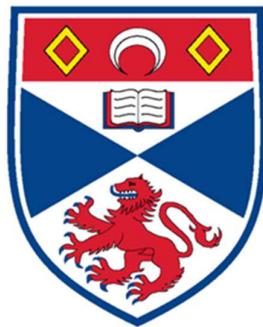


Project report for final year project in Physics & Astronomy

## Investigating AGN Variability in the LSST

Matriculation Number: 200004472



University  
of  
St Andrews

Module: AS5101 MPhys Astrophysics

Science is often highly collaborative. Please clarify here which parts of the work reported have been done by you and which parts include the work of others. For example, this might include provision of simulation or analysis codes which you utilised or modified, provision of experimental data to analyse, collaborations with others in performing new measurements, if you built up things from scratch in your project or started from pre-existing setups, and others. Providing this information helps clarify the scope and focus of your project work.

The overall structure of the synthetic light curve generation pipeline was built from the ground up using standard Python packages (Numpy, SciPy, Pandas, and Matplotlib). The following are existing codes I have modified and utilised for specific sections of the pipeline:

- Modified CAR(1) algorithm from AstroML package to generate stochastic light curve.
- G03\_SMCBar, SMC-like Dust Extinction Model from Astropy to simulate disc spectrum.
- Transfer Function generation code adapted from PyCREAM (Starkey et al. 2016).
- 

In addition to my computer, Juan Hernandez Santisteban's and Rory Brown's desktops were used to run PhoSim simultaneously to produce the simulated data.

For the analysis of the data, the following existing code have been utilised:

- Running optimal average function from PyRoa package (Donnan et al. 2021) for the synthetic light curve processing.
- PyCCF package (Sun et al. 2018) to calculate the cross correlation function and centroid to measure the lag between light curves.

## ABSTRACT

I investigate the feasibility of using the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) to conduct continuum reverberation mapping (CRM) studies of active galactic nuclei (AGN) over a broad range of redshifts and accretion rates. By constructing a physically motivated simulation pipeline, I generate high-fidelity synthetic light curves in the LSST’s ugrizy filter bands by modelling the AGN accretion disc, incorporating reprocessing physics, cosmological redshift, estimated LSST survey cadence, and observational throughput using the Photon Simulator (PhoSim). Time-series analysis is then performed to recover interband lags via cross-correlation techniques, enabling us to assess the reliability of CRM across the parameter space. I demonstrate that lags are recoverable for AGNs up to  $z \sim 1.5$  in the COSMOS field and  $z \sim 0.7$  in lower-cadence Deep Drilling Fields (DDFs). These results imply that, over the 10-year LSST mission, CRM can be applied to hundreds of AGNs with varying physical properties, including high-redshift and low-Eddington ratio sources previously not readily accessible by other studies. The results show that LSST will be capable of probing accretion discs’ inner structure and provide insight into current inconsistencies between theoretical models and realistic observations.

## 1. INTRODUCTION, BACKGROUND SCIENCE, AND MOTIVATION

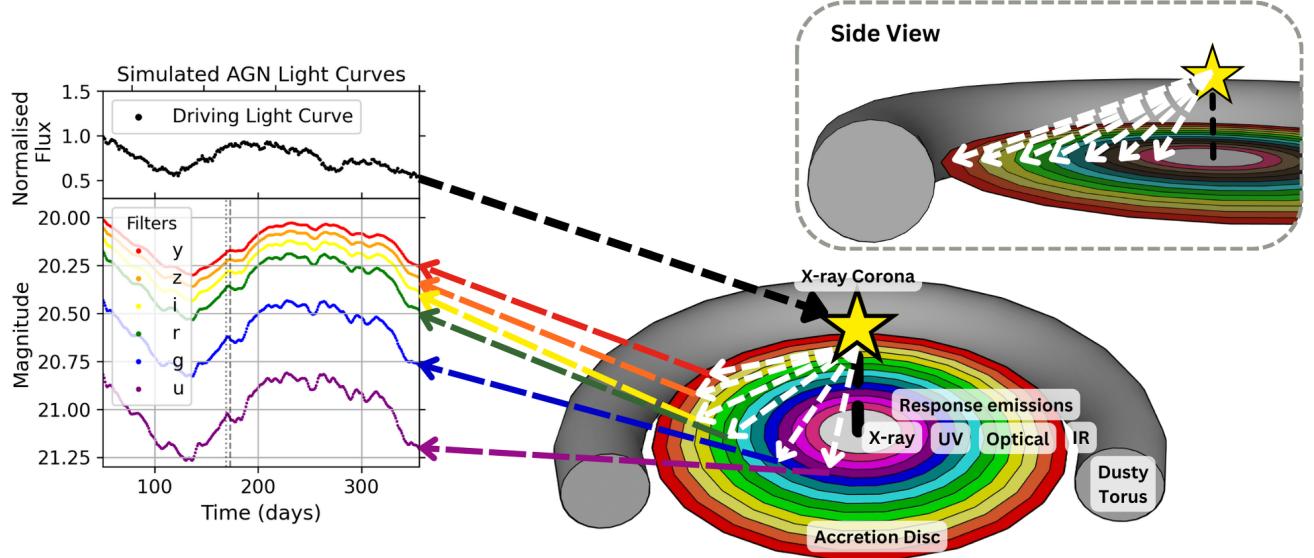
Active Galactic Nuclei (AGNs) are one of the Universe’s most luminous and variable objects. They lie at the centre of active galaxies, are powered by the accretion of matter onto the supermassive black hole (SMBH) (Magorrian et al. 1998; Silk & Rees 1998), and emit across the electromagnetic spectrum. These highly energetic processes affect their host galaxy’s star formation and the intergalactic medium through feedback mechanisms (Silk & Rees 1998; Fabian 2012). Their brightness allows them to be observed over cosmological distances, making them valuable probes for galaxy and black hole co-evolution via the M- $\sigma$  relation (Ferrarese & Merritt 2000).

As such, determining the SMBH mass is key to furthering our understanding of galaxies. To achieve this, we must study the structure and dynamics of the inner regions surrounding the SMBH. However, due to their compact nature, they remain unresolved by conventional observational techniques. Reverberation mapping (Blandford & McKee 1982; Peterson 1993) is an observational technique that utilises temporal resolution to infer the spatial structure of an AGN’s internal components. This method measures the

time delay, or “lag”, between flux variation features in different wavelengths of light curves.

### 1.1. *Studying AGNs with Reverberation Mapping*

Fig. 1 shows a schematic diagram of an AGN demonstrating how different distances between the driving source and the response emission region result in the lag in the light curves described in this section. The current, widely accepted understanding of AGNs uses the Unification Model (Antonucci 1993; Urry & Padovani 1995; Netzer 2015). This model suggests that the different types of AGNs observed are due to the angle at which we observe an AGN. Fig. 1 shows a cross-section of the partial model which only focuses on the central engine. The central engine, which consists of the SMBH with its accretion disc, where the viscosity of the gas heats the accretion disc and causes thermal emission in the UV/optical range. The engine is surrounded by an opaque, torus-shaped structure composed of dust and gas, which absorbs radiation and re-emits it in the infrared. The region inside the torus is filled with fast-moving gas closer to the SMBH and slow-moving gas further away, which is ionised by radiation from the engine. These regions are referred to as the broad lines region (BLR) and narrow lines region (NLR), respectively. A powerful X-



**Figure 1.** A schematic diagram of an AGN demonstrating how the geometry of the AGN leads to the time lag observed during reverberation mapping studies. The left plot shows the stochastic ionising X-ray driving light curve (top panel) and the response emission light curves (bottom panel) for each of the LSST’s filters. The dashed lines in the bottom panel show the lag between the flux variation feature in the  $u$  and  $y$  bands. The bottom right diagram shows a partial Unified Model of an AGN (Antonucci 1993; Urry & Padovani 1995; Netzer 2015), illustrating the accretion disc, torus, and X-ray corona (represented by the yellow star). The white arrows represent the ionising X-ray radiation from the corona travelling to different parts of the disc. The colour arrows represent the response emissions detected by the telescope. The side view diagram (top right) shows clearly the different distances the ionising radiations have to travel from the corona to the response emission regions, resulting in the time lag between flux variations in the light curves. Note that this diagram does not present jets that exist in quasars, blazars, and radio galaxies.

ray corona above the SMBH further irradiated the accretion disc and surrounding gases.

The time lag measured in RM studies is defined as  $\langle \tau \rangle = R/c$ , where  $R$  is the distance the light travels and  $c$  is the speed of light. Reverberation Mapping relies on three key assumptions (Peterson 1993):

1. The driving continuum originates from a single source.
2. The time delay between the driving and response continuum is due entirely to the photon travel times.
3. A simple (linear) relationship exists between the observed and driving ionised continuum.

When the ionising radiation travels from the emitting source to another region of the AGN, it can ionise and cause thermal or line emission at longer

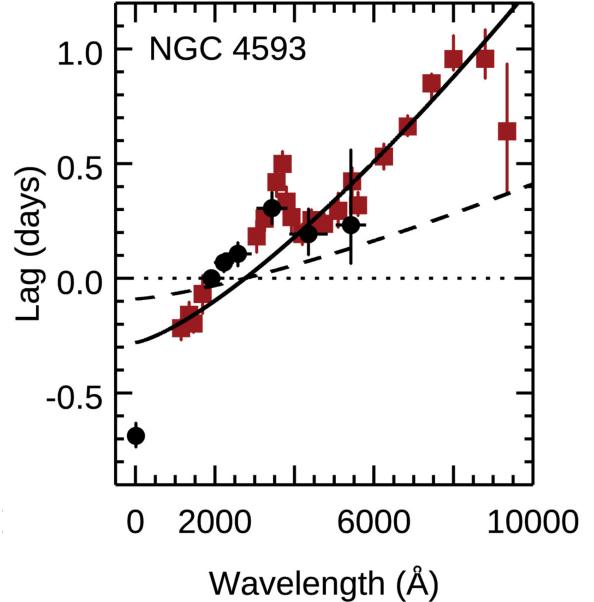
wavelengths. Thus, any flux variation observed in the ionising radiation will also be present in the remitted wavelength, as shown in Fig.1. When using this technique to measure the size  $R_{BLR}$  of the broad line region (BLR) using emission lines, the mass of the SMBH can be estimated (Peterson & Horne 2004) by from the following equation  $M_{BH} = fR_{BLR}\sigma_v^2/G$  where  $G$  is the Gravitational constant,  $f$  is the scaling factor that accounts for the geometry and kinematics of the BLR, and  $\sigma_v$  is the velocity dispersion measured from line broadening. Since the BLR emission lines flux responds to variations in the continuum radiation, by better understanding the continuum-emitting regions, we can more accurately measure the size of the BLR.

Continuum reverberation mapping (CRM) focuses on probing the structure and dynamics of the accretion disc. The thermal reprocessing model (Cackett et al. 2007) theorises that differ-

ent observed wavelengths are associated with different radii of the accretion disc, due to the inner regions of the disc being hotter than the outer regions. Thus, the thermal emission from the inner disc has a shorter wavelength and emits in the UV and gets longer into the optical and IR as we move further out of the disc. By measuring the lag in the response emission light curve, we can recreate the structure of the accretion disc, as different regions of the disc have unique time delays (Sergeev et al. 2005; Fausnaugh et al. 2016). The x-ray corona is the driving ionisation source that causes the response thermal emission in the accretion disc and other parts of the AGN (Edelson et al. 2019). We can predict the time delays by building a model of the disc, utilising its temperature profile and structure. The simplest model that is used in this project and the majority of the literature (Cackett et al. 2007; Starkey 2017; Pozo Nuñez et al. 2023) to generate synthetic response thermal emission light curves is to assume a non-relativistic geometrically thin and optically thick accretion disc (Shakura & Sunyaev 1973) with the X-ray corona as a compact point source above the centre of the disc acting as the 'Lamp-post' that illuminates the disc (Martocchia & Matt 1996). This disc model predicts the temperature profile to follow  $T(R) \propto (M\dot{M})^{-1/4}R^{-3/4}$ , where  $\dot{M}$  is the Eddington ratio. Treating the accretion disc as a blackbody with  $\lambda \propto T^{-1}$ . Relating  $\tau \propto R$  to the temperature profile and the blackbody radiation, we get a relationship between the time delay and wavelength to be  $\tau(\lambda) \propto (M\dot{M})^{-1/4}\lambda^{4/3}$  (Cackett et al. 2021). Observationally, the observed wavelength-dependent lag  $\tau$  can be fitted with the following model:

$$\tau = \tau_0 \left( \left( \frac{\lambda}{\lambda_0} \right)^{4/3} - y_0 \right) \quad (1)$$

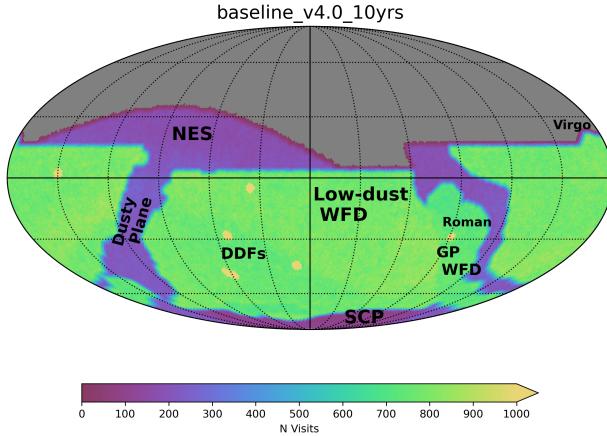
where  $\tau_0$  is the normalisation,  $\lambda_0$  is the reference wavelength, and  $y_0$  is usually set to -1 such that the lag goes through zero at the reference wavelength. However, measured lags from multiple long-term studies (De Rosa et al. 2015; Edelson et al. 2015; Fausnaugh et al. 2016) have shown three inconsistencies:



**Figure 2.** Lag Spectrum of NGC4593 (Fausnaugh et al. 2016) demonstrating the discrepancy between theoretical and observation results. The dashed line shows the theoretical lag, and the solid line shows the  $\tau \propto \lambda^{4/3}$  best-fit. The excess lag in the U-bands around 3500Å and poor x-ray correlation are also exhibited in this AGN. Figure adapted from Cackett et al. (2018)

1. The measured size of the accretion disc is frequently 2-3 times larger than theoretical predictions.
2. There is an excess lag in the u-band.
3. X-ray measurements are inconsistent with the ionising source and poorly correlate with the wavelength-lag relation.

