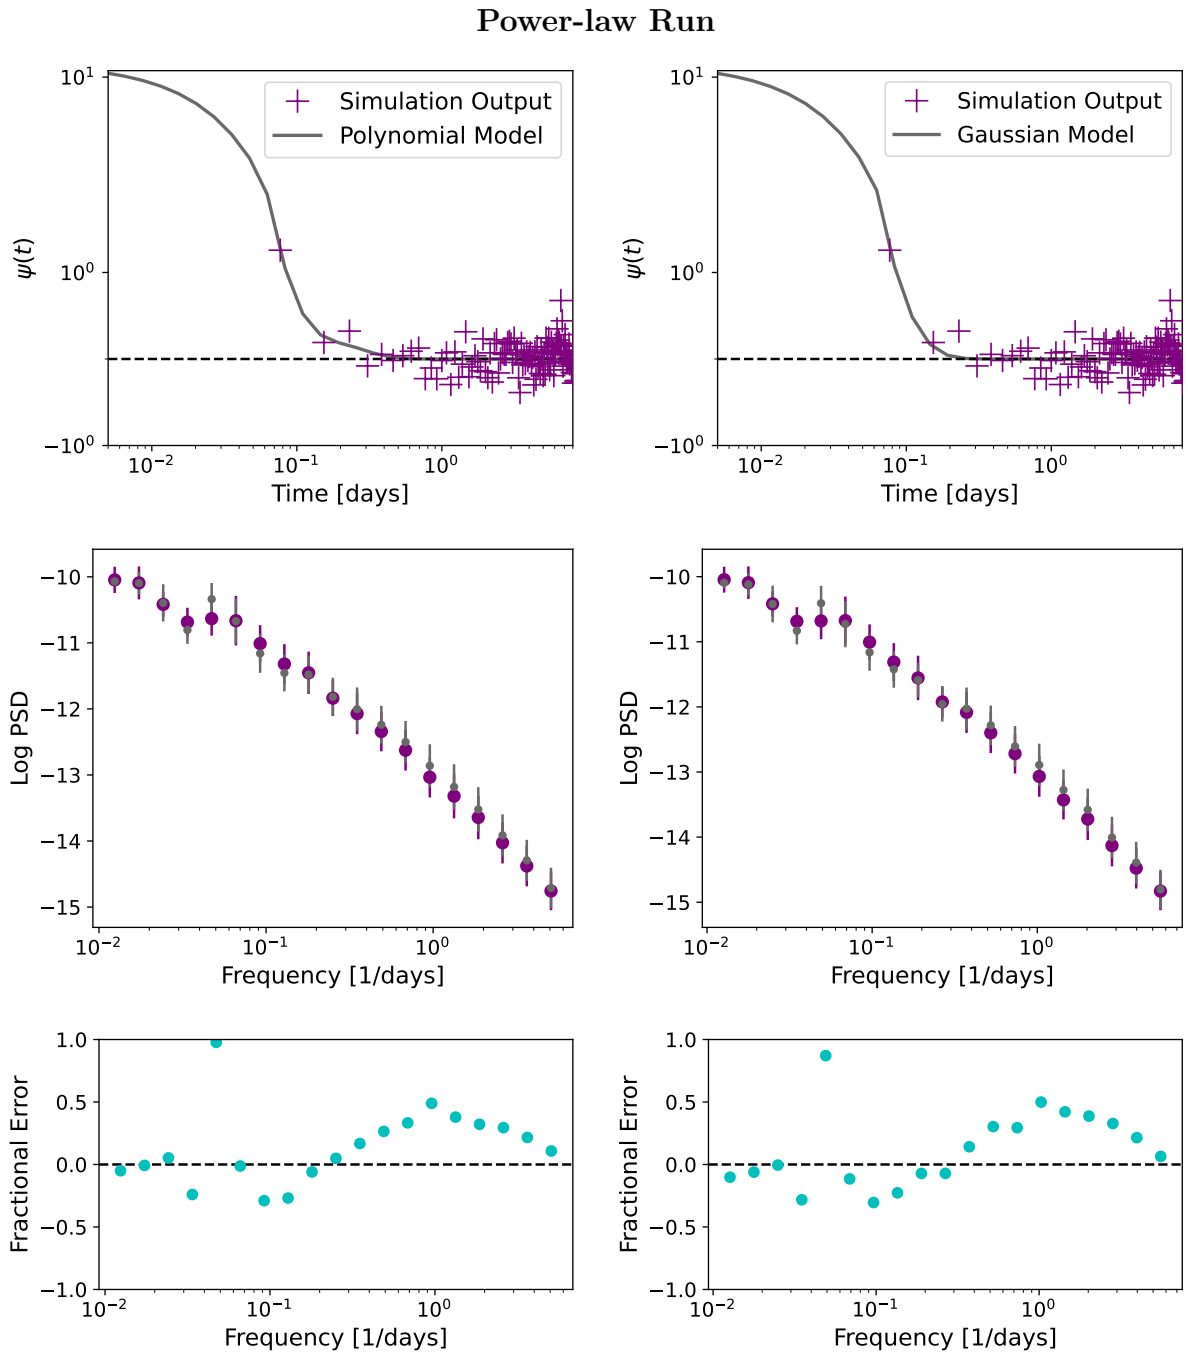


Figure 5.12: For the Power-Law run. **Top panels:** The response function in real space as a function of time, $\psi(t)$, calculated using our simulation data (in purple). The gray line shows our best-fit damped second-order polynomial (left panel) and Gaussian models (right panel) for the simulation output. **Middle panels:** The logarithm of the binned PSD for the simulation data is shown in purple, and the logarithm of the binned PSD of the X-ray light curve reprocessed with the response function shown in the upper panels in gray. **Bottom panels:** The fractional error, (modeled PSD - simulation output PSD)/simulation output PSD, as a function of frequency.

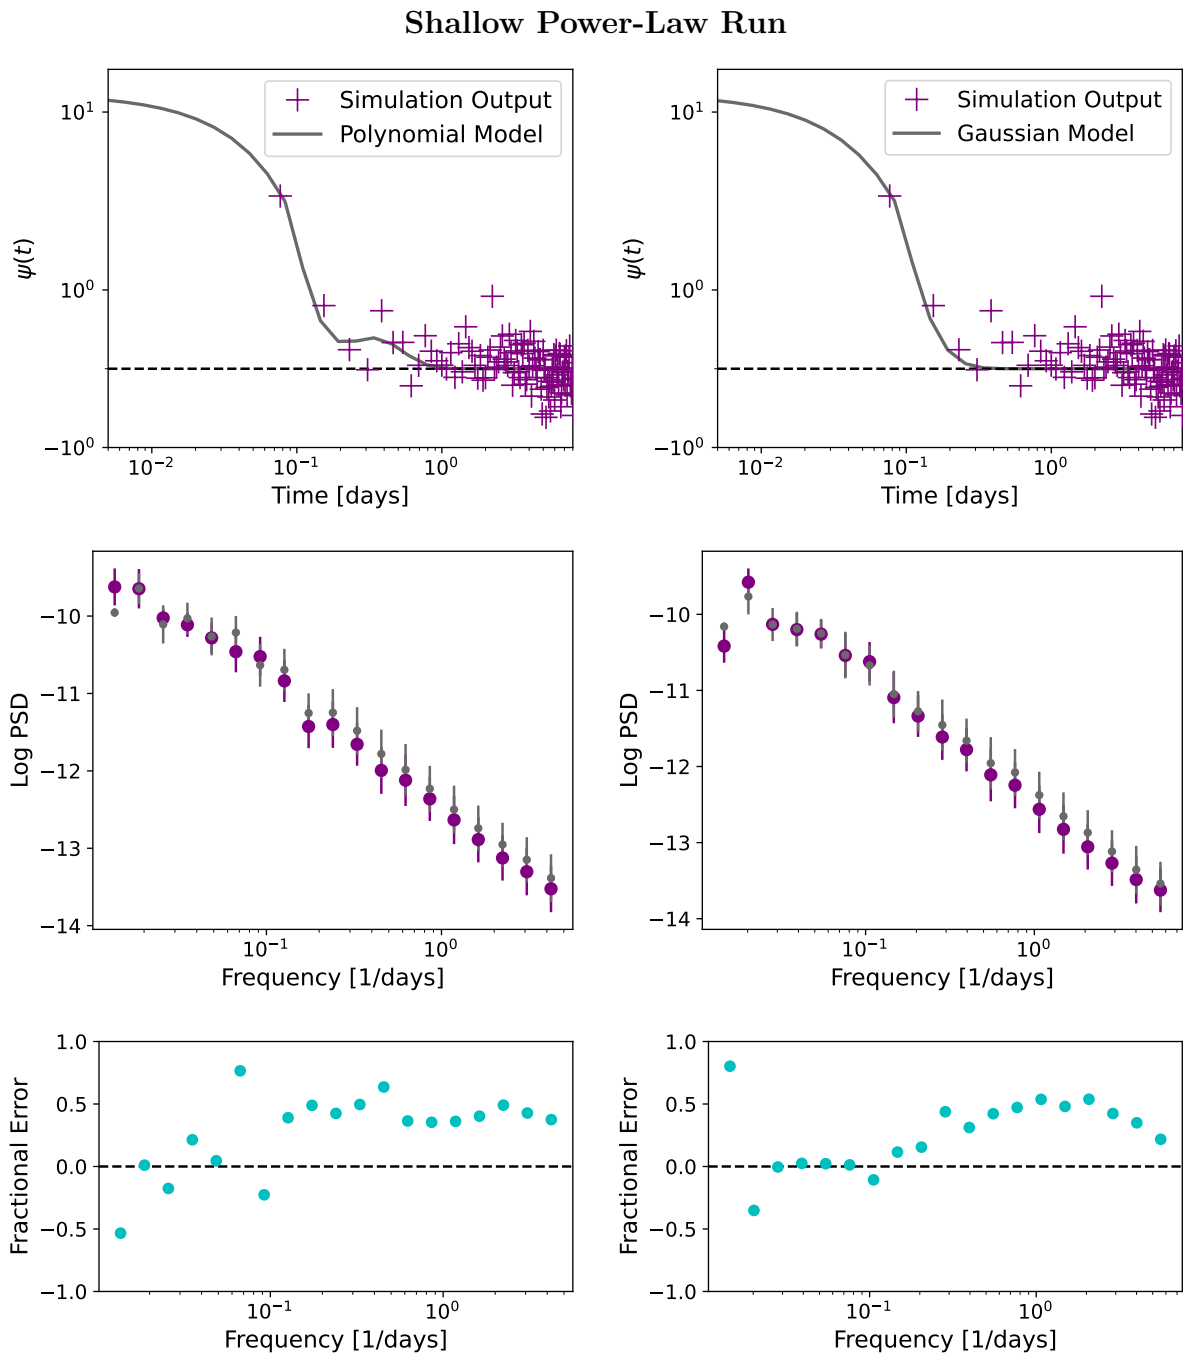


Figure 5.13: For the Shallow Power-Law run. **Top panels:** The response function in real space as a function of time, $\psi(t)$, calculated using our simulation data (in purple). The gray line shows our best-fit damped second-order polynomial (left panel) and Gaussian models (right panel) for the simulation output. **Middle panels:** The logarithm of the binned PSD for the simulation data is shown in purple, and the logarithm of the binned PSD of the X-ray light curve reprocessed with the response function shown in the upper panels in gray. **Bottom panels:** The fractional error, (modeled PSD - simulation output PSD)/simulation output PSD, as a function of frequency.

to solve for the Fourier transform of the response function, $\Psi(f)$. We can then inverse Fourier transform $\Psi(f) \rightarrow \psi(t)$ to derive the response function as a function of time in real space.

We show these response functions for all data points on timescales less than ten days in our simulations in the top panels of Figures 5.10 – 5.13 for our four different runs. We experimented with several functional forms for a model to fit to this raw data and find that the Gaussian response function that is often used is a good fit in most cases. In addition, we also find that a damped second-order polynomial is an equally good fit for all runs. We show the best-fit polynomial and Gaussian models in the left and right panels, respectively, of Figures 5.10 – 5.13.

Using the best-fit damped second-order polynomial and Gaussian models for our data we generate a forward-modeled PSD by transforming our best-fit response function and X-ray light curve to Fourier space and using Equation 5.2 to calculate a reprocessed light curve. We plot the binned PSD of this reprocessed light curve that has been calculated using our best-fit response function (in grey) and the binned PSD of the UV light curve emitted from our simulation (in purple) in the middle panels of Figures 5.10 – 5.13 for comparison. The fractional error, (modeled PSD - simulation output PSD)/(simulation output PSD) for these PSDs is shown in the bottom panels of Figures 5.10 – 5.13. We summarize the best-fit parameters and the reduced-χ^2 values for our modeled PSDs and the PSDs of the UV light curve from our simulation in Table 5.4.

For all runs, reduced-$\chi^2 < 1$ due to the large variance in the PSD of the reprocessed UV light curves we are fitting our models to. The best-fit models are reasonably similar between the different runs, with damping timescales for the polynomial fits ranging from 0.056–0.10 days and slightly shorter Gaussian widths of around 0.03–0.05 days for the Gaussian models. These widths/damping timescales are $\lesssim \Delta t = 0.075$ days, the printing timestep of our simulation. Therefore, it is difficult to determine what is driving these timescales. Currently, we have decreased the printing timestep of our simulation runs, which will help us to examine the shorter timescale structure of the response function for each run.

The best-fit response function for the Fiducial run has the lowest reduced-χ^2 value, likely because the correlation between the X-ray and reprocessed UV light curves is strongest for the Fiducial run, $r = 0.92$. The modeled PSD is roughly consistent at all frequencies with the PSD from the simulation output as can be seen in Figure 5.10. The reduced-χ^2 values are highest for our Low Flux run. In Figure 5.11 the modeled PSD is least consistent with the PSD from the simulation output at low frequencies, where the model underestimates the power in the reprocessed UV light curve. This underestimate is likely due to the fact that there is more power at low frequencies in the reprocessed UV PSD than the X-ray PSD (see Figure 5.7), which suggests that reprocessed X-ray irradiation is not the only driver of low frequency variability in the UV light curve.

Both the damped polynomial and Gaussian models are better fits for the response function for the Power-Law run than for the Shallow Power-Law run. For both power-law runs the forward modeled PSD differs most significantly from the UV PSD at high frequencies $f > 0.2$ days^{-1}. At these frequencies, the forward modeled PSD has more power, suggesting that the model we have fit does not suppress the variability from the X-ray light curves as much as it is smoothed in our simulations.

5.4 Discussion

In Paper I, we used the multi-frequency radiation MHD code described in Jiang (2022) to simulate the reprocessing of high-frequency X-rays into lower-frequency UV light, a key component of disk continuum reverberation mapping, one of the primary tools used to study AGN disks. We found that the X-ray and UV light curves were not strongly correlated due to very low X-ray absorption opacities, making it apparent that it is necessary to model not just the UV opacities using sophisticated opacity models, but also the X-ray opacity. In order to investigate further, in this paper we use tabulated opacity models for our X-ray radiation. We find that depending on the amount of flux and the PSD of the variability of

the injected X-ray light curve, the X-ray irradiation and reprocessed UV light curve can be highly correlated. In the near future, it may be possible to observe line emission using X-ray spectra, which would allow us to further test these X-ray opacity models.

Our Fiducial run here is identical to run A in Paper I, except for the updated opacity model for the X-ray radiation. Instead of the weak correlation of $r = 0.34$ found in Paper I, we find a strong correlation of $r = 0.92$ between the X-ray and UV light curve. This stronger correlation is due to the greater absorption opacity for the X-ray photons which can allow for an energy transfer from the X-ray irradiation to the gas, instead of only the momentum transfer we observed in Paper I. On the other hand, there is nearly no correlation between the inner UV flux below where the X-rays penetrate at $z = \pm 30\ H_{\rm g}$ and the reprocessed UV emitted at the surface of the box.

When we fit a DRW model to the PSD of the X-ray light curve entering the simulation domain and the reprocessed UV light curve leaving the simulation domain, we find $\tau_{\rm damp} \approx$ 8 days for both PSDs. Meanwhile, $\tau_{\rm damp} \approx 3$ days for a DRW model fit to the PSD of the intrinsic UV light curve from the run without injected X-rays in Paper I. In general, the intrinsic UV PSD from the simulation without injected X-rays and the reprocessed UV PSD from the Fiducial run differ significantly, while the reprocessed UV and X-ray PSDs from the Fiducial run are very similar.

Burke et al. (2021) propose that the damping timescale of DRW-modeled UV-optical light curves emitted by AGN disks are set by the thermal timescales of the disk. The best-fit $\tau_{\rm damp} \approx 3$ days for the DRW model we fit to the intrinsic UV PSD is an order of magnitude shorter than the average thermal timescale of our simulations. Instead, it is closer to the thermal timescale at $\pm 10\ H_{\rm g}$ for reasons that are unclear at this time.

In this paper, we also vary the way we model the variability of the injected X-ray light curve using four different runs summarized in Table 5.2. Changing the injected X-ray light curve has an impact on the reprocessed UV light curve, which is highly correlated with the injected X-ray light curve only for the Fiducial and Power-Law runs. However, it is

interesting that in the weakly correlated cases, the Low Flux and Shallow Power-Law runs, the injected X-ray and inner UV light curves are weakly anti-correlated, $r = -0.20$ and $r = -0.31$, respectively. These anti-correlations may be partially responsible for the lower correlations between the X-ray and reprocessed UV light curves, in particular for the Shallow Power-Law run where the PSDs of the X-ray and reprocessed UV light curves are very similar. On the other hand, in the Low Flux case the PSDs of the X-ray and reprocessed UV light curves differ significantly at low frequencies. Therefore it seems that the amount of X-ray flux has a larger affect than the PSD of the injected X-ray light curve on how much it is the X-ray variability driving the variability in the reprocessed UV light curve, despite there being only slightly less disk heating in the Low Flux case.

The source of the weak anti-correlation between the X-ray and inner UV light curves is unknown. While the X-ray irradiation does not penetrate down to $\pm 30\ H_{\rm g}$, it is possible that through disk heating the X-ray variability is translated down to $\pm 30\ H_{\rm g}$ but ends up out of phase. However, it is also possible that these anti-correlations are purely incidental like the anti-correlations we found in Paper I. Running these simulations for longer would help to eliminate incidental correlations or anti-correlations.

For all four runs, the commonly used Gaussian response function is a good fit for our simulation data, although a damped second-order polynomial is also a good fit. For the weakly correlated Low Flux run, both models predict less power at lower frequencies, perhaps because the variability in the reprocessed UV light curve at low frequencies is primarily from intrinsic variability. Conversely, both models predict more variability at high frequencies for both power-law runs. These results indicate that a higher-order model where there is a different ratio of reprocessing at different frequencies may be a better fit to our simulation data than a linear reprocessing model.

Updating the way we model reverberation mapping is crucial to understanding the physical processes behind reprocessing and AGN light curve variability. Reverberation mapping campaigns often compare their empirical results to the Shakura & Sunyaev (1973) optically

thick but geometrically thin standard disk model. However, there is growing evidence that the standard thin disk model does not fully describe an AGN disk (e.g., Antonucci, 1988; Antonucci et al., 1989; Antonucci, 2023). For example, both recent simulations and observations of AGN disks suggest that pressure support from magnetic fields and radiation can lead to puffed up or slim instead of thin disks (e.g., Gaburov et al., 2012; Jiang et al., 2016, 2019; Jiang & Blaes, 2020; Hopkins et al., 2023; Yao et al., 2023; Secunda et al., 2023). In addition, both micro-lensing and reverberation mapping campaigns suggest that the radial extent of the disk is larger than predicted by the standard disk model (e.g., Morgan et al., 2010; Jiang et al., 2017; Guo et al., 2022). A better understanding of the source of AGN light curve variability and reprocessing will allow us to develop new models for AGN disks that go beyond the standard thin disk model.

5.4.1 Next Steps

The results from our Fiducial run are very promising, but work remains to fully resolve the picture provided by the various runs presented here. Our next steps are as follows:

1. Allow the Low Flux, Power-Law, and Shallow Power-Law simulations to continue running until they reach the roughly 800 day time-span that we have for the Fiducial run. A longer run time will allow for better statistics, such as a better DRW or power-law model fitting to the PSDs of the light curves (Figure 5.6–Figure 5.9), less incidental correlations and anti-correlations (Table 5.3), and a more accurate fitting of a response function model (Figure 5.10–Figure 5.13).

2. Run our simulations with a shorter printing cadence so that we can get additional short timescale data which seems very important to more robustly modeling the response function (Figure 5.10–Figure 5.13).

3. Take a closer look at whether a log-normal response function is a good fit to our simulation output in order to make a statement about whether our simulations agree that this very commonly used response function is a good model for X-ray reprocessing.

4. Study the impact of the absorption of X-ray irradiation on the AGN disk in greater detail in Section 5.3.1, such as seeing how the effective temperature of the disk is changed by heating due to this X-ray absorption and why the damping timescale of our light curves appears to be shorter than the average thermal timescale.

Chapter 6

Finding Long Negative Lags in Light Curves from the Zwicky Transient Facility

6.1 Introduction

Despite being among the most luminous objects in the Universe, being an important laboratory for studying accretion physics, and having a large impact on galaxy evolution, active galactic nuclei (AGN) and the accretion disks that surround them and fuel them are still not well understood. A big reason for this lack of understanding is that with the rare exception of two nearby supermassive black holes that have recently been imaged by the Event Horizon Telescope Collaboration (2019) and Gravity Collaboration (2018), it is generally not possible to resolve the microarcsecond scales of an AGN's event horizon, accretion disk, and broadline region. However, there is hope to improve our understanding of these important objects because AGN light curves vary allowing us to substitute temporal resolution for spatial resolution.

One technique that takes advantage of this temporal resolution is reverberation mapping, which allows us to map out the area surrounding the supermassive black hole by observing the echoing of light at different wavelengths emitted from different regions. There are several types of reverberation mapping. The first, is broad line reverberation mapping, first proposed by Blandford & McKee (1982), which uses the time lag between continuum light curve variability and variability in broad emission lines to measure the distance from the accretion disk to the broad line region (see also Peterson et al., 2004; Peterson, 2014; Kaspi et al., 2000; Bentz & Katz, 2015; Grier et al., 2017). There is also traditional disk continuum reverberation mapping, first introduced by Collier et al. (1999), which uses lags between the variability in continuum light curves emitted in different wavebands by different temperature regions of the disk to measure the radial temperature profile and radial extent of the disk (e.g., Sergeev et al., 2005; Cackett et al., 2007, 2018, 2021; De Rosa et al., 2015; Edelson et al., 2015, 2017, 2019; Jiang et al., 2017; Fausnaugh et al., 2016; Starkey et al., 2017; Homayouni et al., 2022). These lags occur on the light crossing timescale of a few days as, according to the lamp post model, time-varying X-rays move outwards along the disk to cooler temperature regions. As the time-varying X-rays move outwards they are absorbed and re-emitted based on the local blackbody temperature resulting in variability in shorter waveband light curves emitted by the inner hotter regions leading variability in longer waveband light curves emitted by the cooler outer regions. We refer to these lags as short lags, here.

A third newer type of reverberation mapping, measures the long negative lag due to the propagation of fluctuations in AGN disks from outer cooler regions to inner hotter regions of the disk on the viscous timescale (Yao et al., 2023; Secunda et al., 2023). These fluctuations, which occur on the thermal timescale, could potentially originate from magneto-rotational instability turbulence (Balbus & Hawley, 1991) or convection due to enhanced opacity in the optical/UV region of the disk (Jiang & Blaes, 2020). Simulations such as those in Jiang & Blaes (2020) and Secunda et al. (2024) show that disk fluctuations can have an impact on the variability of UV-optical light curves emitted from the disk and may help explain

sources of optical-UV variability in AGN light curves that are not present in X-ray light curves (e.g., Arévalo et al., 2009; Neustadt & Kochanek, 2022; Kara et al., 2023; Burke et al., 2021) and account for the observed lack of correlation between X-ray and UV-optical light curves for several AGN (e.g., Edelson et al., 2019; Cackett et al., 2023). As these turbulent perturbations are accreted inwards on the viscous timescale, they will modify the variability of light curves in inner regions of the disk (Lyubarskii, 1997), producing this long negative lag.

We call these lags negative because the fluctuations propagate inwards, causing the variability in long wavebands to lead the variability in short wavebands, which is the opposite of the lags in traditional disk continuum reverberation mapping. These lags are also called a long lag, because they occur on the viscous or inflow timescale, as opposed to the much shorter light crossing timescale of the traditional disk continuum reverberation lags. The Shakura & Sunyaev (1973) optically thick and geometrically thin standard disk model most commonly used to model AGN accretion disks predicts that the viscous timescale would be (Davis & Tchekhovskoy, 2020),

$$t_{\rm in} = \frac{1}{\alpha\Omega}\frac{r^2}{h^2} = 260 \text{ year} \frac{\rm M}{10^8 \rm M_\odot}\frac{0.1}{\alpha}\left(\frac{\rm h}{0.01\rm r}\right)^{-2}\left(\frac{\rm r}{30\rm r_g}\right)^{3/2} \tag{6.1}$$

where α is the Shakura & Sunyaev (1973) stress parameter, $\Omega = (GM/r^3)^{1/2}$ is the orbital frequency, M is the mass of the supermassive black hole, G is the gravitational constant, h is the scale height of the disk, r is the radial position in the disk, and $r_{\rm g}$ is gravitational radii. A lag on this 260 year viscous timescale predicted by the standard disk model is far too long for us to possibly detect.

However, there is recent evidence that the standard thin disk model may not be entirely accurate for AGN disks. For example recent reverberation mapping campaigns and micro-lensing observations have found that the radial extent of the AGN disk is larger than predicted by the standard thin disk model (e.g., Morgan et al., 2010; Jiang et al., 2017; Guo

et al., 2022). Furthermore, recent simulations suggest that pressure support from magnetic fields and radiation can lead to puffed up, or slim not thin, disks (e.g., Jiang et al., 2016, 2019; Jiang & Blaes, 2020; Hopkins et al., 2023). Because the viscous timescale goes as the inverse of the aspect ratio squared, increasing the aspect ratio by an order of magnitude could lead to an inflow timescale on the order of a hundred days instead of a hundred years.

There is also growing evidence for the presence of propagating temperature fluctuations on hundred day timescales in AGN disks from both photometry (Neustadt & Kochanek, 2022; Neustadt et al., 2023) and spectroscopy (Stone & Shen, 2023). Yao et al. (2023) were able to explicitly measure a long negative lag between the *W2*-band and multiple longer wavelength UV and optical wavebands out to the z-band in Fairall 9 light curves from Hernández Santisteban et al. (2020) using several standard methods used in traditional continuum reverberation mapping. They found that the long lag between the *W2*-band and other wavebands increased rapidly as a function of wavelength in the UV, but remained nearly constant between the *W2*-band and the different optical wavebands for a maximum duration of ~ -70 days. Secunda et al. (2023) used these measurements to derive a model for the scale height of the disk as a function of disk radius, because according to the standard thin disk model, the viscous timescale depends on the height of the disk. Their model predicted that not only was the disk thick throughout, the aspect ratio of the disk also increased more than quadratically with radius.