These inconsistencies are demonstrated in Fig. 2 and prevent us from accurately determining the structure of the accretion disc and thus the mass of the SMBH. Contributions from emission lines and the BLR's diffuse continuum emission (Korista & Goad 2001; Chelouche et al. 2019; Netzer 2022), as well as more sophisticated accretion disc models, have been suggested to address these inconsistencies. Ultimately, more AGNs need to be studied to further constrain our model. The relatively short lag time scale and accretion disc's emission in the UV/optical spectrum requires intensive monitoring campaigns for CRM, such as McHardy et al. (2014); Edelson et al.



**Figure 3.** LSST Baseline Survey Visit map showing the coverage of the LSST and the location and expected numbers of visits of each mission. Figure from (Jones et al. 2025)

(2015); Fausnaugh et al. (2016); Hernández Santisteban et al. (2020). However, this makes the high-cadence multi-wavelength photometric surveys such as the LSST potentially suitable for CRM to study large numbers of AGNs (Chan et al. 2020; Kovačević et al. 2022).

## 1.2. LSST

Later in 2025, the Vera C. Rubin Observatory (Collaboration et al. 2009) will begin the Legacy Survey of Space and Time (LSST), a comprehensive 10-year survey of the southern sky, using its 8.4-meter primary mirror, six optical ugrizy filters, and a 3.2-gigapixel camera. Every 2-4 nights, the observatory will capture wide-field images of the entire southern sky; the final coadded image will have an average depth of 26.2 magnitudes and 171 visits per band across the entire survey. Embedded within the wide fields are Deep Drilling Fields, which will receive an average of 23,000 visits. Fig. 3 shows the footprint of the survey as well as the different fields, which will be explained in Sec. 3.1. With over 60 petabytes collected by the end of the survey, it is expected that this observatory will revolutionise our understanding of the Universe.

## 1.3. The Photon Simulator (PhoSim)

The Photon Simulator (PhoSim) (Peterson 2014) is a photon-based Monte Carlo ray-tracing tool originally developed for the Vera C. Rubin

Observatory. It simulates the full photon journey from astrophysical source to the detectors, accounting for atmospheric effects, telescope optics, filter throughput, and instrumental noise. This enables the generation of realistic synthetic images under observing conditions expected from the LSST. PhoSim has been utilised to generate accurate 5-to-10 year data from the LSST for the broader scientific community to test their analysis methods (Sánchez et al. 2020; Abolfathi et al. 2021).

## 1.4. Motivation

The LSST is expected to produce over 20TB/night of data, thus requiring an effective strategy prior to the start of the survey to study AGNs. The LSST's high cadence and depth will allow us to conduct CRM studies on a wide range of AGN populations at varying redshifts. This would provide more evidence of the structure and accretion flow by probing the inner disc region from the rest-frame UV of high-redshift AGN, which is currently possible only with UV telescopes such as the Swift (Burrows et al. 2005). In addition, it may reveal how the structure changes as a function of redshift, which will better constrain our accretion disc model. As such, the scientific aim of this project is to determine the feasibility of studying AGNs with different properties at varying redshifts throughout the 10-year LSST mission by simulating high-fidelity synthetic AGNs light curves as detected by the LSST.

Current literature (Kovačević et al. 2022; Pozo Nuñez et al. 2023) investigation into the feasibility of CRM studies with the LSST focuses on the ability to retrieve lag measurements as a function of cadences, which only involves simulating the light curves, artificially adding uniform noise, and varying the light curve's cadences to represent the LSST's data. This technique is straightforward and quick; however, the output may not accurately represent the true fidelity of the LSST data, as it does not fully account for the physical throughput associated with observing with the LSST.

Beyond simulating a field of astronomical objects, the literature has focused very little on using PhoSim to study specific astronomical transient objects, with one other implementation be-

ing from [Wang et al. \(2015\)](#) who simulated a population of Type 1A supernovae light curves on top of PhoSim. However, despite its efficiency, simulating an entire decade of the LSST, even a small sky region, remains computationally expensive, especially when high temporal resolution is required to resolve short-timescale variability, such as that seen in AGNs. By incorporating PhoSim into our light curve simulation pipeline, this project is bridging the gap between the theoretical models and realistic throughput currently missing in the literature.

A key limitation arises from PhoSim’s input requirements. It expects a static spectral energy distribution (SED) scaled to a given magnitude, which does not inherently account for time-dependent flux variability. Ideally, a unique spectrum would be generated for each epoch, accounting for every filter to match the time-varying flux in the AGN light curve. However, this is beyond the scope of this project. As a workaround, comparison stars are introduced into the field, and differential photometry is used to extract relative light curves from the synthetic images. This method enables us to track variability while compensating for systematic effects. However, it slightly reduces the signal-to-noise ratio due to the lower photon flux per epoch; thus, the results presented will provide an upper limiting case.

### 1.5. Data Production Pipeline

To achieve the goals mentioned above, I developed a light curve simulation pipeline of AGNs’ accretion discs, combining physical models of accretion disc reprocessing, as described in Sec. 2 with the estimated cadence and PhoSim’s realistic image generation describe in Sec. 3. Fig. 4 shows the flowchart of the different components considered to accurately simulate the AGNs’ light curves and the interplay between each component. The pipeline is designed to generate a set of light curves for a specific AGN. Therefore, the pipeline will have to be run multiple times to simulate a population of AGNs. Users provide the AGN’s physical parameters—black hole mass, Eddington ratio, redshifts, and the disc inclination—and the pipeline generates corresponding driving and response light curves, which are then input into PhoSim to produce synthetic LSST observations.

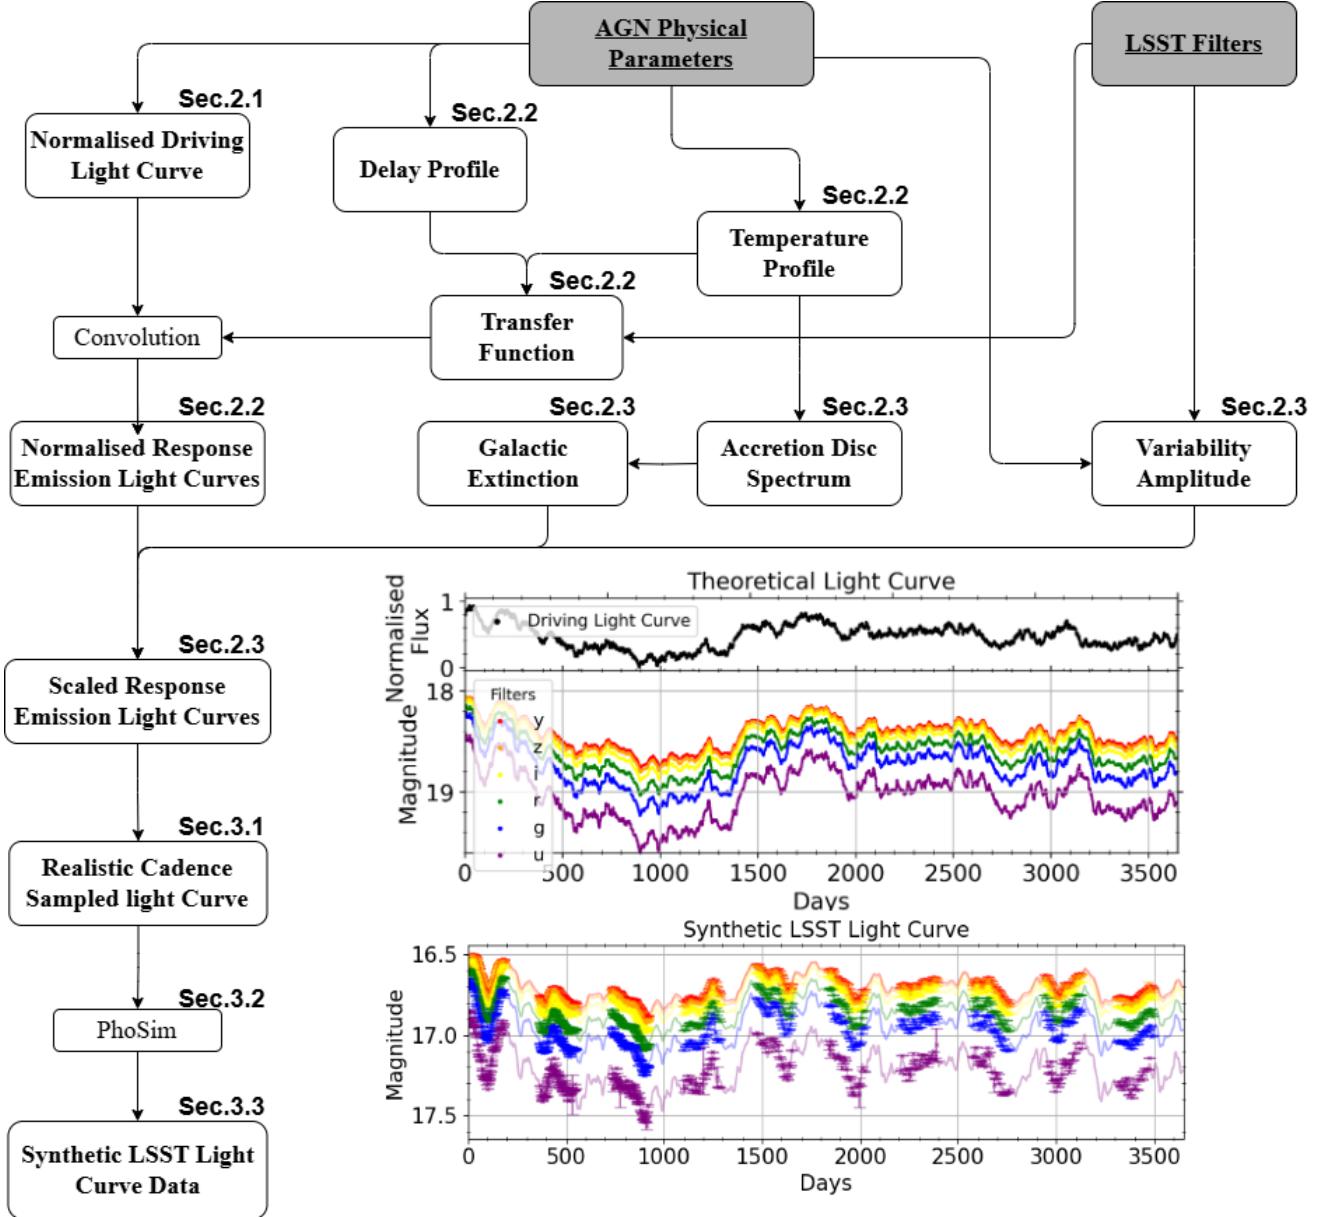
Band	Effective Wavelength (Å)	Limiting Magnitude	Colour in figure
u	3671	23.70	Purple
g	4827	24.97	Blue
r	6223	24.52	Green
i	7546	24.13	Yellow
z	8691	23.56	Orange
y	9712	22.55	Red

**Table 1.** Summary of the LSST’s optical bands, effective wavelengths used in the simulation, the  $5\sigma$  limiting magnitude, and the colours used to represent the light curves in each band. The limiting magnitude assumes a dark sky, zenith case, single 30s exposure visit. Since most light curves rarely overlay each other, the legend description will follow the same order as the light curves being plotted.

This allows us to analyse how different physical parameters affect lag recovery by the LSST. Additionally, the pipeline allows users to modify the temperature profile or input their transfer functions, see Sec. 2.2 for more detail, to incorporate a sophisticated accretion disc model, such as the Bowl model ([Starkey et al. 2023](#)). Appendix A explains the inner workings of PhoSim in further depth and the decisions made based on PhoSim’s performance to optimise the simulation processing time and number of parameters.