In order to determine if this dramatic disk flaring is standard for AGN disks it is crucial to follow this observation with more detections of long negative lags. As laid out in Secunda et al. (2023), where they tested various lag detection methods on mock Vera C. Rubin Legacy Survey of Space and Time (LSST Ivezić et al., 2019) light curves, these long negative lags are difficult to detect. First, it is expected that the signal from short lags, will be much stronger, making it more difficult to detect a second weaker signal. Second, ground-based optical photometry will have large seasonal gaps of around -180 days which are comparable to the predicted length of long negative lags. Finally, because of the long length of these lags,

it is necessary to use multi-year baseline light curves that still have a high enough cadence to make the variability apparent.

An ideal survey to use in our search for long negative lags then, is the Zwicky Transient Facility survey (ZTF, Bellm et al., 2019; Graham et al., 2019; Dekany et al., 2020). ZTF is a multi-waveband (g-, r-, i-bands), optical, time-domain survey with a very wide ~ 47 sq degrees field of view. This wide field of view allows ZTF to scan the entire Northern sky every two days producing very high cadence light curves. Because ZTF has been operating since 2018 it now has AGN light curves with almost a five year baseline. Here, we start with a parent sample of 1568 light curves with a nearly three-year baseline from Ward et al. (2021), and present the seven most promising long negative lag detections. We then test the robustness of these detections with a large suite of mock ZTF light curves and discuss some of the pitfalls of our methods. Finally, we compare our models from Fairall 9 to the results presented here.

6.2 Methods

6.2.1 Lag Recovery Methods

In this chapter we use several different lag detections methods. The first two are JAVELIN (Zu et al., 2011, 2016) and the maximum-likelihood Fourier method (Uttley et al., 2014; Zoghbi et al., 2013; Miller et al., 2010), which are described in detail in Sections 3.4.2 and 3.4.3 of Chapter 3, respectively. JAVELIN is one of the most commonly used methods for performing continuum reverberation mapping by using Monte Carlo Markov chains (MCMCs) to simultaneously fit the DRW parameters of the light curve and the response function between different light curves. We again use the photometric method described in Zu et al. (2016) and model the response function as a top hat with a finite width centered at the time lag. In the following sections we use both the median of the negative part of the JAVELIN distribution and the 10-day width bin with the highest number of MCMC realizations (or

mode) in the negative part of the JAVELIN distribution to describe the long negative lag detected by JAVELIN. In Chapter 3 we found that in general JAVELIN is the most accurate method at recovering long negative lags in mock LSST light curves.

Nonetheless Fourier methods are still very useful because they allow us to recover the lags as a function of frequency, which may make it easier to find the long lag when a short lag is also present. We also found in Chapter 3 that Fourier methods are far less likely to suffer from aliasing of the seasonal gaps. Fourier methods assume a linear reprocessing model where,

$$r(t) = \int_{-\infty}^{+\infty} \psi(\tau) d(t - \tau(d\tau), \tag{6.2}$$

where t is time, τ is the lag, $r(t)$ is the reprocessed light curve, $\psi(\tau)$ is the response function, and $d(t)$ is the driving light curve. In Fourier space this becomes,

$$R(\nu) = \Psi(\nu) D(\nu), \tag{6.3}$$

where $D(\nu)$, $R(\nu)$, and $\Psi(\nu)$ are the Fourier transforms of the driving light curve, reprocessed light curve, and response function, respectively. Using this equation, Fourier methods determine the lags as a function of frequency by deriving the complex cross-spectrum of the two light curves,

$$C(\nu) = D * (\nu) R(\nu) = |D(\nu)|^2 \Psi(\nu), \tag{6.4}$$

The phase lag is then the phase, ϕ, of the the complex cross spectrum, and the time lag can be calculated as,

$$\tau = \phi / 2\pi\nu. \tag{6.5}$$

The maximum-likelihood Fourier method described in Zoghbi et al. (2013) and first presented by Miller et al. (2010), deals with uneven cadence light curves by binning the data in frequency space and maximizing the log-likelihood of first the autocovariance then the cross-covariance to find the best fit parameters for the power spectral density (PSD) and

phase lag, respectively. Here we then sue these best fit parameters as input for an MCMC which gives us the distribution of the phase lag in each frequency bin. For each bin we use the median of the distribution and take the standard deviation as the error and then calculate the lag in each bin using Equation 6.5.

In addition to the maximum-likelihood Fourier method, we also calculate the phase lag using the phase of the complex cross-spectrum by using the Fourier method on light curves that we linearly interpolate to a nightly cadence. It is common to use a linear interpolation in order to calculate the cross-correlation lags between different waveband light curves, with a method known as the interpolated cross-correlation method (Gaskell & Peterson, 1987). However, linearly interpolating light curves in order to use the Fourier method is not as common as the maximum-likelihood Fourier method.

Another, more popular method to handle the uneven cadence of AGN light curves uses Gaussian Processes to model the variability of AGN light curves in order to draw evenly sampled realizations of these light curves (e.g., Wilkins, 2019; Lewin et al., 2022). With this method, frequency-resolved lags are computed from thousands of Gaussian process-modeled light curve realizations from which a final lag distribution is determined. Preliminary tests show that using this Gaussian process method does not produce significantly better results than doing a simple linear interpolation and there are two main reasons why. The first reason is that we are already working with light curves that have been averaged to a nightly cadence, as described in the previous section. The second reason is that there are over 100 day long seasonal gaps in our data, which long exceeds the typically assumed timescales used in the kernel functions for these Gaussian processes. As a result the linearly interpolated and Gaussian process modeled variability in these gaps is non-existent. Nonetheless, future work should test whether Gaussian-process regression leads to more accurate results.

A key feature in the Fourier method is that for a symmetric response function there is a wrapping frequency at $\nu_w \approx 1/2\tau$ (see Chapter 3, Section 2.3.2 for detailed description). Below this frequency we can expect to see a signal from the long lag. Above this frequency,

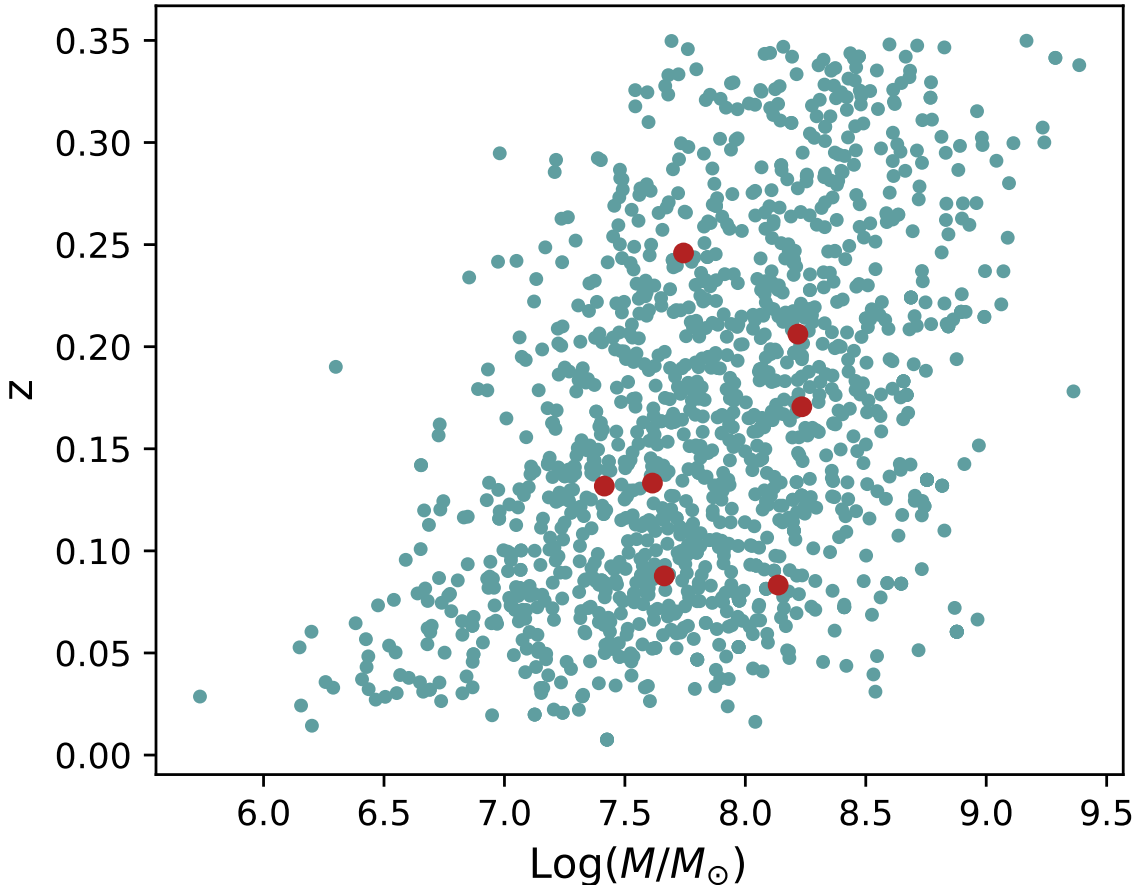


Figure 6.1: The masses and redshift of the AGN in our parent sample, when available. The red points show the masses and redshifts of our seven selected light curves. In general they tend to be of average redshift and mass.

however, we can expect to see either only the signal from the short lag or expect the lags to be consistent with zero if the short lag is not present or not detected. At frequencies near ν_w we can expect to see the lag to oscillate around zero or the short lag value. In the following sections we take the long negative lag to be the largest duration negative lag detected as we did in Chapter 2. Alternatively, we could take the error-weighted mean of all negative lags with longer durations than -10 days as in Chapter 3.

6.2.2 Light Curve Selection

Our initial sample from ZTF includes 1568 light curves from the parent sample in Ward et al. (2024). This parent sample is selected from 5000 variable AGNs identified in ZTF time-domain survey data in Ward et al. (2021) to have low redshifts, $z < 0.4$, an increase in flux of 1.5 mag in either the g- or r-band ZTF difference image photometry, and be classified as "GALAXY AGN BROADLINE" or "'QSO BROADLINE"' in the Sloan Digital Sky Survey Data Release 14 (SDSS DR14) spectroscopic catalog (Blanton et al., 2017). We show the available redshifts and mass for each of these AGN in Figure 6.1.

To make our first round of analysis more tractable, we focus on only the 5–10 light curves with the most promising recovered lags between the g- and i-band. We focus specifically on the most promising recovered lags between the g- and i-band, because as opposed to those between the between the g- and r-band, because the former probes a wider range in wavelength, whereas the latter is a very narrow change in wavelength over which it may be difficult to detect a long negative lag.

We make several cuts to our initial sample in order to end up with less than 10 light curves. First, we eliminate light curves with fewer than 1000 observations in the i-band, leaving 227 light curves. Next, we run JAVELIN and the maximum-likelihood Fourier method on the remaining 227 light curves and cut down on the number of light curves based on these initial results using several criteria. Our first criterion is that for the JAVELIN distribution there must be a strong peak where greater than 60% of the distribution falls between 30 days. The second criterion primarily has to do with eliminating false positives due to the aliasing of seasonal gaps. In Chapter 3, we show that the roughly 180 day seasonal gaps in optical light curves can lead to false detections of lags around this 180 days. While future work should attempt to find ways to eliminate this known signal, currently we avoid these false detections by only selecting out light curves for which the peak of their JAVELIN distribution, l, is either -460 days $< l < -220$ days or $l > -160$ days. We set a lower limit in the first case, because we run JAVELIN from $\pm$500 days and, as we see in our mock light curves in Section 6.3, if

a lag is not recovered it is common for the JAVELIN distribution to peak around the lower limit over which we run JAVELIN.

Our final criterion is to select out AGN for which there is a significant detection with the maximum-likelihood Fourier method. We consider a detection "significant" if the minimum lag detected is less than -50. Combined, these three criteria reduce our first subsample of 227 light curves down to a more manageable 7 light curves. We provide the ZTF names for these AGN in column (1) of Table 6.1 along with the corresponding number 1 through 7 which we use to refer to these light curves throughout the text. We also show these selected AGN in red in Figure 6.1, where they appear to span average masses and redshifts in this sample.

6.2.3 Mock Light Curve Generation

To test the robustness of the long negative lag detections for the seven ZTF light curves that meet the criteria outlined in the previous section, we test our three lag detection methods on mock light curves. For our initial analysis we only study the ability of our methods to detect long negative lags in ZTF light curves between the g- and i-band, because of the larger wavelength difference, which may be necessary to robustly detect long negative lags. Future work should study the ability of our methods to detect long negative lags between the g- and r-band, as well, and also the ability to detect short lags between the g-band and both the r- and i-band.

To generate mock ZTF light curves we follow the same procedure as in Chapter 3. First, we generate a 2000 year-long damped random walk light curve with a time sampling of $\Delta t = 0.001$ days, a damping timescale of $\tau_{\rm damp} = 170$ days, and a structure function at infinity of $\mathrm{SF}_\infty = 2.5$. We then subsample this light curve into 100 damped random walk driving, or g-band, light curves.

Next, we generate three different types of mock reprocessed, or i-band, light curves. The first is a mock light curve with only a long lag signal, the second has both a long and short

lag signal, and the third has no input lag. To create these light curves we reprocess the driving light curves with a Gaussian response as described in Section 3.3.2 of Chapter 3. We use an input long lag of −280 days, which is similar to several of the long lags recovered for the seven light curves studied in detail in Section 6.4.1. For light curves reprocessed with two lags the input short lag is +5 days. The width of each Gaussian is set to be a factor of 5 smaller than the input lag. As in Chapter 3, in cases where we have both a long and short input lag we add two response functions together and scale the relative strength of the two response functions by fixing the amplitude of the long lag response function to the fiducial fraction of the amplitude of the short lag response function in Chapter 3, 0.20.

Once we have a driving and reprocessed light curves, we subsample the light curves with the g- and i-band cadences, respectively, from the ZTF light curves. Instead of creating mocking light curves for all seven AGN, we create mocks sub-sampled with the cadences of three of these light curves by choosing the light curve with the highest number of nightly cadences in the i-band, the lowest number of nightly cadences in the i-band, and the median number of nightly cadences in the i-band, light curves 4, 6, and 3, respectively. We find in Section 6.3 that the Fourier methods perform similarly on the three light curves, but that JAVELIN performs best on the light curve with the highest nightly cadence, and worse on the light curve with the median nightly cadence. Finally, we add noise to the mock light curves using the signal to noise ratio of each measurement from the corresponding ZTF light curve.

6.3 Mock Light Curves

Due to the novelty of using JAVELIN and Fourier methods to detect long negative lags it is crucial to perform both null tests and recovery tests for all methods on mock data. Because we are primarily focused on the g- to i-band lags in this chapter, we only present mock light curves of these two wavebands. Future mock light curves are needed to better understand our results for the g- to r-band lags found in the previous section.

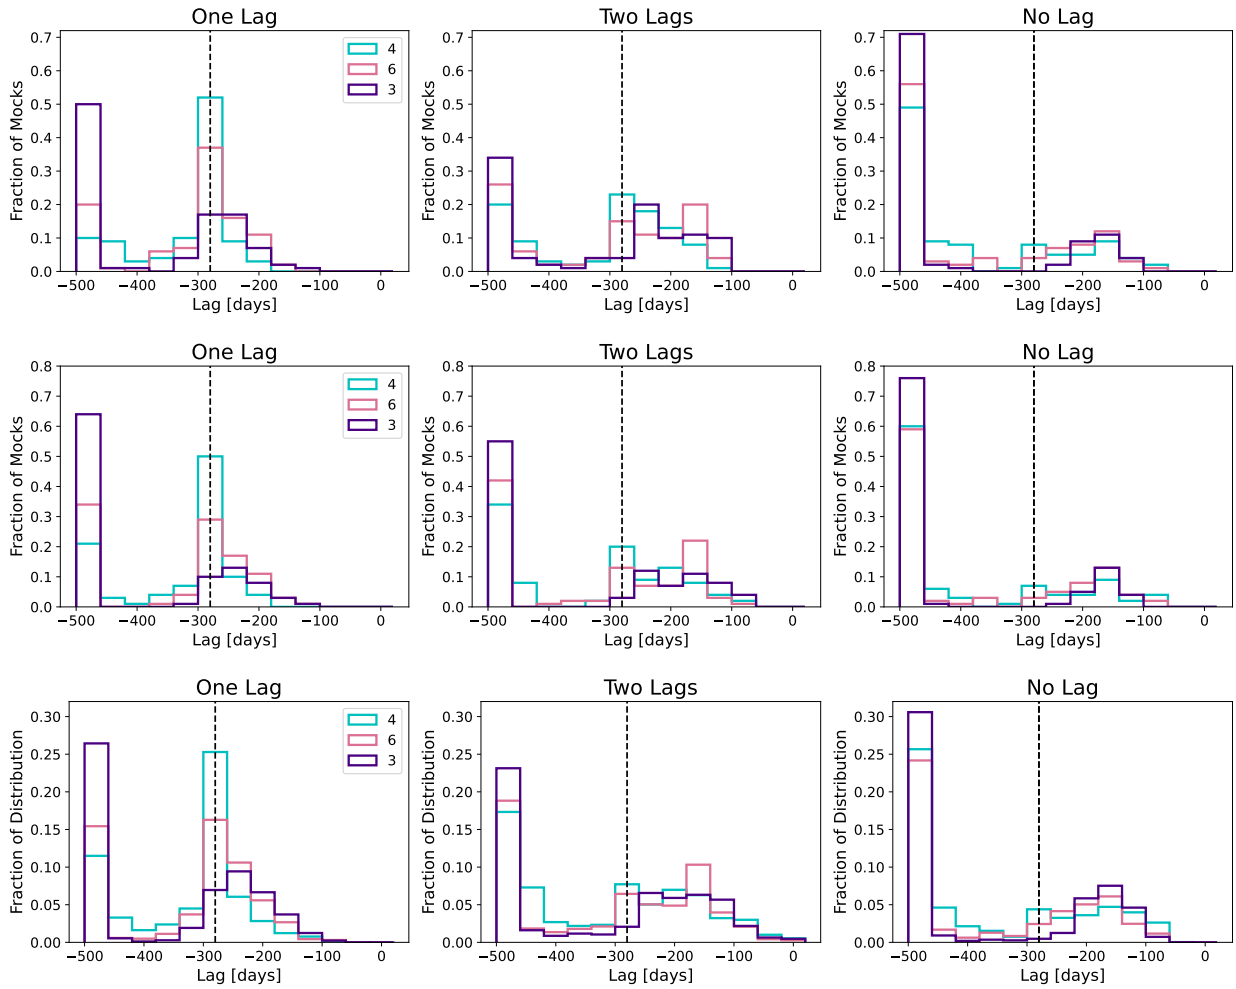


Figure 6.2: Lags recovered by JAVELIN for 100 mock light curves of three of the subsample ZTF AGN with only an input long lag (left panel), an input long and short lag (middle panel), and no input lags (right panel). In the top panel we take the recovered lag as the median of the negative part of the distribution as in Chapters 2 and 3. In the middle panel we take the recovered lag as the mode of the negative JAVELIN distribution. In the bottom panel we stack all of the distributions for each and then calculated the fraction of the distribution within each bin. The dashed vertical line shows the input long lag of -280 days.

As described in Section 6.2.3, we make three kinds of mock light curves, one with only an input long negative lag of −280 days, one with an input long negative lag of −280 days and an input short lag of +5 days, and one with no input lags. We then make 100 mock light curves sub-sampled with the cadences from three of our seven selected light curves. We intentionally choose the light curves with the most (light curve 4), least (light curve 6), and median (light curve 3) amount of nightly cadence in order to probe as wide a range of cadences as possible (see Table 6.1 for details on these light curves). Figure 6.2 shows the median and mode of the negative JAVELIN distribution in the top and middle panels, and the normalized summed negative JAVELIN distributions in the bottom panel, for all 100 mock light curves.

We show our null tests where we have not injected any lag into our mock light curves in the right panels of Figure 6.2. The most prominent peak for all distributions occurs around −500 days or towards the bottom of the range over which we run JAVELIN. Around 50–80% of the medians and modes for all three light curve cadences fall between −480 and −500 days, suggesting that any lag recovered within that range is unreliable. The second, smaller, and broader peak of the null test JAVELIN distribution occurs at around −180 days, or timescales similar to the seasonal gap. This aliasing on timescales corresponding to the length of the seasonal observing gap was also a problem for the mock LSST light curves in Chapter 3.

These null tests aid us in determining false detection zones which informed our selection criteria by only selecting light curves where the lags recovered by JAVELIN are either -460 days $< l < -220$ days or $l > -160$ days. For light curves 3 and 6, only around 10–25% of the negative JAVELIN distributions in these null tests fall outside these false detection zones with never with more than 5% of the the remaining distribution falling within a single 40-day wide bin. For light curve 4, 42% of the medians of the distributions of the 100 mock light curves fall outside of the false detection zones, but most of these medians fall just outside of the false detection zone at just short of −460 days.

For our recovery tests on mocks with input lags, JAVELIN is most accurate when there is only one input long lag. For these mocks, the distributions of the medians and modes of each light curve's JAVELIN distribution and the summed JAVELIN distributions have two peaks, one around the input long lag and one around −500 days. Because this peak near −500 days is a known problem it would be excluded and would not contaminate our results. For light curve 3, the most prominent peak is around −500 days instead of around the input long lag, with over half of the medians and modes falling within the lowest bin and over a quarter of the summed distribution falling in this bin. The next most prominent peak for light curve 3 is centered just short of the input long lag, so it is in some cases possible to detect the long negative lag with JAVELIN with some accuracy for this light curve.

When we add a short lag to our mock light curves the JAVELIN results are less accurate for all light curves, with only 20–30% of the medians and modes accurate to within 10% for light curves 4 and 6, and even fewer for light curve 3. Adding the short lag leads the peak around the input long lag to broaden significantly to include the signal from the seasonal gap of −180 days, especially for light curve 6 whose distributions now have a peak at around −180 days. This shift in the JAVELIN distributions towards zero when we add a short lag is consistent with the mock results in Chapter 3.

We next test out our two Fourier methods on our mock light curves. We show the distribution of minimum lags recovered by the maximum-likelihood Fourier method and the linear interpolation Fourier method for our various types of mock light curves in the top and bottom panel of Figure 6.3, respectively. For our null tests on mock light curves with no input lags, both Fourier methods appear to have a significant peak at around −100 days. For our recovery tests on mock light curves with only an input long lag, the distributions of lags recovered by the maximum-likelihood Fourier method do not peak around the input long lag. Instead, there is a slight peak around −50 days for light curves 4 and 6, which may be related to the peak in the distributions at just short of −100 days in our null tests. The distributions of lags recovered for light curve 3 are mostly flat. For mock light curves with an

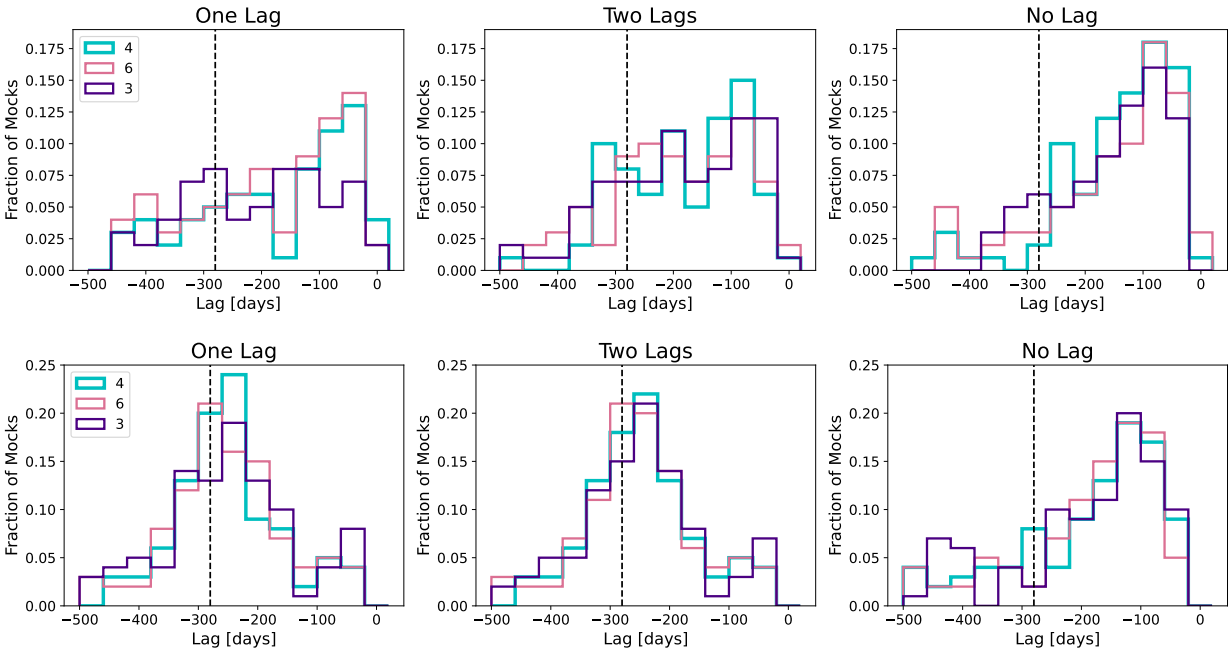


Figure 6.3: Lags recovered by the Fourier methods for 100 mock light curves of three of the subsample ZTF AGN with only an input long lag (left panel), an input long and short lag (middle panel), and no input lags (right panel). In the top panel we use the maximum-likelihood Fourier method. In the bottom panel we use our linear interpolation Fourier method. The dashed vertical line shows the input long lag of −280 days.

input long and short lag, the distribution of long lags recovered by the maximum-likelihood Fourier method also does not have any prominent peaks.

The linear interpolation Fourier method appears to perform much better than the maximum-likelihood Fourier method on our mocks with input lags. For both the one and two lag case, the most prominent peak is at the input long lag or just shorter timescales for all three cadences. In Section 6.2.2 we select the most promising lags based on the JAVELIN and maximum-likelihood Fourier results, because these are more commonly used methods. However, these mocks suggest in the future using the interpolation Fourier method is the better choice.