## 2. THEORETICAL LIGHT CURVE SIMULATIONS

To accurately simulate AGN light curves in different LSST filters, I must consider the physical properties of the AGN and the rest frame wavelengths associated with the filter’s wavelengths. Tab. 1 shows the 6 LSST filter bands, their effective wavelength, and the colours used in the plots. In this project, I will consider a series of AGNs with a mass of  $10^{8.5} M_{\odot}$  with Eddington ratios of 5%, 10%, 50%, and 100% at redshifts 0.5, 1, 2, and 3. These parameters are chosen from an AGN mass and Eddington ratio distribution from the AGN Variability Archive - AVA (JV Hernández Santisteban, Under Development), where the most abundant AGN has a mass of  $10^{8.5} M_{\odot}$  with an Eddington ratio of 10%, which I will treat as the fiducial case. The distribution of AGNs’ redshifts from [De Cicco et al. \(2021\)](#) in



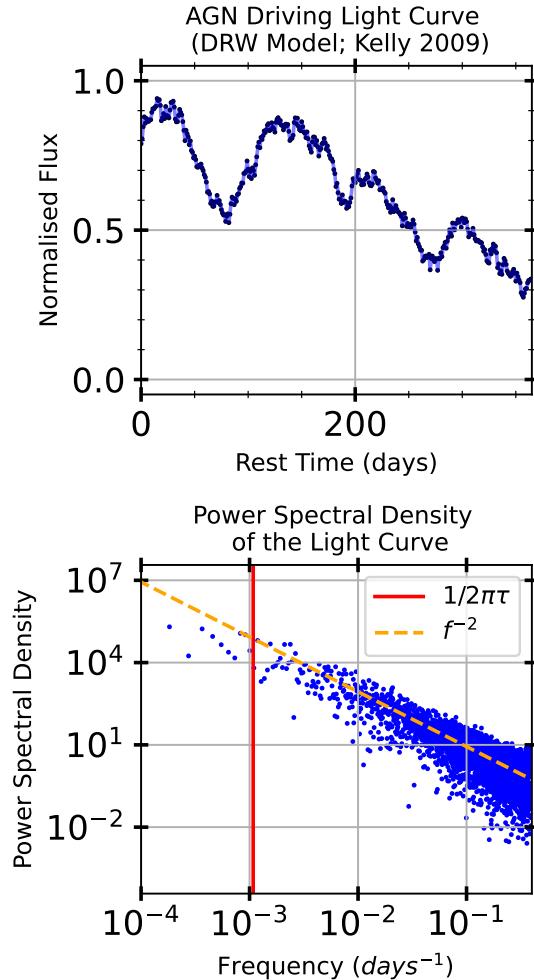
**Figure 4.** Flowchart of the developed pipeline. The pipeline requires the physical parameters of the AGN and the LSST filter's wavelength to produce the light curve. Each section where the different components are described in this dissertation is noted above each element. The theoretical light curve from pipeline [Top Plot] and the synthetic light curve after PhoSim [Bottom Plot] are shown.

the COSMOS field peaks between  $z = 0.5-1.0$  and then tails towards  $z = 4$ ; as such, I have selected the range above to test if a significant fraction of AGNs can be detected, and examine the possibilities of studying  $z > 1$  AGNs throughout the 10-year survey. In total, there are 24 sources simulated by PhoSim, 16 AGN cases, and 8 comparison stars. I will present the results of our

fiducial case throughout the main text from the modelled results within the pipeline. Following literature standards, I use a  $\Lambda$ -CDM model with  $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_\Lambda = 0.7$  and  $\Omega_m = 0.7$  to calculate our luminosity distance (Wright 2006).

### 2.1. Driving Ionisation Light Curves

The stochastic nature of AGN light curves can be modelled using the CAR(1) process as de-



**Figure 5.** The power spectral density (PSD) [Bottom Plot] and the driving light curve [Top Plot] generated using a damped random walk algorithm. The PSD follows the form  $P(\nu) \propto \nu^{-2}$  (dashed line) as measured from observational data (Giveon et al. 1999; Collier & Peterson 2001). The red vertical line indicates the cut-off time scale at which the PSD flattens, highlighting the damped nature of the light curve. The light curve plot displays the flux variation values (points) calculated using the CAR(1) algorithm, with the solid line representing the interpolated data.

scribed by Kelly et al. (2009), which generates the driving light curve in the time domain using a damped random walk algorithm. Note that this method has been used to model emission light curves, not just the driving light curve. Fig. 5 shows the driving AGN light curve and the power spectral density (PSD) generated by this method. It can be seen from the PSD that the variabil-

ity frequencies follow the form  $PSD(\nu) \propto \nu^{-2}$  as expected from AGN light curves (Giveon et al. 1999). This process is advantageous compared to the frequency modelling technique (Timmer & König 1995) as it directly models the variation, allowing us to fit this model to observed light curves. The following equation describes the CAR(1) process

$$dx(t) = -\frac{1}{\tau_{char}}x(t)dt + \sigma\sqrt{dt}\epsilon(t) + bdt \quad (2)$$

Where  $dx(t)$  corresponds to the change in the light curve flux.  $\tau_{char}$  is the characteristic time scale or relaxation time, which describes the time it takes for the long-term variation to return to the light curve's mean.  $\epsilon(t)$  is the white noise described by a Gaussian process. The mean and variance of the light curve are equal to  $b\tau_{char}$  and  $\tau_{char}\sigma^2/2$ , respectively. The  $\sigma$  term describes the long-term variability of the light curve; this value can be calculated from the structure function  $SF_\infty$ , which is the asymptotic rms variability amplitude on a long timescale and is equal to  $SF_\infty = \sqrt{2}\sigma$ . The relation from MacLeod et al. (2010) provides  $SF_\infty$  for an AGN of a given mass, redshift, and the absolute i-band magnitude, which is assumed to be 10% of the total absolute magnitude of the AGN. From this method, I can generate a series of unique AGNs by changing the random seed of the Gaussian process.

## 2.2. Response Emission Light Curve

The normalised response light curve  $F_c(t)$  is simulated by convolving a wavelength-specific transfer function  $\psi(\tau|\lambda)$  over the normalised driving X-ray light curve  $F_x(t)$ ,

$$F_c(\lambda, t) = \int_0^\infty \psi(\tau|\lambda)F_x(t - \tau)d\tau \quad (3)$$

The mean of the transfer function dictates the time lag observed between the driving and response light curve. The transfer function takes into consideration the blackbody spectrum of the disc  $B_\nu$ , the driving luminosity  $L_*$ , the temperature  $T$ , and the delay profile  $\tau(r, \phi, i)$  of the disc. The blackbody spectrum of the disc is calculated by summing up the Planck function across the entire disc.

$$\psi(\tau|\lambda) = \int_{r_{in}}^{r_{out}} \frac{\partial B_\nu}{\partial T} \frac{\partial T}{\partial L_*} \delta(\tau - \tau(r, \phi, i)) d\Omega \quad (4)$$

The temperature profile  $T(r)$  of the disc can be calculated by considering the combined effects of viscous heating due to disc accretion and irradiation heating from the X-ray corona. (Cackett et al. 2007):

$$T^4 = \frac{3GM\dot{M}}{8\pi\sigma r^3} \left(1 - \sqrt{\frac{r_{in}}{r}}\right) + \frac{L_b(1-a)h_x}{4\pi\sigma(r^2 + h_x^2)^{3/2}} \quad (5)$$

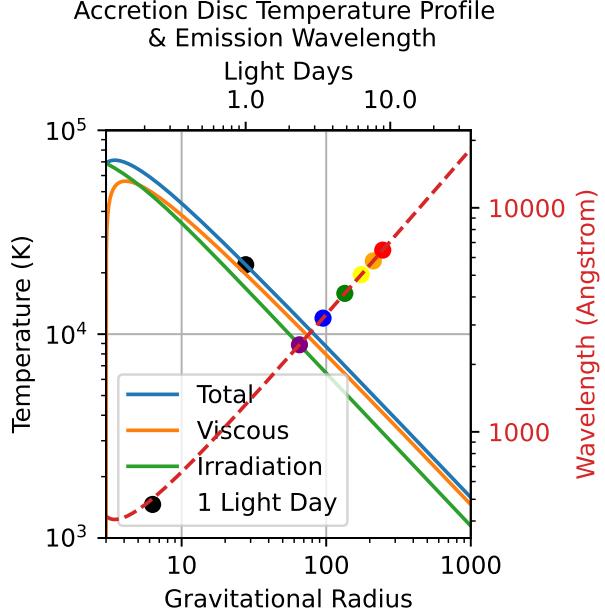
Where  $G$  is the gravitational constant,  $M$  is the SMBH mass,  $\dot{M}$  is the accretion rate,  $\sigma$  is the Stefan-Boltzmann constant,  $L_b$  is the bolometric luminosity of the corona, and  $a$  is the albedo of the disc.  $r_{in}$  is the radius of the most stable orbit, usually measured in units of Gravitational Radii  $r_g$ , where  $r_g$  is defined as  $GM/c^2$ . Here, I set  $r_{in} = 3r_g$  for a Schwarzschild black hole (Starkey 2017). Fig. 6 shows the modelled temperature profile as described in Eq. 5, as well as the peak thermal emission from Wien's law, which shows the radii at which each wavelength is emitted on the accretion disc.

For a simple Keplerian disc with the corona at a height  $h$  above the disc (Sergeev et al. 2005), the delay profile is

$$\tau(r, \phi, i) = \frac{1}{c}(\sqrt{r^2 + h_x^2} + r \sin(i) \cos(\phi) + h_x \sin(i)) \quad (6)$$

Where  $i$  and  $\phi$  are the inclination and azimuthal angles, respectively, and  $r$  is the radius of the disc. Fig. 7 shows the delay profile of the disc plotted at different inclination angles.

A modified PyCream function from Starkey (2017) combines the aforementioned results using the temperature of the disc at one light day radius, the AGN's physical parameters, and the function's default arguments to generate the transfer functions for a particular AGN at specified wavelengths. I now have all the components required to generate the response emission light curves as stated in Eq. 3. Fig. 8 shows the generated transfer functions with respect to the rest frame of the AGN (top panel) and the response thermal emission from convolving the transfer function over the driving light curve from

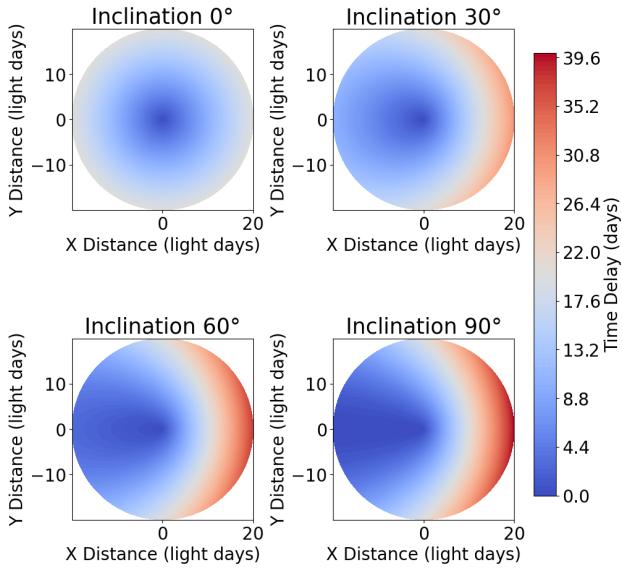


**Figure 6.** Accretion Disc's Temperature Profile and the thermal emission wavelength as a function of the disc's radius in gravitational radii and light days. The solid lines in the plot represent the viscous temperature resulting from accretion, the irradiation temperature due to the corona's X-ray heating, and the disc's total temperature. The dashed red line shows the peak thermal emission wavelength from Wien's law, while the coloured points show the radius of the disc each LSST filter is observing.

Fig. 5 (bottom panel). The transfer function effectively makes the different flux variation features less defined, the wider the transfer function. This is because the response emission from higher wavelengths is coming from larger disc annuli. To account for the time dilation due to cosmological redshift, each time step is multiplied by a factor of  $(1+z)$  to achieve the observed frame lag.

### 2.3. Light Curve Scaling

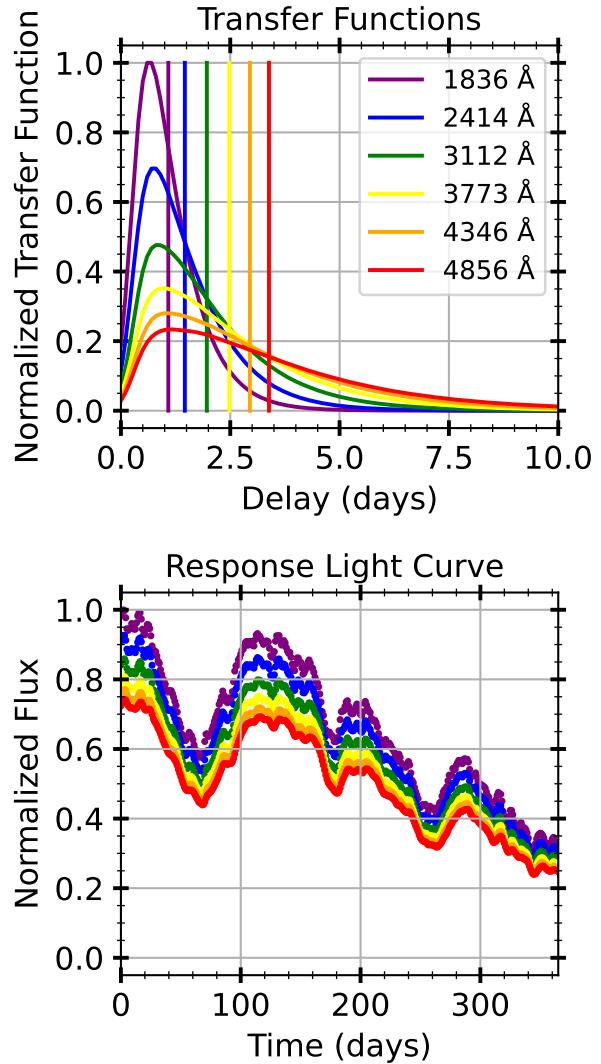
From the previous section, I now have normalised response emission light curves. To simulate the flux measured by our instruments, I need to calculate the background spectrum of the disc to find the mean magnitude of the light curve and the magnitude variability amplitude of the observed light curve. The mean magnitude can be calculated by measuring the flux of the accretion disc spectrum at the rest wavelength. Using



**Figure 7.** Accretion disc's delay profile due to corona source at height  $h$  above the disc from (Sergeev et al. 2005). The effect of inclination on the delay time is clearly shown, where the delay time for the part of the disc inclined away from the observer is larger, as light has to travel to that region and back to the observer. As such, the flux variation in the response emission at a given annulus from the closer side of the disc will arrive before the same flux variation for the further side.