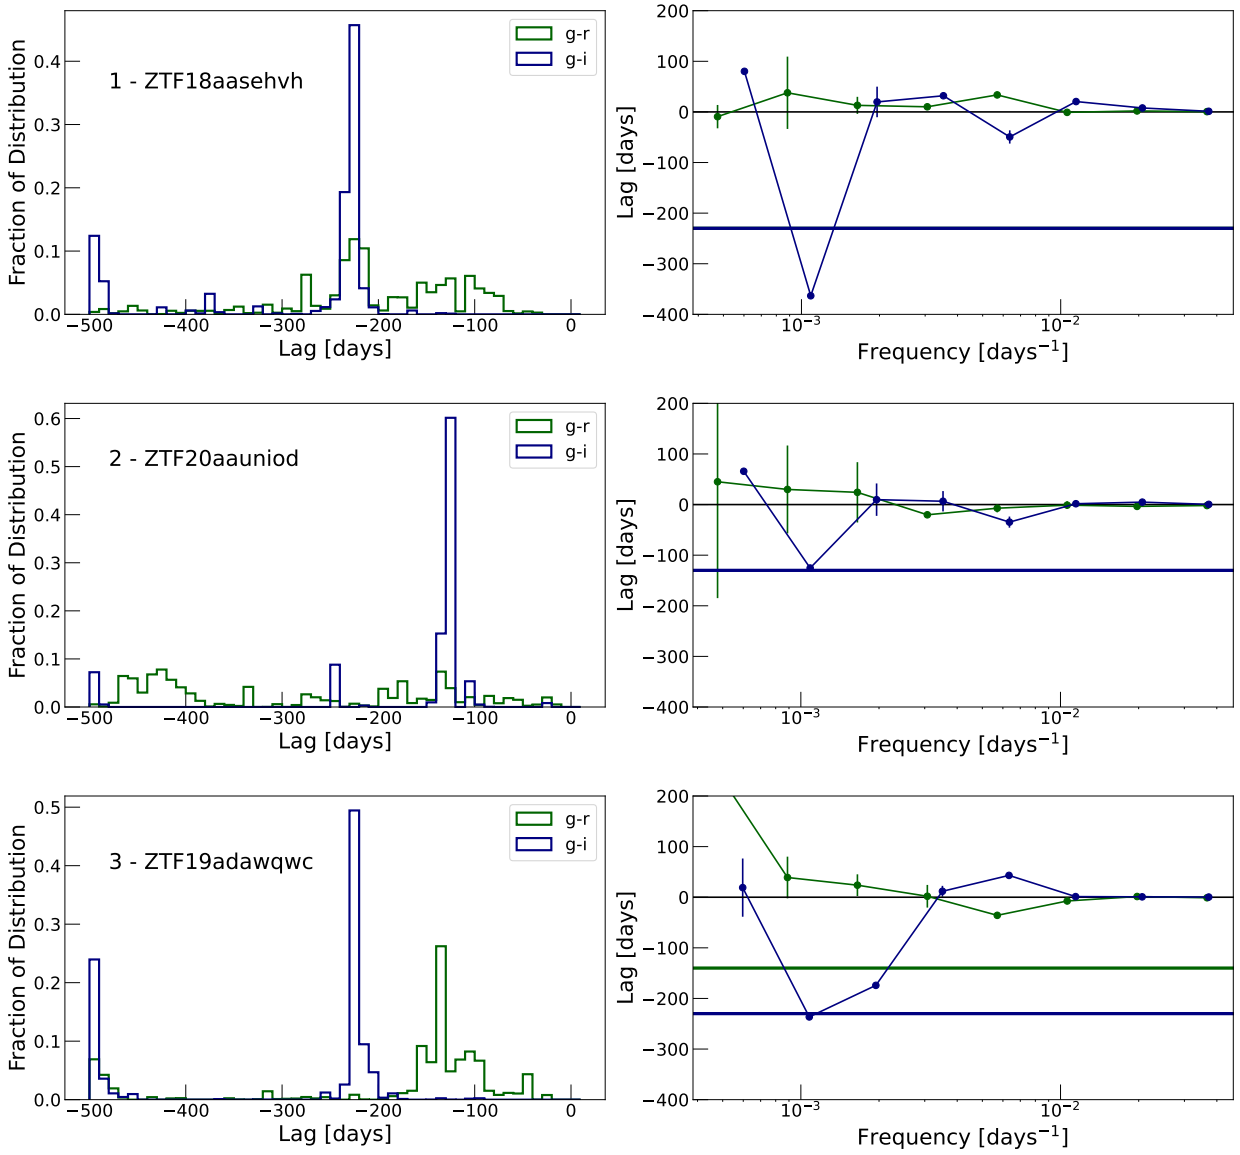


Figure 6.4: The distribution of lags recovered by JAVELIN (left panels) and the lag as a function of frequency recovered by the maximum-likelihood Fourier method (right panels) for the first three of seven ZTF light curves in our sub-sample selected based on our criteria. Lags between the g- and r-band are shown in purple, while lags between the g- and i-band are shown in blue. The horizontal lines in the right panels show half of the peak value of the corresponding JAVELIN distribution.

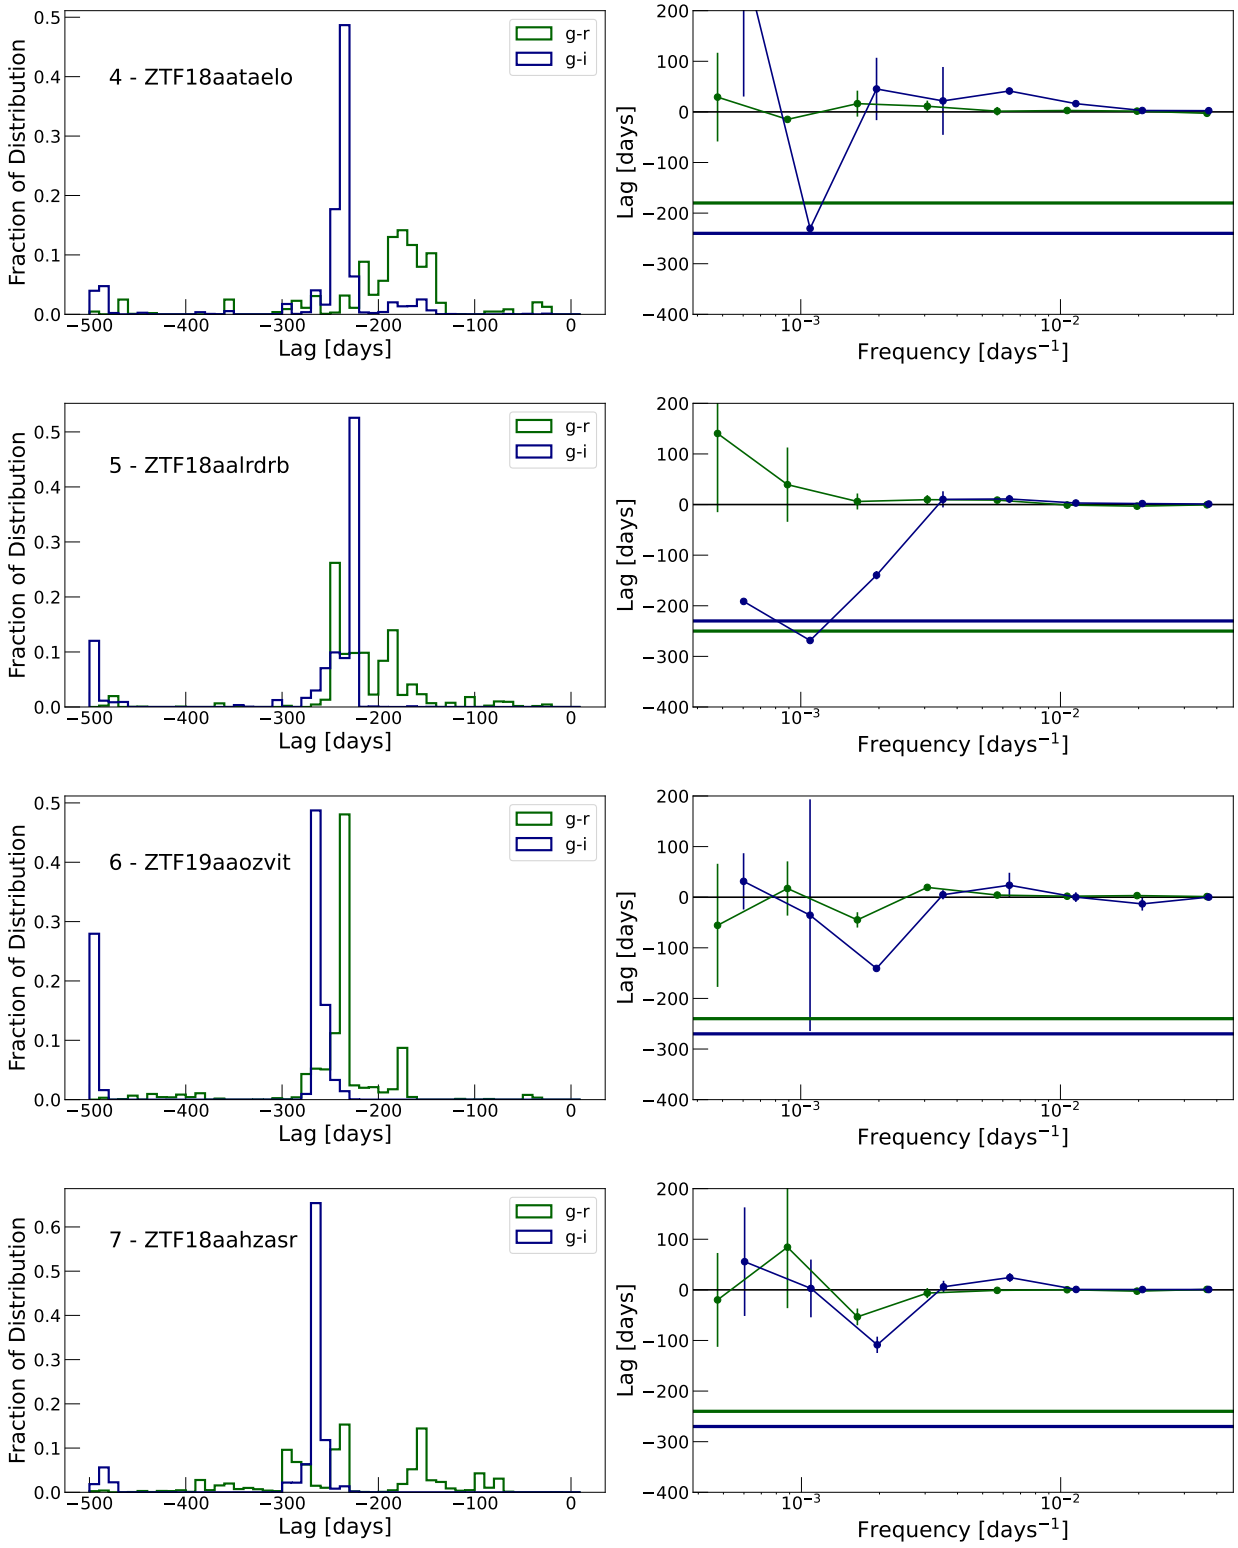


Figure 6.5: Figure 6.4 continued for the next four lags in our sub-sample.

AGN	Median	Mode	Max-Likelihood	Interpolated
1 (ZTF18aasehvh)	−229 (−214)	−230 (−230)	−363 (−9)	−299 (−63)
2 (ZTF20aauniod)	−128 (−334)	−130 (−430)	−125 (−20)	−49 (−471)
3 (ZTF19adawqwc)	−225 (−137)	−230 (−140)	−237 (−36)	−202 (−58)
4 (ZTF18aataelo)	−239 (−179)	−240 (−180)	−230 (−15)	−309 (−42)
5 (ZTF18aalrdrb)	−228 (−222)	−230 (−250)	−269 (−3)	−118 (−481)
6 (ZTF19aaozvit)	−264 (−234)	−270 (−240)	−141 (−56)	−252 (−109)
7 (ZTF18aahzasr)	−266 (−237)	−270 (−240)	−109 (−53)	−79 (−92)

Table 6.1: In each row we give the ZTF name in parentheses and our number system in the first column. We then give the median of the negative part of the JAVELIN distribution, the ten-day wide bin with the largest fraction of the negative JAVELIN distribution, the minimum lag detected with maximum-likelihood Fourier method, and the minimum lag detected with the nightly linearly-interpolated Fourier method in columns 2, 3, 4, and 5, respectively. The values in parentheses are the g-r lags and the values not in parentheses are the g-i lags.

6.4 Results

In Section 6.2.2 we describe our criteria for selecting only the most promising g- to i-band long negative lags recovered by JAVELIN and the maximum-likelihood Fourier method. Here, we present the seven light curves that meet those criteria. In Figures 6.4 and 6.5, we show the JAVELIN distributions (left panels) and the lags as a function of frequency recovered by the maximum-likelihood Fourier method (right panels) between the g- and r-band in green and between the g- and i-band in blue. Each row represents a different AGN and is labeled with both its ZTF name and a number 1 through 7 which we will use to refer to these light curves in the text. Figure 6.6 shows the lags recovered by our linear interpolation Fourier method, with the same numbering and color scheme.

The lags recovered by all of our methods for each AGN and their masses and redshifts are summarized in Table 6.1. It is apparent that with our current selection criteria we are more likely to select long negative lags between −200 and −300 days. Our mock light curves in Section 6.3 suggest that there is no reason that the cadence of ZTF light curves would bias us towards making an erroneous lag detection of this duration.

6.4.1 g- to i-Band Lags

The selection criteria in Section 6.2.2 is based solely on recovered g- to i-band lags, because these two wavebands are farther apart in wavelength than the g- and r-bands. This larger difference in wavelength makes it more likely that the g- and i-band light curves will originate from more distinct radial regions of the disk, which should make it easier to detect a long negative lag. Therefore, in this section we will begin with discussing the g- to i-band lags recovered by our various lag detection methods.

Our first criterion for selecting our sub-sampled AGN light curves, is that over 60% of the negative part of the JAVELIN distribution falls within 30 days (see Section 6.2.2). As a result, the JAVELIN distributions for all seven light curves have a prominent peak. We determine the lag recovered by JAVELIN in two different ways, first with the median of the negative distribution, and second with the mode of the negative distribution, which we determine by binning the data in bins with widths of ten days. Because these JAVELIN distributions have prominent peaks the medians and modes are accurate to within roughly 2% of the lag duration for all of these light curves. Four out of seven of these lags fall around -230 days, with two falling around -270 days. Only one light curve has a lag (-130 days) that falls at $l > -160$, or timescales shorter than those excluded because of the seasonal gap at around 180 days.

The lags as a function of frequency recovered by the maximum likelihood method are shown in the right panel of Figures 6.4 and 6.5 with the lags recovered by JAVELIN shown as a solid horizontal line. For our Fourier methods, we take the long negative lag to be the minimum lag recovered. For roughly half of the light curves there is excellent agreement between the maximum-likelihood Fourier method and JAVELIN. The lags recovered with the maximum-likelihood Fourier method for light curves 6 and 7 are are much shorter than the lags recovered by JAVELIN for these light curves, which may be due to the averaging of the long and short lag together (see Figure 2.11 in Chapter 2). Finally the lag recovered by the

maximum-likelihood Fourier method for light curve 1 does not appear to be at all consistent with the lag recovered by JAVELIN.

In addition to the maximum-likelihood Fourier method, we also use the Fourier method on light curves we have linearly interpolated to a nightly cadence. Figure 6.6 shows the lag as a function of frequency recovered using these interpolated light curves. Only the lags recovered for two of the light curves by the interpolated Fourier method are in good agreement with the lags recovered by JAVELIN, while the lags recovered by the interpolated Fourier method for three light curves are significantly shorter than the lags recovered by JAVELIN. Again these lags recovered by the interpolated Fourier method could be shorter due to the averaging of the long and short lag, but we would expect this averaging to be consistent between the two Fourier methods, which only appears to be the case for light curve 7. Therefore, while there is good agreement between lags recovered by JAVELIN and the maximum-likelihood Fourier methods, the lack of consistency between the two Fourier methods prevents there from being a coherent picture, particularly because we found the interpolated Fourier method to be far more accurate at detecting the input long lags in our mock light curves in Section 6.3.

6.4.2 *g*- to *r*-Band Lags

The previous section described the g- to i-band lags recovered by our various methods for our seven ZTF light curves. We now turn to the g- to r-band lags recovered by these same methods. We note that because the criteria we used to select our sub-sample of ZTF light curves was solely based on the g- to i-band lags recovered, we do not expect these g- to r-band lags to be as robust as the lags in the previous section. Furthermore, physically it is unclear if the narrow wavelength range from the g- to r-band is as reliable for recovering long negative lags as the wider g- to i-band range. Nonetheless recovering a similar, or shorter duration lag between the g- and r-band could be a useful confirmation of the lags recovered between the g- and i-band, especially since we are not selecting our subsample of

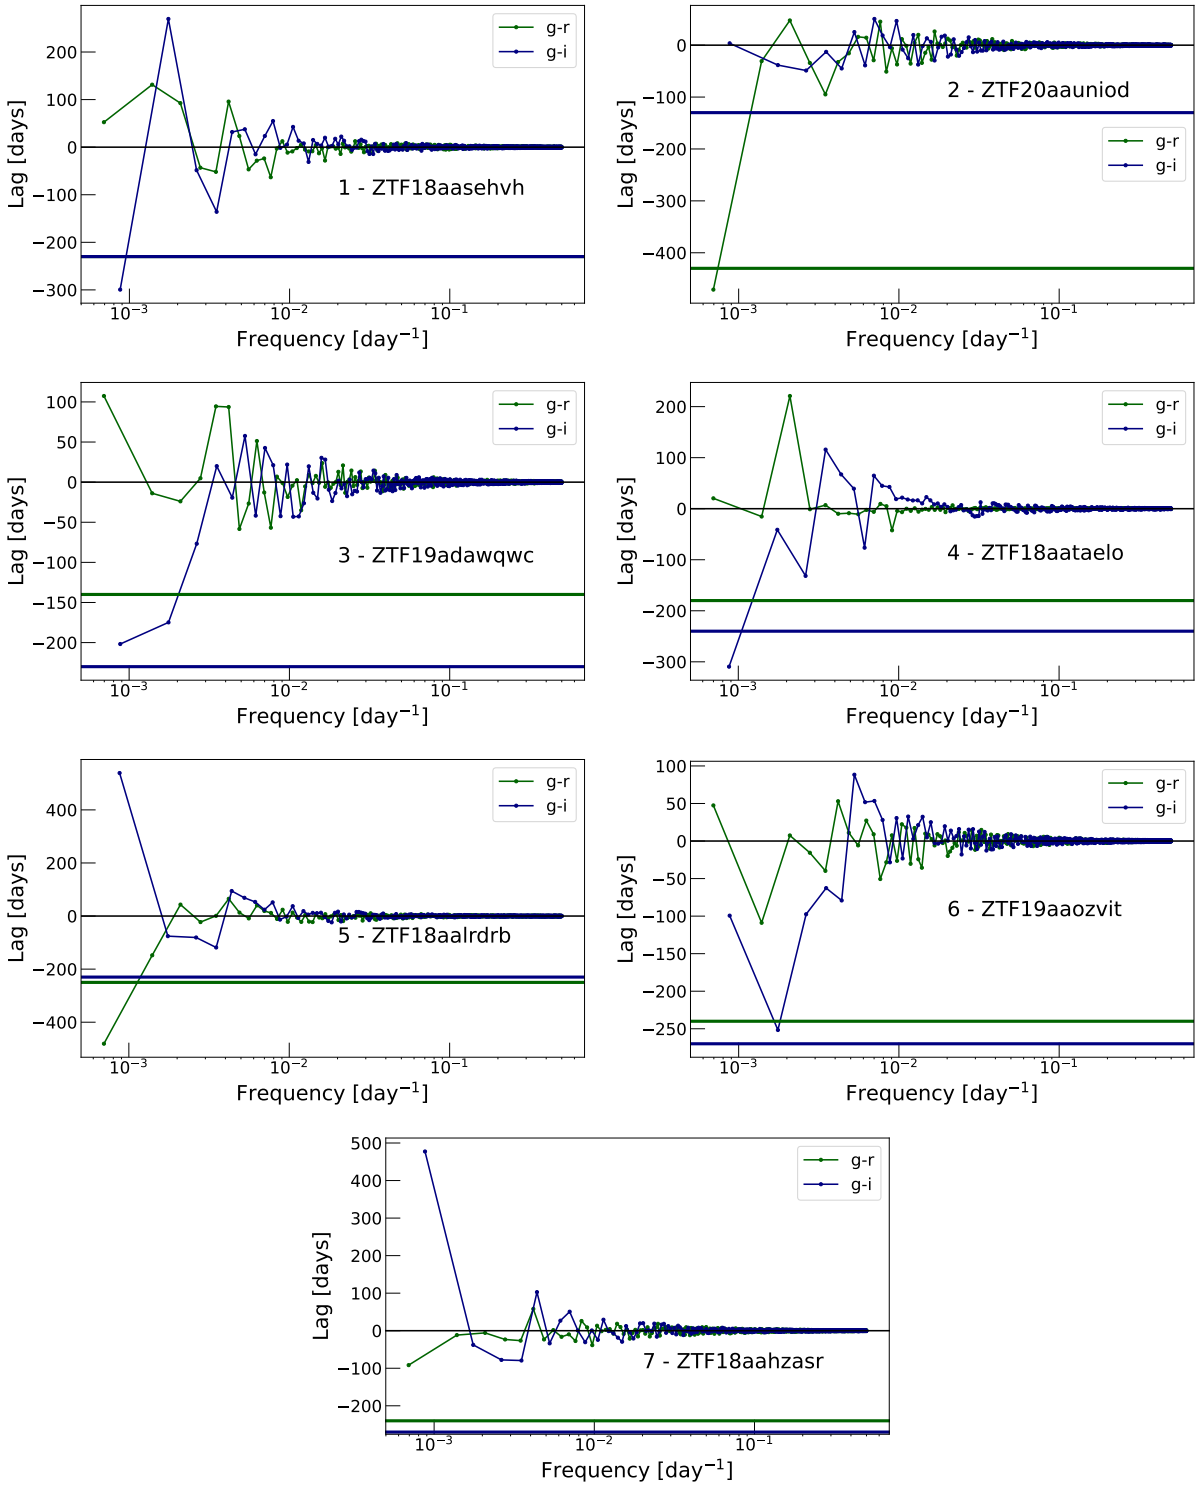


Figure 6.6: The lags as a function of frequency recovered by the linear interpolation Fourier method for the seven AGN light curves in our subsample. The g- to r-band lags are shown in purple and the g- to i-band lags are shown in blue.

light curves using these lags. The g- to r-band lags recovered by our various methods are given in parentheses in the second through fifth column in Table 6.1.

Despite not requiring a prominent peak in the JAVELIN distribution for the g- to r-band lags, we still find good agreement between the median and the mode of the negative JAVELIN distribution to within around 10% for all but light curve 2. For light curve 2, there is no prominent peak in the g- to r-band JAVELIN distribution shown in green in the middle right panel of Figure 6.4, suggesting there is no lag recovered between the g- and r-band. For all of the remaining light curves, the lag recovered by JAVELIN between the g- and r-band is similar to (within roughly 10%) or shorter than the g-i lag. Either result could be explained physically by the smaller distance between the two temperature regions for the narrower difference in wavelength, which would lead to a similar or shorter long negative lag. However, the peaks in the these JAVELIN distributions tend to make up only around 30% of the total negative distribution instead of the over 60% we require for the g- to i-band lag distributions. Only light curve 6, has a peak in the JAVELIN distribution that is over 60% of the distribution for the g- to r-band lags.

For all cases, the maximum-likelihood Fourier method detects long negative lags that are significantly shorter than the lag recovered by JAVELIN for that same light curve. For half of the light curves the minimum lag recovered has a shorter duration than -20 days which could easily be noise. For the other half, the lags recovered by the maximum-likelihood Fourier method are roughly a quarter of the lag recovered by JAVELIN, which could be due to averaging of the long and short lag. The majority of the lags recovered by the interpolated Fourier method are also significantly shorter than the lags recovered by JAVELIN. However, the lags recovered by our two different Fourier methods are not in very close agreement with each other; the lags from the interpolation method are significantly longer than the ones recovered by the maximum-likelihood method for all light curves.

AGN	Log(M)	z	$l/l_{\rm edd}$	Predicted Lag (days)	Local h/r Prediction
1	8.1 (7.6)	0.083	0.0089	-0.53	0.38
2	8.2 (7.9)	0.21	0.069	-3.0	0.85
3	7.7 (7.2)	0.088	0.040	-0.66	0.42
4	8.2 (7.8)	0.17	0.018	-1.3	0.46
5	7.6 (7.6)	0.13	0.042	-0.82	0.41
6	7.7 (7.9)	0.25	0.058	-1.9	0.42
7	7.4 (7.7)	0.13 =	0.067	-0.81	0.38

Table 6.2: For each of our seven AGN we give the log of the mass of in solar masses determined with SDSS DR14 spectra using the width of the $H\alpha$ line, not in parentheses, and the width of the $H\beta$ line, in parentheses, the redshift, and the Eddington ratio of the luminosity calculated using F5100 from the SDSS DR14 spectra. Using these three parameters, in the next two columns we give the redshift-dependent long negative lags predicted by Equation 6.5 and $\beta = 3.59$, the value for the power-law dependence of the scale height found for Fairall 9 (Yao et al., 2023) in day, and the local value of the aspect ratio derived by solving Equation 6.5 for the aspect ratio, h/r, using the lag detected for each AGN by JAVELIN.

6.5 Predictions from Previous Work

In Chapter 3, we derived an equation for the long negative lag using the viscous timescale,

$$\tau = \tau_0 \left[\left(\frac{\lambda}{\lambda_0} \right)^{4/3(7/2-2\beta)} - 1 \right], \tag{6.6}$$

where,

$$\tau_0 = 93 \text{ days} \left(\frac{-4}{7/2-2\beta} \right) \left(\frac{\lambda_0}{\lambda_{\rm W}} \right)^{4/3(7/2-2\beta)} \left(\frac{h_{\rm W}/r_{\rm W}}{0.2} \right)^{-2} \frac{0.1}{\alpha} \left(\frac{l/l_{\rm edd}}{0.1} \right)^{1/2} \left(\frac{M}{10^8 \text{ M}_\odot} \right)^{1/2}, \tag{6.7}$$

λ is the response wavelength (the r- or i-band here), λ_0 is the reference wavelength (the g-band here), $\lambda_{\rm W} = 1928$ Å(the SWIFT *W2*-band used as the reference band in Chapter 2.7), $h_{\rm W}/r_{\rm W}$ is the aspect ratio at the disk radius corresponding to the *W2*-band or the reference aspect ratio, and $l/l_{\rm edd}$ is the Eddington ratio. β represents the power-law scaling of the height of the disk as a function of radius, $h = h_{\rm W}(r/r_{\rm W})^{\beta}$. The parameterization and bands are based on the findings in Chapter 2, where they found $\beta = 3.59 \pm 1.22$, a very

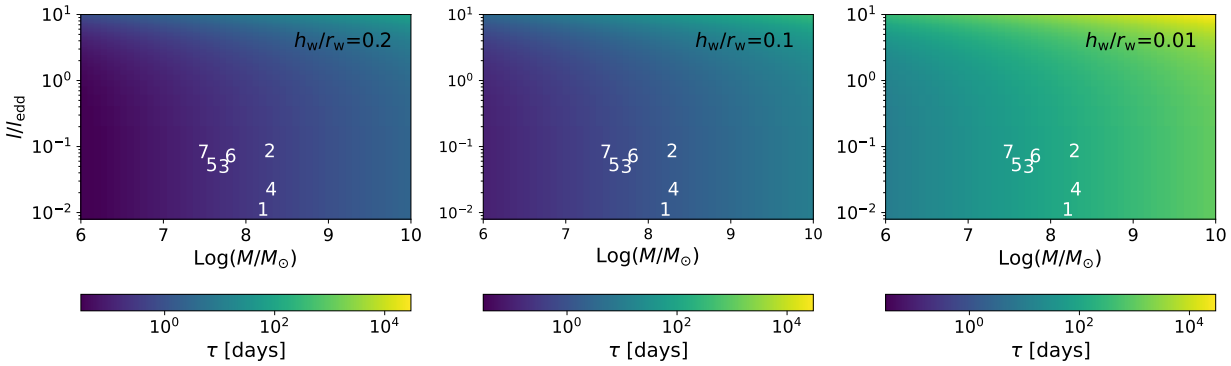


Figure 6.7: The long negative lag as a function of mass and Eddington ratio at redshift zero calculated between the ZTF g- and i-band for three different reference scale heights using the model from Chapter 3 and $\beta = 3.59$ from the fit in Chapter 2. The seven AGN presented here are shown at their mass and Eddington ratio.

steep increase in the scale height as a function of radius. Figure 6.7 shows the lags predicted between the g- and i-band by this model as a function of mass and Eddington ratio for three different reference aspect ratios at $z = 0$.