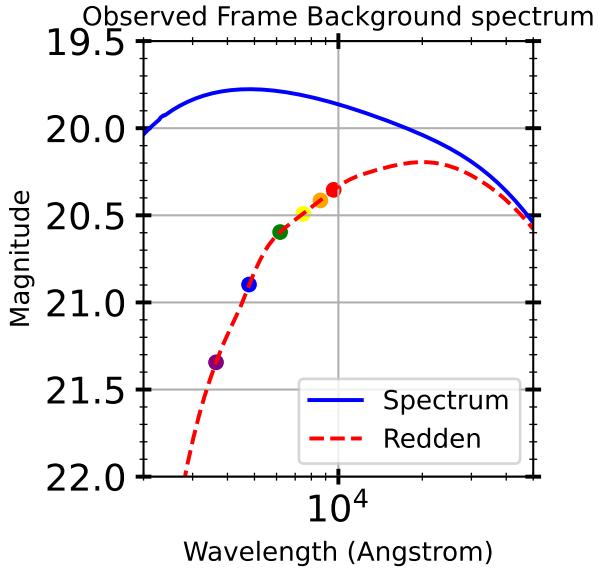
the method described in Sec 2.2, I can generate the accretion disc spectrum in the rest frame of the AGN. The disc spectrum is then redshifted with a factor of  $(1+z)$  accordingly into the observed frame. A SMC-like extinction model (Gordon et al. 2003, 2016) is then used to simulate the host galaxy extinction on the disc spectrum as per the findings from Weaver & Horne (2022). I then measured the flux value at the effective wavelength of each LSST filter. Fig. 9 shows the accretion disc spectrum in the observed frame before and after it has been extinct. The coloured dots represent the effective wavelength of each LSST band and where the flux measurements were made.

Using the fitted relation from MacLeod et al. (2010) which used data from Sloan Digital Sky Survey (Vanden Berk et al. 2004; Gunn et al. 2006) Stripe 82 AGNs (Bramich et al. 2008), I can calculate the structure function  $SF_\infty$  as described in Sec. 2 and convert the structure function to the standard deviation of the expected light curve. I



**Figure 8.** The transfer functions of the accretion disc at the rest wavelength corresponding to the LSST filters [Top Plot] and the response thermal emission light curve from convolving the transfer functions over the driving light curve in Fig 5 [Bottom Plot].

then scale the convolved light curve to have the mean value as the magnitude of the accretion disc spectrum and standard deviation from the fitted relation from MacLeod et al. (2010). Fig. 4 shows the scaled light curves in different bands. At this point, I have successfully simulated the accretion disc response emission light curves. Within the pipeline, the light curves are saved into individual CSV files for further processing to prepare the data for PhoSim's input.



**Figure 9.** Redshifted and SMC-like extinct accretion disc spectrum generated from the temperature profile (solid blue line) and the extinct spectrum (dashed red line), as well as the coloured points showing the location on the spectrum each filter observes. The theoretical spectrum flux follows the  $\lambda^{-1/3}$  relation in the Rayleigh-Jean limit and the  $\lambda^{-2}$  relation following Wein’s displacement law.

### 3. LIGHT CURVE SAMPLING AND REALISTIC THROUGHPUT

#### 3.1. LSST Survey Strategy

The current LSST’s baseline survey strategy (Jones et al. 2025) consists of 4 different missions: the Wide Fast Deep (WFD), Mini and Micro Surveys, Deep Drilling Fields (DDFs), and Targets of Opportunity (ToOs). Fig. 3 shows the map of the observable region of the LSST and the estimated numbers of visits in each field. Each point in the sky will be assigned a 6-month active monitoring period, consistent with when it is observable. The WFD will receive most of the coverage with an expected 800 visits per point across the 10-year survey. Observations will be done in pairs of bands (u+g, u+r, g+r, r+i, i+z, z+y, y+y) every 2-4 days, thus providing us with 14-day cadences per band in the best-case scenario. The single visit per point per night with such a large observing gap would make it unsuited for continuum reverberation mapping. In contrast, the DDFs are expected to receive 10-20 visits per night per

band; however, the cadence between each observing night is yet to be finalised. Tab. 2 summarises the different DDFs, the predicted total numbers of visits as of February 2025 from the Baseline simulation version 4.3<sup>1</sup>, and the average cadence for each DDF considering the average number of 15 visits per band per night and the 6-month active observing period. The COSMOS field will receive double the number of visits, with half of the visits completed within the first three years of the survey as part of a pathfinder and data management program. This gives the COSMOS field the highest cadence of 2 and 5 days for the first 3 and final 7 years, respectively. The ELAIS1, XMM-LSS, and ECDFS should have a cadence of around 7 days, and the EDSF field, split into two “a” and “b” pointings, will have a cadence of around 9 days. Henceforth, I will refer to the other DDFs, excluding the EDSF subfields, as the “low cadence DDFs”. For my simulation, I will sample data points every 2 days in the COSMOS field for the pathfinder period in the first 3 years. I will sample every 6 days for both the nominal period of the COSMOS field and the entire duration of the low cadence DDFs. This allows us to replicate the low cadence DDFs’ cadence by sampling every 3rd data point from the COSMOS field pathfinder period without having to rerun a new simulation for the low cadence DDFs.

Since I am interested in the most extreme case, I only simulated one image per observing night. This would heavily reduce computation time with PhoSim and ensure that if the randomly generated observing condition is poor for an evening, I assume no further observing will be conducted on the same object in the band.

#### 3.2. Synthetic Images

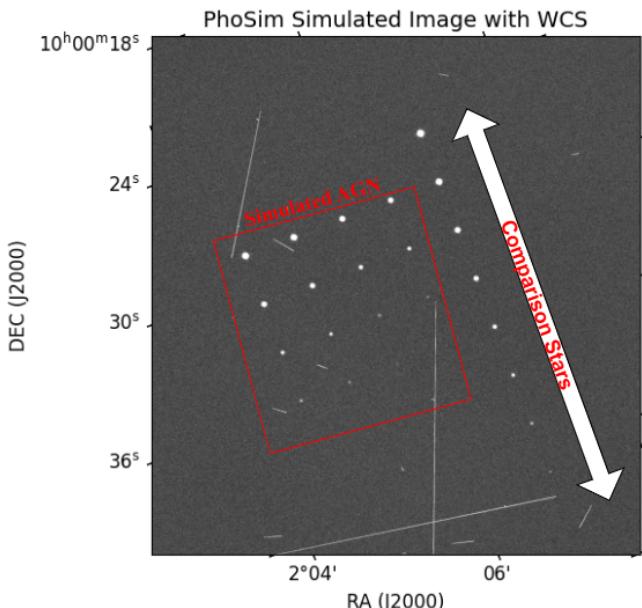
Here, the individual light curves are compiled together with the 6-month seasonal gap masked and sampled according to the above cadence. The data from each epoch is then converted into PhoSim’s input file<sup>2</sup>. The input file requires the user to provide the pointing sky coordinates (Ra and Dec), the exposure time (in seconds), the observing filter (filter index), the number of exposures, and the individual objects to be simulated

<sup>1</sup> <https://survey-strategy.lsst.io/baseline/index.html>

<sup>2</sup> <https://www.phosim.org/documentation/instance-catalog>

Field	RA	Dec	Average Cadence (days)	Estimated Total Visits
COSMOS	150.11	2.23	2 and 5*	45649
ELAIS1	9.45	-44.02	7	22769
XMM-LSS	35.57	-4.82	7	23330
ECDFS	52.98	-28.12	7	23712
EDSF-a	58.9	-49.32	9	11824
EDSF-b	63.6	-47.6	9	11779

**Table 2.** Information of LSST’s Deep Drilling Fields (DDFs) and their respective estimated total visits as predicted by version 4.3 of the LSST’s survey strategy. The average cadence for each DDF is calculated by considering the average 15 visits per band per DDF observing night and a 6-month nominal active observing period. The EDSF is a larger field split into “a” and “b” pointings. This project will not consider these two fields due to the lower cadence. \*The COSMOS field will receive double the visits, with half the visits completed within the first 3 years of the survey. As such, I used two different cadences when simulating this field’s light curve.



**Figure 10.** Example of a synthetic 30s exposure i-band electron image generated from PhoSim. The red square outlines the  $4 \times 4$  grid of simulated AGNs with descending Eddington ratio across the columns, and increasing redshift across the rows. The 8 comparison stars with magnitudes ranging from 16 to 24 are shown next to the white arrows. Other sources in the image are due to the simulated cosmic rays hitting the sensor.

in the field for the software to function properly. If the exact observing time is not provided or an invalid time is provided, PhoSim will generate a random date and time where the object is visible and has an airmass  $<2$ . Unless specified, PhoSim also generates a random observing condition. To replicate the observing strategy, the weather, and to ensure that the required airmass is consistent, I allowed PhoSim to select its own observing time and weather rather than predicting the actual epoch ourselves, which may not be consistent with the survey strategy. This method allows the user to quickly adjust their desired cadence while achieving realistic photometric results.

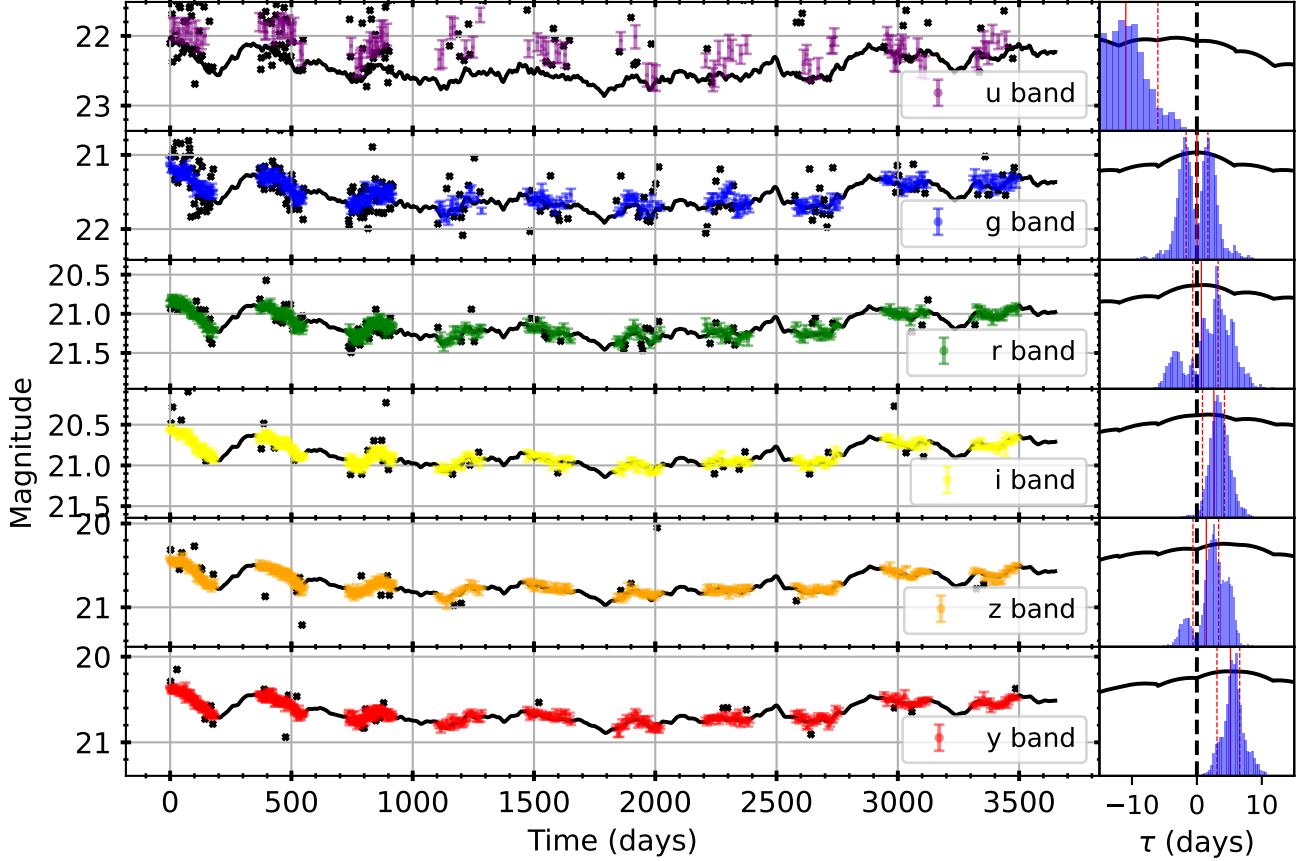
Lastly, PhoSim also accounts for the location of the Sun and Moon and ensures optimal distance between them and the observing target when generating an observing time to ensure as little effect as possible. As such, I set the altitude of the Sun and Moon to  $-90$  deg as the atmospheric glow increases the simulation run time significantly.