Our model predicts that there would be a very short duration long negative lag roughly on the order of one day for $h_W/r_W = 0.2$, the reference aspect ratio derived from Fairall 9 lags in Chapter 2. This duration is very short because of the narrow wavelength difference between the g- and i-band, only 4830 to 7904 Å. In addition, our model is based on the long lag as a function of wavelength detected for Fairall 9 and so assumes there will be a steep disk flaring which leads to a very short viscous timescale at optical-emitting radii. A negative lag shorter than ten days would be very difficult to detect with our current methods. If, however, we decrease our reference aspect ratio to $h_W/r_W \approx 0.01$, and therefore decrease the aspect ratio in the optical-emitting region, we can reach long negative lag timescales on the order of one hundred days, which agrees with the lags detected for the ZTF light curves in this paper.

Alternatively, the disk flaring we modeled in Fairall 9 may not be common to other AGN. Therefore, in Figure 6.8 we show the long negative lag as a function of Eddington ratio and supermassive black hole mass if the scale height is constant as a function of radius, for several different aspect ratios. Without the disk flaring the aspect ratio would need to be large throughout the disk in order for the lag to be on hundred day timescales in the optical region.

The seven AGN presented here are marked at their mass and Eddington ratio in white in Figures 6.2 and 6.8. Eddington ratios for these objects were derived using the black hole masses based on the $H\alpha$ widths and the flux at 5100 Å in the SDSS DR14 spectra. All of the AGN in our sample are low luminosity, falling between roughly 0.01 and 0.1 $l_{\rm edd}$, which according to Equation 6.7 would lead us to expect they might have shorter than average long lags. For the reference scale height and disk flaring found for Fairall 9 in Chapters 2 and 3, we can see in Figure 6.7 that the lags at zero redshift are around two orders of magnitude lower than what we have detected here between the g- and i-bands. It is only for a much thinner disk with $h_{\rm W}/r_{\rm W} \approx 0.01$, that we predict lags on the order of the ones recovered here. Alternatively, if we assume a constant scale height as a function of radius instead of disk flaring as in Figure 6.8, then we require a thicker disk throughout, with the long negative lags approaching thousands of days for our seven selected AGN as the aspect ratio drops below 0.1.

Table 6.2 shows the redshift dependent g-i lag predicted for each of our seven selected light curves with the model assumptions from Chapters 2 and 3 in the second column. These predicted long negative lags vary from -0.53 to -3.0 days, all of which are still much shorter than the lags recovered here, which suggests that the parameters derived from our lag detection in Fairall 9 may not be applicable to our ZTF light curves. We show the prediction for what we would expect the aspect ratio to be if we instead assume a flat disk in the second column of Table 6.2. For most of the seven AGN in this paper the predicted

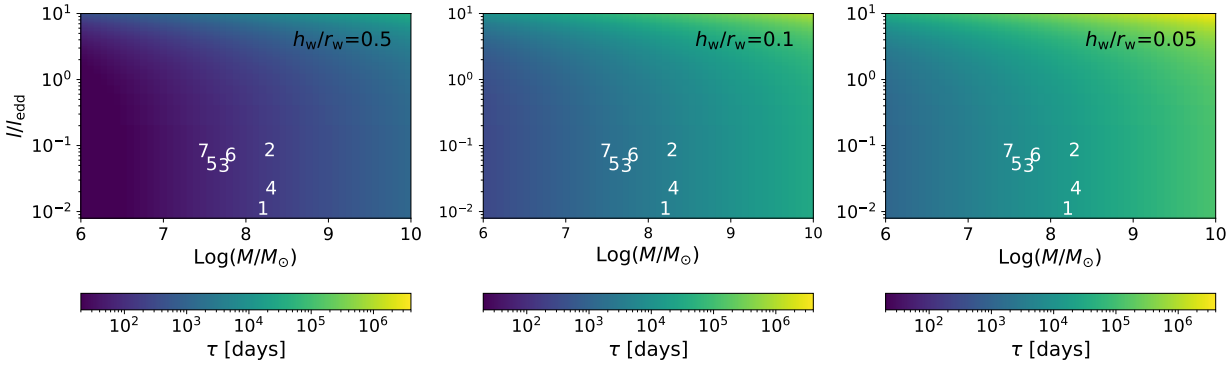


Figure 6.8: The long negative lag as a function of mass and Eddington ratio at redshift zero calculated between the ZTF g- and i-band for three different scale heights using the model from Chapter 3 and $\beta = 0$, i.e. a constant scale height. The seven AGN presented here are shown at their mass and Eddington ratio.

aspect ratio is around 0.4, which is over a factor of ten thicker than values typically predicted in the standard thin disk model.

6.6 Discussion

In this chapter we have summarized the preliminary results in our search for long negative lags in 3-year baseline ZTF light curves. We start with a sample of 227 light curves at $z < 0.4$ with over 1000 i-band visits and reduce the noise from multiple close together nightly observations by calculating an error-weighted nightly average. We then run two lag recovery methods, JAVELIN and the maximum-likelihood Fourier method and select out seven light curves with the the most promising g- to i-band lags according to the criteria described in Section 6.2.2. These criteria do not seem to bias us towards high or low black hole masses nor a particular redshift within our redshift range. We also run a linear nightly interpolation Fourier method on these selected light curves.

The lags recovered by our different methods for these seven light curves are promising. For the g- and i-band lags the JAVELIN distributions for these seven light curves have prominent peaks around the detected long lags. For four of the seven light curves, the lags recovered by the maximum-likelihood Fourier method are in good agreement with the lags recovered by JAVELIN. For two of the seven light curves they are significantly shorter, which could be due to the presence of a short lag (see Chapters 2 and 3). For six of the seven light curves, the lags recovered by JAVELIN between the g- and r-bands are reasonably consistent or shorter than the lags recovered by JAVELIN for the g- to i-band, as expected because the radial distance between the g- and r-band emitting regions will be shorter than the distance between the g- and i-band emitting regions. The two Fourier methods both find lags that are shorter than the lags recovered by JAVELIN for a majority of these g- and r-band light curves, which again could be due to the averaging of the long and short lag. However, both the g- to i- and g to r-band lags recovered by the interpolated Fourier method are generally not consistent with the lags recovered by the maximum-likelihood Fourier method.

We generate mock light curves with the same cadence as three of our seven selected nightly-averaged light curves in order to perform null tests of our methods on mock ZTF light curves with no lag. In our null tests, JAVELIN only recovers a lag near the lower limit of the range we run JAVELIN over (~ -500 days) and around the seasonal gap of ~ -180 days. We can use these false detection zones to inform our selection criteria for robust lag measurements. Our Fourier methods appear to erroneously recover long lags of around -100 days in our mock light curves with no input lag for reasons that are unclear at this time. Fortunately, only one lag recovered by the Fourier methods for the real light curves in Sections 6.4.1 and 6.4.2 falls close to the peak around -100 days. Thus, our null tests do not point to any obvious contamination for a majority of the lags detected in Sections 6.4.1 and 6.4.2 with any method.

We also perform recovery tests by using our methods on two additional types of mock light curves, one with only an input long lag and one with an input long lag and short lag.

For these tests, we find that JAVELIN is reasonably accurate on mocks with one input lag, but less accurate when a short lag is added. In addition, the maximum-likelihood Fourier method does not accurately recover the input long lag for a majority of our mock light curves. Furthermore, we do not find that the lags recovered by JAVELIN and the maximum-likelihood Fourier method agree with each other, which is what we found for a majority of the real light curves presented in Section 6.4.1. Instead, in Section 6.3 we find that the interpolated Fourier method is the most accurate method on our mocks, which was not expected based on our results from the real light curves.

When we use the models derived in Chapter 3 from the lags as a function of wavelength recovered for Fairall 9 in Chapter 2, we calculate a much shorter long negative lag from the g- to i-band than the ones detected here. If we instead assume a flat disk with a constant scale height as a function of radius, we find that for most of the seven AGN presented here the aspect ratio is around 0.5 in the optical region. These results suggest that the model parameters for the vertical structure of AGN disks derived for Fairall 9 are not the same as for these ZTF light curves. Future work should determine whether the disk is thinner throughout the UV-optical region for these AGN or if there is less disk flaring than was found for Fairall 9. Recovering a lag between the g- and r-band in addition to the lag between the g- and i-band would allow us to use our models to derive the scale height as a function of radius instead of only being able to calculate a single measurement of the scale height in one region. In the future, mock light curves can help to determine if detecting a lag between the g- and r-band is possible given that this two wavebands are very similar in wavelength.

A closer examination of each object in this sample of seven is necessary to reconcile the outcomes of our different methods on each and on our mock light curves. For example, measuring the short lag for each AGN will help us to better understand potential averaging affects that may be occurring when we use the Fourier method. Additionally, ZTF provides a much larger sample of AGN that should also be surveyed in order to maximize the number of long negative lags detected. Given the preliminary results in Sections 6.4 and 6.3, it may be

better to select from the wider ZTF sample based on the results from the interpolated Fourier method instead of the criteria we have used here, or jointly select targets with promising g-i and g-r lags.

Improving our methods on light curves from ZTF is excellent preparation for LSST, which will provide hundreds of thousands of AGN light curves, because there are several similarities between the two surveys. ZTF and LSST are both ground-based optical surveys, which means they have overlapping waveband coverage and both will have the large $\sim$ 180 day seasonal gaps. The LSST deep drilling fields, five relatively small (9.6 deg^2) fields which will have deeper coverage and a higher cadence than the rest of the survey, also have a similar cadence to ZTF. There are advantages to LSST, which is anticipated to have better signal to noise ratios, deeper observations, and a broader waveband coverage with six wavebands from the u- to y-band, instead of three. In addition, the planned baseline for LSST is ten years, while the full lifetime of ZTF is unknown and the light curves used here have a duration of only three years. It is most likely for this reason that both JAVELIN and the maximum-likelihood Fourier method appear to perform better on LSST mocks, such as those in Chapter 3, than the mocks in this chapter that have been sub-sampled with ZTF cadences. This result is encouraging, because it suggests that by perfecting our methods on ZTF light curves, we should do even better with LSST light curves in the near future.