To simulate each object, the user must provide the object’s coordinates (Ra and Dec), AB magnitude, SED profile, and spatial modes. Since I am simulating quasars at high redshifts, I treat the quasars’ spatial mode as point sources and set every other parameter, i.e. shear and lensed effects, to none. Additionally, since I have accounted for the redshifts in the light curve simulation, I also set the redshift of the source to zero. As mentioned in Sec. 1.4, I set the input magnitude in each band to be the magnitude value from our simulated light curve at a given epoch. Fig. 10 shows an example of an i-band image simulated by PhoSim. The simulated AGNs are placed into a grid with varying redshifts and Eddington ratios on each axis. A set of 8 comparison stars with magnitudes ranging from 16 to 24 is placed next to the grid for differential photometry.

### 3.3. Image Processing and Data Reduction

The simulated images from each frame were then processed using AstroImageJ (Collins & Kielkopf 2013) for differential photometry, as the WSC values for the images were not accurate enough for processing by AstroPy. Additionally, AstroImageJ allows the aperture to be manually

### Fiducial Case Processed Light Curve with CCFs



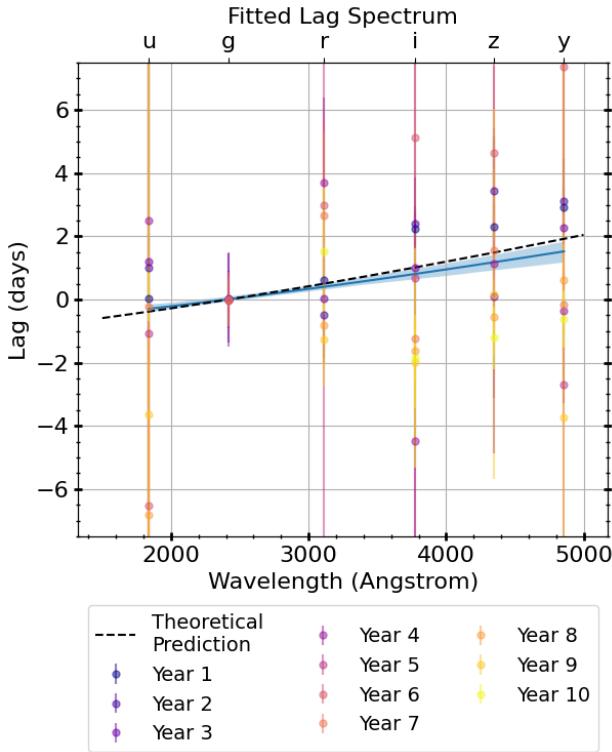
**Figure 11.** Processed fiducial case AGN light curves in the 6 LSST filters [left panels] with the cross correlation functions results (black lines) and centroid distribution (blue) anchored on the g-band [right panels] from year 1. The light curve plot shows all the flux measurement data points with error bars plotted above the theoretical input light curve (black line), with the masked data points plotted in black. See Sec. 1.5 for more details. In this case, the u-band is reaching its limiting magnitude, resulting in sparse data and poor CCF results. The increasing lag trend can be seen as the wavelength of each band increases.

placed over points in the AGN grid, including those that peak-finding functions would not pick up. This enables consistent data output across all bands, facilitating easier processing. The auto variable aperture setting was selected for the differential photometry, the gain setting was set to  $3 \text{ e}^-/\text{count}$  and the CCD readout noise was set to  $5 \text{ e}^-$ ; these are the average values across all the LSST's CCD chips (Collaboration et al. 2009).

#### 4. TIME SERIES ANALYSIS

To use Pyccf (Peterson et al. 1998; Sun et al. 2018) to calculate the interband lags, I must en-

sure that outliers are removed, as the measured lag is highly sensitive to outliers. First, an error cut was performed, removing data points with errors larger than 0.2 mag. Secondly, the median value across each light curve was calculated, and extreme outliers with values less than or more than 4% from the median were removed. This significantly helps with the running optimal average (ROA) clipping, which then removes the smaller outliers. The RunningOptimalAverage function from PyROA (Donnan et al. 2021) was used to calculate the ROA of the light curve. Through trial and error, a window of 0.4% around the ROA



**Figure 12.** Fiducial case fitted lag spectrum. Each lag measurement from each year is plotted with the darker colours associated with earlier years. The dashed black line shows the theoretical spectrum derived from the transfer functions. The fitted model is shown in blue with its  $1-\sigma$  uncertainty window.

was identified to remove small outliers iteratively, and it was found to be effective and efficient for cases with sufficient data. Fig. 12 shows the processed light curves separated in all bands for our fiducial case AGN with their CCF results. Appendix B shows the processed light curves for all other cases. The magnitude is then converted to flux,  $F_\nu = 3631 \times 10^{-0.4mag} [mJy]$  for the PyCCF calculations with their respective errors calculated by multiplying the converted flux with the derivative of the conversion and magnitude error.

The processed light curves are then split into yearly samples and have their interband lags calculated with a window of  $[-100, 100]$  days,  $\text{sigmode}$  of 0.2,  $\text{nsim}$  of 10000, and only using RF sampling ( $\text{mcmode} = 2$ ), following Cackett et al. (2018). The g-band is used as the reference band for redshifts 0.5 and 1, and the r-band for redshifts 2 and 3, as the former approaches its lim-

iting magnitude. The centroid lag, which is the most robust case, is used for the analysis. The observed lag and wavelength are then corrected for their redshift by a factor of  $(1+z)^{-1}$  to retrieve their rest frame lags. The lag and wavelength are then plotted to create wavelength-dependent continuum lags for each case. To measure how well and accurately the lag is retrieved, I fit the lag-wavelength relation in Eq. 1 to retrieve  $\tau$  and compare it against the theoretical  $\tau_0$  calculated from fitting the same relation over the theoretical lags calculated from the mean of the respective transfer functions. Using the emcee package (Foreman-Mackey et al. 2013), I run an MCMC fit on our measured lag, leaving  $\tau$  and  $y_0$  as free parameters. To demonstrate the LSST’s quality of analysis as the survey progresses, I fit the lag-wavelength relation to the cumulative lag measurement year by year. Fig. 12 shows the fitted lag spectrum to the cumulative lag measurements across all 10 years of our fiducial case.

For the COSMOS field, Fig. 13 shows how the quality of the measured  $\tau$  compares with the theoretical  $\tau_0$ , measured by calculating  $\log_{10}(\tau/\tau_0)$  evolves year by year for each Eddington ratio at a fixed redshift. Ideally, I want each case to converge towards zero. Cases where there are breaks in the plot are due to the measured  $\tau$  being negative. Not surprisingly, the low redshifts (0.5 and 1) and high Eddington ratios (50% and 100%) consistently maintained a high-quality fit across all years. This plot suggests that the LSST may struggle to retrieve the lag accurately around  $z = 2$ . For the low cadence DDFs, Fig. 14 shows the same fit comparison as the COSMOS field. In this case, it suggests that the quality of the recovered lags begins to drop for the dimmer cases at  $z = 1$ . Appendix C shows the yearly fits in parameter space for both the COSMOS and low cadence DDFs cases.

Using the AGNs from the SDSS Reverberation Mapping (SDSS-RM) campaign (Shen et al. 2015; Homayouni et al. 2019), which has accurate redshift measurements, I can estimate the density of AGNs up to a certain redshift in a given field by assuming the same density of AGNs in every part of the sky. Tab. 3 shows the density of AGNs per  $\text{deg}^2$  up to a certain redshift from the SDSS-RM campaign. From our lag recovery results, I

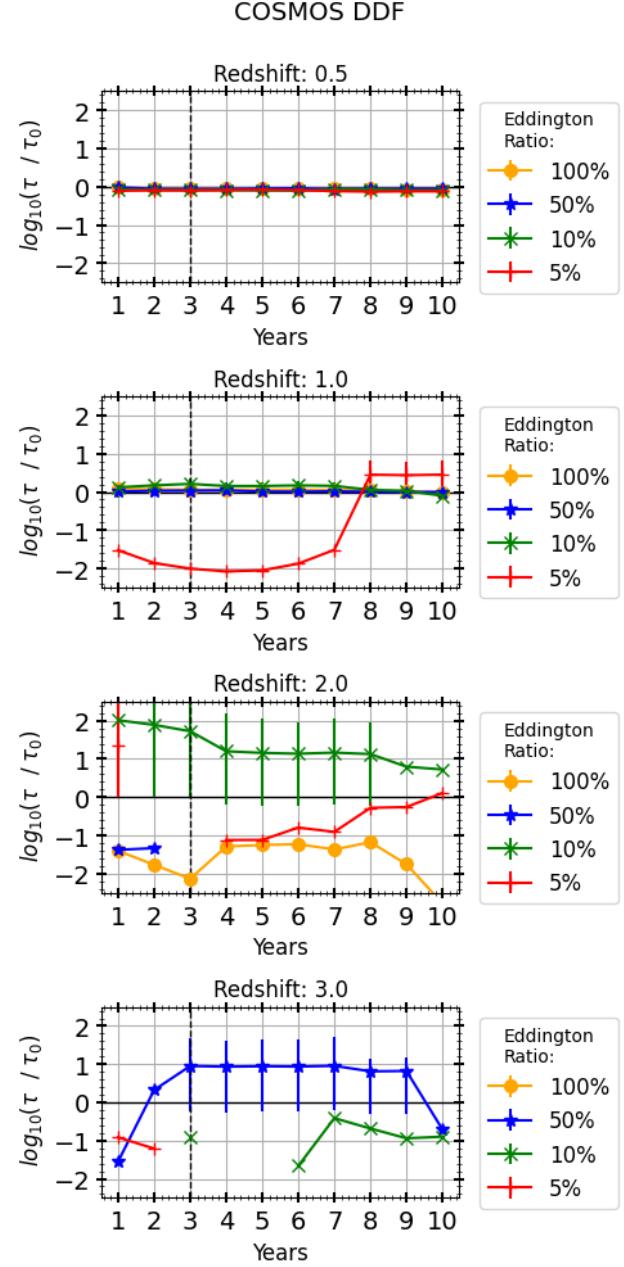
Redshift Range	Numbers per deg <sup>2</sup>	Expected Number of AGNs
0 < $z$ < 0.5	5.8	56
0.5 < $z$ < 0.7	5.5	53
0.7 < $z$ < 1.0	15.7	151
1.0 < $z$ < 1.5	24.9	239
1.5 < $z$ < 2.0	34.4	330

**Table 3.** Number density of AGNs in the SDSS-RM field for each redshift range. By assuming the same AGN distribution in each of the DDFs, the expected number of AGNs with recoverable lags can be predicted.

assume that AGNs at  $z < 1.5$  for the COSMOS field and  $z < 0.7$  for the low cadence DDFs are recoverable. Multiplying the number density by the field of view of the LSST image of  $9.6\text{deg}^2$  I find that the estimated number of AGNs with recoverable lags is 499 and 325 for the COSMOS and the combined low cadence DDFs, respectively. This brings the total number up to 824 AGNs after the 10-years survey. Our prediction of AGNs with recoverable lag in the COSMOS field is consistent with the prediction by Kovačević et al. (2022) of 523 AGNs with recoverable lag, which uses De Cicco et al. (2021) AGNs catalogue taken with the VLT Survey Telescope (VST) (Arnaboldi et al. 1998) in the COSMOS field. Using my results, I could provide further precision to Kovačević et al. (2022) by predicting the exact number of AGNs with recoverable lags for each redshift range using the SDSS-RM AGNs density as shown in Tab. 3. Furthermore, using the latest COSMOS AGN catalogue, with the known Type 2 AGNs removed from De Cicco et al. (2022), I could expect up to 3801 AGNs in the COSMOS field and 5356 AGNs in all the low cadence DDFs combined.

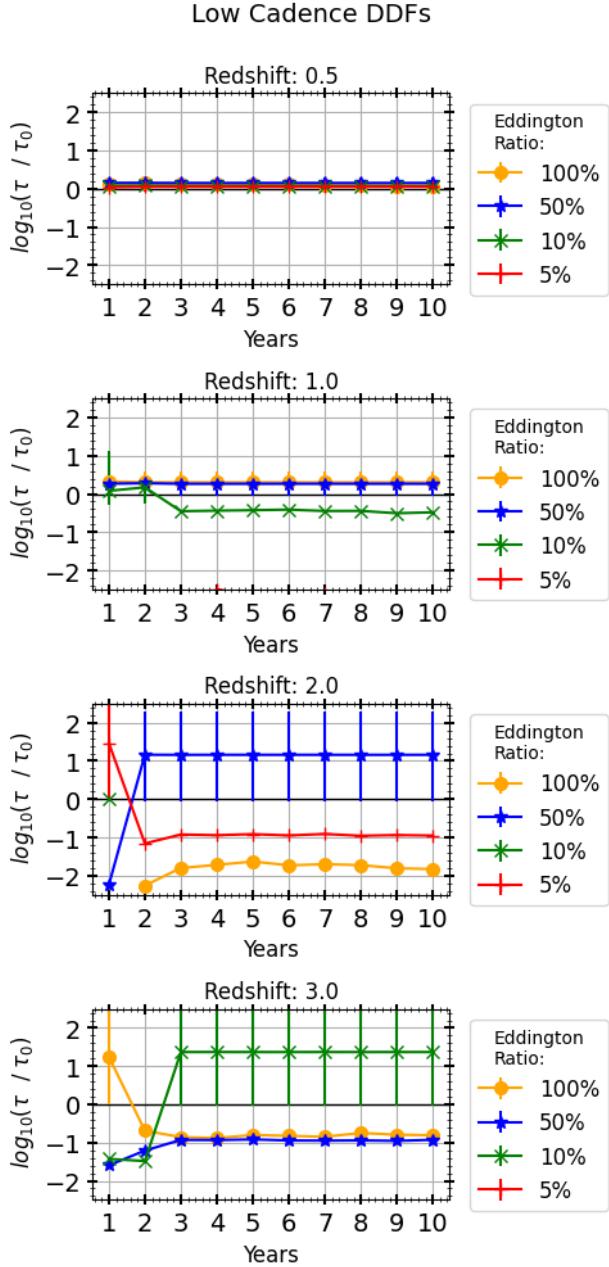
## 5. DISCUSSION

In this work, I have shown that it is possible to generate reasonably high-fidelity synthetic AGN light curves for the LSST and recover the lag. It should be reiterated that the data presented in this project considers the most conservative case, with abnormally high data censoring, a simplistic processing pipeline, and under-estimated flux measurements due to the limitations of the ray tracing software. In reality, the DDFs of the



**Figure 13.** COSMOS Field's year-by-year  $\tau$  measurement at fixed redshift ratio as a function of the Eddington ratios. The black vertical dashed line at year 3 signifies the change in cadence from 2 to 6 days. Recovery of lags begins to struggle at  $z=2$ .

LSST allow for multiple exposures per band per visit, enabling coadded images that significantly improve photometric measurements and increase the limiting magnitude (Jones et al. 2025). This would enable a better signal-to-noise ratio in all



**Figure 14.** ELAIS1, XMM-LSS, and ECDFS Field’s year-by-year  $\tau$  measurement at fixed redshift ratio as a function of the Eddington ratios. Recovery of lags begins to struggle at  $z=1$ .

cases, allowing for more accurate measurement of the flux variation in the dimmer cases that border the limiting magnitude of each band.

The LSST will provide an unprecedented observing depth, allowing us to study most AGNs at redshift 1 at a range of Eddington ratios, which is

only currently possible with very high Eddington ratio AGNs at high redshift. Since the Eddington ratio affects the brightness and structure of the disc (Abramowicz & Fragile 2013), observing low Eddington ratios at high redshifts will enable us to apply our current understanding and models of the accretion disc at low Eddington ratios and test for any changes at varying redshifts, as well as expand our knowledge of how the accretion disc structure changes at high Eddington ratios. Particularly, we will have access to the inner region of these accretion discs from the rest-frame UV emissions, and test if there are any changes at varying redshifts.

The breadth and depth of this project were heavily influenced by the limitations of PhoSim, where the decision was made to focus on a range of parameters rather than simulating a set of light curves. By using only one set of light curves, every AGN case with the same Eddington ratio has the same section of the light curve analysed. This results in a slight bias toward either the positive or negative fit in the redshift 0.5 and 1 cases in Fig. 13 and 14. Additionally, certain trends from the light curve would also appear in the year-by-year fit as seen in the 50% Eddington at redshift 3 from Fig. 13, where the fitted  $\tau$  value increases in the first three years, plateaus, and then decreases. This is due to the long-term variation in the light curve decreasing below the limiting magnitude in year 3. This results in no new lag measurements until the AGN was bright enough for observation again in the final year. These biases and trends are also present in Pozo Nuñez et al. (2023) results for individual light curves. Provided more computational power, a larger set of light curves for the same AGN parameters should be simulated to remove any biases associated with specific light curves.

While this project focuses solely on the light curve from the accretion disc, moving forward, the diffuse continuum emission from the BLR should also be accounted for in future work, as it has been shown to cause the lag measurement to be overestimated (Chelouche et al. 2019). Now that the redshift range is constrained, the host galaxy’s contribution should be considered using techniques such as the FVG method (Choloniewski 1981). These improvements would also allow fur-

ther work on accurately simulating the AGN's spectra, which is required for PhoSim.

Furthermore, the integration of the LSST Operation Simulator (OpSim) (Delgado et al. 2014), which provides the actual observing strategy for a given survey model, and Catalogue Simulator (CatSim) (Connolly et al. 2014), which provides a realistic field of objects including stars and galaxies, would allow for the final output to be even more realistic. These additions could be made at the data sampling and PhoSim input generation stage, and thus would not affect the primary light curve generation pipeline.

Lastly, the nature of PhoSim and the developed pipeline could be used to simulate other large-scale photometric surveys or long-term monitoring campaigns of AGNs, provided the optical system files are available for PhoSim. The optimisation steps taken in this project allowed for a significant reduction in PhoSim's processing time compared to LSST DESC DC1 (Sánchez et al. 2020),  $\sim$ 6 minutes per CCD vs  $\sim$ 66 minutes per CCD or approximately 300 CPU hours vs 3300 CPU hours in total to produce 10 years of data in 6 optical bands.

## 6. CONCLUSION

This dissertation presents the development and implementation of a relatively quick synthetic AGN light curve simulation pipeline designed to test the feasibility of conducting continuum reverberation mapping with the LSST. The study focuses on quantifying how the interplay between black hole mass, redshift, and Eddington ratio affects the ability to recover interband lags from AGN light curves, given LSST's expected cadence and photometric depth.

To achieve this, synthetic light curves are generated by convolving wavelength-dependent transfer functions derived from a thermally reprocessing accretion disc model over a driving ionisation light curve from a damped random walk algorithm. The response light curves are then scaled using empirical structure functions and passed PhoSim to generate synthetic LSST images, which simulate realistic instrumental effects and observing conditions.

Differential photometry is utilised to extract flux measurements from simulated images, and

time-series cross-correlation analysis is used to measure wavelength-dependent continuum lags. The results are compared against the theoretically predicted lag-wavelength relation  $\tau(\lambda) \propto \lambda^{4/3}$ , enabling a quantitative assessment of CRM performance across the parameter space.

Findings indicate that the COSMOS field, with its higher cadence, allows for reliable lag recovery up to  $z = 1.5$ , especially for high Eddington ratio AGNs. Lower cadence DDFs show diminished performance beyond  $z = 0.7$ , but still yield recoverable lags in brighter cases. These results suggests that continuum reverberation mapping will be feasible in LSST's deep drilling fields on hundreds of AGNs, including those at high redshifts, offering a unique opportunity to map the structure of accretion discs and study SMBH growth under varying physical parameters. Future work could include contributions from the diffuse BLR continuum and host galaxy contamination to enhance spectral realism further for PhoSim.

## 7. ACKNOWLEDGEMENTS

First and foremost, I want to thank my project supervisor, Dr Juan V. Hernández Santisteban, with whom I had the privilege of working for the past 3 years. His continuous support, endless patience, valuable advice, and guidance in both academic and pastoral matters have helped me survive this degree.

This semester would have been much more lonely if not for the weekly Echo Mapping Group meetings. I want to express my gratitude to all the group members in no particular order. I want to thank Keith Horne for his insights and comments to make sure my work is up to par. Thank you to Aditya Sinha, Farah (Tori) Wallauer, and Reinosuke (Luke) Kusano for helping me with any and all coding-related and technical issues, as well as our post-meeting PhD applications panic. Thank you to Omar Abdelmonem for being there from the start of university. Finally, thank you to Roberta Vieliute for technical assistance and the weekly badminton sessions.

Lastly, I want to acknowledge the people in my life. I would like to express my gratitude to my parents for their support and to my siblings for encouraging me to visit home. I want to thank all

my friends with whom I spent countless hours at the library, shared pints, and commiserated with. Special shout-out to my flatmate Andrew Chingos for putting up with my unhealthy working habits.

## REFERENCES

Abolfathi, B., Alonso, D., Armstrong, R., et al. 2021, The Astrophysical Journal. Supplement Series, 253, doi: [10.3847/1538-4365/abd62c](https://doi.org/10.3847/1538-4365/abd62c)

Abramowicz, M. A., & Fragile, P. C. 2013, Living Reviews in Relativity, 16, 1, doi: [10.12942/lrr-2013-1](https://doi.org/10.12942/lrr-2013-1)

Antonucci, R. 1993, ARA&A, 31, 473, doi: [10.1146/annurev.aa.31.090193.002353](https://doi.org/10.1146/annurev.aa.31.090193.002353)

Arnaboldi, M., Capaccioli, M., Mancini, D., et al. 1998, The Messenger, 93, 30

Blandford, R. D., & McKee, C. F. 1982, ApJ, 255, 419, doi: [10.1086/159843](https://doi.org/10.1086/159843)

Bramich, D. M., Vidrih, S., Wyrzykowski, L., et al. 2008, MNRAS, 386, 887, doi: [10.1111/j.1365-2966.2008.13053.x](https://doi.org/10.1111/j.1365-2966.2008.13053.x)

Burrows, D. N., Hill, J., Nousek, J., et al. 2005, Space science reviews, 120, 165

Cackett, E. M., Bentz, M. C., & Kara, E. 2021, iScience, 24, 102557, doi: <https://doi.org/10.1016/j.isci.2021.102557>

Cackett, E. M., Chiang, C.-Y., McHardy, I., et al. 2018, ApJ, 857, 53, doi: [10.3847/1538-4357/aab4f7](https://doi.org/10.3847/1538-4357/aab4f7)

Cackett, E. M., Chiang, C.-Y., McHardy, I., et al. 2018, The Astrophysical Journal, 857, 53, doi: [10.3847/1538-4357/aab4f7](https://doi.org/10.3847/1538-4357/aab4f7)

Cackett, E. M., Horne, K., & Winkler, H. 2007, MNRAS, 380, 669, doi: [10.1111/j.1365-2966.2007.12098.x](https://doi.org/10.1111/j.1365-2966.2007.12098.x)

Chan, J. H. H., Millon, M., Bonvin, V., & Courbin, F. 2020, A&A, 636, A52, doi: [10.1051/0004-6361/201935423](https://doi.org/10.1051/0004-6361/201935423)

Chelouche, D., Pozo Nuñez, F., & Kaspi, S. 2019, Nature Astronomy, 3, 251, doi: [10.1038/s41550-018-0659-x](https://doi.org/10.1038/s41550-018-0659-x)

Choloniewski, J. 1981, AcA, 31, 293

Collaboration, L. S., Abell, P. A., Allison, J., et al. 2009, LSST Science Book, Version 2.0. <https://arxiv.org/abs/0912.0201>

Collier, S., & Peterson, B. M. 2001, ApJ, 555, 775, doi: [10.1086/321517](https://doi.org/10.1086/321517)

Collins, K., & Kielkopf, J. 2013, AstroImageJ: ImageJ for Astronomy, Astrophysics Source Code Library, record ascl:1309.001

Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9150, Modeling, Systems Engineering, and Project Management for Astronomy VI, ed. G. Z. Angeli & P. Dierickx, 915014, doi: [10.1117/12.2054953](https://doi.org/10.1117/12.2054953)

De Cicco, D., Bauer, F. E., Paolillo, M., et al. 2021, A&A, 645, A103, doi: [10.1051/0004-6361/202039193](https://doi.org/10.1051/0004-6361/202039193)

—. 2022, A&A, 664, A117, doi: [10.1051/0004-6361/202142750](https://doi.org/10.1051/0004-6361/202142750)

De Rosa, G., Peterson, B. M., Ely, J., et al. 2015, ApJ, 806, 128, doi: [10.1088/0004-637X/806/1/128](https://doi.org/10.1088/0004-637X/806/1/128)

Delgado, F., Saha, A., Chandrasekharan, S., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9150, Modeling, Systems Engineering, and Project Management for Astronomy VI, ed. G. Z. Angeli & P. Dierickx, 915015, doi: [10.1117/12.2056898](https://doi.org/10.1117/12.2056898)

Donnan, F. R., Horne, K., & Hernández Santisteban, J. V. 2021, MNRAS, 508, 5449, doi: [10.1093/mnras/stab2832](https://doi.org/10.1093/mnras/stab2832)

Edelson, R., Gelbord, J. M., Horne, K., et al. 2015, ApJ, 806, 129, doi: [10.1088/0004-637X/806/1/129](https://doi.org/10.1088/0004-637X/806/1/129)

Edelson, R., Gelbord, J., Cackett, E., et al. 2019, ApJ, 870, 123, doi: [10.3847/1538-4357/aaf3b4](https://doi.org/10.3847/1538-4357/aaf3b4)

Fabian, A. C. 2012, ARA&A, 50, 455, doi: [10.1146/annurev-astro-081811-125521](https://doi.org/10.1146/annurev-astro-081811-125521)

Fausnaugh, M. M., Denney, K. D., Barth, A. J., et al. 2016, ApJ, 821, 56, doi: [10.3847/0004-637X/821/1/56](https://doi.org/10.3847/0004-637X/821/1/56)