Bibliography

Abdallah, J. J., & Clark, R. E. H. 1985, TOPS: A Multigroup Opacity Code, Tech. rep., Los Alamos

Alexander, T. 2013, arXiv e-prints, arXiv:1302.1508

Antonucci, R. 1988, in Supermassive Black Holes, ed. M. Kafatos, 26

Antonucci, R. R. J. 2023, Galaxies, 11, 102

Antonucci, R. R. J., Kinney, A. L., & Ford, H. C. 1989, ApJ, 342, 64

Arévalo, P., McHardy, I. M., & Summons, D. P. 2008a, MNRAS, 388, 211

Arévalo, P., & Uttley, P. 2006, MNRAS, 367, 801

Arévalo, P., Uttley, P., Kaspi, S., et al. 2008b, MNRAS, 389, 1479

Arévalo, P., Uttley, P., Lira, P., et al. 2009, MNRAS, 397, 2004

Bai, X.-N., & Stone, J. M. 2013, ApJ, 769, 76

Balbus, S. A., & Hawley, J. F. 1991, ApJ, 376, 214

Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019, PASP, 131, 068003

Bentz, M. C., & Katz, S. 2015, Publications of the Astronomical Society of the Pacific, 127, 67

Blandford, R. D., & McKee, C. F. 1982, ApJ, 255, 419

Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, AJ, 154, 28

Buisson, D. J. K., Lohfink, A. M., Alston, W. N., et al. 2018, MNRAS, 475, 2306

Burke, C. J., Shen, Y., Blaes, O., et al. 2021, Science, 373, 789

Cackett, E. M., Bentz, M. C., & Kara, E. 2021, iScience, 24, 102557

Cackett, E. M., Chiang, C.-Y., McHardy, I., et al. 2018, ApJ, 857, 53

Cackett, E. M., Fabian, A. C., Zogbhi, A., et al. 2013, ApJ, 764, L9

Cackett, E. M., Horne, K., & Winkler, H. 2007, MNRAS, 380, 669

Cackett, E. M., Zoghbi, A., & Ulrich, O. 2022, ApJ, 925, 29

Cackett, E. M., Gelbord, J., Barth, A. J., et al. 2023, arXiv e-prints, arXiv:2306.17663

Caplar, N., Lilly, S. J., & Trakhtenbrot, B. 2017, The Astrophysical Journal, 834, 111

Chelouche, D., Pozo-Nuñez, F., & Zucker, S. 2017, ApJ, 844, 146

Collier, S., Horne, K., Wanders, I., & Peterson, B. M. 1999, MNRAS, 302, L24

Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, in Modeling, Systems Engineering, and Project Management for Astronomy VI, ed. G. Z. Angeli & P. Dierickx, Vol. 9150, International Society for Optics and Photonics (SPIE), 915014

Davis, S. W., & Tchekhovskoy, A. 2020, ARA&A, 58, 407

De Marco, B., Ponti, G., Cappi, M., et al. 2013, MNRAS, 431, 2441

De Rosa, G., Peterson, B. M., Ely, J., et al. 2015, ApJ, 806, 128

Dekany, R., Smith, R. M., Riddle, R., et al. 2020, PASP, 132, 038001

Dexter, J., & Begelman, M. C. 2019, MNRAS, 483, L17

Dexter, J., & Quataert, E. 2012, MNRAS, 426, L71

Edelson, R., Gelbord, J. M., Horne, K., et al. 2015, ApJ, 806, 129

Edelson, R., Gelbord, J., Cackett, E., et al. 2017, ApJ, 840, 41

—. 2019, ApJ, 870, 123

Event Horizon Telescope Collaboration. 2019, ApJ, 875, L1

Fabian, A. C. 2012, ARA&A, 50, 455

Fabian, A. C., Zoghbi, A., Ross, R. R., et al. 2009, Nature, 459, 540

Fausnaugh, M. M., Denney, K. D., Barth, A. J., et al. 2016, ApJ, 821, 56

Friedjung, M. 1985, A&A, 146, 366

Gaburov, E., Johansen, A., & Levin, Y. 2012, ApJ, 758, 103

García, J., Dauser, T., Reynolds, C. S., et al. 2013, ApJ, 768, 146

García, J. A., Kara, E., Walton, D., et al. 2019, ApJ, 871, 88

Gardner, E., & Done, C. 2017, MNRAS, 470, 3591

Gaskell, C. M. 2017, MNRAS, 467, 226

Gaskell, C. M., Anderson, F. C., Birmingham, S. Á., & Ghosh, S. 2023, MNRAS, 519, 4082

Gaskell, C. M., & Peterson, B. M. 1987, ApJS, 65, 1

Gaskell, C. M., & Sparke, L. S. 1986, ApJ, 305, 175

Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, PASP, 131, 078001

Gravity Collaboration. 2018, Nature, 563, 657

Grier, C. J., Trump, J. R., Shen, Y., et al. 2017, The Astrophysical Journal, 851, 21

Grier, C. J., Shen, Y., Horne, K., et al. 2019, The Astrophysical Journal, 887, 38

Guo, H., Barth, A. J., & Wang, S. 2022, arXiv e-prints, arXiv:2207.06432

Hagen, S., & Done, C. 2023, MNRAS, 521, 251

Hall, P. B., Sarrouh, G. T., & Horne, K. 2018, ApJ, 854, 93

Hameury, J. M. 2020, Advances in Space Research, 66, 1004

Hernández Santisteban, J. V., Edelson, R., Horne, K., et al. 2020, MNRAS, 498, 5399

Hirose, S., Krolik, J. H., & Blaes, O. 2009, ApJ, 691, 16

Hogg, J. D., & Reynolds, C. S. 2016, ApJ, 826, 40

Homayouni, Y., Trump, J. R., Grier, C. J., et al. 2020, The Astrophysical Journal, 901, 55

Homayouni, Y., Sturm, M. R., Trump, J. R., et al. 2022, ApJ, 926, 225

Hopkins, P. F., Squire, J., Su, K.-Y., et al. 2023, arXiv e-prints, arXiv:2310.04506

Iglesias, C. A., & Rogers, F. J. 1996, ApJ, 464, 943

Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111

Jiang, Y.-F. 2021, ApJS, 253, 49

—. 2022, ApJS, 263, 4

Jiang, Y.-F., & Blaes, O. 2020, ApJ, 900, 25

Jiang, Y.-F., Blaes, O., Stone, J. M., & Davis, S. W. 2019, ApJ, 885, 144

Jiang, Y.-F., Davis, S. W., & Stone, J. M. 2016, ApJ, 827, 10

Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2013, ApJ, 778, 65

Jiang, Y.-F., Green, P. J., Greene, J. E., et al. 2017, ApJ, 836, 186

Kammoun, E. S., Dovčiak, M., Papadakis, I. E., Caballero-García, M. D., & Karas, V. 2021a, ApJ, 907, 20

Kammoun, E. S., Papadakis, I. E., & Dovčiak, M. 2021b, MNRAS, 503, 4163

Kara, E., Mehdipour, M., Kriss, G. A., et al. 2021, ApJ, 922, 151

Kara, E., Barth, A. J., Cackett, E. M., et al. 2023, ApJ, 947, 62

Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2015, MNRAS, 451, 4328

Kaspi, S., Smith, P. S., Netzer, H., et al. 2000, The Astrophysical Journal, 533, 631

Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, ApJ, 698, 895

Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511

Kovacevic, A. B., Radovic, V., Ilic, D., et al. 2022, arXiv e-prints, arXiv:2208.06203

Kozłowski, S. 2017, ApJ, 835, 250

Kozłowski, S., Kochanek, C. S., Udalski, A., et al. 2010, ApJ, 708, 927

Laor, A., & Davis, S. W. 2014, MNRAS, 438, 3024

Lawrence, A. 2018, Nature Astronomy, 2, 102

Leighly, K. M. 2004, ApJ, 611, 125

Lewin, C., Kara, E., Wilkins, D., et al. 2022, ApJ, 939, 109

Li, J. I.-H., Shen, Y., Brandt, W. N., et al. 2019, The Astrophysical Journal, 884, 119

Lynden-Bell, D. 1969, Nature, 223, 690

Lyubarskii, Y. E. 1997, MNRAS, 292, 679

MacLeod, C. L., Ivezić, Ž., Sesar, B., et al. 2012, ApJ, 753, 106

Mahmoud, R. D., & Done, C. 2020, MNRAS, 491, 5126

Markowitz, A., Edelson, R., Vaughan, S., et al. 2003, ApJ, 593, 96

McHardy, I. M., Arévalo, P., Uttley, P., et al. 2007, MNRAS, 382, 985

McHardy, I. M., Cameron, D. T., Dwelly, T., et al. 2014, MNRAS, 444, 1469

McHardy, I. M., Connolly, S. D., Horne, K., et al. 2018, MNRAS, 480, 2881

Miller, L., Turner, T. J., Reeves, J. N., et al. 2010, MNRAS, 403, 196

Morales, A. M., Miller, J. M., Cackett, E. M., Reynolds, M. T., & Zoghbi, A. 2019, ApJ, 870, 54

Morgan, C. W., Kochanek, C. S., Morgan, N. D., & Falco, E. E. 2010, ApJ, 712, 1129

Mudd, D., Martini, P., Zu, Y., et al. 2018, ApJ, 862, 123

Murray, N., Chiang, J., Grossman, S. A., & Voit, G. M. 1995, ApJ, 451, 498

Mushotzky, R. F., Edelson, R., Baumgartner, W., & Gandhi, P. 2011, ApJ, 743, L12

Neustadt, J. M. M., & Kochanek, C. S. 2022, MNRAS, 513, 1046

Neustadt, J. M. M., Kochanek, C. S., Montano, J., et al. 2023, arXiv e-prints, arXiv:2310.01497

Novikov, I. D., & Thorne, K. S. 1973, in Black Holes (Les Astres Occlus), 343–450

Oknyanskii, V. L. 1993, Astronomy Letters, 19, 416

Oknyansky, V., & Oknyansky, R. 2022, PyMCCF: Python Modernized Cross Correlation Function for reverberation mapping studies, Astrophysics Source Code Library, record ascl:2212.007, ascl:2212.007

Pahari, M., McHardy, I. M., Vincentelli, F., et al. 2020, MNRAS, 494, 4057

Pal, M., Dewangan, G. C., Connolly, S. D., & Misra, R. 2017, MNRAS, 466, 1777

Panagiotou, C., Kara, E., & Dovčiak, M. 2022a, ApJ, 941, 57

Panagiotou, C., Papadakis, I., Kara, E., Kammoun, E., & Dovčiak, M. 2022b, ApJ, 935, 93

Papadakis, I. E., Nandra, K., & Kazanas, D. 2001, ApJ, 554, L133

Pâris, I., Petitjean, P., Aubourg, E., et al. 2018, A&A, 613, A51

Peterson, B. M. 2014, Space Sci. Rev., 183, 253

Peterson, B. M., Wanders, I., Horne, K., et al. 1998, PASP, 110, 660

Peterson, B. M., Ferrarese, L., Gilbert, K. M., et al. 2004, ApJ, 613, 682

Pozo Nuñez, F., Bruckmann, C., Deesamutara, S., et al. 2023, MNRAS, 522, 2002

Pozo Nuñez, F., Gianniotis, N., Blex, J., et al. 2019, MNRAS, 490, 3936

Pringle, J. E., & Rees, M. J. 1972, A&A, 21, 1

Reuter, M. A., Cook, K. H., Delgado, F., Petry, C. E., & Ridgway, S. T. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9911, Modeling, Systems Engineering, and Project Management for Astronomy VI, ed. G. Z. Angeli & P. Dierickx, 991125

Salvesen, G. 2022, ApJ, 940, L22

Sánchez-Sáez, P., Lira, P., Mejía-Restrepo, J., et al. 2018, The Astrophysical Journal, 864, 87

Schimoia, J. S., Storchi-Bergmann, T., Grupe, D., et al. 2015, ApJ, 800, 63

Secunda, A., Greene, J. E., Jiang, Y.-F., Yao, P. Z., & Zoghbi, A. 2023, arXiv e-prints, arXiv:2306.05455

Secunda, A., Jiang, Y.-F., & Greene, J. E. 2024, ApJ, 965, L29

Sergeev, S. G., Doroshenko, V. T., Golubinskiy, Y. V., Merkulova, N. I., & Sergeeva, E. A. 2005, ApJ, 622, 129

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

—. 1976, MNRAS, 175, 613

Shappee, B. J., Prieto, J. L., Grupe, D., et al. 2014, ApJ, 788, 48

Shen, Y., Richards, G. T., Strauss, M. A., et al. 2011, ApJS, 194, 45

Smith, K. L., Mushotzky, R. F., Boyd, P. T., et al. 2018, ApJ, 857, 141

Starkey, D., Horne, K., Fausnaugh, M. M., et al. 2017, ApJ, 835, 65

Stern, D., McKernan, B., Graham, M. J., et al. 2018, ApJ, 864, 27

Stone, Z., & Shen, Y. 2023, MNRAS, 524, 4521

Stone, Z., Shen, Y., Burke, C. J., et al. 2022, MNRAS, 514, 164

Sun, M., Grier, C. J., & Peterson, B. M. 2018, PyCCF: Python Cross Correlation Function for reverberation mapping studies, Astrophysics Source Code Library, record ascl:1805.032, ascl:1805.032

Sun, M., Xue, Y., Guo, H., et al. 2020, ApJ, 902, 7

Tachibana, Y., Graham, M. J., Kawai, N., et al. 2020, ApJ, 903, 54

Timmer, J., & Koenig, M. 1995, A&A, 300, 707

Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, A&A Rev., 22, 72

Uttley, P., Edelson, R., McHardy, I. M., Peterson, B. M., & Markowitz, A. 2003, ApJ, 584, L53

Vincentelli, F. M., McHardy, I., Cackett, E. M., et al. 2021, MNRAS, 504, 4337

Ward, C., Gezari, S., Frederick, S., et al. 2021, ApJ, 913, 102

Ward, C., Gezari, S., Nugent, P., et al. 2024, ApJ, 961, 172

White, R. J., & Peterson, B. M. 1994, PASP, 106, 879

Wilkins, D. R. 2019, MNRAS, 489, 1957

Yao, P. Z., Secunda, A., Jiang, Y.-F., Greene, J. E., & Villar, A. 2023, ApJ, 953, 43

Yu, W., Richards, G. T., Vogeley, M. S., Moreno, J., & Graham, M. J. 2022, ApJ, 936, 132

Yu, Z., Martini, P., Davis, T. M., et al. 2020, ApJS, 246, 16

Yuk, H., Dai, X., Jayasinghe, T., et al. 2022, ApJ, 930, 110

Zoghbi, A., Fabian, A. C., Uttley, P., et al. 2010, MNRAS, 401, 2419

Zoghbi, A., Reynolds, C., & Cackett, E. M. 2013, ApJ, 777, 24

Zu, Y., Kochanek, C. S., Kozłowski, S., & Peterson, B. M. 2016, ApJ, 819, 122

Zu, Y., Kochanek, C. S., Kozłowski, S., & Udalski, A. 2013, ApJ, 765, 106

Zu, Y., Kochanek, C. S., & Peterson, B. M. 2010, JAVELIN: Just Another Vehicle for Estimating Lags In Nuclei, Astrophysics Source Code Library, record ascl:1010.007, ascl:1010.007

—. 2011, ApJ, 735, 80