Ferrarese, L., & Merritt, D. 2000, ApJL, 539, L9, doi: [10.1086/312838](https://doi.org/10.1086/312838)

Foreman-Mackey, D., Conley, A., Meierjurgen Farr, W., et al. 2013, emcee: The MCMC Hammer, Astrophysics Source Code Library, record ascl:1303.002

Giveon, U., Maoz, D., Kaspi, S., Netzer, H., & Smith, P. S. 1999, Monthly Notices of the Royal Astronomical Society, 306, 637, doi: [10.1046/j.1365-8711.1999.02556.x](https://doi.org/10.1046/j.1365-8711.1999.02556.x)

Gordon, K. D., Clayton, G. C., Misselt, K. A., Landolt, A. U., & Wolff, M. J. 2003, ApJ, 594, 279, doi: [10.1086/376774](https://doi.org/10.1086/376774)

Gordon, K. D., Fouesneau, M., Arab, H., et al. 2016, ApJ, 826, 104, doi: [10.3847/0004-637X/826/2/104](https://doi.org/10.3847/0004-637X/826/2/104)

Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, AJ, 131, 2332, doi: [10.1086/500975](https://doi.org/10.1086/500975)

Hernández Santisteban, J. V., Edelson, R., Horne, K., et al. 2020, MNRAS, 498, 5399, doi: [10.1093/mnras/staa2365](https://doi.org/10.1093/mnras/staa2365)

Homayouni, Y., Trump, J. R., Grier, C. J., et al. 2019, ApJ, 880, 126, doi: [10.3847/1538-4357/ab2638](https://doi.org/10.3847/1538-4357/ab2638)

Jones, L., Bianco, F. B., Yoachim, P., & Neilsen, E. 2025, <https://survey-strategy.lsst.io>, v0.1.0, Zenodo, doi: [10.5281/zenodo.15128504](https://doi.org/10.5281/zenodo.15128504)

Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, ApJ, 698, 895, doi: [10.1088/0004-637X/698/1/895](https://doi.org/10.1088/0004-637X/698/1/895)

Korista, K. T., & Goad, M. R. 2001, ApJ, 553, 695, doi: [10.1086/320964](https://doi.org/10.1086/320964)

Kovačević, A. B., Radović, V., Ilić, D., et al. 2022, ApJS, 262, 49, doi: [10.3847/1538-4365/ac88ce](https://doi.org/10.3847/1538-4365/ac88ce)

MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, ApJ, 721, 1014, doi: [10.1088/0004-637X/721/2/1014](https://doi.org/10.1088/0004-637X/721/2/1014)

Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, AJ, 115, 2285, doi: [10.1086/300353](https://doi.org/10.1086/300353)

Martocchia, A., & Matt, G. 1996, Monthly Notices of the Royal Astronomical Society, 282, L53, doi: [10.1093/mnras/282.4.L53](https://doi.org/10.1093/mnras/282.4.L53)

McHardy, I. M., Cameron, D. T., Dwelly, T., et al. 2014, Monthly Notices of the Royal Astronomical Society, 444, 1469, doi: [10.1093/mnras/stu1636](https://doi.org/10.1093/mnras/stu1636)

Netzer, H. 2015, ARA&A, 53, 365, doi: [10.1146/annurev-astro-082214-122302](https://doi.org/10.1146/annurev-astro-082214-122302)

—. 2022, MNRAS, 509, 2637, doi: [10.1093/mnras/stab3133](https://doi.org/10.1093/mnras/stab3133)

Peterson, B. M. 1993, PASP, 105, 247, doi: [10.1086/133140](https://doi.org/10.1086/133140)

Peterson, B. M., & Horne, K. 2004, Astronomische Nachrichten, 325, 248, doi: [10.1002/asna.200310207](https://doi.org/10.1002/asna.200310207)

Peterson, B. M., Wanders, I., Horne, K., et al. 1998, PASP, 110, 660, doi: [10.1086/316177](https://doi.org/10.1086/316177)

Peterson, J. R. 2014in , C04010, doi: [10.1088/1748-0221/9/04/C04010](https://doi.org/10.1088/1748-0221/9/04/C04010)

Pozo Nuñez, F., Bruckmann, C., Deesamutara, S., et al. 2023, Monthly Notices of the Royal Astronomical Society, 522, 2002, doi: [10.1093/mnras/stad286](https://doi.org/10.1093/mnras/stad286)

Sergeev, S. G., Doroshenko, V. T., Golubinskiy, Y. V., Merkulova, N. I., & Sergeeva, E. A. 2005, ApJ, 622, 129, doi: [10.1086/427820](https://doi.org/10.1086/427820)

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

Shen, Y., Brandt, W. N., Dawson, K. S., et al. 2015, ApJS, 216, 4, doi: [10.1088/0067-0049/216/1/4](https://doi.org/10.1088/0067-0049/216/1/4)

Silk, J., & Rees, M. J. 1998, A&A, 331, L1, doi: [10.48550/arXiv.astro-ph/9801013](https://doi.org/10.48550/arXiv.astro-ph/9801013)

Starkey, D. A. 2017, PhD thesis, Saint Andrews University, UK

Starkey, D. A., Huang, J., Horne, K., & Lin, D. N. C. 2023, MNRAS, 519, 2754, doi: [10.1093/mnras/stac3579](https://doi.org/10.1093/mnras/stac3579)

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

Sánchez, J., Walter, C. W., Awan, H., et al. 2020, Monthly Notices of the Royal Astronomical Society, 497, 210, doi: [10.1093/mnras/staa1957](https://doi.org/10.1093/mnras/staa1957)

Timmer, J., & König, M. 1995, A&A, 300, 707

Urry, C. M., & Padovani, P. 1995, PASP, 107, 803, doi: [10.1086/133630](https://doi.org/10.1086/133630)

Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, ApJ, 601, 692, doi: [10.1086/380563](https://doi.org/10.1086/380563)

Wang, Y., Gjergo, E., & Kuhlmann, S. 2015, MNRAS, 451, 1955, doi: [10.1093/mnras/stv1090](https://doi.org/10.1093/mnras/stv1090)

Weaver, J. R., & Horne, K. 2022, Monthly Notices of the Royal Astronomical Society, 512, 899, doi: [10.1093/mnras/stac248](https://doi.org/10.1093/mnras/stac248)

Wright, E. L. 2006, PASP, 118, 1711, doi: [10.1086/510102](https://doi.org/10.1086/510102)

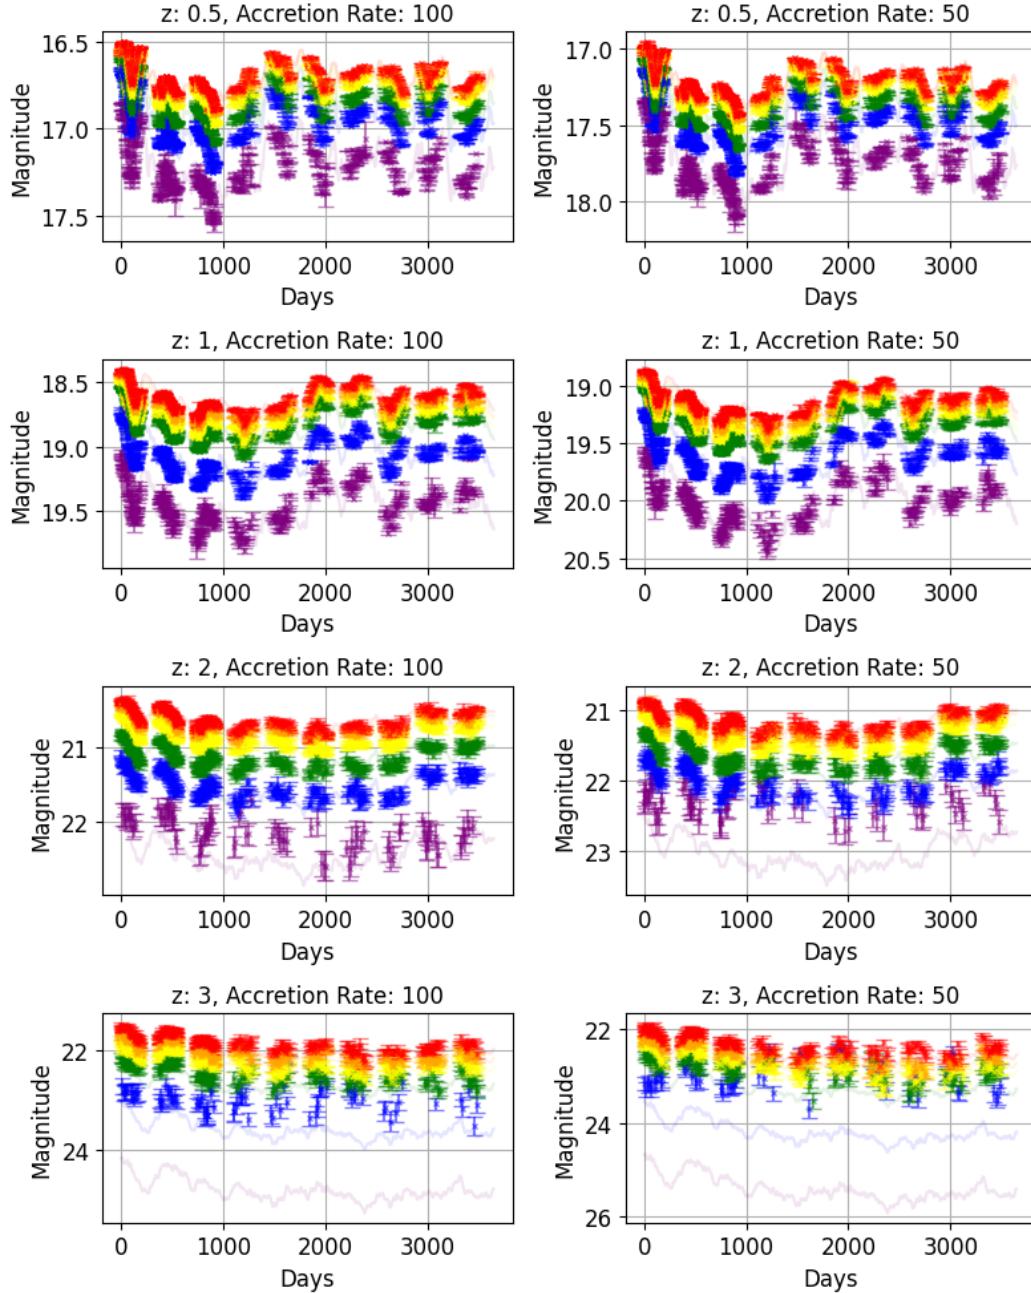
## APPENDIX

## A. PHOSIM LIMITATIONS

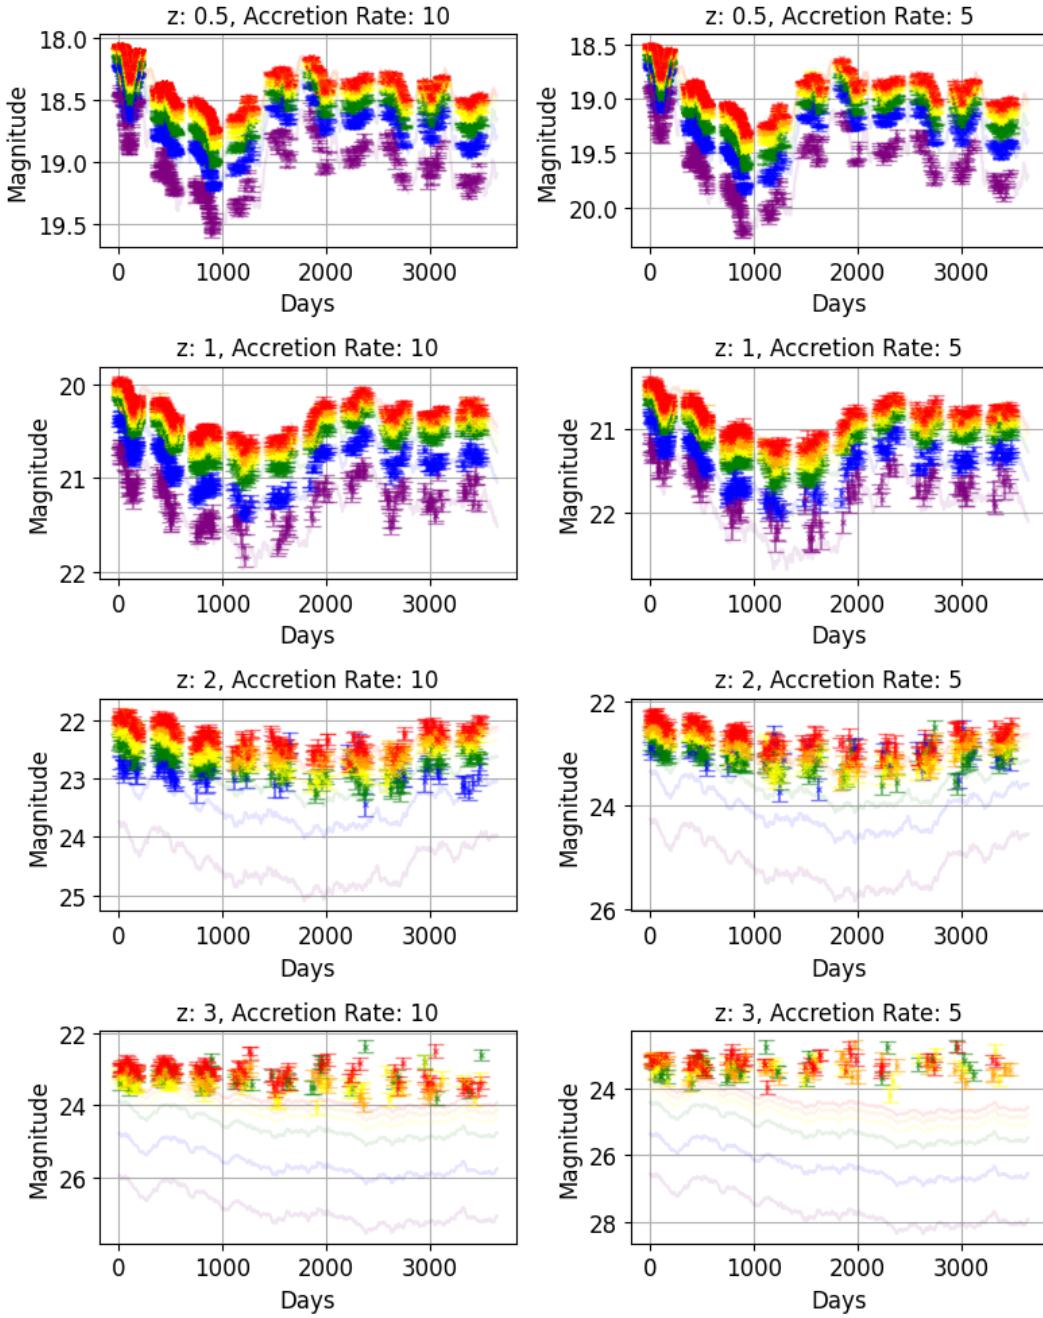
PhoSim offers a significant advantage for time-domain studies, particularly for AGN variability, as it can capture subtle effects introduced by filter throughput, atmospheric conditions, and instrumental noise. However, despite its efficiency, simulating an entire 10 years of the LSST, even a small sky region, remains computationally expensive, especially when high temporal resolution is required to resolve short-timescale variability, such as that seen in AGNs.

PhoSim is designed to operate on the CPU and can only process the images in serial. The software's parallelisation options allow for simulating multiple CCD chips using multiple cores, and simulating multiple light sources in the image using multiple threads simultaneously. To optimise the simulation, we can place all our simulated sources in a grid within one CCD chip and simulate a number of sources corresponding to the multiple number of threads on the CPU. On a 12-thread CPU used to simulate these images, there was a 5% increase in processing time between simulating 24 and 12 sources. As such, we chose to simulate 24 sources, which allows us to provide a 4x4 grid in the parameter space. A significant portion of the simulation runtime is dedicated to simulating the atmosphere, the quick background option is used to simulate our images as it averages the physics on the sensor (e.g. reducing vignetting). This allows for the simulation to still account for realistic atmosphere and optics while reducing the simulation set-up time by a factor of 3. Additionally, by using the “electron images” from the software, we can neglect the need for calibration frames, as simulating these frames on top of our light frame is extremely computationally expensive, with some bias frames taking up to an hour to run.

## B. SYNTHETIC LIGHT CURVES

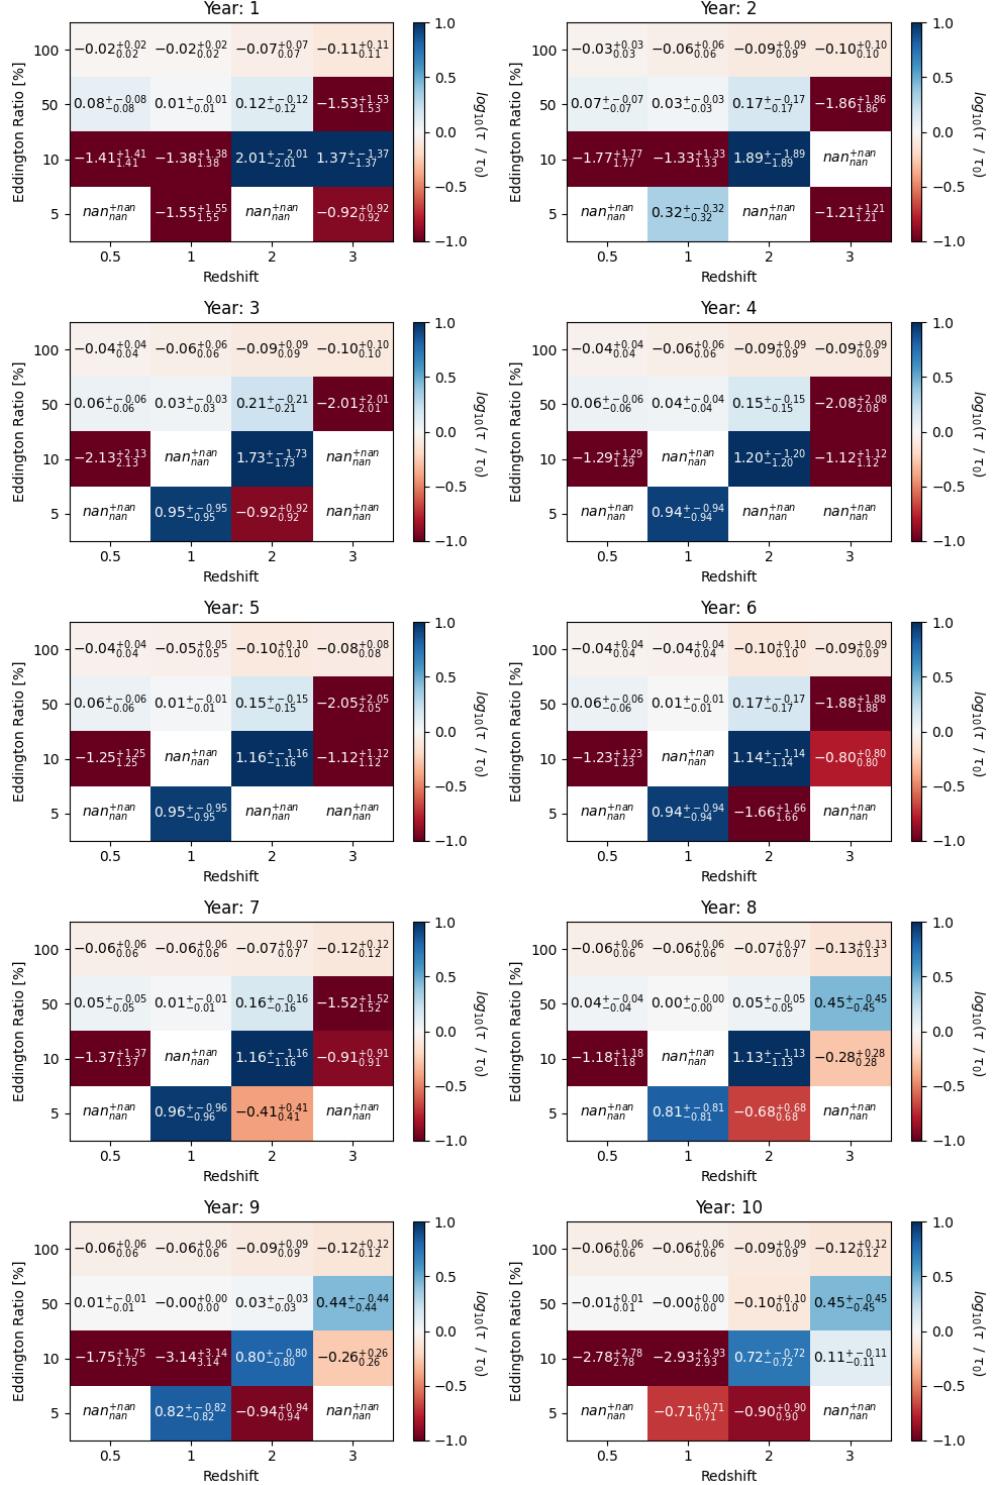


**Figure 15.** Processed Simulated Light Curves data for Eddington ratio of 100% [left column] and 50% [right column]

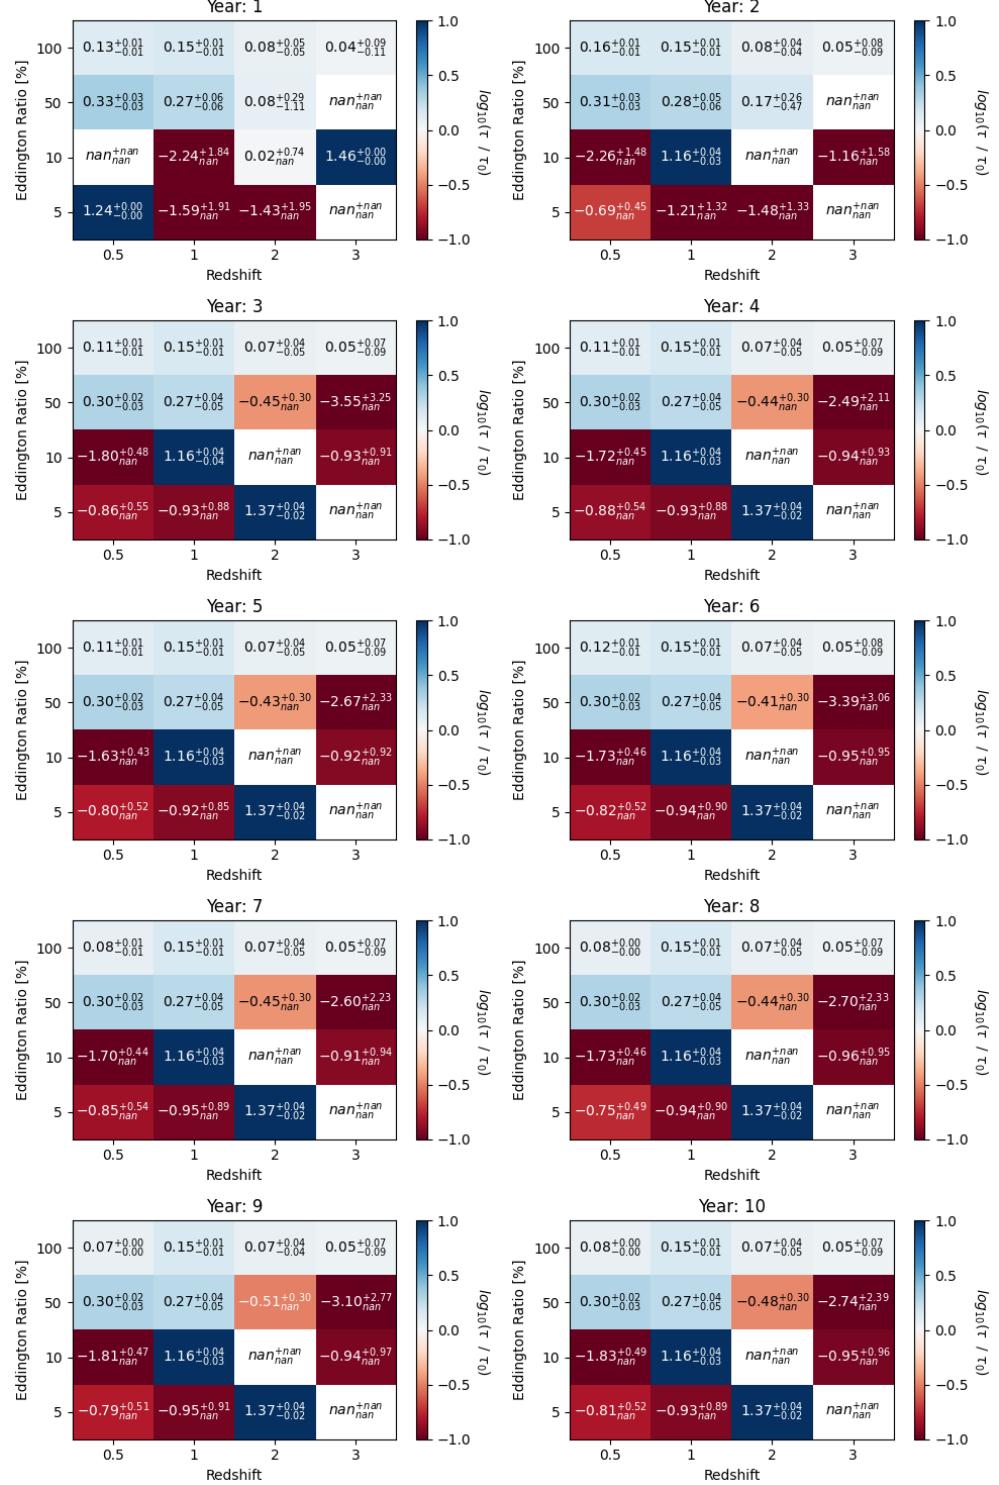


**Figure 16.** Processed Simulated Light Curves data for Eddington ratio of 10% [left column] and 5% [right column]

## C. YEAR BY YEAR PARAMETER SPACE FITS



**Figure 17.** Year by year fits comparing parameter space of the AGNs for the COSMOS Field.



**Figure 18.** Year by year fits comparing parameter space of the AGNs for the low cadence DDFs.